Essays in Information Economics

by

Chen Lyu

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

(Economics)

at the

UNIVERSITY OF WISCONSIN-MADISON

2023

Date of final oral examination: May/01/2023

The dissertation is approved by the following members of the Final Oral Committee:

Marzena Rostek, Professor, Economics

Lones Smith, Professor, Economics

Daniel Quint, Associate Professor, Economics

Zhuoxin Li, Associate Professor, Operations and Information Management

Dedicated to my parents

and to my wife

for their endless love and support

Acknowledgments

I am grateful to Marzena Rostek, Lones Smith, Daniel Quint, Marek Weretka and Bill Sandholm for their invaluable advice and excellent teaching during my PhD study. I would also like to thank all of the amazingly supportive staff members in the economics department at UW-Madison.

Contents

Abstract

In the first chapter, a recommendation platform sequentially collects information on a new product revealed from past consumer trials and uses it to better guide later consumers. Because consumers do not internalize the value of information they bring to others, their incentive for trying out the product can be socially insufficient. Given such a challenge, I study how the platform can maximize the total social surplus generated on it by designing its recommendation policy. In a model with binary product quality and general trial-generated signals, I show that the optimal design features a sequence of time-specific thresholds, which vary in a U-shaped pattern over the product's life. At any time, the platform should recommend the product if, based on its current belief, the probability of the product's quality being high is above the current threshold. This characterization allows me to provide predictions about the optimal recommendation dynamic and study comparative statics regarding the recommendation standards. My analysis also illustrates the potential usefulness of a Lagrangian duality approach for dynamic information design.

The second chapter studies optimal information provision by a search goods seller. While the seller controls a consumer's pre-search information, he cannot control post-search information because the consumer will inevitably learn the product's match after search. A relaxed problem approach is developed to solve the optimal design, which accommodates both continuous value distributions and ex-ante heterogeneous consumers with privately known outside options. The optimal design is shown to crucially depend on the outside option value distribution, and can be implemented by a simple upper-censorship signal under certain regularity conditions. Several applications are provided, including comparing information designs of search goods and experience goods, and studying the effect of competition with a large number of sellers.

The third chapter studies optimal disclosure regulation for entrepreneur public financing with a post-financing moral hazard problem. I show that partial disclosure can improve social

welfare over full disclosure through reducing efficiency loss caused by the moral hazard problem. As a result, a properly designed partial disclosure rule would be optimal without assuming any disclosure cost. This remains true after allowing for endogenous entrepreneur types with adverse selection concerns. With (constrained) Bayesian persuasion tools, the optimal disclosure rule is fully characterized. Although the paper is developed mainly around entrepreneur equity financing, its intuition can be more generally applicable. For instance, I also adapt the basic model to debt financing and provide an application to banking system disclosure.

Chapter 1

Information Design for Social Learning on a Recommendation Platform

1.1 Introduction

Recommendation platforms are quite popular in our daily lives. For examples, people rely on Netflix for what to watch, Yelp for where to eat, and TripAdvisor for where to travel.[1] To provide better recommendations, a common practice of these platforms is to induce a kind of social learning for new products. Namely, they collect information generated from early consumers' trials of a product (e.g., via rating and reviews), and use it to better guide later consumers' decisions. In this process, however, because individual consumers do not internalize the value of information they bring to others, their incentive for trying out the product is typically insufficient. This handicaps learning and can hinder the platform from making better-informed recommendations.

In this paper, I study how a platform facing the above challenge should design its recommendation policy in a dynamic manner, which can potentially "persuade" consumers towards more socially desirable trials of new products. In the model, consumers arrive sequentially over the (finite) lifetime of a product with unknown quality, which can be either high or low, and

[1] For some non-commercial examples, consider FDA for drug uses and Medicare Advantage Star Ratings for Medicare plan choices in the US.

decide whether to consume it. Whenever a consumer consumes the product, a signal about its quality will be generated and privately observed by the platform.[2] Unlike some existing studies surveyed later, I allow such consumption-generated signals to be general and non-conclusive. In each period, based on the signals previously received, the platform can guide the current consumer's choice by providing a recommendation message. Knowing the message and her own arrival time, the consumer then makes her consumption decision in a Bayes-rational way. The platform's design problem is to find a dynamic recommendation policy, to which it can commit ex-ante, in order to maximize the total social surplus generated on it.

Ideally, the platform should recommend the product for trial as long as this is socially desirable, even if consumption is suboptimal for the current consumer based on the current information. With such a policy, however, the expected quality of some recommendations may be too low for the consumers to follow, which renders the design ineffective. An optimal policy therefore must choose when to recommend socially desirable but individually suboptimal consumption most efficiently, subject to the requirement that the consumer in each period will be willing to follow the recommendation. As my results will show, this incentive concern is the central force that shapes the optimal design.

The need to convince consumers to follow the recommendations induces a sequence of incentive-compatibility (abbr., IC) constraints – one for each consumer – in the dynamic design problem, which makes it a *constrained* Markov decision process. Solving such a problem is challenging because the standard dynamic programming technique cannot directly handle those constraints. To overcome the difficulty, I deploy a Lagrangian duality approach. It allows me to partially characterize the shadow values of the IC constraints and finally solve the optimal design. To the best of my knowledge, this is the first paper solving a (non-degenerate) constrained Markov decision process that naturally arises from a dynamic information design problem.

I show that the optimal design features a sequence of time-specific thresholds, which generally vary in a U-shaped pattern over the product's life. At any time, the platform should recommend the product if, based on its current information, the probability of the product's quality being high is above the current threshold. This suggests that the platform should set time-varying standards for recommending the product, which first decrease and then increase as the product ages. Underlying this time-pattern is a tension between the platform's desire

[2] The platform also receives an initial signal about the product before any consumption, which reflects its internal research or data on the previous performance of similar products.

to create information value for future consumers and its need to fulfill the current consumer's IC constraint. When the product is very young, the information value is high due to a long remaining product lifespan, while consumers are "skeptical" about following recommendations because they know that even the platform has not acquired much information about the product yet. This implies a binding IC constraint and necessitates a picky censorship regarding when to recommend the product. As time passes, consumers become easier to convince as they expect that the platform may have gotten better informed by previous signals. The recommendation standard can hence be lowered.[3] This continues until the standard has become sufficiently low such that trials with beliefs further below it are no longer worthwhile. Thereafter, the optimal threshold gradually goes up because the information value of consumption dwindles as the product approaches its end of life.

The results above implies an interesting prediction about the optimal recommendation dynamic – it can feature *temporary* recommendation suspensions following negative consumer feedback for young products. Specifically, following a negative feedback, the platform's belief of high quality can drop below the recommendation threshold, which suspends the recommendation and learning. However, if we are in the early phase of the product's life where the threshold is declining, the threshold can fall below the belief again a few periods later, which restarts recommendation. In practice, it is well-known that temporary suspension or deprioritization of recommendation is often used for punishing the misconduct of a seller or content provider.[4] My finding here suggests another motivation for taking such action, which is to enhance social learning on new products given the inadequacy of individual consumers' incentive.

My characterization of the optimal design also enables a couple of comparative statics. The first analysis considers how the design should be adjusted when consumption becomes more likely to yield non-neutral signals about the product quality (e.g., due to better feedback elicitation designs). I show that the recommendation standards should be lowered uniformly over time when this happens. The second analysis incorporates random consumer arrivals. I show that when we have a thicker market where consumers arrive more frequently, the optimal policy should become more generous in recommending the product.

I note that the binary quality assumption will be relaxed in an extension of my main model. Although a full characterization for the optimal design is not available there, my duality ap-

[3] See Section 1.4.1 for a more concrete intuition behind this.

[4] For a list of real-world examples, see Table 1 in Liang et al. (2020).

proach still helps to reveal certain properties of it. In particular, I show that the optimal policy features a partial-order monotone structure, which can be considered as a generalization of the threshold structure. I will discuss implications of this general result for algorithmic recommendation design following the extension.

The paper is organized as follows. The rest of Section 1 reviews the literature. Section 1.2 presents the main model. Section 1.3 derives the optimal design. Section 1.4 explores dynamic properties of the optimal design. Section 1.5 considers comparative statics. Section 1.6 provides additional discussions. Section 1.7 concludes with some methodological remarks. Appendix 1.A considers the extension with general quality support. All proofs are provided in Appendix 1.B.

Related literature – My study closely relates to Kremer et al. (2014) and Che & Hörner (2018), who also study the optimal recommendation design when early consumption has informational externality to later consumers.[5] These papers have focused on special classes of consumption-generated signals. Specifically, the main model in Kremer et al. (2014) considers fully revealing signals, i.e., the underlying quality will be fully revealed after a single trial. This allows them to reduce the design problem into a decision problem about when to induce the first trial based on the platform's initial information.[6] Che & Hörner (2018) considers a Poisson learning environment with binary quality levels in continuous time. They assume that the platform has either received no news, or has received conclusive news that fully reveals the product's quality. The design problem then boils down to a deterministic control problem about recommendation intensity following the history without news arrival. Unlike these papers, my study accommodates general non-conclusive signals. My characterization of the optimal design is thus about whether to recommend the product in each period based on any current belief of the platform, which goes beyond timing of the first trial or recommendation intensity without previous news. This allows me to interpret my results as regarding the time-varying recommendation standards and necessitates the more general formulation of the problem.

An extension in Kremer et al. (2014) and a strand of subsequent algorithm-oriented research have studied environments more general than in the main model of Kremer et al. (2014),

[5] Lorecchio & Monte (2021) also considers a setting where the designer records previous agents' feedbacks to guide later agents' decisions. However, their designer has state-independent payoff, rely on restricted communication rules, and only focuses on the long-run stationary equilibrium, which makes their paper distinct from mine.

[6] More precisely, the initial information in their paper is about an alternative consumption option, which is always tried by the first consumer with its quality fully revealed since then.

which do allow for non-conclusive consumption-generated signals (Papanastasiou et al., 2018; Mansour et al., 2020).[7] The goal of this literature is to propose algorithms that can achieve better asymptotic performance as the number of consumers coming in sequence goes to infinity, which is often measured by the decay rate of per-consumer welfare loss compared to the full-information first-best benchmark. While such measurement reflects an important aspect of the design performance, it does not take into account the design's finite horizon behavior and can be insensitive to multiplicative changes in the welfare loss.[8] Hence, for the algorithms proposed in the above literature, little is known about their finite-horizon efficiency, and little has been done to improve their finite-horizon performance. My paper complements the literature by solving the finite-horizon optimal design in a stylized setting, which may serve as a short-run performance benchmark for evaluating any algorithm and help to inspire new algorithms with a non-asymptotic focus.[9]

Another growing literature also considers the optimal information provision by a platform to a sequence of short-lived agents (Glazer et al., 2021; Komiyama & Noda, 2021; Küçükgül et al., 2022). In these papers, the agents are either endowed with or are able to acquire private signals, and a central task for the platform is to infer these private signals from the agents' decisions. These papers thus consider very different information sources for the platform and explore design concerns distinct from mine.

More generally, my paper belongs to the broad literature on information design, beginning with Kamenica & Gentzkow (2011) and Rayo & Segal (2010), and especially studies on dynamic designs (e.g., Ely, 2017; Renault et al., 2017; Smolin, 2021; Ely & Szydlowski, 2020; Ball, 2019; Orlov et al., 2020; Lorecchio, 2021). One difference between many of the studies in this literature and mine is that I consider a designer whose private information flow is controlled by the receivers' decisions, rather than being exogenous. My analysis illustrates how such a setting naturally leads to a constrained Markov decision process after simplification with the revelation

[7] Also see, for examples, Bahar et al. (2015), Mansour et al. (2016), Chen et al. (2018), Immorlica et al. (2019) and Bahar et al. (2021) for a variety of extensions.

[8] To see this, notice that an average loss function $L(t)$ is considered to have the same decay rate in t as $\alpha L(t)$ for any $\alpha > 0$.

[9] Papanastasiou et al. (2018) does investigate finite-horizon design in a particular setting, but in that setting the initial information is such that either no exploration can ever happen or consumer IC constraints are never binding, which makes the optimal design obvious. They also propose to formulate the designer's problem as a constrained Markov decision process in a more general setting, but concluded it to be computationally infeasible, and did not derive analytical results from it except for giving a bound on how many belief states need to involve randomization in an optimal design.

principle (Myerson, 1986), and how the Lagrangian duality approach can be useful for solving it.[10]

1.2 The Model

I first describe the model, and then discuss several underlying assumptions in Section 1.2.2.

1.2.1 The Setting and the Design Problem

The model features a platform, a sequence of short-lived Bayes-rational consumers and a product. The product is launched in period 1 and will remain available for consumption over $T < \infty$ discrete time periods. In each period $t = 1, ..., T$, a consumer arrives at the platform and decides whether to consume the product. I denote the consumer's decision as $a_t \in \{0, 1\}$ with $a_t = 1$ meaning consumption occurs. Without consuming the product, the consumer will receive her outside option value, which is normalized to zero. If she consumes the product, the consumer's utility will be equal to $\tilde{\theta}$, which is a random variable taking values in $\{\theta_L, \theta_H\}$, with $\theta_L < 0 < \theta_H$.[11] This $\tilde{\theta}$ measures the underlying quality of the product, which is fixed over time but initially unknown. I assume that the platform and the consumers share a common prior for it.

At the beginning of period 1, the platform receives a signal s_0 about $\tilde{\theta}$, which reflects the platform's initial information based on, for example, its internal research or data about past performance of similar products. Subsequently, an additional signal will be generated for the platform whenever a consumer consumes the product. Let s_i denote the signal from the i'th consumption of the product. Conditional on $\tilde{\theta}$, I assume that $s_1, s_2, ...$ are i.i.d. and are independent from s_0.

In every period, the platform can compute its posterior belief about the product quality based on the previous signals received. I use p_t to denote the platform's belief about $\tilde{\theta} = \theta_H$ at the beginning of period t. Let μ_1 denote the distribution of $p_1 = \mathbb{P}(\tilde{\theta} = \theta_H | s_0)$; let $G(\cdot|\cdot)$

[10] Beutler & Ross (1985) and Beutler & Ross (1986) were the first to use a Lagrangian approach to study constrained Markov decision processes. The method is subsequently developed and applied in many mathematical and engineering papers (see, e.g., Section 1.2 of Altman (1999) for a brief survey). These studies typically only involve a few aggregate constraints corresponding to different design criteria. In contrast, my problem features one constraint for each period, which leads to a novel dynamic aspect of the problem.

[11] As usual, one can interpret $\tilde{\theta}$ as the consumer's expected utility from consumption given the product's true quality, and the actually realized utility can involve an ex-post idiosyncratic shock.

denote the transition kernel of $(p_t)_{t=1}^T$ following one's consumption; let $D(\cdot|p)$ denote the Dirac measure at p. The process of $(p_t)_{t=1}^T$ then follows the following rule:

$$p_1 \sim \mu_1 \tag{1.1}$$

$$p_{t+1}|p_t, a_t \sim a_t G(\cdot|p_t) + (1-a_t) D(\cdot|p_t) \tag{1.2}$$

For later analyses, I define $u(p) := p\theta_H + (1-p)\theta_L$, and define $\bar{p}$ to be the indifference belief for consumers, i.e., $u(\bar{p}) = 0$.

Before the realization of s_0, the platform can commit to an information transmission policy that decides what message to convey to the coming consumer in each period based on the information available at that time. I assume that the consumer can neither observe previous messages, nor observe decisions of earlier consumers. She just observes her arrival time and her own message, and then decides whether to consume the product.

The timeline of the environment is summarized as follows:

1. Before period 1, the platform (publicly) commits to an information transmission policy, and the product quality $\tilde{\theta}$ is secretly realized. Then, the platform privately receives its initial signal s_0 about $\tilde{\theta}$.
2. At the beginning of each period $t = 1, ..., T$, a consumer arrives and receives a message from the platform, which is generated according to the information transmission policy. She then decides whether to consume the product. If she is the n'th consumer who consumes the product, signal s_n will be generated for the platform. The economy then enters into the next period.

The Designer's Problem: I look for the information transmission policy that maximizes the total consumer surplus generated on the platform over the product's lifetime (i.e., $\sum_{t=1}^T \mathbb{E}[a_t\tilde{\theta}]$). By the revelation principle (Myerson, 1986), it suffices to consider *incentive-compatible recommendation policies*, which just decide whether to recommend the product for consumption in every period, subject to the requirement that a Bayes-rational consumer will like to follow the recommendation. Since the belief p_t summarizes all the payoff-relevant information of the platform in period t, standard argument implies that I can further focus on *randomized Markov policies* with respect to the process of $(p_t)_{t=1}^T$.[12] Formally, a randomized Markov rec-

[12] A formal proof for this is available upon request.

ommendation policy is a sequence of measurable mappings $\phi := (\phi_t : t = 1, ..., T)$, where each $\phi_t : [0,1] \to [0,1]$ decides the probability of recommending the product at time t given any $p_t \in [0,1]$. For the rest of the paper, by "policy" I will be referring to a policy of this type.

I impose the following assumption on consumption-generated signals throughout:

Assumption 1.2.1. (i) $\mathbb{P}(\mathbb{E}[\tilde{\theta}|s_0] > 0) > 0$;
(ii) For any $i \geq 0$, we have $\mathbb{P}(s_i \in A|\tilde{\theta} = \theta_L) < \mathbb{P}(s_i \in A|\tilde{\theta} = \theta_H)$ for some (measurable) set A in the realization space of s_i;
(iii) For any $i \geq 0$, we have $\mathbb{P}(s_i \in A|\tilde{\theta} = \theta_L) > 0 \Leftrightarrow \mathbb{P}(s_i \in A|\tilde{\theta} = \theta_H) > 0$ for any (measurable) set A in the realization space of s_i.

Condition (i) implies that, given a certain realization of the platform's initial information, it is optimal for the first consumer to consume.[13] Without such a condition, the first consumer may never want to consume, knowing which the second consumer will never consume either. Induction would then imply that no consumption can ever happen under any design. Condition (i) rules out such a trivial scenario. Condition (ii) simply guarantees that the signals are indeed informative about $\tilde{\theta}$. Condition (iii) implies that no signal realization can conclusively reveal the quality level. It helps to simplify the exposition of certain proof, but is not essential for results in the paper.

1.2.2 Discussion on Model Assumptions

1. **Information v.s. monetary incentive.** The model assumes that the platform cannot directly pay early consumers for trying out the new product. While it can work well in some applications, the use of monetary incentive may be problematic in others. In particular, if consumers are only attracted by the monetary incentive instead of the product itself, they may just *pretend* to consume the product and leave some artificial feedback, especially when the product's pecuniary price is zero (e.g., digital contents on a subscribed platform).[14] I thus focus on the design of information in this paper, which can convince consumers to truly try the product.

[13] To guarantee this, one may instead impose the slightly weaker condition $\mathbb{P}(\mathbb{E}[\tilde{\theta}|s_0] \geq 0) > 0$. I impose the stronger condition for a technical reason when deriving the duality result.

[14] For example, one may play a movie at background without watching it, and then fabricate some feedback to earn the money.

2. **Consumer information on product launch time.** An important assumption of the model is that each consumer can observe when the product was launched. This is reasonable for many products like TV-series, video games, restaurants or amusement parks, which typically have a public release or opening time. For some other products, however, the consumer may only have a rough idea about the launch time. In such cases, my design setting can be considered as being robust to uncertainty in the exact consumer information. Indeed, under the assumption that consumers perfectly observe the product launch time, the optimal design I derive will be incentive-compatible no matter what information consumers actually have about the product's launch time. It thus provides the best guaranteed performance.
3. **Consumer information on early consumer arrivals.** The model above has a feature that each consumer perfectly knows how many consumers have arrived before her, which is unrealistic. Fortunately, this is not truly a concern because my framework can be easily extended to incorporate random consumer arrivals. In such an extension, each consumer can infer how many consumers have arrived earlier based on her own arrival time, but cannot know the number exactly. All of my results will still hold. To ease notation, however, I do not explicitly introduce random consumer arrival until Section 1.5.2, where I will examine how the arrival rate affects the optimal design.

1.3 Characterization for the Optimal Design

1.3.1 The Constrained Markov Decision Process

Let Φ denote the set of all (measurable) policies. Given any $\phi \in \Phi$, I use $\mathbb{P}_\phi$ to denote the probability measure over events of $\big((a_t)_{t=1}^T, (p_t)_{t=1}^T\big)$ provided that consumers follow the recommendations, and use $\mathbb{E}_\phi$ to denote the corresponding expectation operator. Then, the incentive compatibility (IC) constraint for a time-t consumer can be written as:[15]

$$\mathbb{P}_\phi(a_t = 1) > 0 \Rightarrow \mathbb{E}_\phi[u(p_t)|a_t = 1] \geq 0 \tag{1.3}$$

$$\mathbb{P}_\phi(a_t = 0) > 0 \Rightarrow \mathbb{E}_\phi[u(p_t)|a_t = 0] \leq 0 \tag{1.4}$$

These respectively guarantee that the consumer will follow the recommendation when the product is recommended ($a_t = 1$) and when it is not ($a_t = 0$). Since one's consumption of the

[15]Notice $\mathbb{E}_\phi[\tilde{\theta}|a_t] = \mathbb{E}_\phi\big[\mathbb{E}_\phi[\tilde{\theta}|p_t]|a_t\big] = \mathbb{E}_\phi[u(p_t)|a_t]$ since a_t is measurable w.r.t. p_t.

product generally benefits later consumers by yielding information, the designer will never want to recommend $a_t = 0$ when consumption is optimal for the current consumer. This implies that the second constraint above is non-restrictive for the designer and can thus be omitted. For the first constraint, we can more compactly write it as $\mathbb{E}_\phi[u(p_t)|a_t = 1]\mathbb{P}_\phi(a_t = 1) \geq 0$, which is equivalent to $\mathbb{E}_\phi[a_t u(p_t)] \geq 0$. Notice $\mathbb{E}_\phi[a_t u(p_t)]$ is the expected surplus of the time-t consumer when she follows the recommendation under ϕ. The designer's problem can hence be formulated as:[16]

$$\max_{\phi\in\Phi}\Big\{\sum_{t=1}^{T}\mathbb{E}_\phi[a_t u(p_t)]\Big\} \tag{1.5}$$

$$\text{s.t. } \mathbb{E}_\phi[a_t u(p_t)] \geq 0 \quad \forall t = 1, ..., T \tag{1.6}$$

Due to the presence of the expectation operator, each constraint in (1.6) is not just restricting the recommendation decision given a particular realization of p_t, but involves integration over the entire distribution of p_t at a particular time. Such an aggregated constraint arises here because the payoff-relevant process $(p_t)_{t=1}^{T}$ is privately monitored by the platform, and hence the consumer must integrate over its equilibrium distribution when forming her posterior belief given any recommendation message. The presence of such constraints makes the problem a *constrained* Markov decision process, which cannot be directly handled by dynamic programming with p_t being the state variable.[17] To overcome this difficulty, I will provide a dual characterization for it using Lagrangian duality, which allows me to partially reduce the problem to an unconstrained one.

The following lemma reveals the key properties of the belief process needed for later analyses.

Lemma 1.3.1. *The belief process (1.1) – (1.2) satisfies the following conditions:*

(P1) $G(\cdot|p)$ *as a measure-valued function of* p *is weakly continuous.*

(P2) $\int_{p'} u(p')G(dp'|p) = u(p)$.

(P3) $G(\cdot|p)$ *increases in* p *in terms of first-order stochastic dominance.*

(P4) $G\big([0,p)|p\big)$ *and* $G\big((p,1]|p\big)$ *are strictly positive for any* $p \in (0,1)$.

(P5) $\mu_1\big((\bar{p},1]\big) > 0$.

[16] Papanastasiou et al. (2018) first proposed the constrained Markov decision process formulation for this kind of design problem. However, they do not pursue much further analysis with it. See footnote 9 for details.

[17] See Altman (1999) for a textbook treatment to constrained Markov decision processes.

Property (P1) means that small changes in the prior can only lead to small changes in the posterior, which is a technical result that guarantees the existence of the optimal design. Property (P2) is implied by the standard law of iterated expectation. Property (P3) is an inertia property of the belief process, which roughly says that a product looking more promising today is also more likely to look promising tomorrow. It will be important for showing the threshold structure of the optimal policy later. Property (P4) is directly implied by the assumption that signals are informative, which guarantees that the belief process will not stay constant for sure following consumption. Property (P5) is directly implied by Assumption 1.2.1(i), which allows consumption, and hence learning, to occur in period 1. I note that these five properties are all that one needs to know about the belief process for later analyses, which abstract away from other details of the learning process.

The following result guarantees that the designer's problem is well defined.

Proposition 1.3.1. *There exists an optimal solution to the designer's problem (1.5) – (1.6).*

1.3.2 The Dual Characterization

Given any vector of Lagrangian multipliers $\lambda \in \mathbb{R}_+^T$ associated to the IC constraints, I define the Lagrangian function for the designer's problem as:

$$\mathcal{L}(\phi; \lambda) = \sum_{t=1}^{T} \mathbb{E}_\phi[(1 + \lambda_t) a_t u(p_t)] \tag{1.7}$$

Then, we have the following strong-duality result.

Lemma 1.3.2. *Let w^* denote the optimal value of the designer's problem. Then,*

$$w^* = \min_{\lambda \in \mathbb{R}_+^T} \sup_{\phi \in \Phi} \mathcal{L}(\phi; \lambda) \tag{1.8}$$

where the minimum is achieved by some λ^*. Given any such λ^*, a policy ϕ^* is optimal for the designer's problem if and only if:

(i) $\phi^* \in \arg\max_{\phi \in \Phi} \mathcal{L}(\phi; \lambda^*)$

(ii) $\lambda_t^* \mathbb{E}_{\phi^*}[a_t u(p_t)] = 0,\ \forall t = 1, ..., T$

(iii) $\mathbb{E}_{\phi^*}[a_t u(p_t)] \geq 0,\ \forall t = 1, ..., T$

To see how this result is helpful, notice that once we know a solution λ^* to the dual problem

(1.8), which intuitively measures the "shadow values" of the IC constraints, the lemma implies that any optimal policy must solve $\max_{\phi\in\Phi}\mathcal{L}(\phi;\lambda^*)$, which is an unconstrained problem. The optimal design can then be characterized by studying this unconstrained problem with the standard dynamic programming approach.

The problem here, however, is that the value of λ^* is not available. Generally, deriving it requires one to either solve the min-max problem in (1.8) or solve the fixed-point problem defined by conditions (i) – (iii) jointly for (λ^*, ϕ^*), both of which are difficult. Fortunately, as I show below, a property of λ^* can be directly derived from the dual problem, which turns out to suffice for a sharp characterization of the optimal policy.

1.3.3 Main Structures of the Optimal Design

The following lemma is a key result derived from the dual problem (1.8).

Lemma 1.3.3. *There exists* $\lambda^* \in \arg\min_{\lambda\in\mathbb{R}^T_+}\sup_{\phi\in\Phi}\mathcal{L}(\phi;\lambda)$ *such that* $\lambda^*_t \geq \lambda^*_{t+1}$ *for all* $t = 1, ..., T-1$.

To see an intuition behind this result, assume that the dual problem has a unique solution λ^*. As usual, we can interpret λ^*_t as the shadow value of marginally relaxing the time-t IC constraint for the designer's problem. As time passes, two changes happen in the designer's problem. First, as information accumulates over time, we are able to have better selections over the products for recommendation. This makes it possible to obey the consumer's IC constraint with less sacrifice for socially desirable consumption. Second, as the remaining lifetime of the product gets shorter, the dynamic value from having additional myopically suboptimal consumption drops. These both suggest that relaxing later IC constraints is less helpful than relaxing the earlier ones. Hence, the associated shadow values should decrease over time.

Although the argument above is intuitive, it is hard to formalize it into a concrete proof. For the lemma's proof, I directly examine the dual problem and develop an "inter-change" argument. In particular, given any λ^* solving the dual problem, I show that if two adjacent components of it violate the time pattern, then interchanging them will lead to a new solution to the dual problem. Starting with any solution to the dual problem, one can hence construct a new solution satisfying the time pattern by making such interchanges repeatedly.

One immediate implication of Lemma 1.3.3 is that the optimal design will generally feature a two-phase structure. In the first phase, $\lambda^*_t > 0$ and the IC constraints are binding; in the second phase, $\lambda_t^* = 0$ and the IC constraints are essentially non-restrictive.[18] As I will show later, the optimal recommendations corresponding to these two phases exhibit very different dynamic patterns.

The time-pattern of λ^* found in Lemma 1.3.3 also turns out to induce a simple solution structure for the Lagrangian function optimization $\max_{\phi \in \Phi} \mathcal{L}(\phi; \lambda^*)$. To state the result, I define *threshold policies* as follows:

Definition 1.3.1. A time-t policy $\phi_t : [0,1] \to [0,1]$ is called a *threshold time-t policy* if there exists threshold $\eta_t \in [0,1]$ such that $p > \eta_t \Rightarrow \phi_t(p) = 1$ and $p < \eta_t \Rightarrow \phi_t(p) = 0$. A policy ϕ is called a *threshold policy* if ϕ_t is a threshold time-t policy for every t.

Namely, a threshold policy will recommend the product when the current quality belief of $\tilde{\theta} = \theta_H$ is above a threshold, and will not recommend it when the belief is below the threshold. It can also involve arbitrary randomization at the threshold. By applying backward induction on the dynamic programming of $\max_{\phi \in \Phi} \mathcal{L}(\phi; \lambda)$, I show the following result:

Lemma 1.3.4. *Given any non-increasing sequence of* $(\lambda_t)_{t=1}^T$*, any solution to* $\max_{\phi \in \Phi} \mathcal{L}(\phi; \lambda)$ *is almost surely equivalent to a threshold policy. Moreover,* $p_t > \bar{p} \Rightarrow a_t = 1$ *a.s. under such a policy.*[19]

Together with Lemma 1.3.3 and the dual characterization for the optimal design, this directly implies the threshold structure of the optimal design:

Corollary 1.3.1. *Any optimal policy is almost surely equivalent to a threshold policy. Moreover,* $p_t > \bar{p} \Rightarrow a_t = 1$ *a.s. under it for any* t*.*[20]

I note that although threshold policies are intuitively appealing, their optimality is not obvious in my setting. While the myopic value of consumption always increases in p_t, the dynamic informational value of it does not. Given the presence of IC constraints, even measuring such dynamic value is not straightforward, as one not only needs to consider the direct benefit to later consumers, but also needs to consider how better information may help to relax the

[18] It is easy to see that the second phase includes at least the last period, since the optimal policy there will be myopically optimal. The first phase is non-empty as long as the prior on $\tilde{\theta}$ is not sufficiently favorable to support first-best learning.

[19] By saying $A \Rightarrow B$ almost surely (a.s.), I mean that the event in which A happens but B does not happen is of zero probability.

[20] The second statement implies that the product is always recommended when consumption is optimal for the current consumer. Hence the IC constraint (1.4) I ignored before is indeed satisfied.

IC constraints of later consumers, and thereby facilitate more information generation from them. The duality approach I take partly resolves such difficulty by characterizing the shadow values of those IC constraints. Given the monotonicity property of $(\lambda_t^*)_{t=1}^T$, I show that when p_t increases, the positive change in the myopic consumption value always dominates the potentially indeterminate change in the dynamic value of consumption, as is measured in the continuation problem of $\max_{\phi\in\Phi}\mathcal{L}(\phi;\lambda^*)$. The total value of consumption is thus always increasing in p_t, which implies the threshold structure of the optimal design.

Lemma 1.3.4 above also helps to characterize the *dictator's optimal policy*, where by "dictator" I mean a social planner who can dictate consumers' decisions without obeying their IC constraints. Notice that if $\lambda_t = 0$ for all t, the Lagrangian optimization $\max_{\phi\in\Phi}\mathcal{L}(\phi;\lambda)$ is reduced to the dictator's problem. Lemma 1.3.4 then implies that the dictator's problem also features a threshold solution. This solution will be used in the construction of the optimal design below.

1.3.4 The Optimal Policy

Notation: For any vector indexed by time, I will use subscription "$\geq t$" to indicate the sub-vector corresponding to time no earlier than t. For example, $\phi_{\geq t}$ will denote the continuation policy since time t. Notations like $\phi_{>t}$ and $\phi_{<t}$ are similarly defined.

Based on the two-phase structure of the optimal design implied by Lemma 1.3.3 and the threshold structure stated in Corollary 1.3.1, one can explicitly construct the optimal design using a forward induction algorithm. To do so, I define ϕ^d to be the "most conservative" optimal policy for the dictator's problem (i.e., the designer's problem without IC constraints), whose details are provided in Appendix 1.B.5. When the dictator's problem admits multiple solutions, ϕ^d is the most conservative in the sense that it always breaks ties in favor of non-recommendation, and hence it is in favor of the current consumer's surplus.

A candidate optimal threshold policy ϕ^o, together with a cutoff time point $\widehat{t}$, can be inductively defined as follows.

Definition 1.3.2. A policy ϕ^o, a sequence of distributions $(\mu_t^o)_{t=1}^T$ over $[0,1]$ and a time point $\widehat{t}\in\{1,...,T\}$ are defined with the following algorithm.[21]

[21] I note that the algorithm must eventually stop in step 1 because $\int_p \phi_t^d(p)u(p)\mu_t^o(dp)\geq 0$ holds for $t=T$, since ϕ_T^d is myopically optimal. This implies that $\widehat{t}$ is well-defined.

Starting with $t = 1$ and $\mu_1^o = \mu_1$:

- Step 1: If $\int_p [\phi_t^d(p)u(p)]\mu_t^o(dp) \geq 0$, let $\widehat{t} = t$, $\phi_{\geq t}^o = \phi_{\geq t}^d$, and $(\mu_s^o)_{s>t}$ be the (marginal) distributions of $(p_s)_{s>t}$ under $\phi_{\geq t}^d$ given $p_t \sim \mu_t^o$. Otherwise, go to the next step.

- Step 2: Let ϕ_t^o be a threshold time-t policy such that

  (i) $\phi_t^o(p) = 1$ for all $p > \bar{p}$;

  (ii) $\int_p [\phi_t^o(p)u(p)]\mu_t^o(dp) = 0$.

  (See Appendix 1.B.5 for details.) Also let μ_{t+1}^o be the distribution of p_{t+1} under ϕ_t^o given $p_t \sim \mu_t^o$. Then go back to step 1 for time $t + 1$.

Intuitively, ϕ^o in its early phase is a threshold policy just "picky" enough to fulfill the IC constraints. This continues until a time point $\widehat{t}$, at which even the dictator's optimal continuation policy $\phi_{\geq \widehat{t}}^d$ will not be too generous to be incentive-compatible. Then ϕ^o just resumes with $\phi_{\geq \widehat{t}}^d$ later on. The following proposition shows that ϕ^o is indeed an optimal policy and fully characterizes any optimal design.

Proposition 1.3.2. *Any policy ϕ^* is optimal for the designer's problem (1.5) – (1.6) if and only if: (i) $\phi_{<\widehat{t}}^*$ agrees with $\phi_{<\widehat{t}}^o$ almost surely; (ii) given $p_{\widehat{t}} \sim \mu_{\widehat{t}}^o$, $\phi_{\geq \widehat{t}}^*$ satisfies IC constraints for all $t \geq \widehat{t}$ and is optimal for the dictator's continuation problem starting from time $\widehat{t}$. In particular, ϕ^o is optimal.*

The characterization in Proposition 1.3.2 makes it convenient to explore dynamic features of the optimal design and to study how it should be adjusted with changes in market details. I pursue these in the following sections.

1.4 Dynamic Properties of the Optimal Design

1.4.1 Time Pattern of the Recommendation Standards

Threshold policies can be naturally interpreted as policies setting the minimum age-specific standards for a product to qualify for recommendations. Given the dynamic nature of the problem, it is conceivable that such a minimum standard should evolve over the product's life. I explore this time pattern below.

For ease of exposition, I impose the following full-support and atomless assumption on the belief process.

Assumption 1.4.1. The signals (s_0 and $\{s_n\}_{n\geq 1}$) are such that the marginal distributions of $p_1, ..., p_T$ are atomless and have full support over $[0, 1]$ under any policy.[22]

The atomless assumption renders randomization at the threshold irrelevant, so we can solely focus on the thresholds themselves. The full-support assumption guarantees that any deviation in the recommendation threshold matters, which avoids the need to discuss off-path indeterminacy of the optimal policy.

Recall that $\hat{t}$ is defined by the algorithm in Definition 1.3.2. I have the following result.

Proposition 1.4.1. *Under Assumption 1.4.1, the thresholds $(\eta_t^*)_{t=1}^T$ of any optimal threshold policy satisfies: (a) $\eta_t^* > \eta_{t+1}^*$ for $t \leq \hat{t} - 2$; (b) $\eta_t^* < \eta_{t+1}^*$ for $t \geq \hat{t}$. Moreover, $\eta_t^* \leq \bar{p}$ for all t.*

Intuitively, the optimal recommendation standard should first decrease and then increase over the product's life, which correspond to the two phases with binding and non-binding IC constraints respectively. Underlying this result is the tension between our desire to create dynamic informational value for future consumers and the need to fulfill the current consumer's IC constraint. In the early phase, the dynamic value is generally high with a long future to go, while consumers are more "skeptical" about following recommendations, as they know that even the platform has not acquired much information yet. This implies a binding IC constraint and necessitates a picky censorship over the beliefs eligible for recommendations. As time proceeds, consumers become easier to convince and the recommendation criterion can thus be relaxed. This continues until the standard is already sufficiently low such that consumption with beliefs further below it is no longer worthwhile given the remaining time of the product. The IC constraint then turns non-restrictive. Thereafter, the optimal standard gradually goes up, because the dynamic value from myopically suboptimal consumption dwindles as the product approaches its end of life.

Figure 1.1 explains why the optimal threshold drops before time $\hat{t}$ in more detail. The two dots on the left represent two possible realizations of p_t, and the blue bar between them represents the optimal threshold for period t. Since the green dot is above the threshold, it is associated with a recommendation and will hence split in a mean-preserving spread manner, which leads to realizations of p_{t+1} represented by the two new green dots in period $t + 1$. This

[22] The assumption holds, in particular, if the log-likelihood ratios of the signals are continuous random variables with full support over $\mathbb{R}$. That is, for both $i = 0$ and $i \geq 1$, s_i admits density functions f_i^L and f_i^H conditional on $\tilde{\theta} = \theta_L$ and $\tilde{\theta} = \theta_H$ respectively such that $\log\left(\frac{f_i^H(s_i)}{f_i^L(s_i)}\right)$ is a continuous random variable with full support over $\mathbb{R}$.

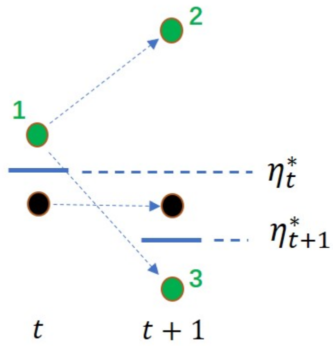

Figure 1.1: Explanation for decreasing recommendation thresholds before time $\widehat{t}$. The dots represent possible realizations of p_t or p_{t+1}. The arrows indicate the evolution of these beliefs. The blue bars represent the optimal recommendation thresholds.

reflects the new information generated by consumption. The black-dot belief in period t does not qualify for a recommendation and is thus carried over into period $t+1$. Now, suppose the designer keeps the threshold unchanged over the two periods. Then in period $t+1$, only the upper green dot (green dot 2) would qualify for a recommendation. Since this belief is more favorable than its predecessor (green dot 1), the consumer's IC constraint in period $t+1$ would turn slack. Intuitively, the better information in period $t+1$ would have induced a more favorable selection for the consumer if the threshold were kept the same as before. This leaves room for the designer to also include the black dot into the recommendation region. When $t+1 < \widehat{t}$, the designer indeed wants to do so since my previous characterization has shown that the consumer's IC constraint should keep binding before time $\widehat{t}$. This implies a lower threshold in period $t+1$.

The U-shaped pattern of recommendation thresholds has interesting implications for the optimal recommendation and learning dynamics, which I turn to in the next subsection.

1.4.2 Optimal Recommendation Dynamic

Figure 1.2 demonstrates an example path of realized recommendations under the optimal policy.[23] The sequence of blue bars represents the age-specific thresholds of the optimal policy, which form a "U" shape by Proposition 1.4.1. The series of crosses and circles tracks a realized path of $(p_t)_{t=1}^T$, where green circles mean that the belief is above the current threshold and the

[23] The figure presents a case where recommendation is eventually abandoned, which is more likely to happen when the product quality is low. However, the property discussed below does not rely on this.

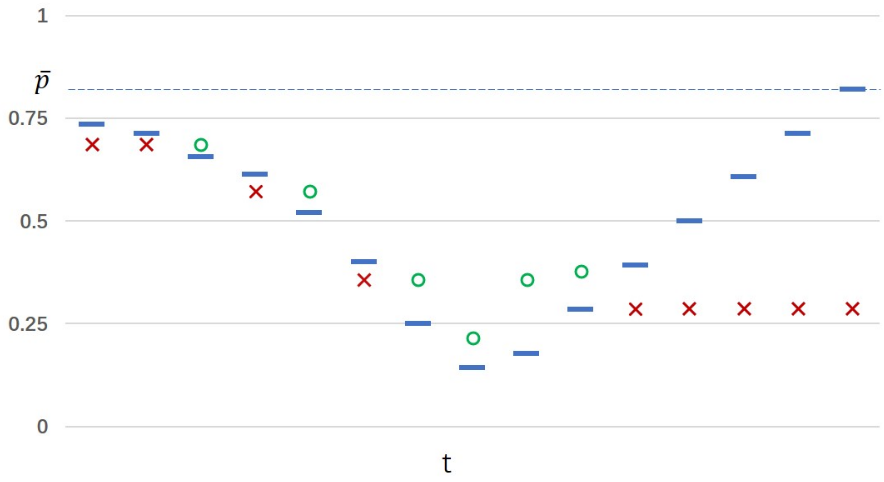

Figure 1.2: A realized path of recommendations under the optimal design. Blue bars represent the thresholds of the optimal policy. The crosses and circles track the platform's belief of $\widehat{\theta} = \theta_H$, where green circles mean that recommendation is made and red crosses mean the opposite. As defined earlier, $\bar{p}$ is the myopically optimal threshold.

product is hence recommended, while red crosses mean the opposite.

The figure highlights an interesting property of the optimal recommendation dynamic – recommendation and learning can be *temporarily* suspended following negative feedback from the last consumption. Such suspension is beneficial because it allows us to support exploration on products looking more promising at the same age without violating the IC constraint. However, the suspension may not last forever when further exploration is still socially desirable. As soon as the threshold for recommendation drops below the current belief, trials for the product will resume. Of course, such restart of recommendations can only happen in the early phase of the product's life, where the IC constraints are binding and the recommendation threshold decreases over time. In the later phase, the threshold goes up, and thus any suspension of recommendation will be permanent.

In practice, *temporary* recommendation suspension (or deprioritization) following negative consumer feedback is often used for punishing misconduct of the product supplier (e.g., a seller or content provider).[24] My finding here suggests another motivation for taking such action, which is to induce social exploration in a more efficient way given the inadequacy of consumer incentive. Compared to punishing supplier misconduct, the temporary recommendation suspensions in my

[24]See, e.g., Table 1 in Liang et al. (2020) for a list of examples.

model have two distinct features. First, they can happen following negative feedback regarding the product's innate features instead of unsatisfactory behavior of the supplier. Second, they mainly happen for young products. Presuming that platforms are in practice indeed trying to enhance social learning about new products by tailoring their recommender systems, these may serve as concrete predictions of the model that can be tested with real data on platform recommendations.

1.5 Comparative Statics

In this section, I provide two comparative static analyses regarding the optimal recommendation standards.

1.5.1 Information Generation Rate

In many applications, having someone try out a product is not guaranteed to generate meaningful information about the product's quality. The consumer may not leave feedback,[25] or may give feedback that is too cursory to be authenticated.[26] These will lead to little post-consumption information generation.

How should the optimal design be adjusted when consumption becomes more likely to yield information, such as with better feedback elicitation designs? To study this, I introduce an information generation rate into my model. Specifically, I assume that following one's consumption, the signal s_i being generated will be a compounded signal. It has probability $\alpha \in (0, 1)$ to be an informative signal and has probability $1-\alpha$ to be uninformative, and the platform can tell which type the signal is.[27] Let $G^I(\cdot|\cdot)$ denote the transition kernel for the platform's belief following an informative signal. Given any α, the transition kernel for p_t following one's consumption

[25] I note that non-feedback is not necessarily uninformative. It will be if there is little self-selection bias, so that the feedback probability is uncorrelated with the product quality, which may be reasonable to assume for some applications.

[26] In particular, platforms like Amazon or Yelp often rely on textual analysis to filter out fake reviews. If a review is not material enough to pass such a test, it will be disregarded or attached with little weight in any recommendation algorithm. This is important to deter fake reviews, which is a topic beyond the scope of this paper.

[27] This compounded signal satisfies requirements in Assumption 1.2.1 as long as the informative component is non-conclusive and indeed informative.

then becomes:

$$G(\cdot|p) = \alpha G^I(\cdot|p) + (1-\alpha)D(\cdot|p) \tag{1.9}$$

The following proposition provides the comparative statics result with respect to α. For simplicity, I still impose the full-support and atomless assumption in Section 1.4.1.

Proposition 1.5.1. *Assume Assumption 1.4.1 holds.*[28] *Given any α, let $(\eta_t^*(\alpha))_{t=1}^T$ denote the thresholds of the optimal threshold policy. Then $\eta_t^*(\alpha)$ weakly decreases in α for all t.*

The proposition suggests that after the information generation rate improves, the optimal recommendation standard should be lowered for all product ages. There are two forces behind this change. First, a higher α implies greater informational value from one's consumption, which motivates more exploration. Second, with a higher α, information from past consumption is accumulated at a faster rate. This enables better-informed recommendations at any time, which makes consumers more willing to follow the recommendations ceteris paribus. We thus have room to lower the recommendation standards without violating the IC constraints. Together, these lead to looser recommendation criteria in the optimal design.

The formal proof of the proposition is technically involved because it requires comparison between two controlled Markov processes corresponding to different α. Central to the proof is a coupling argument, where I explicitly construct the belief processes under the optimal designs corresponding to different α in the same probability space. This allows a direct comparison between them. I refer interested readers to Claim (d) in Appendix 1.B.7.

1.5.2 Random Consumer Arrivals

As is mentioned earlier, my framework easily accommodates random consumer arrivals. I formally illustrate this below and examine how the optimal design should depend on the consumer arrival rate.

I now assume that a consumer arrives in each period with probability $\rho \in (0, 1)$. The arrivals are independent over time and independent from other random objects in the model. Accordingly, I reinterpret a_t as the consumption decision (or the platform's recommendation)

[28] This holds, in particular, if the log-likelihood ratios of s_0 and subsequent *informative* signals are continuous random variables, and the log-likelihood ratio of s_0 has full support over $\mathbb{R}$.

conditional on the consumer's arrival. Then, a signal of quality (i.e., s_i) will be generated following period t if and only if a consumer arrives in that period and $a_t = 1$. The transition rule of the belief process $(p_t)_{t=1}^T$ then becomes:

$$p_{t+1}|p_t, a_t \sim a_t\big[\rho G(\cdot|p_t) + (1-\rho)D(\cdot|p_t)\big] + (1-a_t)D(\cdot|p_t) \tag{1.10}$$

Compared to the transition rule in (1.2), the change is that $G(\cdot|\cdot)$ is now replaced by $\rho G(\cdot|\cdot) + (1-\rho)D(\cdot|\cdot)$. This will be the only change in the design environment.

Another change is needed for the designer's problem. In particular, we should now replace $u(\cdot)$ in the designer's objective function (1.5) with $\rho u(\cdot)$, which reflects the fact that a consumer arrives only with probability ρ. This does not matter for the optimization, however, since it only multiplies the objective function with a strictly positive scalar.

Because Lemma 1.3.1 still holds with $G(\cdot|\cdot)$ replaced by $\rho G(\cdot|\cdot) + (1-\rho)D(\cdot|\cdot)$, all of my previous characterizations for the optimal design will remain valid. The following proposition reveals how the consumer arrival frequency matters for the optimal design.

Proposition 1.5.2. *Assume Assumption 1.4.1 holds.*[29] *Given any arrival rate ρ, let $(\eta_t^*(\rho))_{t=1}^T$ denote the thresholds of the optimal threshold policy. Then $\eta_t^*(\rho)$ weakly decreases in ρ for all t.*

The proposition suggests that when we have a thicker market where consumers arrive more frequently, the recommendation criterion for any product age should be lower. The intuition is again twofold. First, a higher arrival rate means more consumers are likely to come in the future. This increases the informational value from early consumption. Second, with a higher arrival rate, we will in expectation have more consumption and hence more signals accumulated before any given period. This supports more generous recommendations while obeying the IC constraints.

[29] This still holds, in particular, under the conditions in footnote 22, but now there is no need to require $\log\left(\frac{f_i^H(s_i)}{f_i^L(s_i)}\right)$ $(i \geq 1)$ to have full support over $\mathbb{R}$.

1.6 Additional Discussion

1.6.1 Platform's incentive to maximize total consumer surplus

The designer's objective in my model is to maximize the total consumer surplus. One may wonder whether this is in line with a commercial platform's interest. I believe so in at least two scenarios. In the first one, the platform can directly extract consumer surplus by charging a subscription fee (e.g., Netflix). In the second one, the platform does not directly charge consumers, but wants to maximize the user base attracted by its recommendation service. For example, ads-financed search engines want to maximize the quality of their organic recommendations so that more people will use them and see their sponsored ads. In both scenarios, the platform's profit will be increasing in the total consumer surplus generated on it, maximizing which should thus be the primary concern in their recommendation algorithm design.

1.6.2 Non-binary Quality Levels

One restrictive assumption in my main model is that the product quality can only take binary values. As in many other papers on information design or social learning, this allows one to represent the evolving belief with a single-dimensional variable, which significantly simplifies the analysis.[30]

In Appendix 1.A, I extend the model to allow general quality support. Although a full characterization is not available, the duality approach does help to extend certain structures of the optimal design to that general setting. In particular, I show that the optimal design still features the two-phase structure implied by Lemma 1.3.3. Moreover, the optimal policy is more inclined to recommend the product when the platform's current belief about quality is higher in the likelihood-ratio order. This extends the threshold structure to a case with multi-dimensional belief. I will discuss how this result can be helpful for algorithmic recommendation design in the appendix.

1.6.3 Comparison to Previous Studies

As has been mentioned in the introduction, my study is closely related to Kremer et al. (2014) and Che & Hörner (2018), who also study how platform recommendations can improve social

[30] See, e.g., section 2 in Hörner & Skrzypacz (2017) for papers on social experimentation.

learning efficiency when early consumption produces information for later consumers through the platform. Similarly to my paper, they also reveal how past information accrued to the platform enables it to persuade more consumers into socially desirable explorations in the future. Despite this high-level commonality, the three papers yield very different characterizations of the optimal design.

Because Kremer et al. (2014) assumes fully revealing consumption-generated signals, their optimal design is mainly about when to induce the first trial of the product.[31] The main result is that a product that looks better based on the platform's initial information (in its quality relative to an alternative option) should receive the first trial earlier. Che & Hörner (2018) assumes that the platform learns from conclusive news that fully reveals the product quality upon its arrival. The design in their paper is thus about "how much" to recommend the product without news arrival. They show that myopically suboptimal recommendation, given no news arrival, should gain increasing intensity as the product ages until being ceased at some point. In contrast, my study accommodates general non-conclusive consumption-generated signals. My prediction of the optimal design is therefore about whether to recommend the product in each period based on any current belief of the platform. My results suggest that the optimal design features threshold policies with respect to the evolving belief, and the recommendation standard should vary in a U-shaped pattern as the product ages.[32]

Allowing non-conclusive signals also enables richer predictions about the optimal recommendation dynamic. In particular, the phenomenon of temporary recommendation suspensions following negative consumer feedback in Section 1.4.2 cannot exist with conclusive signals, since conclusive negative feedback should necessarily stop recommendation forever. Moreover, both of the previous papers suggest that exploration (i.e., myopically suboptimal trials) should stop after some middle age of the product. With general consumption-generated signals, however, Proposition 1.4.1 implies that exploration can happen until the last period of the product's life, although the belief region for exploration gradually shrinks after some point.

Finally, my study has also provided a couple of comparative statics regarding the optimal recommendation standards in Section 1.5, which do not have counterparts in the previous papers.

[31] After the first trial, quality is fully revealed and recommendations should simply be myopically optimal.

[32] With binary product quality, my result nests that in Kremer et al. (2014) as a special case.

1.7 Conclusion and Methodological Remarks

In this paper, I have studied the optimal design for platform recommendations when early consumption of a product yields informational externality to later consumers through the platform. The optimal design is shown to feature simple threshold policies, and the optimal recommendation standard should vary in a U-shaped pattern over the product's life. An interesting implication about the optimal recommendation dynamic is that recommendation can be temporarily suspended following negative consumer feedback for young products, while such suspension will be permanent for older products. My characterization also enables comparative statics with respect to market details. In particular, I have shown that the recommendation standard should be lowered for all product ages when consumption is more likely to generate informative feedback or when consumers arrive more frequently over time.

My model accommodates non-conclusive consumption-generated signals. Consequently, compared to the literature, it requires the more general formulation of the designer's problem as a *constrained* Markov decision process. I argue that such mathematical formulation can naturally arise in dynamic information design problems when the designer's private information flow is controlled by the receivers' decisions. In such a scenario, if one focuses on direct mechanisms, which are without much loss of generality by the revelation principle, the design problem can be treated as a Markov decision process where the designer decides the receivers' actions and thereby controls his own information flow subject to the receivers' IC constraints. Since the receivers do not observe the designer's information, their IC constraints will involve taking expectations over it. This leads to the aggregated constraints that cannot be directly handled in dynamic programming. I expect that such problem formulation and the Lagrangian duality approach I take will also find applications in other dynamic information design problems with the aforementioned feature.

Appendix

1.A Non-binary Quality Levels

In this appendix, I extend the model to allow for non-binary product quality and generalize certain characterizations of the optimal design.

1.A.1 The General Setting

Consider the same setting as in Section 1.2.1 except that the support of $\tilde{\theta}$ can now be an arbitrary set $\Theta \subset \mathbb{R}$. I assume that the joint distribution of $\tilde{\theta}$ and the signals $(s_0, s_1, ...)$ is such that the platform's posterior belief is always within a family of distributions $\{Q_z\}_{z \in Z}$, where $Z \subset \mathbb{R}^n$ is a countable parameter set.[33] I assume that $\{Q_z\}_{z \in Z}$ admits density functions $\{q_z\}_{z \in Z}$ with respect to some common dominating measure over $\mathbb{R}$, and $q_z(\cdot) > 0$ on Θ for all $z \in Z$. Moreover, I impose the following assumption:

Assumption 1.A.1. (i) $\mathbb{P}(\mathbb{E}[\tilde{\theta}|s_0] > 0) > 0$;
(ii) For any $i \geq 1$, s_i takes values in some set $S \subset \mathbb{R}$. Its distribution conditional on $\tilde{\theta}$ admits a conditional density function $\ell(\cdot|\cdot)$ (w.r.t. some dominating measure over $\mathbb{R}$) such that: $\ell(s|\theta) > 0$ for all $s \in S$ and $\theta \in \Theta$; $\ell(\cdot|\theta)$ increases in θ in the likelihood-ratio order.

Condition (i) plays the same role as its counterpart in Assumption 1.2.1. Condition (ii) implies that higher realizations of s_i suggest that the product is more likely to be of higher quality. This framework is general enough to incorporate many parametric learning models with congruent prior and signals (e.g., the Beta-Binomial model). Moreover, it accommodates any learning model with finite support of $\tilde{\theta}$ that satisfies Assumption 1.A.1.

Let $z_t \in Z$ denote the platform's belief parameter at the beginning of period t. Since z_t (or Q_{z_t}) summarizes all information available to the platform at time t, we can focus on (randomized) Markov recommendation policies w.r.t. $(z_t)_{t=1}^T$. Formally, any policy of this type is a sequence of measurable mappings $\phi := (\phi_t : t = 1, ..., T)$, where each $\phi_t : Z \to [0, 1]$ decides the probability of recommending the product at time t given any belief parameter z_t.

As in Section 1.5.2, I also allow i.i.d. random consumer arrivals and use ρ to denote the arrival probability.

[33] I assume Z to be countable to avoid certain measurability issues about the optimal policy.

1.A.2 Characterizations

Following similar arguments as those in Section 1.3.1 and Section 1.5.2, the designer's problem can be formulated as follows:

$$\max_{\phi\in\Phi}\Big\{\sum_{t=1}^{T}\mathbb{E}_{\phi}[a_t u(z_t)]\Big\}$$
$$\text{s.t. } \mathbb{E}_{\phi}[a_t u(z_t)] \geq 0 \quad \forall t = 1, ..., T$$
$$z_{t+1}|z_t, a_t \sim a_t\big[\rho G(\cdot; z_t) + (1-\rho)D(\cdot; z_t)\big] + (1-a_t)D(\cdot; z_t)$$
$$z_1 \sim \mu_1$$

where $u(z_t) := \int_{\theta\in\Theta} \theta Q_{z_t}(d\theta)$ (i.e., the expected consumption surplus given belief Q_{z_t}), $G(\cdot;\cdot)$ is the transition kernel for z_t following one's consumption, and μ_1 is the distribution of z_1. Compared to the main model, the process of $(z_t)_{t=1}^T$ now replaces the role of $(p_t)_{t=1}^T$. I define the Lagrangian function and the dual problem similarly as those for the main model. The following lemma follows easily from my assumptions and the definition of $(z_t)_{t=1}^T$.

Lemma 1.A.1. *The belief (parameter) process has the following properties*

(P1') $\big[\rho G(\cdot; z) + (1-\rho)D(\cdot; z)\big]$ *as a measure-valued function of* z *is weakly continuous.*[34]

(P2') $\int_{z'} u(z')\big[\rho G(dz'; z) + (1-\rho)D(dz'; z)\big] = u(z)$.

(P5') $\mu_1\big(\{z : u(z) > 0\}\big) > 0$.

These properties are the counterparts to properties (P1), (P2) and (P5) in Lemma 1.3.1. Because Lemmas 1.3.2 and 1.3.3 for my main model only rely on these properties in Lemma 1.3.1, they still hold in the current setting.[35] In particular, we still have the following time pattern of Lagrangian multipliers derived from the dual problem:

Lemma 1.A.2. *There exists* $\lambda^* \in \arg\min_{\lambda\in\mathbb{R}_+^T} \sup_{\phi\in\Phi} \mathcal{L}(\phi;\lambda)$ *such that* $\lambda_t^* \geq \lambda_{t+1}^*$ *for all* $t = 1, ..., T-1$.

One particular implication of this lemma is that the multiplier will stay at zero once dropping to it. As in the main model, this implies that any optimal design features a two-phase structure. In the first phase, IC constraints are binding, and hence the recommendation policy is just

[34] This is trivially true as Z is countable.

[35] The proofs for them remain the same as before except that the role of $(p_t)_{t=1}^T$ is now replaced by $(z_t)_{t=1}^T$.

picky enough for the consumers to follow; in the second phase, the IC constraints become non-restrictive, and the optimal design follows the optimal continuation policy of the dictator.

While I cannot fully pin down the optimal design, the proposition below reveals an important structure of it. Let $\geq_{LR}$ denote dominance in the likelihood-ratio order.

Proposition 1.A.1. *Assume the consumer arrival rate $\rho \in (0,1)$.*[36] *Any optimal policy is almost surely equivalent to some policy ϕ^* such that for all t: if $Q_{z'} \geq_{LR} Q_z$ and $\int_\theta \theta dQ_{z'}(\theta) > \int_\theta \theta dQ_z(\theta)$, then $\phi_t^*(z) > 0 \Rightarrow \phi_t^*(z') = 1$.*

Intuitively, the proposition roughly suggests that any optimal policy should be more inclined to recommend the product when the current belief of quality is higher in the likelihood-ratio order. This naturally extends the threshold structure of the optimal design in my main model to the current setting, where the platform's belief is in a multi-dimensional space endowed only with a partial order.

Although the aforementioned property is intuitively appealing, it is actually violated by many recommendation algorithms proposed in papers that only focus on the asymptotic performance of the design. For example, the algorithm in Mansour et al. (2020) introduces randomized exploration to fulfill the consumers' IC constraints, which necessarily violates the property.[37] My result suggests that modifying their algorithm to be more consistent with this property may help to improve the algorithm's finite-horizon performance. This may be an interesting topic for future algorithm-oriented research.

1.B Proofs

1.B.1 Proof for Lemma 1.3.1

Proof. Property (P2) is implied by the law of iterated expectation; property (P4) is obvious given Assumption 1.2.1(ii) (i.e., the signals are not completely uninformative); property (P5) is directly implied by Assumption 1.2.1(i). I show (P1) and (P3) below.

Fix any $i \geq 1$. Let S denote the signal realization space of s_i. Let $f_L(\cdot)$ and $f_H(\cdot)$ denote s_i's conditional density functions conditional on $\tilde{\theta} = \theta_L$ and $\tilde{\theta} = \theta_H$ respectively, with respect

[36] Although I conjecture that the result should also hold with $\rho = 1$, my current proof requires $\rho < 1$ to avoid some technical subtly.

[37] The algorithm in Mansour et al. (2020) is not Markovian. Hence, more precisely, it is the randomized Markov policy equivalent to their algorithm that does not satisfy the property.

to some dominating measure m over S. Without loss of generality, we can choose m s.t. $f_L(s)$ and $f_H(s)$ are not both equal to zero m-a.s. Since I assume no signal realization fully reveals the value of $\tilde{\theta}$ (i.e., Assumption 1.2.1(iii)), we also have $f_L(s) \neq 0 \Leftrightarrow f_H(s) \neq 0$ m-a.s. Thus $f_L(s)$ and $f_H(s)$ are non-zero m-a.s. Define the log-likelihood ratio $\ell_i = \log\big(f_H(s_i)/f_L(s_i)\big)$, and let J_L and J_H denote its distribution given $\tilde{\theta} = \theta_L$ and $\tilde{\theta} = \theta_H$ respectively.[38]

We have the following observation:

Claim *(a).* For any a, $J_H(a) = \int_{\ell \leq a} e^{\ell} dJ_L(\ell)$.

Proof for Claim (a). The following equalities hold:

$$\begin{aligned}
&\int_{\ell \leq a} e^{\ell} dJ_L(\ell) \overset{①}{=} \mathbb{E}[\mathbb{1}_{\{\ell_i \leq a\}} e^{\ell_i} | \tilde{\theta} = \theta_L] \overset{②}{=} \mathbb{E}\Big[\mathbb{1}_{\{\log \frac{f_H(s_i)}{f_L(s_i)} \leq a\}} \frac{f_H(s_i)}{f_L(s_i)} \Big| \tilde{\theta} = \theta_L\Big] \\
&\overset{③}{=} \int \mathbb{1}_{\{\log \frac{f_H(s)}{f_L(s)} \leq a\}} \frac{f_H(s)}{f_L(s)} f_L(s) m(ds) \overset{④}{=} \int \mathbb{1}_{\{\log \frac{f_H(s)}{f_L(s)} \leq a\}} f_H(s) m(ds) \\
&\overset{⑤}{=} \mathbb{E}[\mathbb{1}_{\{\ell_i \leq a\}} | \tilde{\theta} = \theta_H] \overset{⑥}{=} J_H(a)
\end{aligned}$$

where the first and the last equalities hold by the definitions of J_L and J_H respectively; the second equality holds by the definition of ℓ_i; the third and the fifth equalities hold by the definitions of f_L and f_H respectively; the fourth equality is a trivial identity. □

Now, given any prior belief p about $\tilde{\theta} = \theta_H$, let $\tilde{p}$ denote the posterior belief given s_i. Then by the Bayes rule we have: $\log \frac{\tilde{p}}{1-\tilde{p}} = \log \frac{p}{1-p} + \ell_i$. Let $\mathbb{P}_p$ denote the probability measure given prior p. This then implies that

$$\mathbb{P}_p(\tilde{p} \leq x) = \mathbb{P}_p\big(\ell_i \leq \log \frac{x}{1-x} - \log \frac{p}{1-p}\big) \tag{1.11}$$

$$= pJ_H\big(\log \frac{x}{1-x} - \log \frac{p}{1-p}\big) + (1-p)J_L\big(\log \frac{x}{1-x} - \log \frac{p}{1-p}\big) \tag{1.12}$$

$$= \int_{\ell \leq \log \frac{x}{1-x} - \log \frac{p}{1-p}} [pe^{\ell} + (1-p)] dJ_L(\ell) \tag{1.13}$$

where the last equality holds by Claim (a) above. Now, pick any $p^* \in \mathbb{R}$ and a sequence of $(p_n)_n \to p^*$. When expression (1.13) is continuous in x at $x = x_0$ given $p = p^*$, obviously we must have $J_L(\ell)$ being continuous at $\ell = \log \frac{x_0}{1-x_0} - \log \frac{p^*}{1-p^*}$, which further implies that the expression (1.13) is continuous in p at $p = p^*$ given $x = x_0$. Thus $\mathbb{P}_{p^*}(\tilde{p} \leq x)$ being continuous

[38] Such log-likelihood ratio representation of a signal has been previously used in Smith & Tian (2018).

in x at $x = x_0$ implies $\mathbb{P}_{p_n}(\tilde{p} \leq x_0) \to \mathbb{P}_{p^*}(\tilde{p} \leq x_0)$ as $n \to \infty$. Therefore, the distribution of $\tilde{p}$ given prior p is weakly continuous in p. This proves the weak continuity condition for $G(\cdot|p)$ in (P1).

To check property (P3), we need the following observation:

Claim *(b).* For any a, $\int_{\ell \leq a} e^{\ell} dJ_L(\ell) \leq \int_{\ell \leq a} dJ_L(\ell)$.

Proof for Claim (b). Notice by Claim (a) above, $e^{\ell} dJ_L(\ell)$ just equals to $dJ_H(\ell)$. We can thus treat both $e^{\ell} dJ_L(\ell)$ and $dJ_L(\ell)$ as probability measures over $\mathbb{R}$, with densities e^{ℓ} and 1 respectively w.r.t. the dominating measure $dJ_L(\ell)$. Since e^{ℓ} is increasing in ℓ, we then have $e^{\ell} dJ_L(\ell)$ dominating $dJ_L(\ell)$ in the likelihood-ratio order.[39] This further implies dominance in first-order stochastic dominance and thus $\int_{\ell \leq a} e^{\ell} dJ_L(\ell) \leq \int_{\ell \leq a} dJ_L(\ell)$. □

Now, pick any p_a and p_b s.t. $p_a < p_b$, we have

$$\begin{aligned}
\int_{\ell \leq \log \frac{x}{1-x} - \log \frac{p_b}{1-p_b}} [p_b e^{\ell} + (1-p_b)] dJ_L(\ell) &\leq \int_{\ell \leq \log \frac{x}{1-x} - \log \frac{p_a}{1-p_a}} [p_b e^{\ell} + (1-p_b)] dJ_L(\ell) \\
&\leq \int_{\ell \leq \log \frac{x}{1-x} - \log \frac{p_a}{1-p_a}} [p_a e^{\ell} + (1-p_a)] dJ_L(\ell)
\end{aligned}$$

where the second inequality holds due to Claim (b). Together with equations (1.11)–(1.13), this implies that $\mathbb{P}_p(\tilde{p} \leq x)$ is weakly decreasing in p for any x. Thus we have the property of (P3). *Q.E.D.*

1.B.2 Proof for Proposition 1.3.1

Proof. The proof basically applies Lemma 1(iv) in Feinberg & Piunovskiy (2000) to my setting. Specifically, define

$$\mathcal{V} = \left\{ v \in \mathbb{R}^{T+1} : \exists \phi \in \Phi \text{ s.t. } v_t = \mathbb{E}_{\phi}[a_t u(p_t)] \, \forall t = 1, ...T \text{ and } v_{T+1} = \sum_{t=1}^{T} \mathbb{E}_{\phi}[a_t u(p_t)] \right\}$$

Notice for each admissible policy ϕ, the first T arguments of the corresponding vector v are the values of the IC constraints and the $(T+1)$'th argument of v is just the total surplus in the designer's objective. Lemma 1(iv) in Feinberg & Piunovskiy (2000) implies that $\mathcal{V}$ is a compact

[39] See, e.g., section 1.4 in Müller & Stoyan (2002) for an introduction to such order.

set. This further implies that the set $\mathcal{V} \cap \{v \in \mathbb{R}^{T+1} : v_t \geq 0 \, \forall t = 1, ...T\}$ is compact, and thus when we maximize over its $(T+1)$'th dimension, the supremum is achievable. By the definition of $\mathcal{V}$, this is equivalent to that the designer's problem has its supremum achieved by some ϕ.

Now, it suffices to check that the four conditions of Lemma 1 in Feinberg & Piunovskiy (2000) are indeed satisfied in my setting. Condition 1 holds because my state space $[0,1]$ is closed, and the set of feasible actions $A = \{0,1\}$ is finite and does not vary in time and state. Conditions 2 and 4 hold because the flow payoffs in my setting are bounded and continuous in the pair of action and state, and is non-zero for only finitely many periods. For Condition 3, we just need to show the transition probability $aG(\cdot|p) + (1-a)D(\cdot|p)$ is weakly-continuous in $(a,p) \in \{0,1\} \times [0,1]$. With $\{0,1\}$ endowed with the discrete topology, it suffices to check weak continuity in p when $a = 1$ and $a = 0$ separately. These are respectively implied by the weak continuity of $G(\cdot|p)$ (Property (P1) in Lemma 1.3.1) and $D(\cdot|p)$ in p.[40]

Q.E.D.

1.B.3 Proof for Lemma 1.3.2

Proof. To use the Lagrangian duality theorem, I first transform the designer's problem into a linear program. Throughout, I fix the initial belief state distribution μ_1. Given any policy $\phi \in \Phi$, let m_t^ϕ denote the distribution of (a_t, p_t) under it.[41]

Let $\mathcal{M}$ denote the set of all sequences of such distributions under some ϕ, i.e., $\mathcal{M} = \{(m_t^\phi)_{t=1}^T : \phi \in \Phi\}$. A (standard) characterization for this set is that $(m_t)_{t=1}^T \in \mathcal{M}$ if and only if:

$$m_1(\{0,1\} \times B) = \mu_1(B) \tag{1.14}$$

$$m_{t+1}(\{0,1\} \times B) = \int_{p \in [0,1]} \sum_{a \in \{0,1\}} [aG(B|p) + (1-a)D(B|p)] m_t(a, dp) \, \forall t = 1, ..., T-1 \tag{1.15}$$

for any $B \in \mathcal{B}_{[0,1]}$ (Borel σ-field of $[0,1]$). I use $\widehat{\mathcal{M}}$ to denote the set of $(m_t)_{t=1}^T$ satisfying these conditions. The fact that $\mathcal{M} \subset \widehat{\mathcal{M}}$ is obvious since any $(m_t^\phi)_{t=1}^T$ must be consistent with μ_1 and

[40] To see $D(\cdot|p)$ is weakly-continuous in p, notice its cdf is just $\mathbb{1}_{\{x \geq p\}}$. Given any sequence $(p_n)_n \to p^*$, we have $\mathbb{1}_{\{x \geq p_n\}} \to \mathbb{1}_{\{x \geq p^*\}}$ for any $x \neq p^*$. The weak-continuity is thus implied (see, e.g., Section 3.2 in Durrett (2019) for conditions of weak-continuity).

[41] As is standard, we can construct the underlying measurable space for the process as $(\{0,1\} \times [0,1])^T$ with the Borel σ-field, and treat the corresponding random variables as identity mappings $(\{0,1\} \times [0,1])^T \to (\{0,1\} \times [0,1])^T$. Thus we can treat any distribution for those random variables equivalently as a measure over the underlying measurable space, which is typically how I interpret those distributions.

the transition probabilities, and thus satisfies (1.14) and (1.15).

To see $\widehat{\mathcal{M}} \subset \mathcal{M}$, pick any $(m_t^*)_{t=1}^T \in \widehat{\mathcal{M}}$. Let ϕ^* be a (randomized) Markov policy such that ϕ_t^* is just the conditional probability mass function of a_t given p_t under m_t^*. Formally, for any m_t, treat $m_t(a, dp)$ $(a = 1, 2)$ as a measure over $[0, 1]$ s.t. $m_t(a, B) = m_t(\{a\} \times B)$, $\forall B \in \mathcal{B}_{[0,1]}$. Then ϕ^* is defined as (an arbitrary version of) the Radon-Nikodym derivative of $m_t^*(1, dp)$ w.r.t. $m_t^*(0, dp) + m_t^*(1, dp)$. (Notice $m_t^*(0, dp) + m_t^*(1, dp)$ is just the marginal distribution of m_t^* over $[0, 1]$ and the Radon-Nikodym derivative is by definition measurable.) Then, we can show ϕ^* implements $(m_t^*)_{t=1}^T$ by induction in t. Let $m_t^{\phi^*}$ denote the joint distribution of (a_t, p_t) under ϕ^* for any t. For $t = 1$, we have for all $B \in \mathcal{B}_{[0,1]}$:

$$\begin{aligned} m_1^{\phi^*}(\{1\} \times B) = \int_{p \in B} \phi_1^*(p) \mu_1(dp) &= \int_{p \in B} \phi_1^*(p)[m_1^*(0, dp) + m_1^*(1, dp)] \\ &= \int_{p \in B} m_1^*(1, dp) = m_1^*(\{1\} \times B) \end{aligned}$$

where the second equality holds by condition (1.14) and the third equality holds by the definition of ϕ^*. Since $m_1^{\phi^*}(\{1\} \times B) + m_1^{\phi^*}(\{0\} \times B) = \mu_1(B) = m_1^*(\{1\} \times B) + m_1^*(\{0\} \times B)$, we also have $m_1^{\phi^*}(\{0\} \times B) = m_1^*(\{0\} \times B)$. Thus $m_1^{\phi^*} = m_1^*$.

Now, assume $m_t^{\phi^*} = m_t^*$ and consider the result for $t+1$. Because condition (1.15) holds for m_{t+1}^*, we know that the marginal distribution over $[0, 1]$ under m_{t+1}^* given m_t^* is determined by the same rule as that determines the marginal distribution over $[0, 1]$ under $m_{t+1}^{\phi^*}$ given $m_t^{\phi^*}$. Thus $m_t^* = m_t^{\phi^*}$ implies $m_{t+1}^*(\{0, 1\} \times B) = m_{t+1}^{\phi^*}(\{0, 1\} \times B)$, $\forall B \in \mathcal{B}_{[0,1]}$. This further implies:

$$\begin{aligned} m_{t+1}^{\phi^*}(\{1\} \times B) &= \int_{p \in B} \phi_{t+1}^*(p)[m_{t+1}^{\phi^*}(0, dp) + m_{t+1}^{\phi^*}(1, dp)] \\ = \int_{p \in B} \phi_{t+1}^*(p)&[m_{t+1}^*(0, dp) + m_{t+1}^*(1, dp)] = m_{t+1}^*(\{1\} \times B) \end{aligned}$$

where the second equality holds because the two measures are equal as mentioned right above and the third equality holds by the definition of ϕ^*. Together with $m_{t+1}^*(\{0, 1\} \times B) = m_{t+1}^{\phi^*}(\{0, 1\} \times B)$, this also implies $m_{t+1}^{\phi^*}(\{0\} \times B) = m_{t+1}^*(\{0\} \times B)$. Therefore, $m_{t+1}^* = m_{t+1}^{\phi^*}$. This completes the induction proof for showing that m^* is implemented with ϕ^*.

The above discussion has shown $\mathcal{M} = \widehat{\mathcal{M}}$. We can thus rewrite the designer's problem as

$$\max_{(m_t)_{t=1}^T \in \widehat{\mathcal{M}}} \Big\{ \sum_{t=1}^{T} \Big[\int_{p\in[0,1]} \sum_{a\in\{0,1\}} au(p)m_t(a,dp) \Big] \Big\}$$
$$\text{s.t.} \int_{p\in[0,1]} \sum_{a\in\{0,1\}} au(p)m_t(a,dp) \geq 0 \quad \forall t = 1, ..., T$$

Since conditions (1.14) and (1.15) are affine in $(m_t)_{t=1}^T$, the set $\widehat{\mathcal{M}}$ is a convex subset of $\big\{\text{signed Borel measures on } \{0,1\} \times [0,1]\big\}^T$. The optimization above is thus a linear program over this convex set $\widehat{\mathcal{M}}$.

Let $\widehat{\mathcal{L}}((m_t)_{t=1}^T; \lambda) := \sum_{t=1}^T (1+\lambda_t)\big[\int_{p\in[0,1]} \sum_{a\in\{0,1\}} au(p)m_t(a,dp)\big]$, i.e., the Lagrangian function associated to the linear program. Since $u(\cdot)$ is bounded by Lemma 1.3.1 and $T < \infty$, the optimal value w^* is finite. Standard Lagrangian duality (e.g., Theorem 1 in Section 8.6 of Luenberger (1997)) then implies:[42]

$$w^* = \min_{\lambda\in\mathbb{R}_+^T} \sup_{(m_t)_{t=1}^T \in \widehat{\mathcal{M}}} \widehat{\mathcal{L}}((m_t)_{t=1}^T; \lambda)$$

where the minimum is achieved by some non-negative λ^*. Given any such λ^*, $(m_t^*)_{t=1}^T \in \widehat{\mathcal{M}}$ solves the linear program if and only if:

(i) $(m_t^*)_{t=1}^T \in \arg\max_{(m_t)_{t=1}^T \in \widehat{\mathcal{M}}} \widehat{\mathcal{L}}((m_t)_{t=1}^T; \lambda^*)$

(ii) $\lambda_t^* \int_{p\in[0,1]} \sum_{a\in\{0,1\}} au(p)m_t^*(a,dp) = 0, \forall t = 1, ..., T$

(iii) $\int_{p\in[0,1]} \sum_{a\in\{0,1\}} au(p)m_t^*(a,dp) \geq 0, \forall t = 1, ..., T$

To check the corresponding Slater's condition, notice by properties (P2) and (P5) in Lemma 1.3.1, the consumer's surplus will be strictly positive at all t under the myopically optimal policy, and thus all IC constraints can hold strictly.

Finally, notice $\mathcal{M} = \widehat{\mathcal{M}}$ just means that $(m_t)_{t=1}^T \in \widehat{\mathcal{M}}$ if and only if it is induced by some $\phi \in \Phi$. Thus the above results are equivalent to the statements in the lemma. *Q.E.D.*

[42] Theorem 1 in Section 8.6 of Luenberger (1997) does not directly state the sufficiency of conditions (i) – (iii) for optimality. However, this is obvious as those conditions together imply $(m_t^*)_{t=1}^T$ is feasible and $\sum_{t=1}^T \big[\int_{p\in[0,1]} \sum_{a\in\{0,1\}} au(p)m_t^*(a,dp)\big] = \sum_{t=1}^T (1+\lambda_t)\big[\int_{p\in[0,1]} \sum_{a\in\{0,1\}} au(p)m_t^*(a,dp)\big] = w^*$.

1.B.4 Analyses and Proofs for Section 1.3.3

Preliminaries

Towards using the duality result, I start with examining the Lagrangian function optimization $\max_{\phi\in\Phi}\mathcal{L}(\phi;\lambda)$ given any generic multiplier $\lambda\in\mathbb{R}_+^T$. For this *unconstrained* Markov decision problem, let $V_\lambda(\cdot,t)$ be the value function at time t, which is inductively defined with the Bellman equation:

$$V_\lambda(\cdot,T+1)\equiv 0 \tag{1.16}$$
$$V_\lambda(p,t)=\max\left\{(1+\lambda_t)u(p)+\int_{p'}V_\lambda(p',t+1)G(dp'|p),\,V_\lambda(p,t+1)\right\}\forall t=1,...,T \tag{1.17}$$

where the two arguments in the maximization correspond to the values with and without time-t consumption of the product respectively. I define $H_\lambda(p,t)$ to be the difference between these two values, i.e.,

$$H_\lambda(p,t):=(1+\lambda_t)u(p)+\int_{p'}V_\lambda(p',t+1)G(dp'|p)-V_\lambda(p,t+1)\,\forall t=1,...,T \tag{1.18}$$

Intuitively, $V_\lambda(p,t)$ is the continuation value for the Lagrangian optimization at time t given $p_t=p$; $H_\lambda(p,t)$ measures the net benefit from inducing the time-t consumption of the product given $p_t=p$. A preliminary result needed later is that $H_\lambda(\cdot,t)$ is continuous.

Lemma 1.B.1. *$H_\lambda(\cdot,t)$ is continuous for any t.*

Proof. Since $u(\cdot)$ is continuous by definition, we just need to show $V_\lambda(p,t+1)$ and $\int_{p'}V_\lambda(p',t+1)G(dp'|p)$ are continuous in p. When $t=T$, these hold by the definition of $V_\lambda(\cdot,T+1)$. Given that they hold for time t, by the Bellman equation we know $V_\lambda(p,t)$ is also continuous in p. Furthermore, because $G(dp'|p)$ is weakly continuous in p and $V_\lambda(p,t)$ is bounded due to the boundedness of $u(\cdot)$, this also implies the continuity of $\int_{p'}V_\lambda(p',t)G(dp'|p)$ in p.[43] The proof is thus completed by (backward) induction in t. *Q.E.D.*

We have the standard dynamic programming result:

Lemma 1.B.2. *Given any multiplier λ and initial belief state distribution μ_1, we have:*

[43]Pick any $(p_n)_n\to p^*$, the weak continuity implies $G(dp'|p_n)\overset{w}{\to}G(dp'|p^*)$, which further implies $\int f(p')G(dp'|p_n)\to\int f(p')G(dp'|p^*)$ for any bounded and continuous function f.

(a) $\max_{\phi\in\Phi}\mathcal{L}(\phi;\lambda)=\int_p V_\lambda(p,1)\mu_1(dp)$*;*

(b) A policy $\phi^\lambda\in\arg\max_{\phi\in\Phi}\mathcal{L}(\phi;\lambda)$ *if and only if* $H_\lambda(p_t,t)>0\Rightarrow a_t=1$ *and* $H_\lambda(p_t,t)<0\Rightarrow a_t=0$ *almost surely under it.*

Proof. Pick any $\phi\in\Phi$. Notice that by the definition of V_λ, we have

$$\begin{aligned} V_\lambda(p_t,t) &\geq \phi_t(p_t)\Big[(1+\lambda_t)u(p_t)+\int_{p'}V_\lambda(p',t+1)G(dp'|p_t)\Big]+\big(1-\phi_t(p_t)\big)V_\lambda(p_t,t+1)\\ &\overset{\text{a.s.}}{=} \mathbb{E}_\phi[(1+\lambda_t)a_t u(p_t)|p_t]+\mathbb{E}_\phi[V_\lambda(p_{t+1},t+1)|p_t] \end{aligned} \tag{1.19}$$

where the inequality holds as equality if and only if $H_\lambda(p_t,t)>0\Rightarrow\phi_t(p_t)=1$ and $H_\lambda(p_t,t)<0\Rightarrow\phi_t(p_t)=0$.

Using this repeatedly, we have:

$$\begin{aligned} \int_p V_\lambda(p,1)\mu_1(dp)=\mathbb{E}_\phi[V_\lambda(p_1,1)] &\geq \mathbb{E}_\phi[(1+\lambda_1)a_1u(p_1)]+\mathbb{E}_\phi[V_\lambda(p_2,2)]\\ &\geq \mathbb{E}_\phi[(1+\lambda_1)a_1u(p_1)]+\mathbb{E}_\phi[(1+\lambda_2)a_2u(p_2)]+\mathbb{E}_\phi[V_\lambda(p_3,3)]\\ ...&\geq \mathbb{E}_\phi[(1+\lambda_1)a_1u(p_1)]+...+\mathbb{E}_\phi[(1+\lambda_T)a_Tu(p_T)]+\underbrace{\mathbb{E}_\phi[V_\lambda(p_{T+1},T+1)]}_{=0} \end{aligned}$$

Notice the last line is just $\mathcal{L}(\phi;\lambda)$. This shows $\int_p V_\lambda(p,1)\mu_1(dp)\geq\sup_{\phi\in\Phi}\mathcal{L}(\phi;\lambda)$. Moreover, notice that all these inequalities hold as equalities if and only if the inequality in (1.19) holds as equality for all t almost surely under ϕ. As is mentioned earlier, this is equivalent to $H_\lambda(p_t,t)>0\Rightarrow\phi_t(p_t)=1$ and $H_\lambda(p_t,t)<0\Rightarrow\phi_t(p_t)=0$ almost surely under ϕ. If there is indeed a measurable ϕ satisfying these properties, then we have $\int_p V_\lambda(p,1)\mu_1(dp)=\sup_{\phi\in\Phi}\mathcal{L}(\phi;\lambda)$, the supremum is achieved by such policy, and any other admissible policy is optimal if and only if it also satisfies these properties. Therefore, to prove the lemma, it now suffices to show that there is indeed a measurable policy satisfying $H_\lambda(p_t,t)>0\Rightarrow\phi_t(p_t)=1$ and $H_\lambda(p_t,t)<0\Rightarrow\phi_t(p_t)=0$ almost surely. We can construct such policy by defining $\phi_t(p)=\mathbb{1}_{\{H_\lambda(p,t)\geq 0\}}$, where $\mathbb{1}$ is the indicator function. It is indeed measurable since $H_\lambda(\cdot,t)$ is continuous by Lemma 1.B.1.

Q.E.D.

I now provide some basic properties for the value function V_λ.

Lemma 1.B.3. *Given any* $\lambda\in\mathbb{R}^T_+$*, we have: (a)* $V_\lambda(p,t)$ *is (weakly) increasing in* p*; (b)* $\int_{p'}V_\lambda(p',t)G(dp'|p)\geq V_\lambda(p,t)$ *for any pair of* (p,t)*.*

Proof. The results can be shown by backward induction in t using the Bellman equation. Since $V_\lambda(\cdot, T+1) \equiv 0$, both properties hold trivially for $t = T+1$.

Now, assume property (a) holds for $V_\lambda(\cdot, t+1)$. Then property (P3) in Lemma 1.3.1 implies that $\int_{p'} V_\lambda(p', t+1) G(dp'|p)$ increases in p. Together with the monotonicity of $u(\cdot)$ and the fact that $\lambda_t \geq 0$, we know the RHS of the Bellman equation (1.17) is increasing in p. This shows property (a) also holds for $V_\lambda(\cdot, t)$ and concludes the proof for part (a).

For property (b), assume it holds for all periods later than t and I show $\int_{p'} V_\lambda(p', t) G(dp'|p) \geq V_\lambda(p, t)$. Substituting the Bellman equation in, we know this is equivalent to:

$$\int_{p'} \max\Big\{(1+\lambda_t)u(p') + \int_{p''} V_\lambda(p'', t+1) G(dp''|p'),\, V_\lambda(p', t+1)\Big\} G(dp'|p)$$
$$\geq \max\Big\{(1+\lambda_t)u(p) + \int_{p'} V_\lambda(p', t+1) G(dp'|p),\, V_\lambda(p, t+1)\Big\}$$

It then suffices to check:

$$\int_{p'} \Big((1+\lambda_t)u(p') + \int_{p''} V_\lambda(p'', t+1) G(dp''|p')\Big) G(dp'|p)$$
$$\geq (1+\lambda_t)u(p) + \int_{p'} V_\lambda(p', t+1) G(dp'|p)$$
$$\text{and } \int_{p'} V_\lambda(p', t+1) G(dp'|p) \geq V_\lambda(p, t+1)$$

The second of these inequalities is directly implied by the induction hypothesis. To check the first one, notice by property (P2) in Lemma 1.3.1, we have $\int_{p'} (1+\lambda_t) u(p') G(dp'|p) = (1+\lambda_t)u(p)$. Moreover, by the induction hypothesis we have $\int_{p''} V_\lambda(p'', t+1) G(dp''|p') \geq V_\lambda(p', t+1)$ for any p', and thus $\int_{p'} \int_{p''} V_\lambda(p'', t+1) G(dp''|p') G(dp'|p) \geq \int_{p'} V_\lambda(p', t+1) G(dp'|p)$. These together imply the first inequality above. Thus property (b) holds for period t. This completes the proof by induction.

Q.E.D.

Proof for Lemma 1.3.3

Proof. Pick any $\lambda' \in \mathbb{R}^T_+$ such that $\lambda'_\tau < \lambda'_{\tau+1}$ for some τ. Define λ'' to be equal to λ' except for terms of time τ and $\tau+1$, which are defined as: $\lambda''_\tau = \lambda'_{\tau+1}$ and $\lambda''_{\tau+1} = \lambda'_\tau$. The key to the proof is the following observation:

Claim. Let λ' and λ'' be defined as above. Then $\sup_{\phi\in\Phi} \mathcal{L}(\phi; \lambda') \geq \sup_{\phi\in\Phi} \mathcal{L}(\phi; \lambda'')$.

Proof for the Claim. Since λ' and λ'' agree for $t < \tau$, any policy will lead to the same flow payoffs for periods before τ for both $\mathcal{L}(\phi; \lambda')$ and $\mathcal{L}(\phi; \lambda'')$. It thus suffices to show $V_{\lambda'}(p, \tau) \geq V_{\lambda''}(p, \tau)$ given any p.

Since λ' and λ'' also agree for $t \geq \tau + 2$, we have $V_{\lambda'}(\cdot, \tau + 2) = V_{\lambda''}(\cdot, \tau + 2)$. I can thus let $V_*(\cdot, \tau + 2)$ to denote both of them, i.e., $V_*(\cdot, \tau + 2) := V_{\lambda'}(\cdot, \tau + 2)$ $(= V_{\lambda''}(\cdot, \tau + 2))$. Also let $(y_1, y_2) := (\lambda'_\tau, \lambda'_{\tau+1})$. Then $y_1 < y_2$ and $(\lambda''_\tau, \lambda''_{\tau+1}) = (y_2, y_1)$. By the Bellman equation (1.17), we have:

$$V_{\lambda''}(p, \tau) = \max\left\{(1 + y_2)u(p) + \int_{p'} V_{\lambda''}(p', \tau + 1)G(dp'|p),\ V_{\lambda''}(p, \tau + 1)\right\} \tag{1.20}$$

Consider the following two cases:

- **Case 1:** the maximum in equation (1.20) is achieved with $a_\tau = 0$ (i.e., no consumption).

  In this case, we have:

$$\begin{aligned} V_{\lambda''}(p, \tau) =& V_{\lambda''}(p, \tau + 1) \\ =& \max\left\{(1 + y_1)u(p) + \int_{p'} V_*(p', \tau + 2)G(dp'|p),\ V_*(p, \tau + 2)\right\} \\ \leq& \max\left\{(1 + y_1)u(p) + \int_{p'} V_{\lambda'}(p', \tau + 1)G(dp'|p),\ V_{\lambda'}(p, \tau + 1)\right\} \\ =& V_{\lambda'}(p, \tau) \end{aligned}$$

  where the second equality holds by the Bellman equation for $V_{\lambda''}(p, \tau + 1)$; the inequality holds because the Bellman equation for $V_{\lambda'}(p, \tau + 1)$ implies that $V_*(p, \tau + 2) \leq V_{\lambda'}(p, \tau + 1)$ for any p; the last equality is just the Bellman equation for $V_{\lambda'}(p, \tau)$.

- **Case 2:** the maximum in equation (1.20) is achieved with $a_\tau = 1$.

  In this case, we have:

$$V_{\lambda''}(p, \tau) = (1 + y_2)u(p) + \int_{p'} V_{\lambda''}(p', \tau + 1)G(dp'|p) \tag{1.21}$$

$$\leq (1 + y_1)u(p) + \int_{p'} \big[\underbrace{(y_2 - y_1)u(p') + V_{\lambda''}(p', \tau + 1)}_{=:M(p')}\big]G(dp'|p) \tag{1.22}$$

  where the inequality holds because $y_2 > y_1$ and $u(p) = \int_{p'} u(p')G(dp'|p)$ by property (P2) in Lemma 1.3.1. Using the Bellman equation for $V_{\lambda''}(p', \tau + 1)$, we know the term $M(p')$

satisfies:

$$\begin{aligned} M(p') =&(y_2 - y_1)u(p') + \max\left\{(1+y_1)u(p') + \int_{p''} V_*(p'', \tau+2)G(dp''|p'), V_*(p', \tau+2)\right\} \\ =& \max\left\{(1+y_2)u(p') + \int_{p''} V_*(p'', \tau+2)G(dp''|p'), V_*(p', \tau+2) + (y_2 - y_1)u(p')\right\} \end{aligned} \tag{1.23}$$

Now, consider properties of $M(p')$ in two different scenarios about p':

– Scenario 1: $u(p') > 0$.

  Notice Lemma 1.B.3(b) implies $\int_{p''} V_*(p'', \tau+2)G(dp''|p') \geq V_*(p', \tau+2)$. When $u(p') > 0$, we thus have $(1+y_2)u(p') + \int_{p''} V_*(p'', \tau+2)G(dp''|p') \geq V_*(p', \tau+2) + (y_2 - y_1)u(p')$. Therefore, (1.23) implies:

$$\begin{aligned} M(p') &= (1+y_2)u(p') + \int_{p''} V_*(p'', \tau+2)G(dp''|p') \\ &= \max\left\{(1+y_2)u(p') + \int_{p''} V_*(p'', \tau+2)G(dp''|p'), V_*(p', \tau+2)\right\} \\ &= V_{\lambda'}(p', \tau+1) \end{aligned}$$

  where the second equality holds since the first argument in the bracket is larger when $\int_{p''} V_*(p'', \tau+2)G(dp''|p') \geq V_*(p', \tau+2)$ and $u(p') > 0$.

– Scenario 2: $u(p') \leq 0$.

  In this case, we have $(y_2 - y_1)u(p') \leq 0$. (1.23) thus implies:

$$\begin{aligned} M(p') \leq& \max\left\{(1+y_2)u(p') + \int_{p''} V_*(p'', \tau+2)G(dp''|p'), V_*(p', \tau+2)\right\} \\ =& V_{\lambda'}(p', \tau+1) \end{aligned}$$

  where the equality holds by the Bellman equation.

In both scenarios, we always have $M(p') \leq V_{\lambda'}(p', \tau+1)$. Together with inequality (1.22), this implies that $V_{\lambda''}(p, \tau) \leq (1+y_1)u(p) + \int_{p'} V_{\lambda'}(p', \tau+1)G(dp'|p) \leq V_{\lambda'}(p, \tau)$, where the second inequality is due to the Bellman equation for $V_{\lambda'}(p, \tau)$.

In sum, $V_{\lambda''}(p, \tau) \leq V_{\lambda'}(p, \tau)$ in both cases, which completes the proof for the claim.

□

I now go back to the main proof for Lemma 1.3.3. Pick any $\lambda^0 \in \arg\min_{\lambda\in\mathbb{R}^T_+} \sup_{\phi\in\Phi} \mathcal{L}(\phi;\lambda)$ (whose existence is guaranteed by Lemma 1.3.2). If λ^0_t is already non-increasing in t, we are done; if $\lambda^0_t < \lambda^0_{t+1}$ for some t, then the claim above implies that by interchanging terms λ^0_t and λ^0_{t+1}, we will get a new multiplier still in $\arg\min_{\lambda\in\mathbb{R}^T_+} \sup_{\phi\in\Phi} \mathcal{L}(\phi;\lambda)$. By repeatedly making such interchanges, we can then derive a multiplier $\lambda^* \in \arg\min_{\lambda\in\mathbb{R}^T_+} \sup_{\phi\in\Phi} \mathcal{L}(\phi;\lambda)$ such that $\lambda^*_t \geq \lambda^*_{t+1}$, $\forall t$.[44]

Q.E.D.

Properties of H_λ

Next, I provide some properties of H_λ, which are key to the proof for Lemma 1.3.4.

Lemma 1.B.4. *H_λ satisfies the following properties:*

(a) $p > \bar{p} \Rightarrow H_\lambda(p,t) > 0$, $\forall t$;

(b) If λ_t is non-increasing in t, then $H_\lambda(p,t)$ is (weakly) increasing in p for any t;

(c) If λ_t is non-increasing in t, then for any x, y s.t. $x < y$ and $H_\lambda(y,t) \leq 0$, we have $H_\lambda(x,t) < H_\lambda(y,t)$ (thus $H_\lambda(\cdot,t) = 0$ has at most one root);

(d) If $\lambda_t = \lambda_{t+1}$, then $H_\lambda(p,t) \leq 0 \Rightarrow H_\lambda(p,t+1) < 0$.

Proof. **Part (a):** Because $\int_{p'} V_\lambda(p',t+1)G(dp'|p) - V_\lambda(p,t+1) \geq 0$ according to Lemma 1.B.3, part (a) is directly implied by the definition of H_λ.

Part (b): I prove (b) by backward induction in t. By definition, $H_\lambda(p,T) = (1+\lambda_T)u(p)$ is strictly increasing in p and thus the monotonicity property holds for $H_\lambda(p,T)$. Now, assuming

[44] More specifically, starting with $n = 0$, we can run the following algorithm:

1. Let $\tau_n = \inf\{t : \lambda^n_t < \lambda^n_{t+1}\}$. If $\tau_n = +\infty$, end the algorithm and out-put λ^n; otherwise, go to the next step.
2. Let $s = \max\left\{0, \sup\{t < \tau_n : \lambda^n_t \geq \lambda^n_{\tau_n+1}\}\right\}$ and define

$$\lambda^{n+1}_t = \begin{cases} \lambda^n_{\tau_n+1} & \text{if } t = s+1; \\ \lambda^n_{t-1} & \text{if } t = s+2, ..., \tau_n+1; \\ \lambda^n_t & \text{elsewhere} \end{cases}$$

Then, repeat the procedures with n replaced by $n+1$.

Intuitively, in step 2 of the algorithm we advance the first term in λ^n greater than its predecessor to an earlier position such that the first $\tau_n + 1$ terms will be in descending order. It is then easy to see that this algorithm will end in finite time and the vector it delivers will be non-increasing over t. Moreover, to derive λ^{n+1} from λ^n in step 2, one can just interchange the (τ_n+1)'th term with the τ_n'th term, then interchange the (new) τ_n'th term with the (τ_n-1)'th term,..., and finally interchange the (new) $(s+2)$'th term with the $(s+1)$'th term. In each of these steps, we interchange two adjacent terms with the latter greater than the former. By the claim proved above, this keeps each of the (intermediate) vector within $\arg\min_{\lambda\in\mathbb{R}^T_+} \sup_{\phi\in\Phi} \mathcal{L}(\phi;\lambda)$. Thus the multiplier we derive in the end remains in $\arg\min_{\lambda\in\mathbb{R}^T_+} \sup_{\phi\in\Phi} \mathcal{L}(\phi;\lambda)$. This completes the proof.

it holds for $H_\lambda(p,t+1)$, I show it also holds for $H_\lambda(p,t)$. Notice the following equations hold:

$$\begin{aligned} H_\lambda(p,t) =&(1+\lambda_t)u(p) + \int_{p'} V_\lambda(p',t+1)G(dp'|p) - V_\lambda(p,t+1) \\ =&(1+\lambda_t)u(p) + \int_{p'} \Big(\max\{H_\lambda(p',t+1),0\} + V_\lambda(p',t+2)\Big)G(dp'|p) \\ &- \Big(\max\{H_\lambda(p,t+1),0\} + V_\lambda(p,t+2)\Big) \\ =&(1+\lambda_t)u(p) + \int_{p'} V_\lambda(p',t+2)G(dp'|p) - V_\lambda(p,t+2) \\ &+ \int_{p'} \max\{H_\lambda(p',t+1),0\}G(dp'|p) - \max\{H_\lambda(p,t+1),0\} \\ =&(1+\lambda_t)u(p) - (1+\lambda_{t+1})u(p) + H_\lambda(p,t+1) \\ &+ \int_{p'} \max\{H_\lambda(p',t+1),0\}G(dp'|p) - \max\{H_\lambda(p,t+1),0\} \\ =&(\lambda_t - \lambda_{t+1})u(p) + \min\{H_\lambda(p,t+1),0\} + \int_{p'} \max\{H_\lambda(p',t+1),0\}G(dp'|p) \end{aligned} \tag{1.24}$$

where the second equality is because $V_\lambda(p,t+1) = \max\{H_\lambda(p,t+1),0\} + V_\lambda(p,t+2)$ according to the Bellman equation (1.17); the first and the fourth equalities are directly implied by the definition of H_λ; the other two are just trivial identities. Recall that: $\lambda_t \geq \lambda_{t+1}$ by assumption; $H_\lambda(p,t+1)$ weakly increases in p by the induction hypothesis; and $G(\cdot|p)$ increases in first-order stochastic dominance in p by property (P3) of Lemma 1.3.1. These imply that all of the three terms in the last expression are (weakly) increasing in p. Thus $H_\lambda(p,t)$ is (weakly) increasing in p. This completes the proof for (b).

Part (c): I still prove by induction. The result holds obviously for $H_\lambda(p,T) = (1+\lambda_T)u(p)$. Now, assuming it holds for period $t+1$, I show it also holds for period t. In particular, with any $x < y$ in $[0,1]$, we want to show $H_\lambda(y,t) \leq 0 \Rightarrow H_\lambda(x,t) < H_\lambda(y,t)$. Given result (b) and equation (1.24) derived above, $H_\lambda(x,t) < H_\lambda(y,t)$ obviously holds when $\lambda_t > \lambda_{t+1}$, since $u(p) = \theta_H p + \theta_L(1-p)$ is strictly increasing in p. It thus suffices to assume $\lambda_t = \lambda_{t+1}$. In this case, I have the following observation:

Claim. If $\lambda_t = \lambda_{t+1}$ and $H_\lambda(y,t) \leq 0$, then $H_\lambda(y,t+1) \leq 0$.

Proof for the claim. Given $\lambda_t = \lambda_{t+1}$ and $H_\lambda(y,t) \leq 0$, suppose $H_\lambda(y,t+1) > 0$. Then equation (1.24) implies that $\int_{p'} \max\{H_\lambda(p',t+1),0\}G(dp'|y) = H_\lambda(y,t) \leq 0$. However, by the law of iterated expectation condition on the belief process (property (P2) in Lemma 1.3.1), we must have $G\big([y,1]|y\big) > 0$. Moreover, by the monotonicity of $H_\lambda(\cdot,t+1)$ proved in part (b), we know

$H_\lambda(p', t+1) \geq H_\lambda(y, t+1)$ for any $p' \geq y$. Together with the hypothesis $H_\lambda(y, t+1) > 0$, these then imply $\int_{p'} \max\{H_\lambda(p', t+1), 0\} G(dp'|y) > 0$, which contradicts with the previous conclusion. Thus we must have $H_\lambda(y, t+1) \leq 0$. □

Now, go back to the main proof for part (c). Notice when $\lambda_t = \lambda_{t+1}$ and $H_\lambda(y, t) \leq 0$, the following holds:

$$
\begin{aligned}
H_\lambda(y, t) &= \min\{H_\lambda(y, t+1), 0\} + \int_{p'} \max\{H_\lambda(p', t+1), 0\} G(dp'|y) \\
&= H_\lambda(y, t+1) + \int_{p'} \max\{H_\lambda(p', t+1), 0\} G(dp'|y) \\
&> H_\lambda(x, t+1) + \int_{p'} \max\{H_\lambda(p', t+1), 0\} G(dp'|x) = H_\lambda(x, t)
\end{aligned}
$$

The first equality is just by equation (1.24) with $\lambda_t = \lambda_{t+1}$. The second equality holds because the claim proved above implies $H_\lambda(y, t+1) \leq 0$. The strict inequality holds because: $H_\lambda(y, t+1) \leq 0$ further implies $H_\lambda(x, t+1) < H_\lambda(y, t+1)$ by the induction hypothesis; $G(\cdot|y)$ first order stochastic dominates $G(\cdot|x)$; and $H_\lambda(\cdot, t+1)$ is increasing by part (b). The last equality holds also by equation (1.24) and the fact that $H_\lambda(x, t+1) < 0$. This completes the proof by induction.

Part (d): Suppose $\lambda_t = \lambda_{t+1}$, $H_\lambda(p, t) \leq 0$, but $H_\lambda(p, t+1) \geq 0$. Equation 1.24 would imply

$$
H_\lambda(p, t) = \int_{p'} \max\{H_\lambda(p', t+1), 0\} G(dp'|p)
$$

I now argue that the RHS above must be strictly positive. Notice $H_\lambda(p, t) \leq 0$ obviously imply $p < 1$. Property (P4) in Lemma 1.3.1 then implies that $G((p, 1]|p) > 0$. Moreover, since $H_\lambda(p, t+1) \geq 0$, parts (b) and (c) proved earlier imply $H_\lambda(p', t+1) > 0$ for any $p' > p$. These together imply $\int_{p'} \max\{H_\lambda(p', t+1), 0\} G(dp'|p) > 0$. This contradicts with $H_\lambda(p, t) \leq 0$ given the equation above. Thus the result in part (d) holds.

Q.E.D.

Proof for Lemma 1.3.4

The proof for Lemma 1.3.4 easily follows from Lemma 1.B.4.

Proof. For each t, I construct the threshold η_t as follows:

- Case 1: $\{p : H_\lambda(p,t) = 0\} = \emptyset$.

  In this case, define $\eta_t = \inf\{p : H_\lambda(p,t) > 0\}$.

- Case 2: $\{p : H_\lambda(p,t) = 0\} \neq \emptyset$.

  In this case, Lemma 1.B.4(c) implies that $\{p : H_\lambda(p,t) = 0\}$ contains a single element. Define η_t to be this element.

By the monotonicity property of $H_\lambda(\cdot,t)$ in Lemma 1.B.4(b), in both cases we have $p_t > \eta_t \Rightarrow H_\lambda(p_t,t) > 0$ and $p_t < \eta_t \Rightarrow H_\lambda(p_t,t) < 0$. Together with Lemma 1.B.2(b), this implies that under any solution to $\max_{\phi\in\Phi} \mathcal{L}(\phi;\lambda)$, we have $p_t > \eta_t \Rightarrow a_t = 1$ a.s. and $p_t < \eta_t \Rightarrow a_t = 0$ a.s. Hence any solution is almost surely equivalent to a threshold policy with thresholds being $(\eta_t)_{t=1}^T$.

Moreover, notice Lemma 1.B.4(a) implies $u(p_t) > 0 \Rightarrow H_\lambda(p_t,t) > 0$. Together with Lemma 1.B.2(b), this then implies $u(p_t) > 0 \Rightarrow a_t = 1$ a.s. under any optimal solution to $\max_{\phi\in\Phi} \mathcal{L}(\phi;\lambda)$.

Q.E.D.

1.B.5 Proofs for Section 1.3.4

Definition of ϕ^d

As is mentioned in the main text, I define ϕ^d as the "most conservative" optimal policy for the dictator's problem. Formally, for any $t = 1, ..., T$:

$$\phi_t^d(p) := \begin{cases} 1 & \text{if } H_{\mathbf{0}}(p,t) > 0; \\ 0 & \text{otherwise} \end{cases} \tag{1.25}$$

where $H_{\mathbf{0}}$ is as defined in (1.18) in Appendix 1.B.4 with $\lambda = \mathbf{0}$. By Lemma 1.B.2, it is easy to see that we not only have ϕ^d being optimal for the dictator's problem, but also have $\phi_{\geq t}^d$ to be optimal for the dictator's continuation problem starting from time t regardless of the belief distribution at t. Moreover, ϕ^d (or $\phi_{\geq t}^d$) is conservative in the sense that it breaks any tie in favor of non-recommendation (i.e., $H_{\mathbf{0}}(p,t) = 0 \Rightarrow \phi_t^d(p) = 0$). This makes it most favorable to the current consumer among all dictator's optimal (continuation) policies.

Details in the Construction of ϕ^o and a Uniqueness Property

I first construct a threshold time-t policy ϕ_t^o satisfying the requirements in step 2 of the algorithm in Definition 1.3.2. Given μ_t^o, define the policy's threshold as

$$\eta_t^o = \inf\{x \in [0,1] : \int_{p>x} u(p)\mu_t^o(dp) > 0\}$$

and define the recommendation probability at the threshold as

$$\phi_t^o(\eta_t^o) = -\frac{\int_{p>\eta_t^o} u(p)\mu_t^o(dp)}{u(\eta_t^o)\mu_t^o(\{\eta_t^o\})}$$

When the denominator above is zero, I define $\phi_t^o(\eta_t^o) = 0$ for simplicity. Now, I check the ϕ_t^o such defined indeed satisfies the desired properties.

Claim. ϕ_t^o above is well-defined and satisfies the properties in step 2 of the algorithm.

Proof. First notice properties (P2) and (P5) in Lemma 1.3.1 together imply $\mu_t^o((\bar{p},1]) > 0$ and thus $\int_{p>\bar{p}} u(p)\mu_t^o(dp) > 0$. This implies $\eta_t^o \le \bar{p}$ and thus the first property is satisfied.

Now, I argue that $\int_{p>\eta_t^o} u(p)\mu_t^o(dp) \ge 0$. To see this, notice by the definition of η_t^o, there exists a sequence $\{x_n\} \downarrow \eta_t^o$ such that $\int_{p>x_n} u(p)\mu_t^o(dp) > 0$ for all n. Since $\mathbb{1}_{\{p>x_n\}} \to \mathbb{1}_{\{p>\eta_t^o\}}$ and $u(\cdot)$ is bounded, we must have $\int_{p>\eta_t^o} u(p)\mu_t^o(dp) \ge 0$ by the dominated convergence theorem.

Next, I argue that $\int_{p\ge\eta_t^o} u(p)\mu_t^o(dp) \le 0$. As the algorithm has not been ended in step 1, we must have $\int_{p\in[0,1]} u(p)\mu_t^o(dp) < 0$.[45] Thus the argument is true when $\eta_t^o = 0$. When $\eta_t^o > 0$, notice by the definition of η_t^o, there exists a sequence $\{x_n\} \uparrow \eta_t^o$ such that $\int_{p>x_n} u(p)\mu_t^o(dp) \le 0$ for all n. Since $\mathbb{1}_{\{p>x_n\}} \to \mathbb{1}_{\{p\ge\eta_t^o\}}$, by the dominated convergence theorem we then must have $\int_{p\ge\eta_t^o} u(p)\mu_t^o(dp) \le 0$.

Now, consider two cases:

- Case 1: $\mu_t^o(\{\eta_t^o\}) = 0$. In this case, the above arguments imply $\int_{p>\eta_t^o} u(p)\mu_t^o(dp) = 0$. Therefore the second property is satisfied.
- Case 2: $\mu_t^o(\{\eta_t^o\}) > 0$. In this case, the above arguments imply $\int_{p>\eta_t^o} u(p)\mu_t^o(dp) \ge 0$ and $\int_{p>\eta_t^o} u(p)\mu_t^o(dp) + u(\eta_t^o)\mu_t^o(\{\eta_t^o\}) \le 0$. These imply $0 \ge \frac{\int_{p>\eta_t^o} u(p)\mu_t^o(dp)}{u(\eta_t^o)\mu_t^o(\{\eta_t^o\})} \ge -1$ whenever $u(\eta_t^o)\mu_t^o(\{\eta_t^o\}) \ne 0$. Thus $\phi_t^o(\eta_t^o)$ defined above is a valid probability. By the definition of $\phi_t^o(\eta_t^o)$, we have $\int_p \phi_t^o(p)u(p)\mu_t^o(dp) = \int_{p>\eta_t^o} u(p)\mu_t^o(dp) + \phi_t^o(\eta_t^o)u(\eta_t^o)\mu_t^o(\{\eta_t^o\}) = 0$. Thus

[45] Notice $\phi_t^d(p) = 1$ for any $p > \bar{p}$ by the definition of ϕ^d and Lemma 1.B.4(a).

ϕ_t^o satisfies the second desired property.

Q.E.D.

Due to the possible existence of off-path belief states, there can also be other forms of ϕ_t satisfying the desired properties in step 2 of the algorithm. However, the following lemma implies that any such policy must μ_t^o-a.e. agree with ϕ_t^o.

Lemma 1.B.5. *Given any probability measure μ over $[0,1]$ such that $\mu((\bar{p},1]) > 0$,*[46] *any threshold time-t policies satisfying $p > \bar{p} \Rightarrow \phi_t(p) = 1$ (μ-a.e.) and $\int_p \phi_t(p)u(p)\mu(dp) = 0$ must agree μ-a.e.*

Proof. For $i = 1, 2$, let ϕ_t^i be a threshold time-t policy with threshold η_t^i, which satisfies $\int_p \phi_t^i(p)u(p)\mu(dp) = 0$ and $p > \bar{p} \Rightarrow \phi_t^i(p) = 1$ (μ-a.e.). Without loss of generality, assume $\eta^1 \leq \eta^2$. Notice under the assumption $\mu((\bar{p},1]) > 0$, $\int_p \phi_t^i(p)u(p)\mu(dp) = 0$ implies that we must have $\phi_t^i(p) > 0$ for some p with $u(p) < 0$. Thus the threshold structure implies $\phi_t^i(p) = 1$ for all p such that $u(p) \geq 0$, which holds for both $i = 1, 2$. Now, notice we have:

$$0 = \int_p \phi_t^1(p)u(p)\mu(dp) - \int_p \phi_t^2(p)u(p)\mu(dp) = \int_{p:u(p)<0} \big(\phi_t^1(p) - \phi_t^2(p)\big)u(p)\mu(dp)$$

If $\eta^1 < \eta^2$, then $\phi_t^1(p) \geq \phi_t^2(p)$ for all p because of the threshold structure. Supposing the policies do not agree μ-a.e. , which can only happen when $u(p) < 0$, then $\phi_t^1(p) > \phi_t^2(p)$ for a positive μ-measure set of p with $u(p) < 0$. This implies that the last expression above is strictly negative, which is a contradiction.

If $\eta^1 = \eta^2 =: \eta$, then the two policies can only differ at $p = \eta$ with $u(\eta) < 0$. Supposing they do not agree μ-a.e, we must have $\mu(\{\eta\}) > 0$ and $\phi_t^1(\eta) \neq \phi_t^2(\eta)$. These imply that the last expression above equals to $[\phi_t^1(\eta) - \phi_t^2(\eta)]u(\eta)\mu(\{\eta\}) \neq 0$, which is a contradiction. Thus the policies must agree μ-a.e. *Q.E.D.*

Proof for Proposition 1.3.2 and Related Results

We need to first prove some properties of ϕ^d.

Lemma 1.B.6. *Set $\{p : \phi_t^d(p) = 1\}$ shrinks in set inclusion order as t increases.*

[46] Notice properties (P2) and (P5) in Lemma 1.3.1 together imply that this holds for time-t belief distribution μ_t under any policy.

Proof. Since by construction $\phi_t^d(p) = 1$ if and only if $H_{\mathbf{0}}(p,t) > 0$, the result is directly implied by Lemma 1.B.4(d) in Section 1.B.4. *Q.E.D.*

This observation leads to the following result:

Lemma 1.B.7. *Given any time-t belief μ_t, we have $\int_p \phi_t^d(p)u(p)\mu_t(dp) \le \int_p \phi_{t+1}^d(p)u(p)\mu_{t+1}(dp)$, where μ_{t+1} is the period $t+1$ belief distribution under ϕ_t^d given μ_t. (That is, the consumer's expected payoff is weakly higher in period $t+1$ than in period t under ϕ^d.)*

Proof. The result is proved by the following arguments:

$$
\begin{aligned}
&\int_{p'} \phi_t^d(p')u(p')\mu_t(dp') \\
=&\int_p \int_{p'} u(p')\phi_{t+1}^d(p')[\phi_t^d(p)G(dp'|p) + (1-\phi_t^d(p))D(dp'|p)]\mu_t(dp) \\
=&\int_p \int_{p'} u(p')\phi_{t+1}^d(p')G(dp'|p)\phi_t^d(p)\mu_t(dp) \ge \int_p u(p)\phi_t^d(p)\mu_t(dp)
\end{aligned}
$$

The first equality holds by the transition rule for p. The second equality holds because $\phi_t^d(p) = 0 \Rightarrow \phi_{t+1}^d(p) = 0$ by Lemma 1.B.6, and thus $\int_{p'} u(p')\phi_{t+1}^d(p')(1-\phi_t^d(p))D(dp'|p) = u(p)\phi_{t+1}^d(p)(1-\phi_t^d(p)) = 0$. The last inequality holds because $\int_{p'} u(p')\phi_{t+1}^d(p')G(dp'|p) \ge \int_{p'} u(p')G(dp'|p) = u(p)$, where the "$\ge$" is due to $u(p') > 0 \Rightarrow \phi_{t+1}^d(p') = 1$ and the "$=$" is implied by property (P2) in Lemma 1.3.1.

Q.E.D.

To ease notation, let IC_t denote the IC constraint for time-t consumer. Repeated use of Lemma 1.B.7 implies that given any μ_t, if $\phi_{\ge t}^d$ satisfies IC_t, then it satisfies all later IC's. We are now ready to prove the proposition.

Proof for Proposition 1.3.2. First consider the "only if" part. Let ϕ^{opt} be any optimal policy for the designer and let μ_t^{opt} be the distribution of p_t under it for any t. Due to Corollary 1.3.1, I can assume ϕ^{opt} is a threshold policy without loss of generality. Let λ^* be a Lagrangian multiplier solving the dual problem that is non-increasing over t (which exists by Lemma 1.3.3). I first check condition (i):

***Claim** (a).* $\phi_{<\widehat{t}}^{opt}$ agrees with $\phi_{<\widehat{t}}^{o}$ a.s.

Proof for Claim (a). I check by forward induction in t for $t < \widehat{t}$. For $t = 1$, $t < \widehat{t}$ implies that ϕ_1^d violates IC_1 given the initial state distribution μ_1. In this case, we must have $\lambda_1^* > 0$ and thus

IC_1 is binding under ϕ^{opt}. (Suppose not. Then $\lambda_t^* = 0$ for all $t \geq 1$ since it is non-increasing in t, and thus the Lagrangian problem $\max \mathcal{L}(\phi; \lambda^*)$ would coincide with the dictator's problem. By Lemma 1.3.2, this implies that ϕ^{opt} must also solve the dictator's problem. However, by construction ϕ^d provides the highest time-1 expected consumer surplus given μ_1 among all optimal policies for the dictator's problem. Thus when ϕ_1^d violates IC_1, so does ϕ_1^{opt}, which is a contradiction.) Also notice the optimality of ϕ^{opt} implies $p > \bar{p} \Rightarrow \phi_1^{opt}(p) = 1$ (μ_1-a.e.). Thus ϕ_1^{opt} satisfies the properties of ϕ_1^o in step 2 of the algorithm. By Lemma 1.B.5, we then must have ϕ_1^{opt} and ϕ_1^o agree μ_1-a.e.

Now, assume ϕ^{opt} and ϕ^o agree a.s. for all periods before t and $t < \widehat{t}$. Then, they induce the same distribution for p_t, which is just μ_t^o constructed in the algorithm. I want to show ϕ^{opt} must also satisfy $p > \bar{p} \Rightarrow \phi_t^{opt}(p) = 1$ μ_t^o-a.s. and $\int_p \phi_t^{opt}(p)u(p)\mu_t^o(dp) = 0$. The former is directly implied by the optimality of ϕ^{opt}. For the latter, notice $t < \widehat{t}$ implies that ϕ^d violates IC_t given $p_t \sim \mu_t^o$. We then must have $\lambda_t^* > 0$ and thus IC_t is binding under ϕ^{opt}. (Suppose not. Then $\lambda_{t'}^* = 0$ for all $t' \geq t$, and thus the continuation Lagrangian problem starting with time t coincides with the corresponding dictator's continuation problem. The optimality of ϕ^{opt} then implies that $\phi_{\geq t}^{opt}$ must be optimal for this dictator's continuation problem given $p_t \sim \mu_t^o$. However, among all such policies, $\phi_{\geq t}^d$ delivers the highest time-t expected consumer surplus. Thus when ϕ_t^d violates IC_t, so does ϕ_t^{opt}, which is a contradiction.) Thus $\int_p \phi_t^{opt}(p)u(p)\mu_t^o(dp) = 0$. Again by Lemma 1.B.5, we must have ϕ_t^{opt} and ϕ_t^o agree μ_t^o-a.e. The proof is then completed by induction. □

Now, since $\phi_{<\widehat{t}}^{opt}$ and $\phi_{<\widehat{t}}^o$ agree almost surely, they lead to the same distribution for $p_{\widehat{t}}$, which is just $\mu_{\widehat{t}}^o$. Hence $\phi_{\geq \widehat{t}}^{opt}$ must satisfy all IC constraints after time $\widehat{t}$ given $p_{\widehat{t}} \sim \mu_{\widehat{t}}^o$. The following claim checks the rest of condition (ii) in the proposition.

Claim *(b).* $\phi_{\geq \widehat{t}}^{opt}$ is optimal for the dictator's continuation problem since time $\widehat{t}$ given $p_{\widehat{t}} \sim \mu_{\widehat{t}}^o$.

Proof for Claim (b). By the definition of $\widehat{t}$, $\phi_{\widehat{t}}^d$ satisfies $IC_{\widehat{t}}$ given $p_{\widehat{t}} \sim \mu_{\widehat{t}}^o$ and thus $\phi_{\geq \widehat{t}}^d$ satisfies all later IC constraints by Lemma 1.B.7 given $p_{\widehat{t}} \sim \mu_{\widehat{t}}^o$. This implies that if we deviate from $\phi_{\geq \widehat{t}}^{opt}$ to $\phi_{\geq \widehat{t}}^d$ since period $\widehat{t}$, no IC constraint will be violated. Also notice such deviation can only improve the total surplus since $\phi_{\geq \widehat{t}}^d$ is optimal for the dictator's continuation problem, which is more relaxed than the original continuation problem. For such deviation to be unprofitable, we then need $\phi_{\geq \widehat{t}}^{opt}$ to achieve the same value as $\phi_{\geq \widehat{t}}^d$ for that continuation problem and thus $\phi_{\geq \widehat{t}}^{opt}$ is also optimal for it. □

Now, I turn to the "if" part. Given the existence of optimal policy (guaranteed by Proposition 1.3.1), it suffices to see that all policies satisfying the two conditions (i) and (ii) are feasible and yield the same total payoff for the designer. For the total payoff, this is obviously true since: for periods before $\widehat{t}$, all such policies agree a.s. and lead to $p_{\widehat{t}} \sim \mu^o_{\widehat{t}}$; for periods $t \geq \widehat{t}$, all such policies achieve the same total payoff as that under $\phi^d_{\geq \widehat{t}}$ given $p_{\widehat{t}} \sim \mu^o_{\widehat{t}}$. For feasibility, it suffices to check $\phi^o_{<\widehat{t}}$ is feasible. This is true since $\phi^o_{<\widehat{t}}$ satisfies all IC constraints for $t < \widehat{t}$ as equalities by construction. This completes the proof for the "if" part.

Finally, notice by construction ϕ^o does satisfy conditions (i) and (ii). In particular, Lemma 1.B.7 implies that $\phi^d_{\geq \widehat{t}}$ satisfies all IC constraints after time $\widehat{t}$ given $p_{\widehat{t}} \sim \mu^o_{\widehat{t}}$, which implies the feasibility of ϕ^o following time $\widehat{t}$. Thus ϕ^o is optimal.

Q.E.D.

1.B.6 Proof for Proposition 1.4.1

Proof. Let ϕ^* denote an optimal policy, let $(\eta^*_t)_{t=1}^T$ denote its thresholds, and let μ^*_t denote the distribution of p_t under it. Notice that under the full support condition in Assumption 1.4.1, Corollary 1.3.1 implies that we must have $\eta^*_t \leq \bar{p}$ for all t. (Recall that $\bar{p}$ is the myopic threshold, i.e., $u(\bar{p}) = 0$.) Moreover, under the atomless condition in Assumption 1.4.1, randomization at the thresholds does not matter.

For part (a), I first show the following observation:

Claim. For any $\eta \in (0, \bar{p}]$, we have $\int_{p \geq \eta} \left[\int_{p' \geq \eta} u(p') G(dp'|p) \right] \mu^*_t(dp) > \int_{p \geq \eta} u(p) \mu^*_t(dp)$.

Proof for the claim. Given any $\eta \in (0, \bar{p}]$, recall that property (P4) in Lemma 1.3.1 implies $G([0, \eta)|\eta) > 0$. By the weak continuity of $G(\cdot|p)$ on p (i.e., property (P1) in Lemma 1.3.1), this further implies that there exists $\delta > 0$ s.t. $G([0, \eta)|p) > 0 \, \forall p \in [\eta, \eta + \delta]$.[47] Together with the full support assumption on μ^*_t, we then have $\int_{p \geq \eta} G\big([0, \eta)|p\big) \mu^*_t(dp) > 0$. Since $\eta \leq \bar{p}$, $u(p) < 0$ for all $p < \eta$. Thus $\int_{p \geq \eta} \int_{p' < \eta} u(p') G(dp'|p) \mu^*_t(dp) < 0$. This then implies $\int_{p \geq \eta} \int_{p' \geq \eta} u(p') G(dp'|p) \mu^*_t(dp) > \int_{p \geq \eta} \int_{p'} u(p') G(dp'|p) \mu^*_t(dp) = \int_{p \geq \eta} u(p) \mu^*_t(dp)$, where the equality holds by property (P2) in Lemma 1.3.1. □

[47] See Theorem 3.2.11 in Durrett (2019) (equivalence between conditions (i) and (ii)).

Now, I argue that the following holds given any $t \leq \widehat{t} - 2$:

$$
\begin{aligned}
&\int_{p \geq \eta_t^*} u(p)\mu_{t+1}^*(dp) \\
&= \int_p \int_{p' \geq \eta_t^*} u(p')\left[\mathbb{1}_{\{p \geq \eta_t^*\}} G(dp'|p) + \mathbb{1}_{\{p < \eta_t^*\}} D(dp'|p)\right] \mu_t^*(dp) \\
&= \int_{p \geq \eta_t^*} \left[\int_{p' \geq \eta_t^*} u(p')G(dp'|p)\right] \mu_t^*(dp) + \int_{p \geq \eta_t^*} \mathbb{1}_{\{p < \eta_t^*\}} u(p)\mu_t^*(dp) \\
&= \int_{p \geq \eta_t^*} \left[\int_{p' \geq \eta_t^*} u(p')G(dp'|p)\right] \mu_t^*(dp) \\
&> \int_{p \geq \eta_t^*} u(p)\mu_t^*(dp) = 0
\end{aligned}
$$

The first equality holds by the transition rule of p_t; the second equality is trivial identity; the third equality holds because the second term in line 3 is obviously zero; the last expression equals to zero because the IC constraint is binding for any $t < \widehat{t}$ by Proposition 1.3.2. To see the inequality holds, notice that $t \leq \widehat{t} - 2$ necessarily implies $\eta_t^* > 0$, since otherwise ϕ_t^d would be feasible at time t and the algorithm in Definition 1.3.2 would have stopped in step 1 at time t. The desired inequality is then directly implied by the claim proved above.

For part (b), notice under the atomless assumption in Assumption 1.4.1, Lemma 1.B.4(c) (see Appendix 1.B.4) implies that $H_{\mathbf{0}}(p_t, t)$ is non-zero almost surely under any policy. Lemma 1.B.2(b) then implies that any optimal policy for the dictator must almost surely agree with ϕ^d. By Proposition 1.3.2, this further implies that $\phi_{\geq \widehat{t}}^*$ must almost surely agree with $\phi_{\geq \widehat{t}}^d$ (given $p_t \sim \mu_{\widehat{t}}^*$). Under the full support assumption in Assumption 1.4.1, this then requires that $\phi_{\geq \widehat{t}}^*$ and $\phi_{\geq \widehat{t}}^d$ share the same thresholds. It thus suffices to prove the desired property for ϕ^d.

Let $(\eta_t^d)_{t=1}^T$ denote the sequence of thresholds of ϕ^d. Recall that $\phi_t^d(p) = \mathbb{1}_{\{H_{\mathbf{0}}(p,t) > 0\}}$ (as is defined in Appendix 1.B.5). By the continuity of $H_{\mathbf{0}}(\cdot, t)$ (Lemma 1.B.1), we then have $H_{\mathbf{0}}(\eta_t^d, t) = 0$ for any t.[48] By Lemma 1.B.4(d), this further implies $H_{\mathbf{0}}(\eta_t^d, t+1) < 0$. Since $H_{\mathbf{0}}(\cdot, t+1)$ is increasing (Lemma 1.B.4(b)), we thus must have $\eta_{t+1}^d > \eta_t^d$ for any t. This completes the proof for part (b).

Finally, $\eta_t^* \leq \bar{p}$ for all t is directly implied by Corollary 1.3.1 under the full support assumption in Assumption 1.4.1.

Q.E.D.

[48] Notice it is easy to see that $H_{\mathbf{0}}(0, t) < 0$ and $H_{\mathbf{0}}(1, t) > 0$.

1.B.7 Proof for Proposition 1.5.1

Proof. The designer's problem is formally written as:

$$\max_{\phi\in\Phi}\Big\{\sum_{t=1}^{T}\mathbb{E}_\phi[a_t u(p_t)]\Big\}$$

$$\text{s.t. } \mathbb{E}_\phi[a_t u(p_t)] \geq 0 \quad \forall t=1,...,T$$

$$p_{t+1}|p_t,a_t \sim a_t\big[\alpha G^I(\cdot|p_t)+(1-\alpha)D(\cdot|p_t)\big]+(1-a_t)D(\cdot|p_t)$$

$$p_1 \sim \mu_1$$

Pick α_a and α_b with $\alpha_a < \alpha_b$. Corresponding to these two information generation rates respectively, let G^a and G^b be the transition kernels of p_t following one's consumption, as is defined in equation (1.9); let $V^a_{\mathbf{0}}$ and $V^b_{\mathbf{0}}$ be the value functions for the dictator's problem (i.e., with $\lambda=\mathbf{0}$), as is defined in Section 1.B.4; let $H^a_{\mathbf{0}}$ and $H^b_{\mathbf{0}}$ be the associated H-functions (with $\lambda=\mathbf{0}$) as in equation (1.18); let $\widehat{t}^a$ and $\widehat{t}^b$ denote the critical time points defined in Definition 1.3.2; let ϕ^a and ϕ^b be the optimal threshold policies and denote their sequences of thresholds as $(\eta^a_t)^T_{t=1}$ and $(\eta^b_t)^T_{t=1}$. I show $\eta^a_t \geq \eta^b_t$ for all t with a sequence of claims below.

First, since higher α is beneficial, we have the following non-surprising result for $V^a_{\mathbf{0}}$ and $V^b_{\mathbf{0}}$:

Claim *(a).* $V^b_{\mathbf{0}}(p,t) \geq V^a_{\mathbf{0}}(p,t)$ for any pair of (p,t).

Proof for Claim (a). I show by backward induction on t. For $t=T+1$, the result holds trivially. Assuming it holds for all $t'>t$, I now consider time t.

By the Bellman equation, we have:

$$\begin{aligned}V^b_{\mathbf{0}}(p,t)-V^a_{\mathbf{0}}(p,t) =& \max\Big\{u(p)+\int_{p'}V^b_{\mathbf{0}}(p',t+1)G^b(dp'|p),\,V^b_{\mathbf{0}}(p,t+1)\Big\}\\ &-\max\Big\{u(p)+\int_{p'}V^a_{\mathbf{0}}(p',t+1)G^a(dp'|p),\,V^a_{\mathbf{0}}(p,t+1)\Big\}\end{aligned}$$

By the induction hypothesis, we know $V^b_{\mathbf{0}}(p,t+1) \geq V^a_{\mathbf{0}}(p,t+1)$. It thus suffices to check

$\int_{p'} V_{\mathbf{0}}^b(p', t+1) G^b(dp'|p) \geq \int_{p'} V_{\mathbf{0}}^a(p', t+1) G^a(dp'|p)$. Notice the following relations hold:

$$\begin{aligned}
&\int_{p'} V_{\mathbf{0}}^b(p', t+1) G^b(dp'|p) - \int_{p'} V_{\mathbf{0}}^a(p', t+1) G^a(dp'|p) \\
=&\alpha_b \int_{p'} V_{\mathbf{0}}^b(p', t+1) G^I(dp'|p) + (1-\alpha_b) V_{\mathbf{0}}^b(p, t+1) \\
&- \alpha_a \int_{p'} V_{\mathbf{0}}^a(p', t+1) G^I(dp'|p) - (1-\alpha_a) V_{\mathbf{0}}^a(p, t+1) \\
\geq&\alpha_a \int_{p'} V_{\mathbf{0}}^b(p', t+1) G^I(dp'|p) + (1-\alpha_a) V_{\mathbf{0}}^b(p, t+1) \\
&- \alpha_a \int_{p'} V_{\mathbf{0}}^a(p', t+1) G^I(dp'|p) - (1-\alpha_a) V_{\mathbf{0}}^a(p, t+1)
\end{aligned}$$

where the equality is by the definition of G^a and G^b. To see the inequality holds, recall that Lemma 1.B.3(b) implies $\int_{p'} V_{\mathbf{0}}^b(p', t+1) G^b(dp'|p) \geq V_{\mathbf{0}}^b(p, t+1)$, which further implies $\int_{p'} V_{\mathbf{0}}^b(p', t+1) G^I(dp'|p) \geq V_{\mathbf{0}}^b(p, t+1)$ since $G^b(\cdot|p)$ is a weighted average of $G^I(\cdot|p)$ and $D(\cdot|p)$. The above inequality thus holds given $\alpha_b > \alpha_a$. Now, notice the induction hypothesis implies that the last expression above is indeed non-negative. We thus have the desired result. □

An important implication of the above claim is:

***Claim** (b).* $H_{\mathbf{0}}^b(p, t) \leq 0 \Rightarrow H_{\mathbf{0}}^a(p, t) \leq 0$ for all pairs of (p, t).
Proof for Claim (b). When $H_{\mathbf{0}}^b(p, t) \leq 0$, by Lemma 1.B.4(d) we know $H_{\mathbf{0}}^b(p, t') \leq 0$ for all $t' > t$. With $\alpha = \alpha_b$, it is thus optimal for the dictator to stop recommendation from time t on given $p_t = p$, which leads to the optimal continuation value being zero. Hence $V_{\mathbf{0}}^b(p, t) = 0$. By Claim (a), this implies $V_{\mathbf{0}}^a(p, t) \leq 0$ and it is thus also optimal for the dictator to stop recommendation at (p, t) given $\alpha = \alpha_a$. This then implies $H_{\mathbf{0}}^a(p, t) \leq 0$. □

With Claim (b), we can now prove the desired result for $t \geq \hat{t}^b$.

***Claim** (c).* $\eta_t^a \geq \eta_t^b$ for all $t \geq \hat{t}^b$.
Proof for Claim (c). By Proposition 1.3.2, we know that $\phi_{\geq \hat{t}^b}^b$ is optimal for the dictator's continuation problem with $\alpha = \alpha_b$ given the distribution of $p_{\hat{t}^b}$ under ϕ^b. Under Assumption 1.4.1, this together with Lemma 1.B.2(b) implies that $H_{\mathbf{0}}^b(p, t) \leq 0$ for Lebesgue-a.e. $p < \eta_t^b$ for all $t \geq \hat{t}^b$.

Now, suppose $t \geq \hat{t}^b$ but $\eta_t^a < \eta_t^b$. Then the above conclusion together with Claim (b) implies that $H_{\mathbf{0}}^a(p, t) \leq 0$ for Lebesgue-a.e. $p \in (\eta_t^a, \eta_t^b)$. By Lemma 1.B.4, we know $H_{\mathbf{0}}^a(\cdot, t)$ is weakly increasing and has at most one zero point. Thus we must have $H^a_{\mathbf{0}}(p,t) < 0$ for all $p \in (\eta^a_t, \eta^b_t)$.

The above result further implies that given $\alpha = \alpha_a$, the most conservative optimal policy for the dictator (i.e., ϕ^d as is defined in Section 1.3.4) has a threshold larger than η^a_t at time t and it is hence incentive compatible at time t provided that we have been following ϕ^a before. By the algorithm in Definition 1.3.2 and Proposition 1.3.2, we then must have $t \geq \widehat{t}^a$. This in turn implies the optimality of $\phi^a_{\geq t}$ for the dictator's continuation problem given $\alpha = \alpha_a$ and hence under Assumption 1.4.1 we must have $H^a_{\mathbf{0}}(p,t) < 0 \Rightarrow p < \eta^a_t$ (Lebesgue-a.e.) by Lemma 1.B.2(b). This contradicts with the result $H^a_{\mathbf{0}}(p,t) < 0$ for all $p \in (\eta^a_t, \eta^b_t)$ above. Thus we must have $\eta^a_t \geq \eta^b_t$ when $t \geq \widehat{t}^b$. □

Now it suffices to show $\eta^a_t \geq \eta^b_t$ for $t < \widehat{t}^b$. For $k = a, b$, let F^k_t denote the cdf of p_t under ϕ^k for any t given $\alpha = \alpha_k$. Under Assumption 1.4.1, notice that F^a_t and F^b_t are continuous and have full support over $[0,1]$ for any t. The following observation is the key part of the proof.

Claim *(d).* Let τ be any fixed time. If $\eta^a_t \geq \eta^b_t$ for all $t < \tau$, then F^b_τ is a mean-preserving spread of F^a_τ and $F^a_\tau(p) \leq F^b_\tau(p)$ for any $p \leq \min_{t<\tau}\{\eta^b_t\}$.

Proof for Claim (d). To prove the claim, I construct two belief processes (truncated at time τ) $(p^a_t)^\tau_{t=1}$ and $(p^b_t)^\tau_{t=1}$ on the same probability space, where $(p^a_t)^\tau_{t=1}$ follows the transition rule decided by ϕ^a given α_a and $(p^b_t)^\tau_{t=1}$ follows the transition rule decided by ϕ^b given α_b.

Specifically, fix a probability space on which a Markov process $(x_n)^\infty_{n=0}$ and a sequence of i.i.d. random variables $(\xi_t)^\tau_{t=1}$ with $\xi_t \sim \text{Uniform}[0,1]$ are defined. $(x_n)^\infty_{n=0}$ is independent from $(\xi_t)^\tau_{t=1}$ and satisfies:

$$x_0 \sim \mu_1$$
$$x_{n+1}|x_n \sim G^I(\cdot|x_n),\ \forall n$$

Intuitively, one can interpret x_n as the value that the platform's belief will take after receiving the n'th informative signal from consumers; ξ_t will serve as a randomization device deciding whether an informative signal will be generated after consumption at time t. I define a filtration of σ-fields $(\mathcal{F}_n)^\infty_{n=0}$ such that $\mathcal{F}_n = \sigma((\xi_t)^\tau_{t=1}, x_0, ..., x_n)$ for all n. Then obviously $(x_n)^\infty_{n=0}$ is a martingale w.r.t. $(\mathcal{F}_n)^\infty_{n=0}$.

For $k = a, b$, I now define process $(p^k_t)^\tau_{t=1}$ together with an auxiliary process $(n^k_t)^\tau_{t=1}$ by the following rule:

$$n_1^k = 0; \quad p_1^k = x_0$$

$$n_{t+1}^k = n_t^k + \mathbb{1}_{\{p_t^k > \eta_t^k\}} \mathbb{1}_{\{\xi_t < \alpha_k\}}; \quad p_{t+1}^k = x_{n_{t+1}^k} \quad \forall t$$

Intuitively, under the scenario with (α_k, ϕ^k), n_t^k tracks how many informative signals have been recorded at the beginning of time t. It is added by 1 after each period if and only if consumption has been made in that period (i.e., $p_t^k > \eta_t^k$) and an informative signal is generated (which is assumed to happen when $\xi_t < \alpha^k$).[49] Given that n_t^k informative signals have been received, p_t^k just equals to $x_{n_t^k}$, which reflects the posterior belief given those signals. It is easy to check that $(p_t^k)_{t=1}^{\tau}$ indeed satisfy the initial distribution and the transition rule of the belief process under policy ϕ^k given response rate α_k and thus $p_t^k \sim F_t^k$ for all $t \leq \tau$.

Now, I notice that n_τ^k is a bounded stopping time w.r.t. $(\mathcal{F}_n)_{n=0}^{\infty}$. The boundedness is obvious since $n_\tau^k < \tau$. To show it is a stopping time, notice by the construction of $(n_t^k)_{t=1}^{\tau}$, whether $\{n_\tau^k \leq n\}$ happens is solely determined by $(x_0, ..., x_n)$ together with $(\xi_t)_{t=1}^{\tau}$. (With $(x_0, ..., x_n)$ and $(\xi_t)_{t=1}^{\tau}$, we can perfectly predict when the $(n+1)$'th informative signal will come.) Thus $\{n_\tau^k \leq n\} \in \mathcal{F}_n$ and n_τ^k is a stopping time w.r.t. $(\mathcal{F}_n)_{n=0}^{\infty}$ by definition. Moreover, since $\eta_t^a \geq \eta_t^b$ for all $t < \tau$ and $\alpha_b > \alpha_a$, it is easy to see that $n_\tau^b \geq n_\tau^a$ for sure.[50] By the Doob's optional sampling theorem, we then have $\mathbb{E}[x_{n_\tau^b} | \mathcal{F}_{n_\tau^a}] = x_{n_\tau^a}$ and thus $\mathbb{E}[x_{n_\tau^b} | x_{n_\tau^a}] = x_{n_\tau^a}$.[51] This further implies $\mathbb{E}[p_\tau^b | p_\tau^a] = p_\tau^a$ and therefore F_τ^b is a mean-preserving spread of F_τ^a.

Now, pick any $p \leq \min_{t<\tau}\{\eta_t^b\}$. Notice that because $p \leq \eta_t^b \leq \eta_t^a$ for all $t < \tau$, by construction the processes of $(p_t^a)_{t=1}^{\tau}$ and $(p_t^b)_{t=1}^{\tau}$ will stop once they fall into $[0, p]$. Thus when $p_\tau^b > p$, we must have $x_n > p$ for all $n \leq n_\tau^b$. Because $n_\tau^a \leq n_\tau^b$, we then must have $p_\tau^a = x_{n_\tau^a} > p$. This implies $p_\tau^b > p \Rightarrow p_\tau^a > p$ for sure and therefore $F_\tau^a(p) \leq F_\tau^b(p)$.

□

We are now ready to complete the last piece of the proof:

***Claim** (e).* $\eta_t^a \geq \eta_t^b$ for all $t < \widehat{t}^b$.

Proof for Claim (e). It suffices to assume $\widehat{t}^b > 1$, otherwise the result is vacuous. I prove the

[49]Under Assumption 1.4.1, what happens when $p_t^k = \eta_t^k$ does not matter since it has zero probability to occur.

[50]By construction, whenever $n_t^b = n_t^a$, we have $p_t^b = p_t^a$ and hence $n_{t+1}^a = n_t^a + 1 \Rightarrow n_{t+1}^b = n_t^b + 1$. Thus the sequence of $(n_t^a)_t$ can never surpass $(n_t^b)_t$.

[51]See, for example, theorem 10.11 in Klenke (2020) (third edition).

claim by induction in $t = 1, ..., \widehat{t}^b$. For $t = 1$, $t < \widehat{t}^b$ implies that $\int_{p \geq \eta_1^b} u(p)\mu_1(dp) = 0$ (by Proposition 1.3.2). Feasibility of ϕ^a also implies $\int_{p \geq \eta_1^a} u(p)\mu_1(dp) \geq 0$. Since $\eta_1^a, \eta_1^b \leq \bar{p}$ by the optimality of ϕ^a and ϕ^b, these obviously imply $\eta_1^a \leq \eta_1^b$ under Assumption 1.4.1.

Now, assuming the result holds for all periods $t < \tau$, I show it for period τ. First, notice $\tau < \widehat{t}^b$ implies that η_t^b has been decreasing over time up to time τ by Proposition 1.4.1. Thus if $\eta_\tau^a > \eta_t^b$ for some $t < \tau$, we must have $\eta_\tau^a > \eta_\tau^b$ to hold. Therefore, it suffices to consider the case where $\eta_\tau^a \leq \min_{t<\tau} \eta_t^b$. In this case, I argue that the following holds:

$$\begin{aligned}\int_{\eta_\tau^a}^1 u(p)dF_\tau^b(p) &= u(1) - u(\eta_\tau^a)F_\tau^b(\eta_\tau^a) - (u_H - u_L)\int_{\eta_\tau^a}^1 F_\tau^b(p)dp \\ &\geq u(1) - u(\eta_\tau^a)F_\tau^a(\eta_\tau^a) - (u_H - u_L)\int_{\eta_\tau^a}^1 F_\tau^a(p)dp \\ &= \int_{\eta_\tau^a}^1 u(p)dF_\tau^a(p) \geq 0\end{aligned}$$

where the two equalities are just by integration by parts,[52] and the last inequality is due to the incentive compatibility of ϕ^a. To check the first inequality, notice under the induction hypothesis and the assumption that $\eta_\tau^a \leq \min_{t<\tau} \eta_t^b$, Claim (d) implies $F_\tau^b(\eta_\tau^a) \geq F_\tau^a(\eta_\tau^a)$ and F_τ^b is a mean-preserving spread of F_τ^a, the latter of which further implies $\int_{\eta_\tau^a}^1 F_\tau^b(p)dp \leq \int_{\eta_\tau^a}^1 F_\tau^a(p)dp$.[53] Together with the fact that $u(\eta_\tau^a) \leq 0$,[54] these imply the desired inequality.

The above discussion has shown $\int_{\eta_\tau^a}^1 u(p)dF_\tau^b(p) \geq 0$. Since $\tau < \widehat{t}^b$, by Proposition 1.3.2 we also know $\int_{\eta_\tau^b}^1 u(p)dF_\tau^b(p) = 0$. Thus we must have $\eta_\tau^b \leq \eta_\tau^a$ under Assumption 1.4.1. This completes the proof by induction.

□

Q.E.D.

[52] Notice $u(p) = u_L + (u_H - u_L)p$ in the current application, and F_τ^a and F_τ^b are continuous under Assumption 1.4.1.

[53] See, e.g., Theorem 3.A.1(a) in Shaked & Shanthikumar (2007). In their notation, we have $F_\tau^b \geq_{cx} F_\tau^a$. The theorem implies $\int_{\eta_\tau^a}^\infty [1 - F_\tau^b(p)]dp \geq \int_{\eta_\tau^a}^\infty [1 - F_\tau^a(p)]dp$, which implies my result since $F_\tau^a(p) = F_\tau^b(p) = 1$ for $p > 1$.

[54] Otherwise, there would be a non-empty interval of p in which $p > \bar{p}$ but $p < \eta_\tau^a$. This would violate the optimality of ϕ^a under Assumption 1.4.1.

1.B.8 Proof for Proposition 1.5.2

Proof. Based on my discussion in the main text, the designer's problem is equivalent to:

$$\max_{\phi\in\Phi}\Big\{\sum_{t=1}^{T}\mathbb{E}_\phi[a_t u(p_t)]\Big\}$$

$$\text{s.t. } \mathbb{E}_\phi[a_t u(p_t)] \geq 0 \quad \forall t = 1, ..., T$$

$$p_{t+1}|p_t, a_t \sim a_t\big[\rho G(\cdot|p_t) + (1-\rho)D(\cdot|p_t)\big] + (1-a_t)D(\cdot|p_t)$$

$$p_1 \sim \mu_1$$

With ρ replaced by α and G replaced by G^I, this is equivalent to the designer's problem studied in Appendix 1.B.7. Thus the effect of an increment in ρ here is equivalent to the effect of an increment in α there. The result is hence directly implied by Proposition 1.5.1. *Q.E.D.*

1.B.9 Proof for Proposition 1.A.1

In the following proof, I assume $\{Q_z\}_{z\in Z}$ and the conditional distributions of s_i $(i \geq 1)$ conditional on $\tilde{\theta}$ are all continuous distributions, so the dominating measure for their densities is chosen as the Lebesgue measure. In the general case, the proof remains the same with Lebesgue measure replaced by proper dominating measures on $\mathbb{R}$ (e.g., counting measure for discrete distributions).

Proof. For any $z \in Z$ and $s \in S$, I define $\psi_{(z,s)}$ as a probability density over $\mathbb{R}$ such that

$$\psi_{(z,s)}(\theta) = \frac{q_z(\theta)\ell(s|\theta)}{\int q_z(\theta)\ell(s|\theta)d\theta}$$

That is, $\psi_{(z,s)}$ is the density function of the posterior about $\tilde{\theta}$ computed from Bayes rule given prior Q_z and post-consumption signal realization s.

Claim *(a).* For any $x, y \in Z$ and $s_a, s_b \in S$, we have $Q_y \geq_{LR} Q_x$ and $s_b \geq s_a$ together imply $\psi_{(y,s_b)} \geq_{LR} \psi_{(x,s_a)}$.

Proof for Claim (a). Assume $Q_y \geq_{LR} Q_x$ and $s_b \geq s_a$. By the definition of ψ, we have:

$$\frac{\psi_{(y,s_b)}(\theta)}{\psi_{(x,s_a)}(\theta)} = \frac{q_y(\theta)}{q_x(\theta)} \cdot \frac{\ell(s_b|\theta)}{\ell(s_a|\theta)}$$

Since $Q_y \geq_{LR} Q_x$ and $\ell(\cdot|\theta)$ increases in likelihood-ratio order in θ, both fractions on the right-hand-side are increasing in θ. Thus $\psi_{(y,s_b)} \geq_{LR} \psi_{(x,s_a)}$. □

Now, I show the following observation:

***Claim** (b).* Assume λ_t is non-increasing over t. Then, for any $x, y \in Z$, we have $Q_y \geq_{LR} Q_x \Rightarrow H_\lambda(y,t) \geq H_\lambda(x,t)$ for all t.

Proof for Claim (b). I show by backward induction in t. The result holds with $t = T$ since $H_\lambda(z,T) = (1+\lambda_T)\int \theta dQ_z(\theta)$. Now, assuming the result holds for all periods since time $t+1$, I show it for period t. Recall that equation (1.24) derived in Section 1.B.4 implies

$$\begin{aligned}
H_\lambda(z,t) =&(\lambda_t - \lambda_{t+1})u(z) + \min\{H_\lambda(z,t+1),0\} \\
&+ \int_{z'\in Z} \max\{H_\lambda(z',t+1),0\}[\rho G(dz';z) + (1-\rho)D(dz';z)] \\
=&(\lambda_t - \lambda_{t+1})u(z) + \rho \min\{H_\lambda(z,t+1),0\} + (1-\rho)H_\lambda(z,t+1) \\
&+ \rho \int_{z'\in Z} \max\{H_\lambda(z',t+1),0\} G(dz';z)
\end{aligned}$$

(Since we have random consumer arrivals with arrival rate ρ, the transition kernel G in equation (1.24) is replaced with $\rho G + (1-\rho)D$.)

Pick any $x, y \in Z$ s.t. $Q_y \geq_{LR} Q_x$. We obviously have $(\lambda_t - \lambda_{t+1})u(y) \geq (\lambda_t - \lambda_{t+1})u(x)$ given the assumption that λ_t is non-increasing in t. Moreover, the induction hypothesis implies $\min\{H_\lambda(y,t+1),0\} \geq \min\{H_\lambda(x,t+1),0\}$ and $(1-\rho)H_\lambda(y,t+1) \geq (1-\rho)H_\lambda(x,t+1)$. It then suffices to show $\int_{z'\in Z} \max\{H_\lambda(z',t+1),0\}G(dz';y) \geq \int_{z'\in Z} \max\{H_\lambda(z',t+1),0\}G(dz';x)$ below.

Notice in the current setting, state z matters only through the belief it represents. With slight abuse of notation, I write $H_\lambda(q_z,t+1) = H_\lambda(z,t+1)$. Then, we have:

$$\begin{aligned}
\int_{z'\in Z} \max\{H_\lambda(z',t+1),0\}G(dz';y) &= \int_\theta \Big[\int_s \max\{H_\lambda(\psi_{(y,s)},t+1),0\}\ell(s|\theta)ds\Big] q_y(\theta)d\theta \\
&\geq \int_\theta \Big[\int_s \max\{H_\lambda(\psi_{(x,s)},t+1),0\}\ell(s|\theta)ds\Big] q_y(\theta)d\theta \\
&\geq \int_\theta \Big[\int_s \max\{H_\lambda(\psi_{(x,s)},t+1),0\}\ell(s|\theta)ds\Big] q_x(\theta)d\theta \\
&= \int_{z'\in Z} \max\{H_\lambda(z',t+1),0\}G(dz';x)
\end{aligned}$$

where the two equalities hold by the definition of $\psi_{(z,s)}$. The first inequality holds due to the induction hypothesis and that Claim (a) above implies $\psi_{(y,s)} \geq_{LR} \psi_{(x,s)}$. To see the second inequality, notice Claim (a) implies that $\psi_{(x,s)}$ increases in likelihood-ratio order in s. Together with the induction hypothesis, this implies that $\max\{H_\lambda(\psi_{(x,s)}, t+1), 0\}$ increases in s, which further implies that $\int_s \max\{H_\lambda(\psi_{(x,s)}, t+1), 0\}\ell(s|\theta)ds$ increases in θ since $\ell(\cdot|\theta)$ increases in likelihood-ratio order in θ. The inequality is hence implied by $q_y \geq_{LR} q_x$. □

Now, I slightly strengthen both the condition and the conclusion in Claim (b).

Claim *(c).* Assume λ_t is non-increasing over t. Then, for any $x, y \in Z$, we have $Q_y \geq_{LR} Q_x$ and $\int_\theta \theta dQ_y(\theta) > \int_\theta \theta dQ_x(\theta)$ together imply $H_\lambda(y,t) > H_\lambda(x,t)$ for all t.

Proof for Claim (c). I show by backward induction in t. The result holds with $t = T$ since $H_\lambda(z,T) = (1+\lambda_T)\int \theta dQ_z(\theta)$. Now, assuming the result holds for all periods since time $t+1$, I show it for period t. By the same argument as in the proof of Claim (b), we have

$$\begin{aligned}H_\lambda(z,t) =&(\lambda_t - \lambda_{t+1})u(z) + \rho\min\{H_\lambda(z,t+1),0\} + (1-\rho)H_\lambda(z,t+1)\\ &+\rho\int_{z'\in Z}\max\{H_\lambda(z',t+1),0\}G(dz';z)\end{aligned}$$

and that for any x and y satisfying the conditions in the claim: (i) $(\lambda_t - \lambda_{t+1})u(y) \geq (\lambda_t - \lambda_{t+1})u(x)$; (ii) $\min\{H_\lambda(y,t+1),0\} \geq \min\{H_\lambda(x,t+1),0\}$; (iii) $\int_{z'\in Z}\max\{H_\lambda(z',t+1),0\}G(dz';y) \geq \int_{z'\in Z}\max\{H_\lambda(z',t+1),0\}G(dz';x)$. Moreover, the induction hypothesis directly imply that $(1-\rho)H_\lambda(y,t+1) > (1-\rho)H_\lambda(x,t+1)$ for $\rho < 1$. These together imply $H_\lambda(y,t) > H_\lambda(x,t)$ as is desired. □

Now, define functions V_λ and H_λ in the same way as in Appendix 1.B.4, but with p_t replaced with z_t. Then the dynamic programming result – Lemma 1.B.2 – still applies to the current setting, because its proof only relies on property (P1) in Lemma 1.3.1, which has its counterpart in Lemma 1.A.1. Given the result of Claim (c) and Lemma 1.A.2, Lemma 1.B.2(b) implies that any solution to the Lagrangian optimization $\max_\phi \mathcal{L}(\phi;\lambda^*)$ (with λ^* solves the dual problem) is almost surely equivalent to some ϕ^* satisfying the property specified in Proposition 1.A.1. The proposition hence holds by the duality result in Lemma 1.3.2.

Q.E.D.

Chapter 2

Information Design for Selling Search Goods and the Effect of Competition

2.1 Introduction

Modern information technology has significantly improved firms' ability in communicating product information with potential customers. For examples, they can provide virtual demonstrations online, make personalized recommendations to individual customers (e.g., via targeted email ads or mobile pushes), or have key opinion leaders evaluate their products for specific consumer groups. Given this ease and flexibility of information provision, what and how much information to provide has become an important design question for firms.

Most existing studies on consumer information design have focused on *experience goods*, which are products whose matches cannot be uncovered by consumers on their own before purchase. For this type of products, consumers have to make purchase decisions solely based on the information provided by the seller. However, many products in practice are better understood as *search goods*, whose match values are naturally revealed to consumers after a *search* step before purchase.[1] Classical examples include clothing, home supplies and furniture, for which a consumer can easily tell how she likes the product upon visiting the seller. More generally, a product that allows returns may also be better understood as a search good because

[1] These concepts on goods classification are introduced by Nelson (1970).

its consumers can almost learn their match values before committing on purchase. Treating a search good as an experience good ignores the post-search information revealing and may thus exaggerate the seller's control over consumer information.

The consumer's ability in learning her true match value before purchase seems to make the seller's information design irrelevant. In many situations, however, the search step required before one can learn the product's match incurs a *search cost*. This can include the cost of visiting a local store in off-line shopping or non-refundable shipping costs in online shopping. Although a search goods seller's information is irrelevant after search, it decides whether the consumer will search the product in the first place. By providing proper pre-search information, the seller can attract as many as possible consumers who will purchase once having sunk their search costs into searching.

In this paper, I study optimal pre-search information provision by search goods sellers. In the main model, a monopoly seller of a search good can design a general pre-search signal (à la Bayesian persuasion) for a representative consumer on her match with the product, based on which the consumer decides whether to search the product.[2] After search, the consumer will fully learn her match value and then decide whether to make a purchase. Since in practice a consumer often knows her best alternative to the product, which is unknown to the seller, I allow the consumer to have private information on her outside option. This induces a more general model with richer applications.

The post-search revelation of product match, as the identifying property of search goods, significantly complicates the seller's information design problem compared to that for experience goods. In particular, the first moment of the consumer's posterior belief on match value no longer suffices for determining the purchase outcome.[3] Together with the existence of the consumer's private information on her outside option, this makes it difficult to derive the optimal design with existing tools.[4] To overcome this challenge, I propose a relaxed problem of the seller. Despite being much simpler than the original problem, this relaxed problem turns out to be sufficient for solving the optimal design under certain regularity conditions.

When the consumer's outside option value is unimodal (i.e., has a quasi-concave density),

[2] Section 2.2.2 provides some concrete interpretations for the seller's information provision.

[3] If the product's match value has a chance to be really high, even if its mean is low, it can still be optimal for the consumer to search and learn the actual match. This suggests that higher moments of the consumer's posterior belief must matter.

[4] A discussion about the existing tools is provided in Appendix 2.A.1.

I show that it is optimal for the seller to provide an *upper-censorship signal*, which is the unique optimal signal up to outcome-irrelevant modifications when the unimodality is strict. Under this signal, consumers with match values below a threshold will be fully informed of their matches, while others will only learn that their match values are above the threshold. The characterization also shows that the optimal design crucially depends on the curvature of the consumer's outside option value distribution. In particular, when the distribution is convex over a relevant region, the optimal signal will be fully revealing; when it is concave, the optimal signal will be completely uninformative.

Besides solving the optimal design, my relaxed problem approach also enables a straightforward comparison between the information design problems of search goods and experience goods. Specifically, the search goods seller's relaxed problem can be treated as the experience goods seller's problem with just one more constraint. On one hand, the additional constraint formalizes the intuition that the seller's control over consumer information is weaker with search goods than with experience goods, which is critical in shaping the optimal information provision. On the other hand, despite being more constrained, the relaxed problem does share an important structure with the design problem of experience goods. This explains analogies between many results of these two types of goods and has enabled me to extend results from one to the other.

My characterization of the optimal design enables several applications. The first one studies the effect of policies turning experience goods into search goods.[5] It is shown that while these policies directly benefit consumers by forcing post-search information revealing, under certain conditions they may reduce the seller's incentive in providing pre-search information and thereby lead to more inefficient searches. The overall effect on consumer welfare can be negative. This suggests that for such policies to benefit consumers, additional efforts may be needed to maintain sellers' incentives in providing better pre-search information.

In the second application, I consider the possibility that the seller can partially observe consumers' outside option values and tailor information accordingly. I show that the seller optimally provides better information to those who are expected to have higher outside option values, which forms a kind of discrimination in information provision. This discrimination affects different consumers heterogeneously and its effect on total consumer welfare is generally

[5]Examples of such policies include product labeling laws and regulations forcing sellers to accept product returns. See my discussion in Section 2.4 for more details.

indeterminate. This highlights a non-price discrimination channel through which the design of consumer privacy can influence consumer welfare.

In the third application, I examine the effect of changes in the search cost. When the product price is exogenously fixed,[6] the model predicts that the consumer welfare will unambiguously increase when the search cost decreases, although this gain will be partially offset by coarser pre-search information provided by the seller. This holds not only for overall consumer welfare, but also for each individual consumer given her realized outside option value.

My last application is to study the effect of competition among multiple sellers with horizontally differentiated products. Although I am not able to solve the equilibrium in general, my approach suffices for pinning down an equilibrium when the number of sellers is sufficiently large under certain regularity conditions. In particular, I show that as competition becomes stronger, there is a sequence of equilibria with pre-search information converging to full information. This extends the corresponding result in Hwang et al. (2019) from experience goods to search goods.

Related literature – Possibly due to its technical difficulty, the economic literature on information design for search goods is relatively scarce. To the best of my knowledge, the only major study on seller's information provision for search goods is Anderson & Renault (2006).[7] They consider a similar setting as mine, but the consumers in their paper are ex-ante homogeneous without private information. In Section 2.7.1, I will compare their optimal design with mine and explain how their result is *partially* extended in my setting. I note that most of my findings in the applications above cannot be made in their setting because the consumer's private information plays important roles in those findings.

My paper also relates to several other strands of literature. The first strand studies sellers' provision of real product information without consumer search (e.g., Meurer & Stahl II, 1994; Lewis & Sappington, 1994; Johnson & Myatt, 2006; Ivanov, 2013; Boleslavsky et al., 2017; Hwang et al., 2019). Most relatedly, Ivanov (2013) and Hwang et al. (2019) consider competitive sellers and show that the equilibrium information converges to full information when the number of sellers goes to infinity. In particular, the equilibrium characterization in Hwang et al. (2019)

[6] This mainly applies when the seller is a salesperson or information intermediary, who does not set price but just tries to induce purchase via providing match information.

[7] Also see Anderson & Renault (2013), which extends Anderson & Renault (2006) to also encompass vertical information.

also plays a key role in my analysis. As mentioned earlier, the corresponding result of mine extends their results from experience goods to search goods.

The second strand of literature studies various information design questions in consumer search environments. Unlike mine, most of these papers focus on post-search information design instead of pre-search information design (Bar-Isaac et al., 2012; Board & Lu, 2018; Dogan & Hu, 2018; Au, 2018; Whitmeyer, 2018; Whitmeyer, 2020). Since their consumers cannot receive additional information than that provided by the seller, these papers can still be considered as being about experience goods. Two exceptions are Choi et al. (2019) and Hinnosaar & Kawai (2020). Choi et al. (2019) studies the consumer-optimal pre-search information; Hinnosaar & Kawai (2020) studies the seller-worst pre-search information in a robust mechanism design problem. Unlike these papers, I study the seller-optimal pre-search information provision. Another difference is that I allow continuous value distributions, while the two papers above assume binary values.

More closely related, Wang (2017) considers an environment where consumers can engage in costly search for additional information after receiving the seller's signal. With ex-ante homogeneous consumers, the paper concludes that the seller's optimal signal should deter the consumer from searching for more information. Matyskova (2018) also derives a similar result in a more abstract Bayesian persuasion setting. Importantly, while search is optional in these papers, it is necessary for making purchase in my model. Thus their main topic of search deterrence is not relevant in my study. Moreover, their analyses heavily rely on the ex-ante homogeneity of consumers (as in Anderson & Renault (2006)). In contrast, I allow ex-ante heterogeneous consumers with privately known outside option values.

My paper greatly benefits from recent developments in the consumer search literature. In particular, I draw upon a characterization for the consumer's purchase outcome given any belief on the product match value, which is discovered in several papers (Kleinberg et al., 2016; Armstrong, 2017; Choi et al., 2018).

Finally, the paper relates to the general Bayesian persuasion literature (Rayo & Segal, 2010; Kamenica & Gentzkow, 2011). In particular, the relaxed problem I propose is analogous to the optimization in Dworczak & Martini (2019) and Section 4.3 of Kolotilin (2018), but features one more constraint. The optimality conditions I provide essentially extend Theorem 1 in Dworczak & Martini (2019) to accommodate that additional constraint. Kolotilin et al. (2017) and Guo &

Shmaya (2019) also consider a receiver with private information. However, they do not consider the design for selling search goods, whose complexity motivates my relaxed problem approach.

The paper is organized as follows: Section 2.2 introduces the main model; Section 2.3 develops the relaxed problem and characterizes the optimal design; Section 2.4 compares the designs of search goods and experience goods; Section 2.5 considers comparative statics and related applications; Section 2.6 considers information provision by competing sellers; Section 2.7 provides further discussions. All proofs are provided in the appendix.

2.2 The Model

2.2.1 The Setup

The model features a seller, a representative consumer and a single product. The consumer's match value with the product is denoted as U, which is initially unknown to both agents and has a continuous distribution F_U with finite mean and compact support $[\underline{u}, \bar{u}]$. At the beginning of the game, the seller can design a pre-search signal (statistical experiment) to provide information about U to the consumer. Following Anderson & Renault (2006) and the Bayesian persuasion literature, I do not make any restriction on the signal structure.[8] I use S to denote the signal's realization, and use $\phi(\cdot; S)$ to denote the consumer's posterior belief on U given S.

If the consumer does not consume the product, she will consume her outside option, whose value is the consumer's private information and is denoted as U_0. I assume U_0 is drawn from a distribution J, and is independent from U.

I will consider two kinds of sellers. A *non-pricing seller* treats the product's price as exogenous, and just wants to maximize the consumer's purchase probability by providing information;[9] a *pricing seller* also sets the product's price p, and tries to maximize the expected profit. I normalize the seller's marginal cost to zero.

The game goes as follows:

1. The seller designs the pre-search signal about U. A pricing seller also chooses the product's

[8]Formally, a signal consists of a measurable realization space $\mathcal{S}$ and a transition kernel $\pi : [\underline{u}, \bar{u}] \to \Delta(\mathcal{S})$ that maps any realized match value to a probability distribution over $\mathcal{S}$, according to which the signal realization S will be drawn.

[9]In practice, non-pricing sellers can be salesmen, brokers or information intermediaries, who typically do not set price but just try to induce purchase by providing information.

price p. These are observed by the consumer.[10]

2. The match value U is (secretly) realized, and the pre-search signal realization S is generated to the consumer.
3. After learning S and her outside option value U_0, the consumer decides whether to *search* the product. If not, she consumes the outside option and receives utility U_0.
4. If the consumer chooses to search, she will incur a search cost $c > 0$ and learn U. She then makes her purchase decision. With purchase, her final utility will be $U - p - c$; without purchase, her final utility will be $U_0 - c$.

The assumption that the consumer can fully learn her match value after search is the identifying property of search goods. Formally, I adopt the following dichotomy in Nelson (1970).

Definition 2.2.1. A product is a *search good (abbr., SG)* if its match value will be fully revealed to the consumer after search; it is an *experience good (abbr., EG)* if the consumer cannot learn additional information after search before purchase.

Notice that I assume search is a necessary step towards purchase even if the product is an experience good. This ensures the two types of goods have the same total purchase cost so that they are comparable. I will call the seller in the model described above an *SG seller*, and call a seller *EG seller* if he faces the same problem except that the consumer will learn no additional information after search. In Section 2.4, I will compare the information design problems of these two types of sellers, and highlight the key similarity and dissimilarity between them.

2.2.2 Interpretations for the Seller's Information Provision

The seller's information provision admits two general interpretations.

In the first interpretation, which is adopted by Anderson & Renault (2006), the seller has a single product and faces a large population of potential consumers, whose matches with the product are independently drawn. To provide pre-search information, the seller advertises selected product characteristics to all consumers.[11] With these characteristics, each consumer can privately update her belief about the product's match value based on her own taste. By

[10] I assume that the price set by a pricing seller is observed before search, which avoids the hold-up problem in the Diamond's Paradox (Diamond, 1971). This is innocuous as long as the seller can provide price information along with the pre-search match information.

[11] Only *horizontally differentiating* characteristics are considered here, which better informs a consumer of her individual match without vertically shifting the aggregate demand.

selecting different characteristics to advertise, the seller then imparts different information to the consumers.[12]

In the second interpretation, which is more popular in the Bayesian persuasion literature, the seller interacts with each individual consumer repeatedly over time. Each time, the seller offers one *issue* of his product, whose match with the consumer is drawn independently from those of the other issues. Knowing both features of the product issues and the consumer's preference (e.g., revealed by demographics and browsing history), the seller can predict the consumer's match with each issue and provide information about it (e.g., via recommendation messages). Since the seller repeatedly interacts with the consumer, it is conceivable that he is able to commit on a particular information provision rule to maximize the long-run profit, and the consumer can correctly interpret the messages she receives.[13] This justifies the seller's commitment power over the signal structure being used.[14]

2.3 Solving the Seller's Optimal Design

I first consider a non-pricing seller with exogenous price p, and focus on the information design problem. The pricing seller's problem will be studied in Section 2.3.4.

2.3.1 Preliminaries

To characterize the consumer's optimal search behavior, given any posterior belief ϕ on U, let z_ϕ denote the corresponding Pandora's index.[15]

That is, z_ϕ solves:

$$\int [(x - z_\phi)_+ - c]\phi(dx) = 0 \tag{2.1}$$

where $(y)_+ := \max\{y, 0\}$. Given any pre-search signal, I define the random variable Z as

[12] In this interpretation, the seller's flexibility in designing information is limited by the richness of the available product characteristics. See Section III and Appendix A in Anderson & Renault (2006).

[13] Several recent papers have studied how repeated interaction can support commitment in strategic communication. See, e.g., Mathevet et al. (2019) and Best & Quigley (2020).

[14] I note that for this second interpretation to fit the model with a pricing seller, the seller cannot price-discriminate the consumer based on her realized match values over time. In practice, this can be reasonable for multiple reasons. First, consumers typically have different match values with the same issue and it may be difficult to set personalized prices for them. Second, as is argued in Ichihashi (2020), committing to non-discrimination can relax consumers' privacy concerns and encourages them to share preferences with the seller.

[15] This index is called "reservation price" in the original paper of Weitzman (1979). It is easy to show its existence and uniqueness when $c > 0$.

the Pandora's index of the consumer's posterior belief conditional on the pre-search signal realization, i.e., $Z := z_{\phi(\cdot;S)}$. Then, according to the Pandora's rule in Weitzman (1979), the consumer will search the product if and only if $Z \geq U_0 + p$. Since the consumer will purchase after search if and only if the revealed $U \geq U_0 + p$, the product will be finally sold if and only if $U \wedge Z \geq U_0 + p$, where $x \wedge y$ is a shorthand for $\min\{x, y\}$.[16] I will call this key statistic $U \wedge Z$ the product's *effective-search-value*.[17]

Let G denote the CDF of $U \wedge Z$ induced by the pre-search signal. Let J_p denote the CDF of $U_0 + p$, i.e., $J_p(x) := J(x - p)$. The consumer's purchase probability then equals to $\mathbb{P}(U_0 + p \leq U \wedge Z) = \mathbb{E}[J_p(U \wedge Z)] = \int J_p(x)dG(x)$. The non-pricing seller's problem can then be formulated as:

$$\max_G \int J_p(x)dG(x) \tag{2.2}$$

$$\text{s.t. } G \text{ is a feasible distribution of } U \wedge Z \tag{2.3}$$

Unfortunately, I cannot handle this optimization directly because a full characterization of its feasible set will be too complicated to handle in such an optimization. To get things simplified, I will hence propose a relaxed problem of it, which is based on two lemmas below.

Given any search cost c, I will call $U - c$ the consumer's *net-match-utility* and let F_{U-c} denote its CDF. The following lemma provides a key necessary condition for the constraint (2.3) to hold.

Lemma 2.3.1. *Given F_U and c, the distribution of $U \wedge Z$ under any signal is a mean-preserving contraction (MPC) of F_{U-c}. As a special case, $U \wedge Z = U - c$ under the fully revealing signal.*

The lemma suggests that we can relax the seller's problem by replacing constraint (2.3) with the constraint that G is a MPC of F_{U-c}. However, this will lead to an optimization that is often too relaxed for its solution to be feasible for the seller. Actually, as I will discuss in Section 2.4, such an optimization is equivalent to one faced by an EG seller, with which the difference between the two types of goods will be neglected. Therefore, to derive a more useful relaxed problem, some additional characterization for the constraint (2.3) is needed. The next lemma provides such a result.

[16] I assume the consumer will search and buy the product when being indifferent.

[17] This characterization of the consumer's purchase outcome has been proposed by several papers (Kleinberg et al., 2016; Armstrong, 2017; Choi et al., 2018).

Lemma 2.3.2. *Under any signal, $U \wedge Z \leq U$ and thus the distribution of $U \wedge Z$ is first-order stochastic dominated by F_U.*

While being obvious, the lemma highlights an important restriction faced by an SG seller. Namely, no matter what pre-search information is provided, the consumer's effective-search-value $U \wedge Z$ is always bounded by the true match value U. This limits the seller's ability in manipulating the consumer's purchase behavior through inducing mean-preserving contraction of G by withholding information.

2.3.2 A Relaxed Problem

Let $\preceq_{MPS}$ denote the mean-preserving spread order and let $\preceq_{FOD}$ denote the first-order stochastic dominance order. Based on Lemmas 2.3.1 and 2.3.2, a *Relaxed Problem* of optimization (2.2) – (2.3) can be formulated as:

$$\max_G \int J_p(x)dG(x) \tag{2.4}$$

$$\text{s.t. } G \preceq_{MPS} F_{U-c} \tag{2.5}$$

$$G \preceq_{FOD} F_U \tag{2.6}$$

For any distribution F, I denote its support as *supp*$\{F\}$. The following theorem provides the optimality conditions for solving the linear program above. It extends Theorem 1 in Dworczak & Martini (2019) to accommodate constraint (2.6).

Theorem 2.3.1. *A distribution G solves problem (2.4) – (2.6) if there exists functions $v(\cdot)$ and $\rho(\cdot)$ such that:*

(C1) $v(\cdot)$ is convex over $[\underline{u} - c, \bar{u} - c]$ and $\rho(\cdot)$ is weakly increasing over $[\underline{u} - c, \bar{u}]$.

(C2) $v(x) + \rho(x) \geq J_p(x)$ for all $x \in [\underline{u} - c, \bar{u} - c]$, with equality holding for any $x \in supp\{G\}$.

(C3) $\int v(x)dG(x) = \int v(x)dF_{U-c}(x)$; $\int \rho(x)dG(x) = \int \rho(x)dF_U(x)$.

(C4) G satisfies constraints (2.5) and (2.6).

Moreover, if there exist G, $v(\cdot)$ and $\rho(\cdot)$ satisfying the above conditions, then another distribution $\hat{G}$ also solves the problem if and only if it satisfies conditions (C1) – (C4) with the same $v(\cdot)$ and $\rho(\cdot)$.

In general, the Relaxed Problem does not necessarily admit a solution that is feasible for the seller to induce. However, as I will show in the next subsection, it does under some mild conditions on the outside option value distribution.

2.3.3 The Optimal Signal with Unimodal Outside Option

I impose the following unimodal assumption on the distribution of $U_0 + p$.

Assumption 2.3.1 (Unimodal J_p)**.** J_p admits a continuous quasi-concave density over $[\underline{u} - c, \bar{u} - c]$.

Let j_p denote the density of J_p. The assumption simply requires that over the support of $U - c$, j_p is first increasing and then decreasing, which implies J_p to be first convex and then concave. It is weaker than requiring U_0 to admit a log-concave density, which is satisfied by many common distributions.[18] Under the assumption, I will use $[r_p^{min}, r_p^{max}]$ to denote the mode (interval) of J_p, i.e.,

$$[r_p^{min}, r_p^{max}] := \underset{x \in [\underline{u}-c, \bar{u}-c]}{\arg\max}\ j_p(x) \tag{2.7}$$

Following Kolotilin et al. (2021), I define an *upper-censorship signal* as follows:

Definition 2.3.1. A signal is an upper-censorship signal if there is a threshold $\eta \in [\underline{u} - c, \bar{u} - c]$ such that the signal fully reveals any net-match-utility $U - c$ below η, and pools all net-match-utilities above η together.

Given any upper-censorship signal with threshold η, I will use G_η to denote the distribution of $U \wedge Z$ it induces, and use $z(\eta)$ to denote the Pandora's index of the consumer's posterior belief after learning $U - c \geq \eta$.[19] It is easy to see that $z(\eta) \geq \eta$ and is strictly increasing in η. Figure 2.1 illustrates how $U - c$ maps to $U \wedge Z$ under an upper-censorship signal. We can see that compared to the full revelation case, where $U \wedge Z$ always equals to $U - c$, upper-censorship leads to a contraction for $U \wedge Z$ over the region above η, which by Lemma 2.3.1 must be mean-preserving. In particular, consumers with $U - c \in (\eta, z(\eta))$ will have their effective-search-values increased to either U or $z(\eta)$, whichever is smaller; those with $U - c > z(\eta)$ will have their effective-search-values decreased to $z(\eta)$.

[18]For a list of distributions with log-concave density, see Bagnoli & Bergstrom (2005) Table 1.

[19]I provide detailed properties of $z(\cdot)$ and the formula of G_η in Appendix 2.B.3.

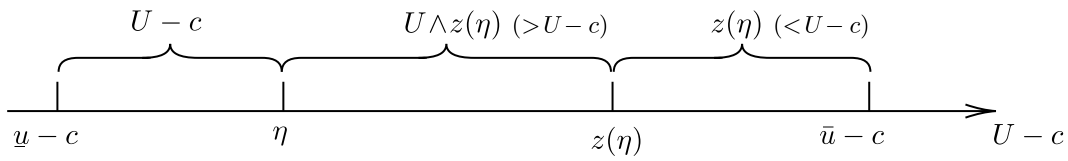

Figure 2.1: Values of $U \wedge Z$ given different values of $U - c$ under an upper-censorship signal with threshold η. The axis represents different values of $U - c$ and the expressions above the brackets indicate the value of $U \wedge Z$ given $U - c$ in each region. We can see that $U \wedge Z > U - c$ when $U - c \in \big(\eta, z(\eta)\big)$ and $U \wedge Z < U - c$ when $U - c > z(\eta)$.

To state the main result below, given any J_p satisfying the unimodal assumption, I define:

$$\eta_0 := \inf\{\eta \in [\underline{u} - c, \bar{u} - c] : z(\eta) \geq r_p^{min}\} \tag{2.8}$$

$$\Gamma(\eta) := \frac{J_p\big((\eta + c) \wedge z(\eta)\big) - J_p(\eta)}{(\eta + c) \wedge z(\eta) - \eta} - j_p\big(z(\eta)\big), \; \forall \eta \in [\underline{u} - c, \bar{u} - c) \tag{2.9}$$

Intuitively, η_0 is the smallest η such that $z(\eta)$ will fall in the region where $J_p(\cdot)$ is concave; $\Gamma(\eta)$ measures the difference between the average slope of $J_p(\cdot)$ over $[\eta, (\eta + c) \wedge z(\eta)]$ and its slope at $z(\eta)$. An optimal signal is characterized by the following proposition.

Proposition 2.3.1. *Under Assumption 2.3.1, for a non-pricing seller:*

(a) If $r_p^{max} = \bar{u} - c$, the upper-censorship signal with threshold $\eta^* = \bar{u} - c$ (i.e., full disclosure) is optimal.

(b) If $r_p^{max} < \bar{u} - c$, there exists $\eta^* \in [\eta_0, r_p^{max}]$ such that either of the following conditions hold:

(i) $\eta^* > \underline{u} - c$ and $\Gamma(\eta^*) = 0$;

(ii) $\eta^* = \underline{u} - c$ and $\Gamma(\eta^*) \geq 0$.

For any such η^*, the upper-censorship signal with threshold η^* is optimal. Moreover, G_{η^*} solves the Relaxed Problem (2.4) – (2.6).

To gain some intuition about why upper-censorship is optimal under the unimodal assumption, consider two events with equal probabilities: in event A, $U - c = a$; in event B, $U - c = b > a$. If they are separately revealed, then $U \wedge Z = U - c$ in both events. Now, suppose the seller instead pools these events. Then the benefit for him is that $U \wedge Z$ will be raised to some $a' > a$ in event A, which increases the sale probability by $\frac{J_p(a') - J_p(a)}{2}$; the cost is that $U \wedge Z$ will be decreased to some $b' < b$ in event B, which decreases the sale probabil-

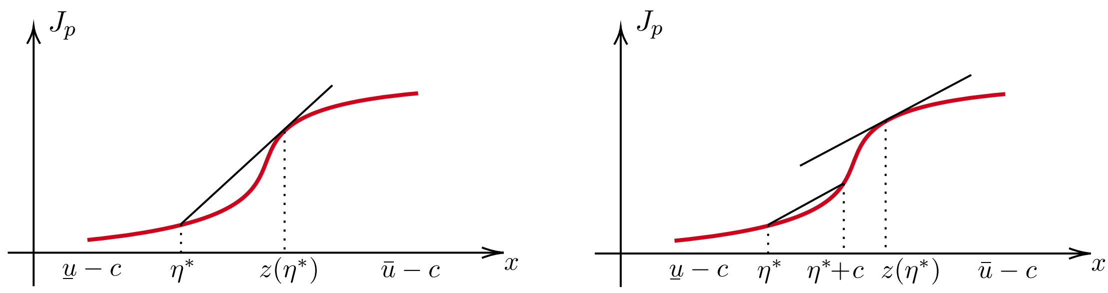

(a) Case 1: $z(\eta^*) \leq \eta^* + c$

(b) Case 2: $z(\eta^*) > \eta^* + c$

Figure 2.2: Graphical illustration for the optimality condition $\Gamma(\eta^*) = 0$. The red curve is $J_p(\cdot)$. The condition implies that the secant of $J_p(\cdot)$ over $[\eta^*, (\eta^* + c) \wedge z(\eta^*)]$ is parallel to the tangent of $J_p(\cdot)$ at $z(\eta^*)$.

ity by $\frac{J_p(b)-J_p(b')}{2}$. Notice the mean-preserving condition implies $\frac{a'-a}{2} = \frac{b-b'}{2}$. We thus must have $\frac{J_p(a')-J_p(a)}{2}$ to be less (resp., greater) than $\frac{J_p(b)-J_p(b')}{2}$ when J_p is convex (resp., concave) over $[a, b]$. This suggests that the optimal signal should reveal better information over a region where J_p is convex, and tends to withhold information over a region where J_p is concave. Under Assumption 2.3.1, this is right in accordance with upper-censorship.[20]

In part (b) of Proposition 2.3.1, an interior optimal threshold can be pinned down by the condition $\Gamma(\eta^*) = 0$. Graphically, this means that the secant of J_p over $[\eta^*, (\eta^* + c) \wedge z(\eta^*)]$ has the same slope as J_p at $z(\eta^*)$ (see Figure 2.2). To understand this condition, suppose that the seller adds some additional consumers with $U - c$ right below η^* to the pooling region, whose total mass is dm. The benefit is that these consumers will have their $U \wedge Z$ increased from η^* to $(\eta^* + c) \wedge z(\eta^*)$, which increases the total sale probability by $\left[J_p\big((\eta^* + c) \wedge z(\eta^*)\big) - J_p(\eta^*)\right] dm$. The cost is that consumers originally with $U \wedge Z = z(\eta^*)$, I denote whose mass as M, will have their effective-search-values marginally decreased from $z(\eta^*)$ by some $|dz|$ to obey the mean-preserving condition. This decreases the total sale probability by $J_p'\big(z(\eta^*)\big)|dz|M$. The seller's net benefit thus equals to $\left[J_p\big((\eta^* + c) \wedge z(\eta^*)\big) - J_p(\eta^*)\right] dm - J_p'\big(z(\eta^*)\big)|dz|M$, which is proportional to:

$$\frac{\left[J_p\big((\eta^* + c) \wedge z(\eta^*)\big) - J_p(\eta^*)\right]}{|dz|M/dm} - J_p'\big(z(\eta^*)\big)$$

[20]This does not imply that the optimal threshold coincides with the mode of J_p, since $J_p(a')/2 - J_p(a)/2 > J_p(b)/2 - J_p(b')/2$ can also hold when J_p is first convex and then concave over $[a, b]$, in which case pooling is more profitable. Proposition 2.3.1 only implies $\eta^* \leq r_p^{max}$.

By the mean-preserving condition, we must have $[(\eta^* + c) \wedge z(\eta^*) - \eta^*]dm = |dz|M$. The above expression is thus equal to $\Gamma(\eta^*)$. Therefore, the graphical condition guarantees that any marginal deviation from η^* will not be profitable for the seller.[21]

Proposition 2.3.1 suggests that the seller's optimal design crucially depends on the curvature of the consumer's outside option value distribution. The following corollary of it illustrates this with the extreme cases.

Corollary 2.3.1. *For a non-pricing seller, we have:*

(a) If $J_p(\cdot)$ is convex over $[\underline{u} - c, \bar{u} - c]$, then the fully revealing signal is optimal.

(b) If $J_p(\cdot)$ is concave over $[\underline{u} - c, \bar{u} - c]$, then the fully pooling signal is optimal.

In general, the seller's optimal design is not unique. For example, suppose an optimal upper-censorship signal is fully revealing for $U - c \in [a, b]$ and J_p is affine over this region. Then another signal pooling over this region (otherwise identical to the original signal) will also be optimal. However, under a strict version of the unimodal assumption on J_p, we do have a uniqueness result.

Assumption 2.3.2 (Strictly Unimodal J_p)**.** Assumption 2.3.1 holds with j_p being strictly quasi-concave over $[\underline{u} - c, \bar{u} - c]$.

The assumption implies that j_p is first strictly increasing and then strictly decreasing over $[\underline{u} - c, \bar{u} - c]$. When it holds, I will use r_p to denote the unique mode of J_p, i.e., $r_p = \arg\max_{x \in [\underline{u} - c, \bar{u} - c]} j_p(x)$.

Proposition 2.3.2. *Under Assumption 2.3.2, all optimal signals induce the same joint distribution of (U, Z), and the optimal upper-censorship signal is unique.*

Notice that the pair of (U, Z) determines a consumer's search and purchase decisions as well as her ex-post utility given any U_0. Thus the proposition implies that under Assumption 2.3.2, the equilibrium search-purchase outcome and consumer surplus are both unique. This makes it convenient to study any comparative statics.

2.3.4 The Optimal Design for a Pricing Seller

Given the results for a non-pricing seller, characterizing the optimal design for a pricing seller is straightforward. By the same argument as in Section 2.3.1, the pricing seller's problem can

[21] For corner solution $\eta^* = \underline{u} - c$, the threshold cannot be further lowered, so we only need to rule out profitable upward deviations from it. Therefore, we only need to require $\Gamma(\eta^*) \geq 0$.

be written as:

$$\max_{G,p}\{p\int J(x-p)dG(x)\} \tag{2.10}$$

$$\text{s.t. } G \text{ is a feasible distribution of } U \wedge Z \tag{2.11}$$

The following assumption allows me to use the simple characterization of the non-pricing seller's optimal design in Section 2.3.3.

Assumption 2.3.3 (Unimodal J)**.** *J admits a continuous quasi-concave density.*

The assumption guarantees that for any price p, the distribution of $U_0 + p$ (i.e., J_p) satisfies Assumption 2.3.1 and hence there exists an optimal upper-censorship signal. Since an upper-censorship signal is solely pinned down by its threshold η, we can transform the pricing seller's problem into an optimization over (p, η), which leads to the following result.

Proposition 2.3.3. *Under Assumption 2.3.3, a price p^* and an upper-censorship signal with threshold η^* are optimal for the pricing seller if and only if:*

$$(p^*, \eta^*) \in \operatorname*{arg\,max}_{p\geq 0;\ \eta\in[\underline{u}-c,\bar{u}-c]} \{p\int J(x-p)dG_\eta\}$$

where the formula of G_η is given by equation (2.23) in Appendix 2.B.3. Moreover, if $J(\cdot)$ is log-concave, then an optimal pair of (p^*, η^*) exists.

Proposition 2.3.3 has simplified the original infinite-dimensional optimization (2.10) – (2.11) into a two-dimensional problem. It will allow me to study several applications of the model in later sections.

2.4 Search Goods vs. Experience Goods

In practice, whether a product is a search good (SG) or an experience good (EG) not only depends on its own property, but also depends on the shopping environment and related consumer protection policies. For examples, product labeling laws require sellers to provide detailed product information on packages, which enables consumers to learn product characteristics (e.g., food nutrition) that they would otherwise not know before (or even after) purchase; consumer protection laws in many countries require online sellers to accept returns without any reason within certain time period.[22] These policies allow consumers to better learn their matches with a product after some search step but before committing on purchase, and can thus approximately transform experience goods into search goods.[23]

In this section, I investigate how such policies can change the equilibrium outcomes and compare the two types of goods from an information design perspective. For simplicity, I will focus on non-pricing sellers with exogenous product price unless otherwise stated. Some numerical analyses will be provided in Appendix 2.A.2 for the case of pricing sellers.

2.4.1 Comparison between the Design Problems

When the product is an experience goods, the consumer will receive no additional information after the search step. She will hence buy the product if and only if her *posterior mean net-match-utility* $\mathbb{E}[U-c|S]$ exceeds U_0+p. The statistic $\mathbb{E}[U-c|S]$ thus replaces the role of $U \wedge Z$ in determining the purchase outcome. Let H denote the distribution of $\mathbb{E}[U-c|S]$ to be induced. It is well known that H is feasible if and only if it is a MPC of F_{U-c}.[24] A non-pricing EG seller's problem can thus be written as:

$$\max_{H} \int J_p(x)dH(x) \tag{2.12}$$

$$\text{s.t. } H \preceq_{MPS} F_{U-c} \tag{2.13}$$

This kind of optimization has been well studied in the Bayesian persuasion literature (e.g., Kolotilin (2018) and Dworczak & Kolotilin (2019)).

Comparing optimization (2.12) – (2.13) with optimization (2.4) – (2.6), one can see some key similarity and dissimilarity between the problems of the two types of goods. In terms of similarity, both problems involve the same MPC constraint. This roots from the fact that both $U \wedge Z$ and $\mathbb{E}[U-c|S]$ must be a MPC of the true net-match-utility $U-c$. This commonality

[22] For instances, China requires online sellers to fully refund no-reason returns (except for special products) within 7 days of the sale, excluding any shipping cost; European Union has a similar policy with the cooling period being 14 days and the refund there includes initial (standard) shipping charges. See `http://lawinfochina.com/display.aspx?id=23187&lib=law` and `https://europa.eu/youreurope/citizens/consumers/shopping/guarantees-returns`.

[23] To interpret the search goods model in the situation of online shopping with returnable products, one should consider c to also include any return cost of the consumer, and consider U to be the product's consumption utility plus the return cost. Then the consumer's final utility with search without purchase is indeed $U_0 - c$, and her utility with purchase is indeed $U - c - p$.

[24] See, e.g., Proposition 2 in Kolotilin (2018).

allows many results to be carried over between the two types of goods. In particular, under the unimodal assumption on J_p, the optimal design of an EG seller also features upper-censorship signals. Moreover, the results in Corollary 2.3.1 equally hold for both types of goods.

The key difference between the two problems is that an SG seller faces the additional constraint (2.6) requiring $G \preceq_{FOD} F_U$, which is absent in the EG seller's problem. This is because $U \wedge Z$ must be bounded by U, but $\mathbb{E}[U - c|S]$ needs not be. In this sense, it is harder to induce MPC for $U \wedge Z$ by withholding information than for $\mathbb{E}[U - c|S]$, which makes the SG seller's problem more "constrained". An obvious implication of this is the following, which holds for both pricing and non-pricing sellers.

Proposition 2.4.1. *The seller's profit is (weakly) lower when the product is a search good than when it is an experience good.*

In the following subsections, I will further compare the two types of goods in terms of the equilibrium information provision and consumer welfare. For convenience, I first present a useful lemma here.

Lemma 2.4.1. *For any belief ϕ on U, we have $z_\phi \geq \mathbb{E}_{U\sim\phi}[U - c]$, which holds as equality if and only if* $\inf(supp\{\phi\}) \geq \mathbb{E}_{U\sim\phi}[U - c]$.

The lemma implies that given any posterior belief ϕ of the consumer, the corresponding Pandora's index z_ϕ is no less than the posterior mean of $U - c$. This is intuitive since even if the consumer always ignores the post-search information, it is optimal for her to search when $U_0 + p \leq \mathbb{E}_{U\sim\phi}[U - c]$. Thus z_ϕ must be no less than $\mathbb{E}_{U\sim\phi}[U - c]$ for the Pandora's rule to be optimal. The lemma also implies that these two values are equal when c is sufficiently large such that even the smallest value in the support of ϕ is greater than $\mathbb{E}_{U\sim\phi}[U - c]$. Intuitively, when the search cost is very large, the post-search information will not make a difference to the consumer's decisions because it comes only after the consumer has sunk a significant cost. The range of U_0 that makes it optimal for the consumer to search should hence be the same regardless of the product's type, which implies $z_\phi = \mathbb{E}_{U\sim\phi}[U - c]$.

2.4.2 Comparison for the Equilibrium Information Provision

I first define some notations. Let $\mu(\eta) := \mathbb{E}[U - c | U - c \geq \eta]$ and

$$\Gamma^E(\eta) := \frac{J_p\big(\mu(\eta)\big) - J_p(\eta)}{\mu(\eta) - \eta} - j_p\big(\mu(\eta)\big),\ \forall \eta \in [\underline{u} - c, \bar{u} - c) \tag{2.14}$$

Intuitively, $\mu(\eta)$ is the posterior mean of $U - c$ after observing the pooling signal realization of an upper-censorship signal with threshold η; $\Gamma^E(\eta)$ measures the difference between the average slope of J_p over $[\eta, \mu(\eta)]$ and its slope at $\mu(\eta)$. I note that these functions are fully determined by J_p and F_{U-c}. They do not separately depend on the search cost c once the distribution of $U - c$ is fixed.

I will focus my analysis on the case where the following assumption holds.

Assumption 2.4.1. J_p satisfies Assumption 2.3.2 with $r_p < \bar{u} - c$ and $\Gamma^E(\underline{u} - c) < 0$.

Here, the conditions $r_p < \bar{u} - c$ and $\Gamma^E(\underline{u} - c) < 0$ guarantee that the EG seller's optimal signal is neither fully revealing nor fully pooling. These are not essential for the analysis, but avoid separate discussions of "corner" solutions. Under Assumption 2.4.1, the EG seller's optimal signal is characterized by Kolotilin et al. (2021). For convenience, I summarize its properties in the following proposition.

Proposition 2.4.2. *Assume Assumption 2.4.1 holds. For a non-pricing EG seller, an upper-censorship signal with threshold η_E^* is optimal, where η_E^* is the unique solution to $\Gamma^E(\eta) = 0$. Moreover, all optimal signals are outcome-equivalent to it.*

The optimality condition $\Gamma^E(\eta_E^*) = 0$ implies that the secant of J_p over $[\eta_E^*, \mu(\eta_E^*)]$ is tangent to J_p at $\mu(\eta_E^*)$ (see Figure 2.3). Like the optimality condition $\Gamma(\eta^*) = 0$ for search goods, this condition guarantees that a marginal deviation of the threshold is not profitable for the seller.

For search goods and experience goods respectively, let η_S^* and η_E^* denote the thresholds of the non-pricing seller's optimal upper-censorship signals. Under Assumption 2.4.1, comparing the equilibrium information provisions boils down to comparing these two thresholds. Notice that by Proposition 2.4.2, η_E^* can be equivalently defined as the unique solution to $\Gamma^E(\eta) = 0$. Given the value of it, the following proposition provides conditions for comparing η_S^* with η_E^*:

Proposition 2.4.3. *Assume Assumption 2.4.1 holds. We have:*

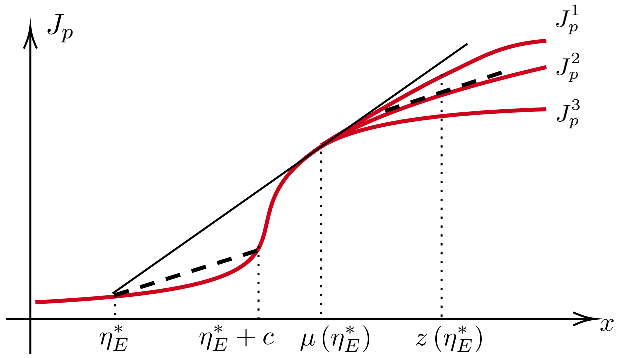

Figure 2.3: An illustration for the cases in Proposition 2.4.3(b). The red curves represent three versions of J_p. Regardless of the version, the solid black line is tangent to J_p, which implies $\Gamma^E(\eta_E^*) = 0$ and thus η_E^* is the EG seller's optimal threshold. The two dashed black segments are parallel, which implies $\Gamma(\eta_E^*) = 0$ when $J_p = J_p^2$. As J_p shifts from J_p^1 to J_p^3, its slope at $z(\eta_E^*)$ decreases, and hence $\Gamma(\eta_E^*)$ shifts from negative to positive.

(a) If $c \geq \mu(\eta_E^*) - \eta_E^*$, then $\eta_S^* = \eta_E^*$ and the equilibrium outcomes of the two types of goods are the same.

(b) If $c < \mu(\eta_E^*) - \eta_E^*$, then we have: (i) $\Gamma(\eta_E^*) < 0 \Rightarrow \eta_S^* > \eta_E^*$; (ii) $\Gamma(\eta_E^*) = 0 \Rightarrow \eta_S^* = \eta_E^*$; (iii) $\Gamma(\eta_E^*) > 0 \Rightarrow \eta_S^* < \eta_E^*$.

Since $\mu(\cdot)$ and $\Gamma^E(\cdot)$ depend on c only through F_{U-c}, so does $\mu(\eta_E^*) - \eta_E^*$. Thus given any fixed F_{U-c}, part (a) of the proposition applies when c is sufficiently large, while part (b) applies otherwise.

Part (a) of Proposition 2.4.3 implies that if c is sufficiently large, then the equilibrium outcome will be the same for the two types of goods. This is intuitive since when c is large, the post-search information revealing will make no difference to the consumer's optimal decisions as I have discussed below Lemma 2.4.1. The discrepancy between the two types of goods should hence disappear.

When $c < \mu(\eta_E^*) - \eta_E^*$, the post-search information revealing does make a difference. In this case, part (b) of Proposition 2.4.3 shows that the order between η_S^* and η_E^* is decided by the sign of $\Gamma(\eta_E^*)$. To understand this, notice that according to my discussion in Section 2.3.3, $\Gamma(\eta_E^*)$ reflects the SG seller's profit from marginally changing the threshold from η_E^*. If $\Gamma(\eta_E^*) > 0$, a decrease will be profitable; if $\Gamma(\eta_E^*) < 0$, an increase will be profitable. The proof of Proposition 2.4.3(b) shows that this local analysis can be extended globally, which implies the result.

Figure 2.3 illustrates each case of Proposition 2.4.3(b). Notice when $c < \mu(\eta_E^*) - \eta_E^*$, Lemma

2.4.1 implies $z(\eta_E^*) > \mu(\eta_E^*) > \eta_E^* + c$.[25] Thus $\Gamma(\eta_E^*)$ equals to the difference between the slope of the secant of J_p over $[\eta_E^*, \eta_E^* + c]$ and the slope of J_p at $z(\eta_E^*)$, which generally differs from $\Gamma^E(\eta_E^*)$ $(= 0)$. In the figure, we can see that if J_p shifts from J_p^1 to J_p^2 and then to J_p^3, $\Gamma(\eta_E^*)$ will change from negative to zero and then to positive, which correspond to the three cases in Proposition 2.4.3(b). This in particular suggests that given the portion of J_p over $[\eta_E^*, \mu(\eta_E^*)]$ fixed, if its slope (density) turns low sufficiently fast after $\mu(\eta_E^*)$, then we will have $\eta_S^* < \eta_E^*$. In this case, the post-search information revealing of a search good will "crowd out" the seller's pre-search information provision.

2.4.3 Comparison for the Consumer Welfare

Above analyses have illustrated that the pre-search information in equilibrium may get either better or worse after an experience good is turned into a search good. If it gets better (i.e., $\eta_S^* > \eta_E^*$), all consumers will certainly be better-off due to better information both before and after search. If it gets worse (i.e., $\eta_S^* < \eta_E^*$), however, the welfare impact is unclear.

When Assumption 2.4.1 holds and $\eta_S^* < \eta_E^*$, Figure 2.4 illustrates that consumers with different outside option values will be heterogeneously affected when we change an experience good into a search good. It shows that consumers with relatively large U_0, including those with $U_0 + p \in \big[\mu(\eta_E^*), z(\eta_S^*)\big)$, will be better-off.[26] However, the opposite is true for those with relatively small U_0, including those with $U_0 + p \in (\eta_S^*, \eta_E^* + c]$. For these consumers, the poorer pre-search information will lead to too many additional inefficient searches, which makes them worse-off despite the better post-search information available.

Now, a natural question to ask is whether it is possible for the total consumer welfare to decrease when the product is turned into a search good. Proposition 2.4.4 below provides a sufficient condition for this to happen. For tractability, I still focus on strictly unimodal J_p and use r_p to denote its mode over $[\underline{u} - c, \bar{u} - c]$. I also assume the model primitives (F_U, c, r_p) satisfy the following assumption.

Assumption 2.4.2. (F_U, c, r_p) satisfy:

(1) F_U has a log-concave density f_U over $[\underline{u}, \bar{u}]$.

---

[25] Under the EG seller's optimal signal, the lowest value of U leading to the pooling signal realization is $\eta_E^* + c$, which is strictly less than $\mu(\eta_E^*)$ here. Thus $z(\eta_E^*) > \mu(\eta_E^*)$ by Lemma 2.4.1.

[26] I note that in the current situation, although $\eta_S^* < \eta_E^*$, we must have $z(\eta_S^*) > \mu(\eta_E^*)$ for the seller's design to be optimal. This means that search given the pooling signal realization must be more attractive when the product is a search good. Its formal proof is provided in Appendix 2.C.4.

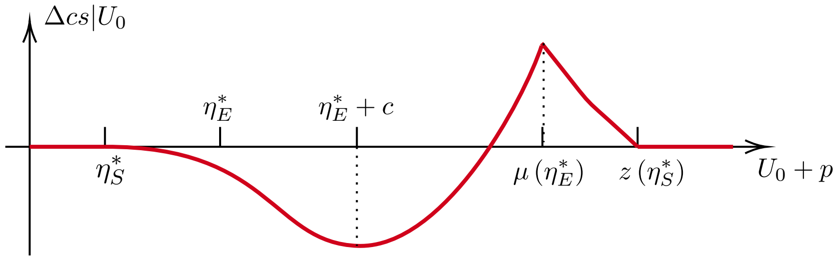

Figure 2.4: When Assumption 2.4.1 holds and $\eta_S^* < \eta_E^*$, the red curve represents the surplus change of consumers with different $U_0 + p$ when the product is turned from an EG to an SG. See Appendix 2.C.4 for the proof of its main qualitative properties.

(2) $c < \mu(r_p) - r_p$.

(3) $z(\underline{u} - c) < r_p < \bar{u} - c$

Condition (1) is a mild condition on F_U that is satisfied by many common distributions. Condition (2) rules out the case of Proposition 2.4.3(a) so that the discrepancy between the two types of goods truly matters. Condition (3) simplifies the discussion by ruling out corner solutions for η_E^* and η_S^*.

Let cs^E and cs^S denote the equilibrium consumer surplus with experience goods and search goods respectively. Given any (F_u, c, r_p) satisfying Assumption 2.4.2, the following proposition provides a condition regarding J_p's behavior on the two sides of its mode r_p that guarantees $cs^S < cs^E$. As before, I define η_0 as the smallest η such that $z(\eta)$ surpasses r_p, i.e., $\eta_0 := \inf\{\eta \in [\underline{u} - c, \bar{u} - c] : z(\eta) \geq r_p\}$.

Proposition 2.4.4. *For any (F_u, c, r_p) satisfying Assumption 2.4.2 and constant $\kappa > 0$, there exists $\nu > 0$ such that: $cs^S < cs^E$ if J_p satisfies Assumption 2.3.2 with its mode being r_p and (i) $j_p(r_p + \nu) < j_p(\eta_0)$; (ii) $j_p(\eta_0) > \kappa \frac{J_p(\bar{u}-c) - J_p(\underline{u}-c)}{\bar{u}-\underline{u}}$.*

The proposition implies that given any (F_u, c, r_p) satisfying Assumption 2.4.2, we can find a pair of κ and ν such that $cs^S < cs^E$ as long as J_p is strictly unimodal and satisfies conditions (i) and (ii). To understand these conditions, notice $\eta_0 < r_p < r_p + \nu$. Thus condition (i) holds when j_p drops sufficiently fast to the right of r_p versus to the left of r_p. This guarantees that the shape of J_p is more aligned with J_p^3 in Figure 2.3 than with J_p^1, which leads to $\eta_S^* < \eta_E^*$. Actually, when ν is chosen to be small enough, the condition ensures that η_S^* will be considerably less than η_E^*, which makes the deterioration of pre-search information significant. Condition (ii) holds when j_p is sufficiently "fat" to the left of its mode. It guarantees that there will be a significant

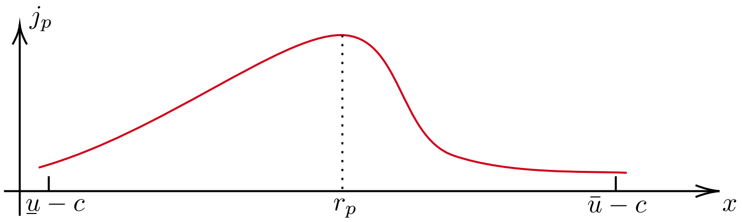

Figure 2.5: Shape of j_p tending to induce $cs^S < cs^E$.

proportion of consumers with $U_0 + p$ lying in the region of $[\eta_S^*, \eta_E^* + c]$. As I have shown in Figure 2.4, such consumers have a net loss when the product changes into a search good. We thus have $cs^S < cs^E$ when their proportion is large. Overall, Proposition 2.4.4 suggests that the total consumer surplus tends to be lower in the search goods case when j_p has a moderately sloped "left-hillside" and a steep "right-hillside" relative to its mode, as is illustrated in Figure 2.5.

My analyses have shown that policies turning experience goods into search goods may crowd out the seller's pre-search information provision and unintentionally reduce the total consumer welfare. For such policies to truly benefit consumers, we may thus need to provide extra incentives to sellers for offering pre-search information. In the context of online shopping, for example, one possibility is to require the sellers to afford some shipping cost for returned products.[27] This makes wasteful searches also costly for the sellers and can thus incentivize them to provide better pre-search information.

2.5 Comparative Statics

2.5.1 Discriminatory Information Provision

In this subsection, I consider comparative statics about the seller's optimal information provision with respect to the consumer's outside option value distribution. The analysis reveals how the seller may want to tailor different pre-search signals to different consumer groups, which forms a kind of discriminatory information provision. Because I will maintain the strictly unimodal assumption on the distribution of U_0, Proposition 2.3.2 implies that I can focus on comparative statics regarding the threshold of the seller's optimal upper-censorship signal, which I denote as η^*. The following proposition provides such a result for a non-pricing seller.

Proposition 2.5.1. *Consider a non-pricing seller and assume J_p remains in the family of*

[27] European Union indeed has such a policy (see footnote 22).

distributions satisfying Assumption 2.3.2. Then η^* (weakly) increases when J_p increases in the likelihood-ratio order.

The proposition suggests that it is optimal for a non-pricing seller to provide better pre-search information when the distribution of $U_0 + p$ is higher (in the likelihood-ratio order). If the seller can distinguish consumers or consumer groups with different outside option value distributions, this implies that he tends to impart different information to them accordingly.[28]

The analysis can also be extended to pricing sellers in a more concrete setting. Specifically, assume that the consumer's outside option value now consists of two parts: $U_0 = W + \epsilon$, where W and ϵ are independent random variables. Let f_W and f_ϵ denote the densities of W and ϵ respectively. If the seller can observe W and tailor pre-search information and price accordingly, then we have the following result on the seller's discriminatory information provision:

Proposition 2.5.2. *Assume f_ϵ is strictly log-concave. For both pricing and non-pricing sellers, if he can observe W, then η^* (weakly) increases in W.*[29]

The proposition suggests that it is optimal for the seller to provide better information to consumers with higher average outside option value. For an intuition on this, notice that conditional on higher $W + p$, the distribution of $U_0 + p$ is located more to the right and is thus convex over a larger portion of $[\underline{u} - c, \bar{u} - c]$, which makes the trade-off mentioned below Proposition 2.3.1 favor a higher η^*. This implies the stated result for a non-pricing seller. For a pricing seller, the proof shows that although the seller may want to charge a lower p when W is higher, optimality requires $W + p$ to be increasing in W. We thus also have η^* to be increasing in W.

The corollary below shows how discriminatory information provision affects consumer welfare. For simplicity, it only considers the case of a non-pricing seller. This shuts down the traditional effect of price discrimination and allows us to solely focus on informational discrimination.

Corollary 2.5.1. *Consider a non-pricing seller and assume:*

(A1) f_W is log-concave and f_ϵ is strictly log-concave;

---

[28] Another implication of Proposition 2.5.1 is that a non-pricing seller tends to provide better information when p is higher. Formally, if U_0 has a log-concave density, then $U_0 + p$ will increase in p in the likelihood-ratio order. The proposition then implies η^* to be increasing in p.

[29] This is in the strong set order if we have a pricing seller and the optimal design is not unique.

(A2) $\arg\max_x\{\int f_W(x-y)f_\epsilon(y)dy\}$ *is a singleton.*[30]

Then there exists w^* such that when the seller can discriminate based on W, all consumers with $W < w^*$ becomes (weakly) worse-off and all consumers with $W > w^*$ becomes (weakly) better-off, compared to the case without discrimination.

The corollary suggests that the interests of consumers with different W are generally not aligned. Those with higher W are more likely to benefit from discriminatory information provision, while the opposite is true for those with lower W. A particular implication of this is that a consumer with high W may have incentive to voluntarily share it with the seller (e.g., by choosing a low privacy setting). Since this will help the seller to also identify those with lower W, it can lead to unraveling of W and harm consumers with lower average outside option values.[31]

As the discriminatory information provision has different welfare implications for different consumers, a natural question is then what its effect on total consumer welfare is. The answer to this is ambiguous and generally depends on the curvatures of the value distributions. If the CDF of $W+\epsilon$ is convex over the support of $U-c$, the seller will provide full information without discrimination. Hence allowing discrimination will merely harm consumers with low W. If the CDF of $W+\epsilon$ is concave over $[\underline{u}-c, \bar{u}-c]$, the seller will provide no information without discrimination. Then discrimination will only help by allowing some consumers with high W to receive better information.

I conclude this subsection with a note on the related literature. Since higher outside option value is equivalent to lower inside option value, the result in Proposition 2.5.2 can also be interpreted as that a seller with lower product quality tends to provide better pre-search information. Anderson & Renault (2013) draw a similar lesson in their Appendix B.[32] My result extends and refines theirs in two ways. First, I allow for ex-ante heterogeneous consumers; second, I provide a detailed comparative statics on the optimal signal threshold, while they only concern whether any threshold information will be provided.

[30] These are satisfied, for example, when W and ϵ are normally distributed.

[31] Such an effect is well-recognized for price discrimination (e.g., Acquisti et al. (2016) page 453).

[32] Also related are earlier papers including Lewis & Sappington (1994), Sun (2011) and Bar-Isaac et al. (2010). These papers only consider experience goods and restricted information structures.

2.5.2 The Effect of Changing Search Cost

Technological developments can significantly reduce a consumer's cost in accessing products. For examples, modern means of transportation make it easier to visit physical stores; advanced logistics systems make product shipping cheaper in online shopping; faster internet connection reduces the time needed for loading digital contents. These all help to reduce a consumer's search cost in various contexts.

The following proposition describes the effect of changing search cost with a non-pricing seller.

Proposition 2.5.3. *Consider a non-pricing seller and assume the following:*

(A1) J_p *admits a continuous density that is strictly quasi-concave over* $(-\infty, \bar{u}]$.[33]

(A2) $1 - F_u$ *is log-concave.*

Then, when the search cost c *decreases, we have:*

(a) The seller's optimal upper-censorship signal will become less informative.

(b) A consumer with any outside option value will become better-off.

When the product price is fixed and the regularity conditions hold, the proposition shows that the consumer welfare will unambiguously increase when the search cost drops, although this gain from lower search cost will be partially offset by coarser pre-search information. This holds not only for consumers as a whole, but also for each individual consumer given her realized outside option value.

A natural question is how the result will change if we have a pricing seller who can adjust his price accordingly when the search cost changes. The general answer to this is ambiguous, which depends on the range of search cost being considered and the detailed shapes of the value distributions. In particular, lower search cost may lead to better pre-search information but lower total consumer welfare because the seller charges a higher price.[34]

[33] If one considers search costs less than some value $\bar{c}$, this only needs to hold over $[\underline{u} - \bar{c}, \bar{u}]$.

[34] As has been noticed by Anderson & Renault (2006), when the search cost decreases, the seller's optimal price can increase faster than how the search cost changes. With ex-ante heterogeneous consumers in my setting, this especially hurts the "infra-marginal" consumers with low outside option values and can thus lead to lower consumer surplus. A concrete example is available upon request.

2.6 Information Provision by Competing Sellers

In this section, I turn to consider multiple sellers and study the effect of competition on equilibrium information provision.

2.6.1 The Model

Consider a discrete choice model with horizontally differentiated products. There are N sellers, each of whom sells a product with zero marginal cost. I call the product sold by seller i as product i. There is a (representative) consumer, whose match value with product i is denoted as U_i. Assume that $\{U_i\}_{i=1}^N$ are independently drawn from a common distribution F_U with compact support $[\underline{u}, \bar{u}]$. The consumer can consume one product and is assumed to have no outside option for simplicity.

Because it is not harder to consider pricing sellers than non-pricing sellers, I will assume all sellers are pricing sellers. The analysis for non-pricing sellers is analogous and will lead to the same qualitative result. The game proceeds as follows:

1. Each seller i decides his price p_i and designs a pre-search signal about U_i. These are done by all sellers simultaneously and observed by the consumer.
2. $\{U_i\}_{i=1}^N$ are (secretly) realized and the pre-search signal realizations $\{S_i\}_{i=1}^N$ are generated to the consumer.
3. After observing $\{S_i\}_{i=1}^N$, the consumer starts a sequential search process as in Weitzman (1979). Specifically, at any time she decides whether to keep searching. If so, she chooses one of the products to search, pay the search cost c, and then learns the product's match value. If not, she chooses one of the products that have been searched before for purchase, and the game ends.[35]

I will call the setting above the SG-environment. As in Definition 2.2.1, I will say a product is an experience good (EG) if the consumer will receive no additional information after searching it. The environment is an EG-environment if all products are experience goods and everything else is the same as above. In particular, the consumer still needs to pay the search cost c in order to buy any product.

[35] As is standard in the literature, I assume that a product must be searched before purchase and it is free for the consumer to recall an earlier searched product for purchase.

2.6.2 Equilibrium Definitions

For any i, I use Z_i to denote the Pandora's index of the consumer's posterior belief on U_i after observing S_i. The following lemma characterizes the consumer's consumption outcome, where part (a) has been discovered in several papers (Kleinberg et al., 2016; Armstrong, 2017; Choi et al., 2018) and part (b) is trivial to see.

Lemma 2.6.1. *We have the following results:*

(a) In the SG-environment, the consumer will buy from a seller with the highest $U_i \wedge Z_i - p_i$. Her expected surplus is $\mathbb{E}[\max_i\{U_i \wedge Z_i - p_i\}]$.

(b) In the EG-environment, the consumer will buy from a seller with the highest $\mathbb{E}[U_i - c|S_i] - p_i$. Her expected surplus is $\mathbb{E}[\max_i\{\mathbb{E}[U_i - c|S_i] - p_i\}]$.

Let G_i and H_i denote the distributions of $U_i \wedge Z_i$ and $\mathbb{E}[U_i - c|S_i]$ respectively. The lemma implies that in the SG-environment (resp., EG-environment), $\{(G_i, p_i)\}_{i=1}^N$ (resp., $\{(H_i, p_i)\}_{i=1}^N$) is sufficient in determining the equilibrium outcomes. We can thus equivalently consider each seller i's decision as choosing the pair of (G_i, p_i) in the SG case and choosing (H_i, p_i) in the EG case.

Given $\{(G_k, p_k)\}_{k\neq i}$ with $\{G_k : k \neq i\}$ being continuous, seller i's best response problem in the SG-environment can be written as:

$$\max_{G_i, p_i}\{p_i \int \prod_{k\neq i} G_k(x - p_i + p_k) dG_i(x)\} \tag{2.15}$$

$$\text{s.t. } G_i \text{ is a feasible distribution of } U_i \wedge Z_i \tag{2.16}$$

Notice the integration in the objective function measures $\mathbb{P}(U_i \wedge Z_i - p_i \geq U_k \wedge Z_k - p_k, \forall k)$, which equals to the sale probability of seller i by Lemma 2.6.1.

Similarly, given $\{(H_k, p_k)\}_{k\neq i}$ with $\{H_k : k \neq i\}$ being continuous, seller i's best response problem in the EG-environment can be written as:

$$\max_{H_i, p_i}\{p_i \int \prod_{k\neq i} H_k(x - p_i + p_k) dH_i(x)\} \tag{2.17}$$

$$\text{s.t. } H_i \preceq_{MPS} F_{U-c} \tag{2.18}$$

I define the equilibria with the two types of goods as follows.

Definition 2.6.1 (SG-equilibrium)**.** A pair of (G^*, p^*) is a *symmetric SG-equilibrium* if G^* is continuous and (G^*, p^*) solves optimization (2.15) – (2.16) given $(G_k, p_k) = (G^*, p^*)$ for all $k \neq i$.

Definition 2.6.2 (EG-equilibrium)**.** A pair of (H^*, p^*) is a *symmetric EG-equilibrium* if H^* is continuous and (H^*, p^*) solves optimization (2.17) – (2.18) given $(H_k, p_k) = (H^*, p^*)$ for all $k \neq i$.

I note that in the above definitions, I have focused on equilibria where G^* or H^* is continuous. This avoids the discussion of ties in the consumer's choices, and also suffices for presenting my results below.[36]

2.6.3 Equilibrium with a Large Number of Sellers

I assume the following regularity condition holds:

Assumption 2.6.1. F_U admits a log-concave density f_U, which is continuously differentiable with bounded derivative over $[\underline{u}, \bar{u}]$. Moreover, f_U is strictly positive over $[\underline{u}, \bar{u}]$.[37]

Under Assumption 2.6.1, Theorem 3 in Hwang et al. (2019) shows that there is a unique symmetric EG-equilibrium. Given any number of competing sellers N, I denote this equilibrium as (H_N^{ex}, p_N^{ex}). My main result in this section is the following:

Proposition 2.6.1. *Under Assumption 2.6.1, there exists $N^* < \infty$ such that for all $N \geq N^*$, $(G^*, p^*) = (H_N^{ex}, p_N^{ex})$ is a symmetric SG-equilibrium when there are N competing sellers.*

The proposition suggests that when the number of competing sellers is sufficiently large, the distribution of $U_i \wedge Z_i$ in the SG-environment will be the same as the distribution of $\mathbb{E}[U_i - c|S_i]$ in the EG-environment, and the product prices in the two environments will also be the same. By Lemma 2.6.1, this implies that the purchase outcome, sellers' profits and consumer welfare will all be the same in the two environments. Actually, the proposition's proof shows that when N is large, any equilibrium profile of signals in the EG-environment also serves as an equilibrium profile of signals in the SG-environment. Under these signals, $Z_i \leq U_i$ almost surely for all i,

[36] Requiring H^* to be continuous in the EG-equilibrium is actually without loss of generality because Hwang et al. (2019) shows that it holds in any symmetric equilibrium with EG.

[37] The assumption that f_U is strictly positive over $[\underline{u}, \bar{u}]$ is not crucial for my results. However, it allows me to directly refer to certain results in Hwang et al. (2019), who make such kind of assumption. Details on this are available upon request.

and thus the consumer will always purchase from the first seller being searched by the Pandora's rule. The discrepancy between the two types of environments therefore vanishes when there is sufficient competition among sellers.

In the EG-environment, Hwang et al. (2019) shows that the equilibrium information will converge to full information when the number of competing sellers goes to infinity. The following corollary of Proposition 2.6.1 extends this result to the SG-environment.

Corollary 2.6.1. *Assume Assumption 2.6.1 holds. For any* $\epsilon > 0$*, there exists* $N_\epsilon < \infty$ *such that for all* $N \geq N_\epsilon$*: there exists a symmetric SG-equilibrium in which* $G^*(x) = F_{U-c}(x)$ *for all* $x \leq \bar{u} - c - \epsilon$*.*

The corollary implies that when the number of competing sellers goes to infinity, the equilibrium distribution of $U_i \wedge Z_i$ will converge to F_{U-c} and hence the equilibrium pre-search information will converge to full information. This suggests that strong competition leads to informational efficiency also for search goods.

I note that Proposition 2.6.1 and Corollary 2.6.1 only consider one symmetric equilibrium for search goods. Although I conjecture that the symmetric equilibrium will be unique when N is large, showing this formally is beyond the scope of this paper.[38] The main difficulty is that without a complete and simple characterization for the feasible set of G_i, it is hard to derive meaningful necessary optimality conditions for the best response problem (2.15) – (2.16), which are needed to establish the equilibrium uniqueness.

2.6.4 Discussion: Equilibrium with a Small Number of Sellers

I have studied the case with a monopoly seller and the case with a large number of competing sellers. The remaining question is what will happen with a small number of sellers. One may guess that the equilibrium signal informativeness will be in the middle of the two extreme cases. Unfortunately, however, a formal analysis for this is difficult.

Similar to what has been done in the proof of Proposition 2.6.1, one tentative approach is to first consider a "relaxed game", where each seller chooses G_i only subject to the two constraints in the Relaxed Problem (2.4) – (2.6). After deriving the equilibrium for this relaxed game, one can check whether the equilibrium G_i can be induced by some signal. If so, this will indeed

[38] I have the conjecture because of the similarity between the two types of goods and that we do have equilibrium uniqueness in the EG-environment.

be an equilibrium of the original game. The main difficulty is that there is no easy sufficient condition to show a distribution is inducible for $U_i \wedge Z_i$. Throughout the paper, I have been showing this by explicitly constructing the underlying signal. However, such a construction is typically difficult when G_i has a degree of complexity that one must face in the competing environment. I thus leave this analysis for future studies.

2.7 Additional Discussions

2.7.1 Comparison with Anderson & Renault (2006)

As is mentioned in the introduction, Anderson & Renault (2006) studies a similar setting as mine while assuming that the consumers are ex-ante homogeneous without private information. Under this assumption, they solve the optimal signal and show that it suffices for the seller to provide *threshold information*, which just informs the consumer whether her match value is above certain threshold.

My analyses in Section 2.3 partially extend their result that the optimal design involves threshold information to the situation where the consumers' outside option values are private and admit a unimodal distribution. Unlike in Anderson & Renault (2006), however, the optimal signal in my setting typically also requires fully revealing match values below the threshold, especially when the distribution of U_0 is strictly unimodal. To understand this discrepancy, let U^* denote the threshold match value under the optimal design. In Anderson & Renault (2006), U_0 is common for all consumers and we have $U_0 > U^* - c - p$ for sure. Thus once a consumer has learned her match value is below U^*, she will not search the product anyway. Providing additional below-threshold information is therefore useless. In my setting, however, U_0 has a non-degenerate distribution and typically some consumers will have $U_0 < U^* - c - p$. Providing additional below-threshold information allows these consumers to tell how far their match values are below U^*, and can thereby alter their decisions. If the CDF of U_0 is strictly convex below $U^* - c - p$, which is always the case under the strict unimodal assumption on U_0, this will indeed increase these consumers' total purchase probability in light of my discussion right below Proposition 2.3.1. Thus finer below-threshold information can attract more consumers with relatively low outside option values into purchase.

2.7.2 Incentive for Multi-Stage Information Provision

I have been studying situations where the seller controls the consumer's pre-search information, while many earlier studies reviewed in the introduction have focused on cases where the seller controls the consumer's post-search information. A natural question is then what if the seller can control information at both stages.

There is a simple answer to this question: the seller never benefits from multi-stage information provision. As is mentioned in Section 2.4, an SG seller's pre-search design problem is more constrained than the EG seller's. This implies that the seller is weakly better-off when the consumer cannot receive additional information after search. Thus even if the seller can control both pre-search and post-search information, it is optimal for him to only provide information at the pre-search stage. The equilibrium outcome will hence be identical to that with an EG seller.

2.7.3 Incentive to Subsidize Search

I have been assuming that the seller cannot incentivize consumer search by directly subsidizing it. An interesting question is how the equilibrium outcome may change if we allow such subsidization. When the distribution of outside option value is unimodal and the seller is a pricing seller, a short answer to this is available: the seller has no incentive to subsidize search, and thus the equilibrium outcome will remain unchanged. A detailed analysis for this is provided in Appendix 2.A.3.

2.7.4 Multimodal Outside Option

Most of my analyses have relied on the unimodal assumption on the consumer's outside option value distribution. What if we go beyond this assumption? Although my relaxed problem approach is not guaranteed to find the seller's optimal design in that case, one can still try the following procedure: (1) solving the Relaxed Problem (2.4) – (2.6); (2) constructing a signal that implements the Relaxed Problem's solution. If the second step is feasible, then the signal constructed will be optimal. In Appendix 2.A.4, I discuss when this procedure works and provide an analytical example. It is shown that the optimal signal can involve multiple pooling and revealing intervals, which is more complex than under the unimodal assumption.

2.8 Concluding Remarks

In this paper, I have examined the optimal pre-search information provision by seller(s) for search goods. The fact that the consumer can inevitably learn her match value after search significantly complicates the design problem compared to its experience goods counterpart. To overcome the challenge, a relaxed problem approach has been developed, which not only solves the optimal design under certain regularity conditions while allowing for ex-ante heterogeneous consumers, but also brings out the key similarity and dissimilarity between information designs for the two types of goods.

The optimal design is fully characterized when the consumer's outside option value distribution is unimodal, which features a simple upper-censorship structure. This partially extends the result in Anderson & Renault (2006) that the optimal design involves threshold information, but also contrasts with it by showing that the optimal design typically also requires fully revealing below-threshold information. Based on the main characterization, several applications are further considered. These include studies on (1) comparison between search goods and experience goods, (2) discriminatory information provision, (3) the effect of reduction in search cost and (4) competition by multiple sellers. In some of these applications, my approach has allowed existing results of experience goods to be qualitatively extended to search goods.

Many research questions still remain open. Examples include how to solve the optimal design with multimodal outside option in general, and how to characterize the equilibrium with a small number of competing sellers. These probably require sharper characterizations for the feasible set of the effective-search-value distribution. My lemmas in Section 2.3.1 and the relaxed problem approach can be considered as the starting point for this more general research agenda.

The paper's setting can also be generally understood as having a sender persuading a receiver to take certain action that involves two costly steps, after the first of which the receiver will inevitably learn additional information and can choose to quit. This kind of environment is very common in practice. For example, one can consider an entrepreneur persuading an investor to finance a project involving two phases, before the second of which the investor can choose whether to continue funding based on the information revealed in the first phase. One may find my approach also useful in such applications.

Appendix

2.A Additional Results and Discussions

2.A.1 Tentative existing approaches to the seller's problem

As is mentioned in the introduction, the first moment of the consumer's posterior belief does not suffice for deciding the purchase outcome. In fact, with continuous U_0, the sender's objective function must depend on infinitely many moments of the posterior belief. This makes the approaches in Dworczak & Martini (2019) and Dworczak & Kolotilin (2019) hard to apply. Another tentative approach is the concavification method in Kamenica & Gentzkow (2011). As is well known, this method is hardly tractable unless U takes no more than three values, which is very restrictive in the current context. In particular, the clean signal structures derived in this paper and Anderson & Renault (2006), like the threshold structure, would not be available if U only takes a few discrete values. A third approach is the linear programming approach in Kolotilin (2018), which handles receiver private information. With this approach, one will have to consider an optimization over the joint distribution of (U, Z), which turns out to be very challenging. A fourth approach is to draw upon a result in Guo & Shmaya (2019) and transform the problem into a nonlinear optimization over some thresholds as in their Section 4. Unless one assumes U_0 takes only finite values, the optimization will be over infinite-dimensional functions. Even if U_0 takes finite values, it is still hard to derive semi-analytical solution for such a non-linear optimization.

2.A.2 SG vs. EG with a Pricing Seller

In Section 2.4, I have shown that with a non-pricing seller, turning an experience good into a search good may unintentionally crowd-out the seller's pre-search information provision and decrease the total consumer welfare. With a numerical example, I show that these can also happen with a pricing seller below.

Assume $c = 0.1$, $U \sim \text{Uniform}[0, 1]$ and U_0 has an *approximate skew-normal distribution* (Ashour & Abdel-hameed, 2010). Specifically, U_0 has density:

$$j(x) = 2\phi\big(\frac{x-\xi}{0.3}\big)F\big(-8 \times \frac{x-\xi}{0.3}\big)$$

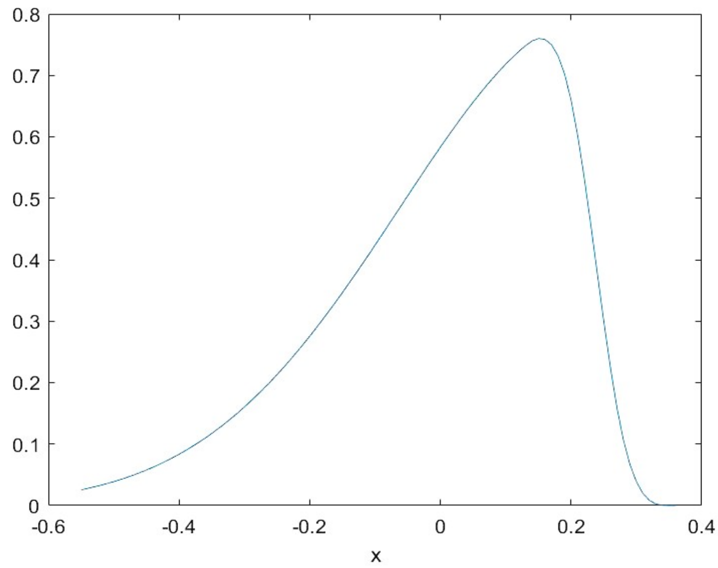

Figure 2.6: $j(\cdot)$ in the example of Appendix 2.A.2.

Product type	η^*	**Equilibrium Price**	**Consumer Surplus**
Experience Goods	0.284	0.387	0.0400
Search Goods	0.273	0.414	0.0366
% Change	-3.89	6.93	-8.33

Table 2.1: Equilibrium outcomes in the example of Appendix 2.A.2.

where ϕ is the standard normal density; F is defined in equation (2.1) in Ashour & Abdel-hameed (2010); ξ is chosen such that the mean of U_0 is normalized to zero. The density of U_0 is plotted in Figure 2.6.

The comparison for equilibrium outcomes of the two types of goods is summarized in Table 2.1. As one can see, when the product turns into a search good, the threshold of the seller's optimal upper-censorship signal (η^*) and the total consumer surplus both drop. Moreover, the consumer welfare decreases due to not only poorer pre-search information, but also higher product price.

2.A.3 Incentive for Subsidizing Search

In this appendix, I show a pricing seller has no incentive to subsidize the consumer's search. Given any original product price, consider two strategies of the seller:

A. Offer a subsidy $y > 0$ per search.

B. Give no search subsidy, but offer a price discount y for the product.

Without loss of generality, assume y is less than the original product price. Then, we have the following observation:

Proposition 2.A.1. *Assume J is unimodal (i.e., Assumption 2.3.3 holds). A pricing seller's maximal profit under strategy A is lower than that under strategy B.*[39]

The proposition implies that subsidizing for search is always an inferior option for the seller than directly lowering the product price. Two intuitions are underlying this result. First, given any purchase probability, strategy A is more costly to implement than strategy B since it offers subsidy not only to those who purchase, but also to those who search without purchase. Second, by making search cheaper, strategy A weakens a consumer's reliance on pre-search information. This makes it less effective to manipulate consumer behavior by controlling the pre-search information and thus decreases the maximal purchase probability.

2.A.4 Multimodal Outside Option

When J_p is multimodal over $[\underline{u}-c, \bar{u}-c]$, my results about the seller's optimal design in Section 2.3.3 does not apply. In this case, one can try the following approach for solving the non-pricing seller's optimal signal: (1) solving the Relaxed Problem (2.4) – (2.6); (2) constructing a signal that implements the Relaxed Problem's solution. If the second step is possible, then the signal constructed will be optimal.

Let G^*_{RP} denote a solution to the Relaxed Problem. Let $\hat{z}(a,b)$ be the Pandora's index of the posterior belief on U after learning $U-c \in [a,b]$. An important case where the above procedure works is when G^*_{RP} features a *monotone-partitionable* structure.

Definition 2.A.1. Distribution G is called monotone-partitionable if there exists $(\{s_i\}_{i=0}^n, \{\eta_i\}_{i=0}^n)$ such that $\underline{u}-c = s_0 < ... < s_n = \bar{u}-c$; $s_i \le \eta_i \le s_{i+1}$ $\forall i$; and

$$G(x) = \begin{cases} F_{U-c}(x) & s_i \le x < \eta_i \\ F_{U-c}(\eta_i) & \eta_i \le x < (\eta_i + c) \wedge \hat{z}(\eta_i, s_{i+1}) \\ F_U(x) & (\eta_i + c) \le x < \hat{z}(\eta_i, s_{i+1}) \\ F_{U-c}(s_{i+1}) & \hat{z}(\eta_i, s_{i+1}) \le x \le s_{i+1} \end{cases} \tag{2.19}$$

[39] This holds strictly as long as there will be some search without purchase under strategy A.

With a similar argument as in the proof of Lemma 2.B.2 (in Appendix 2.B.3), it is easy to see that if G^*_{RP} is monotone-partitionable, it can be induced by a *monotone-partitioning signal* that fully reveals $U - c \in [s_i, \eta_i]\,\forall i$, and pools $U - c \in [\eta_i, s_{i+1}]$ separately for each i. Intuitively, such a signal first partitions net-match-utilities into subintervals $\{[s_i, s_{i+1}]\}_{i=1}^n$, and then imposes an upper-censorship signal on each of these subintervals. Thus, if the Relaxed Problem turns out to have a monotone-partitionable solution, one can conclude that it is indeed inducible and the optimal signal can be the monotone-partitioning signal described above.

In the rest of this appendix, I consider a particular situation where the Relaxed Problem indeed features a monotone-partitionable solution. This both illustrates how the procedure above works and shows that when J_p is not unimodal, the optimal signal can involve more complicated structure than upper-censorship.[40]

Assume the following assumption holds:

Assumption 2.A.1. J_p admits a continuous density j_p over $[\underline{u} - c, \bar{u} - c]$. There exist r_1 and r_2 ($r_1 < r_2$) in $(\underline{u} - c, \bar{u} - c)$ such that $j_p(\cdot)$ is strictly increasing on $[\underline{u} - c, r_1]$ and $[r_2, \bar{u} - c]$, and is strictly decreasing on $[r_1, r_2]$.

Further assume for now that there exists $\eta^* \in [\underline{u} - c, r_1]$ and $\xi^* \in [r_2, \bar{u} - c]$ such that $\hat{z}(\eta^*, \xi^*) \in [r_1, r_2]$ and:

$$\frac{J_p\big((\eta^* + c) \wedge \hat{z}(\eta^*, \xi^*)\big) - J_p(\eta^*)}{(\eta^* + c) \wedge \hat{z}(\eta^*, \xi^*) - \eta^*} = j_p\big(\hat{z}(\eta^*, \xi^*)\big) = \frac{J_p(\xi^*) - J_p(\hat{z}(\eta^*, \xi^*))}{\xi^* - \hat{z}(\eta^*, \xi^*)} \tag{2.20}$$

Figure 2.7 plots J_p satisfying Assumption 2.A.1 with (η^*, ξ^*) satisfying condition (2.20). Intuitively, the condition requires that the two black secants of J_p plotted in the graph have the same slope, which is equal to the slope of J_p at $\hat{z}(\eta^*, \xi^*)$.

Then a solution to the Relaxed Problem (G^*_{RP}) turns out to be monotone-partitionable as in Definition 2.A.1, with $n = 2$, $s_0 = \underline{u} - c$, $\eta_1 = \eta^*$, $s_1 = \xi^*$ and $\eta_2 = s_2 = \bar{u} - c$. As the discussion earlier suggests, this G^*_{RP} can be induced by a monotone-partitioning signal that fully reveals net-match-utilities in $[\underline{u} - c, \eta^*] \cup [\xi^*, \bar{u} - c]$ and pools the rest in the middle. Thus G^*_{RP} is indeed feasible for the seller, and the signal inducing it is optimal.

To check that G^*_{RP} suggested above indeed solves the Relaxed Problem, one can apply Theorem 2.3.1. Namely, one can find a convex function $v(\cdot)$ and an increasing function $\rho(\cdot)$

[40] More examples are available upon request.

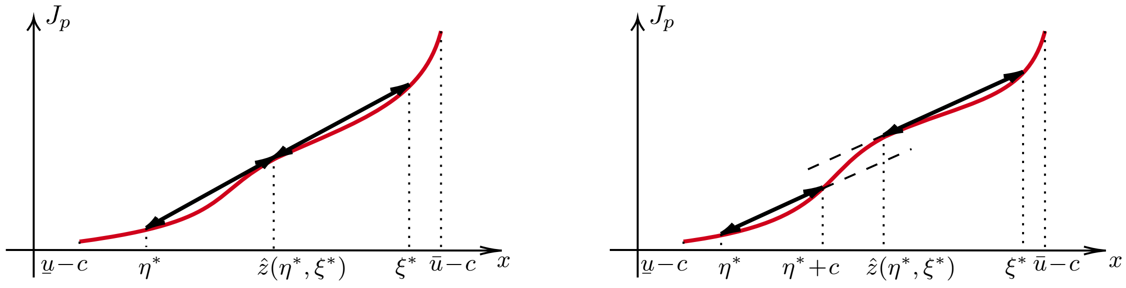

(a) Case 1: $\hat{z}(\eta^*,\xi^*) \leq \eta^* + c$ (b) Case 2: $\hat{z}(\eta^*,\xi^*) > \eta^* + c$

Figure 2.7: $J_p(\cdot)$ satisfying Assumption 2.A.1 with (η^*,ξ^*) satisfying (2.20). The red curve is J_p; the black segments are secants of J_p with common slope $j_p(\hat{z}(\eta^*,\xi^*))$.

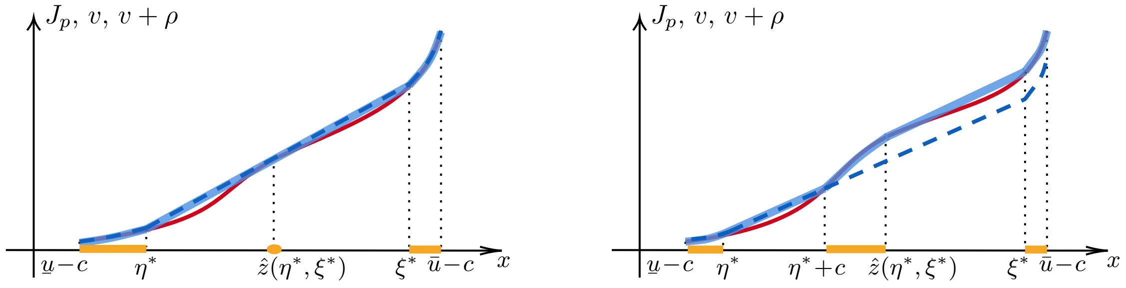

(a) Case 1: $\hat{z}(\eta^*,\xi^*) \leq \eta^* + c$ (b) Case 2: $\hat{z}(\eta^*,\xi^*) > \eta^* + c$

Figure 2.8: Graphical Check of Optimality Conditions for (G^*_{RP}, v, ρ). The red curve is J_p; the solid blue curve is $v + \rho$; the dashed blue curve is v; the orange area on x-axis is the support for G^*_{RP}.

such that (G^*_{RP}, v, ρ) satisfies conditions (C1) – (C4) in Theorem 2.3.1. Figure 2.8 gives a graphical illustration for this. The fact that conditions (C1) and (C2) hold is evident from the graph (notice $\rho(\cdot)$ is represented by the difference between the solid blue curve and the dotted blue curve). (C4) holds since G^*_{RP} is indeed inducible and thus satisfies the constraints. To check (C3), notice that: (1) over regions where $v(\cdot)$ is strictly convex, G^*_{RP} equals to F_{U-c}; (2) over the region where $\rho(\cdot)$ is strictly increasing, G^*_{RP} equals to F_U. Given that G^*_{RP} satisfies the constraints, these imply that (C3) indeed holds. (A detailed proof for this is similar to that for Proposition 2.3.1. See Observation (4) in Appendix 2.B.3 (Step 3).)

Up to now, I have been assuming the existence of (η^*,ξ^*) such that $\hat{z}(\eta^*,\xi^*) \in [r_1, r_2]$ and (2.20) holds. Actually, when such η^* and ξ^* do not exist, one can still derive the optimal signal by solving the Relaxed Problem, but just ends up with a simpler signal structure, which can be thought of as a corner solution. In particular, the optimal signal can feature either upper-censorship or lower-censorship (i.e., it pools values below a threshold and fully reveals values above the threshold). In general, as long as Assumption 2.A.1 holds, the Relaxed Problem always has a monotone-partitionable solution, and the seller's optimal signal can be found with the above procedure.

2.B Proofs for Section 2.3

2.B.1 Proof for Lemma 2.3.1

Proof. First notice by the definition of Pandora's index (2.1), we have $\mathbb{E}[(U-Z)_+ - c|S] = 0$ for any signal realization S. Since $(U-Z)_+ = U - U \wedge Z$, we have $\mathbb{E}[U-c|S] = \mathbb{E}[U \wedge Z|S]$ (i.e., the posterior mean of net-match-utility and the effective-search-value are equal).[41]

Next, I show that conditioning on any signal realization S, the conditional distribution of $U \wedge Z$ single crosses that of $U-c$ from below. To see this, given any S, let z_S be the Pandora's index of posterior belief $\phi(\cdot; S)$. Then, for any $u < z_S$, we have:

$$\mathbb{P}(U \wedge Z \leq u|S) = \mathbb{P}(U \wedge z_S \leq u|S) = \mathbb{P}(U \leq u|S) \leq \mathbb{P}(U - c \leq u|S)$$

For any $u \geq z_S$, we have:

$$\mathbb{P}(U \wedge Z \leq u|S) = \mathbb{P}(U \wedge z_S \leq u|S) = 1 \geq \mathbb{P}(U - c \leq u|S)$$

Together with the observation $\mathbb{E}[U-c|S] = \mathbb{E}[U \wedge Z|S]$ above, this single crossing property implies that $U \wedge Z$ is MPC of $U-c$ conditional on S. Then the result that $U \wedge Z$ is MPC of $U-c$ follows because MPC is preserved under mixture. (To see this, just notice for any convex function $v(\cdot)$, $\mathbb{E}[v(U \wedge Z)|S] \leq \mathbb{E}[v(U-c)|S] \Rightarrow \mathbb{E}[v(U \wedge Z)] \leq \mathbb{E}[v(U-c)]$.)

To see $U \wedge Z = U - c$ under fully revealing signal, take $\phi = \delta_U$ into (2.1), where δ_U is the Dirac measure at U. It is easy to derive $z_{\delta_U} = U - c$ then. *Q.E.D.*

2.B.2 Proof for Theorem 2.3.1

Proof. For the first claim (the part before "Moreover"), let distribution G, convex function v and increasing function ρ satisfy all the conditions. Let $\hat{G}$ be another distribution that satisfies

[41] Throughout my proofs, the (in)equalities involving conditional expectation or Z are understood as holding almost surely when appropriate.

constraints (2.5) and (2.6). We have:

$$
\begin{aligned}
&\int J_p(x)dG(x) \\
&= \int [J_p(x) - v(x) - \rho(x)]dG(x) + \int v(x)dG(x) + \int \rho(x)dG(x) \\
&= \int [J_p(x) - v(x) - \rho(x)]dG(x) + \int v(x)dF_{U-c}(x) + \int \rho(x)dF_U(x) \\
&\geq \int [J_p(x) - v(x) - \rho(x)]d\hat{G}(x) + \int v(x)d\hat{G}(x) + \int \rho(x)d\hat{G}(x) \\
&= \int J_p(x)d\hat{G}(x)
\end{aligned}
$$

The second equality is because G satisfies the complementary-slackness condition (C3). For the inequality in fourth line, notice it holds for each term on the two sides. Specifically,

(1) Since $\hat{G}$ is MPC of F_{U-c}, its support is a subset of $[\underline{u} - c, \bar{u} - c]$. Then by condition (C2), we know $\int [J_p(x) - v(x) - \rho(x)]d\hat{G}(x) \leq 0 = \int [J_p(x) - v(x) - \rho(x)]dG(x)$.

(2) v being convex over $[\underline{u} - c, \bar{u} - c]$ and $\hat{G} \preceq_{MPS} F_{U-c}$ imply that $\int v(x)dF_{U-c}(x) \geq \int v(x)d\hat{G}(x)$.

(3) ρ being increasing on $[\underline{u} - c, \bar{u}]$ and $\hat{G} \preceq_{FOD} F_U$ imply that $\int \rho(x)dF_U(x) \geq \int \rho(x)d\hat{G}(x)$.

Thus, the inequality above holds. Since $\hat{G}$ is an arbitrary feasible distribution, I conclude that G is optimal.

For the second claim (the part after "Moreover"), notice for $\hat{G}$ to be another optimal solution, we need the inequality above to hold as equality. Based on observations (1) – (3) above, this requires equality to hold for each term. Specifically, we need:

First, $\int [J_p(x) - v(x) - \rho(x)]d\hat{G}(x) = \int [J_p(x) - v(x) - \rho(x)]dG(x) = 0$. Since $J_p(x) - v(x) - \rho(x) \leq 0$ on the support of $\hat{G}$, this implies condition (C2) holds for $\hat{G}$.

Second, $\int v(x)d\hat{G}(x) = \int v(x)dF_{U-c}(x)$ and $\int \rho(x)d\hat{G}(x) = \int \rho(x)dF_U(x)$. This means that condition (C3) holds for $\hat{G}$.

Thus, the triple $(\hat{G}, v, \rho)$ satisfies (C1) – (C4). *Q.E.D.*

2.B.3 Proof for Proposition 2.3.1

Since $r_p^{max} = \bar{u} - c$ implies that $J_p(\cdot)$ is convex over $[\underline{u} - c, \bar{u} - c]$, part (a) is directly implied by Lemma 2.3.1. I prove part (b) with the following three steps.

Step 1: preliminaries

First, we need the following lemma regarding $z(\cdot)$:

Lemma 2.B.1. *(a)* $\bar{u} - c > z(\eta) > \eta$ *for* $\eta \in [\underline{u} - c, \bar{u} - c)$*;* $z(\eta) = \eta = \bar{u} - c$ *for* $\eta = \bar{u} - c$.

(b) $z(\cdot)$ *is continuous and strictly increasing on* $[\underline{u} - c, \bar{u} - c]$.

Proof. (a) The result for $\eta = \bar{u} - c$ holds obviously by the definition of the Pandora's index. When $\eta < \bar{u} - c$, the definition of the Pandora's index implies $\mathbb{E}[(U - z)_+ - c|U - c \geq \eta]|_{z=z(\eta)} = 0$. In contrast, $\mathbb{E}[(U - z)_+ - c|U - c \geq \eta]|_{z=\eta} = \mathbb{E}[(U - \eta)_+|U - \eta \geq c] - c = \mathbb{E}[U - \eta|U - \eta \geq c] - c > 0$, where the inequality is strict because $\mathbb{P}[U - \eta > c] > 0$; $\mathbb{E}[(U - z)_+ - c|U - c \geq \eta]|_{z=\bar{u}-c} = \mathbb{E}[(U - (\bar{u} - c))_+ - c|U - c \geq \eta] < \mathbb{E}[(\bar{u} - (\bar{u} - c))_+ - c|U - c \geq \eta] = 0$, where the inequality is strict because $\mathbb{P}(\bar{u} > U \geq \eta + c) > 0$. Since $\mathbb{E}[(U - z)_+ - c|U - c \geq \eta]$ is decreasing in z, these imply $\bar{u} - c > z(\eta) > \eta$.

(b) Suppose $\eta_1, \eta_2 \in [\underline{u} - c, \bar{u} - c]$ are such that $\eta_1 < \eta_2$ but $z(\eta_1) \geq z(\eta_2)$. Then by the result of (a), we must have $\eta_1 < \eta_2 < z(\eta_2) \leq z(\eta_1) < \bar{u} - c$. By the definition of the Pandora's index, we have:

$$\mathbb{E}[(U - z(\eta_1))_+ - c|U - c \geq \eta_1] = 0 \tag{2.21}$$

$$\mathbb{E}[(U - z(\eta_2))_+ - c|U - c \geq \eta_2] = 0 \tag{2.22}$$

The later implies that $\mathbb{E}[\big((U - z(\eta_2))_+ - c\big)\mathbb{1}_{\{U-c\geq\eta_2\}}] = 0$, where $\mathbb{1}$ denotes the indicator function. This further implies that $\mathbb{E}[\big((U - z(\eta_2))_+ - c\big)\mathbb{1}_{\{U-c\geq\eta_1\}}] < 0$ because the event $\{\eta_2 > U - c > \eta_1\}$ has strictly positive probability and $(U - z(\eta_2))_+ - c < 0$ on it. Thus, we have $\mathbb{E}[\big((U - z(\eta_2))_+ - c\big)|U - c \geq \eta_1] < 0$. Since $z(\eta_2) \leq z(\eta_1)$, this further implies $\mathbb{E}[\big((U - z(\eta_1))_+ - c\big)|U - c \geq \eta_1] < 0$, which contradicts with (2.21). Therefore, we must have $\eta_1 < \eta_2 \Rightarrow z(\eta_1) < z(\eta_2)$.

To show the continuity of $z(\cdot)$, define $\psi(z; \eta) := \int_{\eta+c}^{\bar{u}}[(x - z)_+ - c]dF_U(x)$. Then by the definition of $z(\cdot)$, we must have $\psi(z(\eta); \eta) = 0$. Moreover, it is easy to see that the value of z satisfying $\psi(z; \eta) = 0$ is unique when $\eta < \bar{u} - c$. Thus,

$$z(\eta) = \underset{z\in[\underline{u}-c,\bar{u}-c]}{\arg\max} -(\psi(z; \eta))^2$$

for any $\eta < \bar{u} - c$. Because $\psi(z, \eta)$ is continuous in (z, η) due to the assumption that F_U is continuous, this implies that $z(\cdot)$ is continuous at any $\eta \in [\underline{u}-c, \bar{u}-c)$ by the Maximum Theorem. To show it is also continuous at $\eta = \bar{u}-c$, just notice the results in (a) implies $\lim_{\eta \nearrow \bar{u}-c} z(\eta) = \bar{u}-c = z(\bar{u}-c)$.

Q.E.D.

Next, I provide the detailed formula of G_η, the distribution of $U \wedge Z$ under an upper-censorship signal with threshold η.

Lemma 2.B.2.

$$G_\eta(x) = \begin{cases} F_{U-c}(x) & x < \eta \\ F_{U-c}(\eta) & \eta \leq x < (\eta + c) \wedge z(\eta) \\ F_U(x) & (\eta + c) \leq x < z(\eta) \\ 1 & z(\eta) \leq x \end{cases} \tag{2.23}$$

(The third piece vanishes if $\eta + c \geq z(\eta)$.)

Proof. Under the upper-censorship signal with threshold η, all match values with $U - c < \eta$ are fully revealed while the others are pooled together. Therefore, when $U - c < \eta$, we have $Z = U - c$ and thus $U \wedge Z = U - c$;[42] when $U - c \geq \eta$, we have $Z = z(\eta)$ and thus $U \wedge Z = U \wedge z(\eta) \geq (\eta + c) \wedge z(\eta)$. Therefore, we have:

- For $x < \eta$, $U \wedge Z \leq x$ if and only if $U - c \leq x$. Thus $G_\eta(x) = F_{U-c}(x)$.
- For $x \in \big[\eta, (\eta + c) \wedge z(\eta)\big)$, $U \wedge Z \leq x$ if and only if $U - c \leq \eta$. Thus $G_\eta(x) = F_{U-c}(\eta)$.
- For $x \in \big[\eta + c, z(\eta)\big)$, $U \wedge Z \leq x$ if and only if either $U - c < \eta$ or $U \wedge z(\eta) \leq x$. Since $z(\eta) > x \geq \eta + c$ here, this is equivalent to having $U \leq x$. Thus $G_\eta(x) = F_U(x)$.
- For $x \geq z(\eta)$, obviously $U \wedge Z \leq x$ for sure. Thus $G_\eta(x) = 1$.

Q.E.D.

Step 2: existence of such η^* in Proposition 2.3.1(b)

Lemma 2.B.3. *Under the assumptions in Proposition 2.3.1(b), there exists $\eta^* \in [\eta_0, r_p^{max}]$ such that one of conditions (i) and (ii) is satisfied. Moreover, $z(\eta^*) \geq r_p^{min}$ for any such η^*.*

Proof. I first show some basic properties of the function Γ and the set $[\eta_0, r_p^{max}]$.

[42] To check this, notice $z = u - c$ solves equation (2.1) when ϕ places probability 1 at u.

By definition, $\Gamma(\eta) = \frac{J_p\big((\eta+c)\wedge z(\eta)\big)-J_p(\eta)}{(\eta+c)\wedge z(\eta)-\eta} - j_p\big(z(\eta)\big)$. Since $J_p(\cdot)$, $j_p(\cdot)$ and $z(\cdot)$ are all continuous and $z(\eta) > \eta$ (thus the denominator is non-zero) for any $\eta < \bar{u} - c$ (by Lemma 2.B.1(a)), $\Gamma(\cdot)$ is continuous on $[\underline{u} - c, r_p^{max}]$.

By definition, $\eta_0 = \inf\{\eta \in [\underline{u} - c, \bar{u} - c] : z(\eta) \geq r_p^{min}\}$. Because Lemma 2.B.1(a) implies $z(r_p^{min}) \geq r_p^{min}$, we have $\eta_0 \leq r_p^{min}$ and thus the set $[\eta_0, r_p^{max}]$ is non-empty. Moreover, since $z(\cdot)$ is continuous and increasing (Lemma 2.B.1 (b)), we have $z(\eta) \geq r_p^{min}$ for any $\eta \geq \eta_0$. This verifies that the desired η^* must satisfy $z(\eta^*) \geq r_p^{min}$.

Now, I assume that no $\eta^* \in [\eta_0, r_p^{max}]$ satisfies condition (ii) and show that there then must exist $\eta^* \in [\eta_0, r_p^{max}]$ satisfying condition (i). For this purpose, I show $\Gamma(\eta_0) \leq 0$ and $\Gamma(r_p^{max}) > 0$.

- Check $\Gamma(\eta_0) \leq 0$:

  If $\eta_0 = \underline{u} - c$, then the assumption that no $\eta^* \in [\eta_0, r_p^{max}]$ satisfies condition (ii) directly implies that $\Gamma(\eta_0) < 0$.

  If $\eta_0 > \underline{u} - c$, then by the continuity of $z(\cdot)$ we must have $z(\eta_0) = r_p^{min}$ and thus J_p is convex over $[\eta_0, z(\eta_0)]$. This convexity implies $\Gamma(\eta_0) \leq 0$.[43]
- Check $\Gamma(r_p^{max}) > 0$:

  Since $r_p^{max} < (r_p^{max} + c) \wedge z(r_p^{max}) \leq z(r_p^{max})$ by Lemma 2.B.1(a), $\Gamma(r_p^{max}) > 0$ is implied by the concavity of $J_p(\cdot)$ over $[r_p^{max}, \bar{u} - c]$ and the strict concavity of it over $[r_p^{max}, r_p^{max} + \epsilon]$ for some $\epsilon > 0$.

These observations together with the continuity of $\Gamma(\cdot)$ imply that there exists $\eta^* \in [\eta_0, r_p^{max}]$ s.t. $\Gamma(\eta) = 0$ by the Intermediate Value Theorem. Since condition (ii) does not hold for this η^* by my assumption earlier, it must satisfy condition (i). *Q.E.D.*

Step 3: optimality of upper-censorship signal with threshold η^*

Proof. Pick any $\eta^* \in [\eta_0, r_p^{max}]$ satisfying one of the conditions (i) and (ii) in the proposition (existence has been shown in Step 2). Then $\Gamma(\eta^*) \geq 0$, where the inequality is strict only if $\eta^* = \eta_0 = \underline{u} - c$. It suffices to show G_{η^*} solves the Relaxed Problem (2.4) – (2.6).

To ease notation, fixing an η^*, I denote $z(\eta^*)$ as z^* for short. Moreover, I use $\ell(\,\cdot\,; b, (x_0, y_0))$ to denote an affine function with slope b passing point (x_0, y_0).

[43] To see this, notice $\Gamma(\eta_0) = \int_{\eta_0}^{(\eta_0+c)\wedge z(\eta_0)} j(x)dx/\big((\eta_0 + c) \wedge z(\eta_0) - \eta_0\big) - j\big(z(\eta_0)\big) \leq \int_{\eta_0}^{(\eta_0+c)\wedge z(\eta_0)} j\big(z(\eta_0)\big)dx/\big((\eta_0 + c) \wedge z(\eta_0) - \eta_0\big) - j\big(z(\eta_0)\big) = 0$, where the inequality is because $j(x)$ is increasing on $[\eta_0, z(\eta_0)]$.

Define

$$v(x) = \begin{cases} J_p(x) & x < \eta^* \\ \ell\Big(x; j_p(z^*), \big((\eta^*+c)\wedge z^*, J_p((\eta^*+c)\wedge z^*)\big)\Big) & x \geq \eta^* \end{cases}$$

By construction, v follows J_p when $x < \eta^*$, is affine with slope $j_p(z^*)$ when $x \geq \eta^*$, and intersects with J_p at $x = (\eta^*+c)\wedge z^*$.

Define

$$\rho(x) = \begin{cases} 0 & x < (\eta^*+c)\wedge z^* \\ J_p(x) - v(x) & (\eta^*+c)\wedge z^* \leq x \leq z^* \\ J_p(z^*) - v(z^*) & x > z^* \end{cases}$$

Notice if $\eta^* + c \geq z^*$, then the second piece vanishes and $\rho \equiv 0$ ($J_p(z^*) - v(z^*) = 0$ in this case by definition of v). If $\eta^* + c < z^*$, ρ equals to the difference between J_p and v over interval $[\eta^*+c, z^*]$, and is constantly extended to regions above z^* or below $\eta^* + c$.

Now, it suffices to check that G_{η^*} (whose formula is given in Lemma 2.B.2), v and ρ satisfy the optimality conditions in Theorem 2.3.1. This is done by showing four observations below. While the detailed proofs are abstract and tedious, I plot the key functions in Figure 2.9, which serves as a graphical check for the optimality conditions. I highlight that by construction and the conclusion in Lemma 2.B.3, we have $\eta^* \leq r_p^{max}$ and $z^* \geq r_p^{min}$ (i.e., η^* is always on the convex portion of J_p and z^* is always on the concave portion of J_p). This fact is used repeatedly below.

Observation (1). v is convex over $[\underline{u}-c, \bar{u}-c]$. Moreover, $v(\eta^*) \geq J_p(\eta^*)$, which holds as equality if $\eta^* > \underline{u} - c$.

Subproof. To show $v(\eta^*) \geq J_p(\eta^*)$, just notice that

$$\begin{aligned} v(\eta^*) =& \ell\Big(\eta^*; j_p\big(z^*\big), \big((\eta^*+c)\wedge z^*, J_p((\eta^*+c)\wedge z^*)\big)\Big) \\ =& J_p((\eta^*+c)\wedge z^*) - j_p\big(z^*\big)\big((\eta^*+c)\wedge z^* - \eta^*\big) \\ \geq& J_p((\eta^*+c)\wedge z^*) - \Big(J_p((\eta^*+c)\wedge z^*) - J_p(\eta^*)\Big) = J_p(\eta^*) \end{aligned}$$

where the inequality holds due to the condition $\frac{J_p\big((\eta^*+c)\wedge z^*\big) - J_p(\eta^*)}{(\eta^*+c)\wedge z^* - \eta^*} \geq j_p\big(z^*\big)$, which is implied by either of condition (i) and condition (ii). If $\eta^* > \underline{u} - c$, it must hold as equality (i.e., condition (i) holds) and thus $v(\eta^*) = J_p(\eta^*)$.

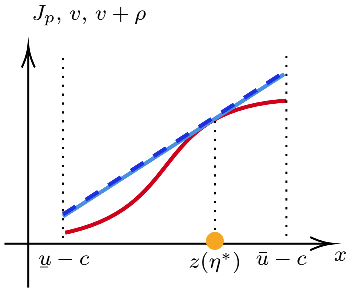

(a) $\eta^* = \underline{u} - c$ and $z(\eta^*) \leq \eta^* + c$

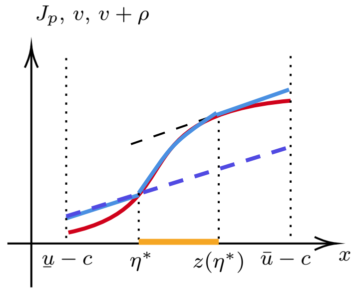

(b) $\eta^* = \underline{u} - c$ and $z(\eta^*) > \eta^* + c$

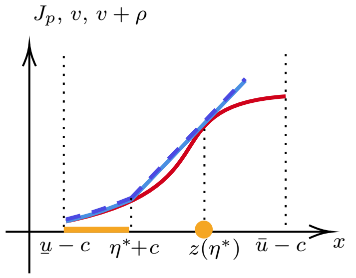

(c) $\eta^* > \underline{u} - c$ and $z(\eta^*) \leq \eta^* + c$

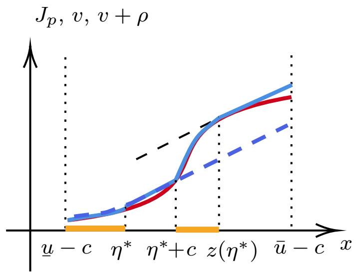

(d) $\eta^* > \underline{u} - c$ and $z(\eta^*) > \eta^* + c$

Figure 2.9: Graphical Check of Optimality Conditions for (G_{η^*}, v, ρ). The red curve is J_p; the solid blue curve is $v + \rho$; the dashed blue curve is v; the orange area on x-axis is the support for G_{η^*}.

Now I show convexity of v. If $\eta^* = \underline{u} - c$, convexity trivially holds since v is affine then. Thus we only need to consider the case where $\eta^* > \underline{u} - c$ and condition (i) holds (i.e., $\Gamma(\eta^*) = 0$) . In this case, notice that both pieces of v are separately convex, and that $v(\eta^*) = J_p(\eta^*)$ shown above implies continuity of v at η^*. Thus it suffices to show convexity around η^*. Specifically, we need to show $v'_+(\eta^*) \geq v'_-(\eta^*)$, where v'_+ and v'_- denote v's right and left derivatives respectively. Consider two cases:

- $(\eta^*+c) \wedge z^* \leq r_p^{max}$. In this case, notice $v'_+(\eta^*) = j_p(z^*) = \frac{J_p\big((\eta^*+c)\wedge z^*\big) - J_p(\eta^*)}{(\eta^*+c)\wedge z^* - \eta^*} \geq j_p(\eta^*) = v'_-(\eta^*)$, where the inequality is by convexity of J_p over $[\underline{u} - c, r_p^{max}]$.
- $(\eta^* + c) \wedge z^* > r_p^{max}$. In this case, $v'(x) = j_p(z^*) \leq j_p(x)\ \forall x \in [r_p^{max}, z^*)$ due to the concavity of J_p over $[r_p^{max}, z^*]$. Since the affine part of v intersects J_p at $(\eta^* + c) \wedge z^*$ by construction, this implies that $v(x) \geq J_p(x)$ for all $x \in [r_p^{max}, (\eta^* + c) \wedge z^*)$ and, in particular, $v(r_p^{max}) \geq J_p(r_p^{max})$.[44] Since it has been shown that $v(\eta^*) = J_p(\eta^*)$, this implies that $v'_+(\eta^*) \geq \frac{J_p(r_p^{max}) - J_p(\eta^*)}{r_p^{max} - \eta^*} \geq j_p(\eta^*) = v'_-(\eta^*)$, where the second inequality is due to convexity of J_p over $[\underline{u} - c, r_p^{max}]$.

[44] If J_p is strictly concave over $[r_p^{max}, \bar{u} - c]$, then the inequalities hold strictly.

□

Observation (2). ρ is (weakly) increasing.

Subproof. It suffices to consider the case $\eta^*+c < z^*$ and show that ρ is increasing over $[\eta^*+c, z^*]$. Notice for all $x \in \left[\max\{r_p^{min}, (\eta^*+c)\}, z^*\right]$, we have

$$\rho'(x) = j_p(x) - v'(x) = j_p(x) - j_p(z^*) \geq 0$$

where the inequality holds because J_p is concave over this region. Moreover, for all $x \in [\eta^* + c, r_p^{min}]$ (supposing $r_p^{min} > \eta^* + c$), we have

$$j_p(x) \geq \frac{J_p\big(\eta^*+c\big) - J_p(\eta^*)}{c} \geq j_p(z^*) = v'(x)$$

where the first inequality is due to convexity of J_p over $[\underline{u} - c, r_p^{min}]$ and the second inequality holds because $\Gamma(\eta^*) \geq 0$ (recall that in the current case $z^* > \eta^* + c$). In sum, $\rho' = j_p - v'$ is non-negative over $[\eta^* + c, z^*]$, and thus ρ is increasing.[45] □

Observation (3). $v(x)+\rho(x) \geq J_p(x)$ over $[\underline{u}-c, \bar{u}-c]$, where equality holds for $x \in [\underline{u}-c, \eta^*) \cup [(\eta^*+c) \wedge z^*, z^*]$. If $\eta^* > \underline{u} - c$, equality also holds for $x = \eta^*$.

Subproof. The fact that equality holds when $x \in [\underline{u} - c, \eta^*) \cup [(\eta^* + c) \wedge z^*, z^*]$ is obvious by construction of v and ρ. If $\eta^* > \underline{u} - c$, we also have $v(\eta^*) + \rho(\eta^*) = J_p(\eta^*)$ because $\rho(\eta^*) = 0$ and $v(\eta^*) = J_p(\eta^*)$ by Observation (1).

For $x > z^*$, by construction we have $v(x) + \rho(x) = v(x) + J_p(z^*) - v(z^*) = j_p(z^*)(x - z^*) + J_p(z^*) \geq J_p(x)$, where the last inequality is due to concavity of J_p over $[r_p^{min}, \bar{u} - c]$.[46]

For $x \in \big[\eta^*, (\eta^* + c) \wedge z^*\big)$, $\rho(x) = 0$ and thus it suffices to show $v(x) \geq J_p(x)$ over this interval. $v(\eta^*) \geq J_p(\eta^*)$ has been shown in Observation (1). Thus it suffices to consider $x \in \big(\eta^*, (\eta^* + c) \wedge z^*\big)$. Consider two possibilities:[47]

- $(\eta^* + c) \wedge z^* \leq r_p^{max}$. In this case, J_p is convex and v is affine over the interval. Thus the desired result is implied by $v(\eta^*) \geq J_p(\eta^*)$ and $v\big((\eta^* + c) \wedge z^*\big) = J_p\big((\eta^* + c) \wedge z^*\big)$.

---

[45] If J_p is strictly convex over $[\underline{u} - c, r_p^{min}]$ and is strictly concave over $[r_p^{min}, \bar{u} - c]$, then the corresponding inequalities hold strictly and thus ρ is strictly increasing over $[\eta^* + c, z^*]$.

[46] The last inequality holds strictly if J_p is strictly concave over $[r_p^{min}, \bar{u} - c]$.

[47] In both cases, if J_p is strictly convex over $[\underline{u} - c, r_p^{max}]$ and is strictly concave over $[r_p^{max}, \bar{u} - c]$, then the same argument shows $v(x) > J_p(x)$ over $\big(\eta^*, (\eta^* + c) \wedge z^*\big)$. When referring to the proof of Observation (1), also refer to footnote 44.

- $(\eta^* + c) \wedge z^* > r_p^{max}$. In this case, for $x \in [r_p^{max}, (\eta^* + c) \wedge z^*)$, it has been shown that $v(x) \geq J_p(x)$ (see the last paragraph in the proof of Observation (1)). For $x \in (\eta^*, r_p^{max})$, the result holds because J_p is convex over $[\eta^*, r_p^{max}]$, $v(\eta^*) \geq J(\eta^*)$ and $v(r_p^{max}) \geq J(r_p^{max})$.

□

Observation (4). G_{η^*}, v and ρ satisfy condition (C3) in Theorem 2.3.1.

Subproof.

- $\int v(x) dG_{\eta^*}(x) = \int v(x) dF_{U-c}(x)$ holds because G_{η^*} agree with F_{U-c} over $[\underline{u} - c, \eta^*)$ (see Lemma 2.B.2) and v is affine over $[\eta^*, \bar{u} - c]$. (Recall that $G_{\eta^*} \preceq_{MPS} F_{U-c}$.)
- When $\eta^* + c \geq z^*$, $\int \rho(x) dG_{\eta^*}(x) = \int \rho(x) dF_U(x)$ holds trivially since $\rho \equiv 0$. When $\eta^* + c < z^*$, we have:

$$
\begin{aligned}
&\int \rho(x) dG_{\eta^*}(x) - \int \rho(x) dF_U(x) \\
=&\int_{[z^*, \bar{u}]} \rho(x) dG_{\eta^*}(x) - \int_{[z^*, \bar{u}]} \rho(x) dF_U(x) \\
=&\big(J_p(z^*) - v(z^*)\big)\Big(\big(1 - \lim_{x \nearrow z^*} G_{\eta^*}(x)\big) - \big(1 - \lim_{x \nearrow z^*} F_U(x)\big)\Big) = 0
\end{aligned}
$$

where the first equality holds because $\rho(x) \equiv 0$ for $x \leq \eta^* + c$ and G_{η^*} agrees with F_U over $(\eta^* + c, z^*)$; the second equality is because $\rho(x) \equiv J_p(z^*) - v(z^*)$ when $x \geq z^*$; the last line equals to zero because $G_{\eta^*} = F_U$ on a left neighborhood of z^* when $\eta^* + c < z^*$ (see Lemma 2.B.2).

□

To sum up, we have:

- (v, ρ) satisfies (C1) by Observations (1) and (2);
- G_{η^*} supports within $[\underline{u} - c, \eta^*] \cup [(\eta^* + c) \wedge z^*, z^*]$ when $\eta^* > \underline{u} - c$ and supports within $[(\eta^* + c) \wedge z^*, z^*]$ when $\eta^* = \underline{u} - c$ (see Lemma 2.B.2). Thus (G_{η^*}, v, ρ) satisfies (C2) by Observation (3).
- (G_{η^*}, v, ρ) satisfies (C3) by Observation (4).
- G_η^* satisfies (C4) since it is induced by a particular feasible signal (upper-censorship signal with threshold η^*).

Thus, G_{η^*} is indeed optimal for the Relaxed Problem, and the corresponding upper-censorship signal is optimal.

Q.E.D.

2.B.4 Proof for Corollary 2.3.1

Proof. Part (a) is equivalent to Proposition 2.3.1(a). For part (b), if $r_p^{max} = \bar{u} - c$, then J_p is also convex and thus affine on $[\underline{u} - c, \bar{u} - c]$. Then the result is directly implied by Lemma 2.3.1. If $r_p^{max} < \bar{u} - c$, it suffices to check that $\eta^* = \underline{u} - c$ satisfies the conditions in Proposition 2.3.1(b). Indeed, concavity of J_p implies $r_p^{min} = \underline{u} - c$ and thus $\eta_0 = \underline{u} - c$, so $\underline{u} - c$ is within $[\eta_0, r_p^{max}]$. Moreover, concavity of J_p directly implies $\Gamma(\underline{u} - c) \geq 0$ and thus condition (ii) holds. *Q.E.D.*

2.B.5 Proof for Proposition 2.3.2

Part 1: uniqueness of the optimal distribution of $U \wedge Z$ and the optimal upper-censorship signal

Proof.

Case 1: $r_p = \bar{u} - c$

In this case, j_p is strictly increasing over $[\underline{u} - c, \bar{u} - c]$ and $G = F_{U-c}$ is optimal for the Relaxed Problem. Then, for any optimal distribution G^* of the effective-search-value, we must have:

$$
\begin{aligned}
0 &= \int_{(\underline{u}-c,\bar{u}-c]} J_p(x)dG^*(x) - \int_{(\underline{u}-c,\bar{u}-c]} J_p(x)dF_{U-c}(x) \\
&= \Big[J_p(x)\big(G^*(x) - F_{U-c}(x)\big)\Big]\Big|_{\underline{u}-c}^{\bar{u}-c} - \int_{\underline{u}-c}^{\bar{u}-c} [G^*(x) - F_{U-c}(x)]j_p(x)dx \\
&= -\int_{\underline{u}-c}^{\bar{u}-c} [G^*(x) - F_{U-c}(x)]j_p(x)dx \\
&= -\Big[\Big(\int_{\underline{u}-c}^{x} [G^*(t) - F_{U-c}(t)]dt\Big)j_p(x)\Big]\Big|_{\underline{u}-c}^{\bar{u}-c} + \int_{(\underline{u}-c,\bar{u}-c]} \Big(\int_{\underline{u}-c}^{x} [G^*(t) - F_{U-c}(t)]dt\Big)dj_p(x) \\
&= \int_{(\underline{u}-c,\bar{u}-c]} \Big(\int_{\underline{u}-c}^{x} [G^*(t) - F_{U-c}(t)]dt\Big)dj_p(x)
\end{aligned}
$$

(For any function h, $h(x)|_a^b := h(b) - h(a)$.) In the first equality, point $\underline{u} - c$ is omitted from the integrals because F_{U-c} is atom-less and because of this, G^* also puts zero probability at $\underline{u} - c$ to be a MPC of F_{U-c}. The 2nd and 4th equalities are by integration by parts.[48] The 3rd equality holds because $G^* \preceq_{MPS} F_{U-c} \Rightarrow supp\{G^*\} \subset supp\{F_{U-c}\}$ and thus $G^*(x) = F_{U-c}(x)$ for $x = \underline{u} - c$ or $\bar{u} - c$. The 5th equality holds because $G^* \preceq_{MPS} F_{U-c}$ implies that

[48]Integration by parts holds in the two equalities respectively because $J_p(x)$ and $\int_{\underline{u}-c}^{x}[G^*(t) - F_{U-c}(t)]dt$ are continuous in x. See Folland (1999) Theorem 3.36.

$\int_{\underline{u}-c}^{x}[G^*(t)-F_{U-c}(t)]dt = 0$ for $x=\bar{u}-c$.

Now, notice term $\int_{\underline{u}-c}^{x}[G^*(t)-F_{U-c}(t)]dt \leq 0$ for all x since $G^* \preceq_{MPS} F_{U-c}$. As j_p is increasing, dj_p is a positive measure over $(\underline{u}-c,\bar{u}-c]$ and the above result implies $\int_{\underline{u}-c}^{x}[G^*(t)-F_{U-c}(t)]dt = 0$ for $dj_p - a.e.$ x in $(\underline{u}-c,\bar{u}-c]$.

Suppose $G^*(t_0) > F_{U-c}(t_0)$ for some $t_0 \in [\underline{u}-c,\bar{u}-c)$. Then by right-continuity we have $G^*(t) > F_{U-c}(t)$ over some right-neighborhood of t_0, denoted as I. Then over interval I, $\int_{\underline{u}-c}^{x}[G^*(t)-F_{U-c}(t)]dt$ is strictly increasing in x. This implies that $\int_{\underline{u}-c}^{x}[G^*(t)-F_{U-c}(t)]dt \neq 0$ over some subinterval $I' \subset I$. Then we must have j_p constant over I', which violates the assumption that j_p is strictly increasing. Similar argument also rules out the possibility that $G^*(t_0) < F_{U-c}(t_0)$. Thus $G^*(t) = F_{U-c}(t)$ for all $t \in [\underline{u}-c,\bar{u}-c)$. We also have $G^*(\bar{u}-c) = 1 = F_{U-c}(\bar{u}-c)$ since $supp\{G^*\} \subset supp\{F_{U-c}\}$. In conclusion, $G^* = F_{U-c}$ is the unique optimal solution.

Case 2: $r_p < \bar{u}-c$

Let η^* be an optimal threshold characterized in Proposition 2.3.1(b), and let v and ρ be the corresponding functions defined in Section 2.B.3 (Step 3). I have shown that (G_{η^*}, v, ρ) satisfies all conditions in Theorem 2.3.1. Thus, by Theorem 2.3.1, if another distribution G^* also solves the Relaxed Problem (which is now necessary for it to be optimal), then (G^*, v, ρ) must also satisfy all optimality conditions in Theorem 2.3.1. Then we have following observations:

Observation (1). G^* has zero probability over $\big(\eta^*, (\eta^*+c)\wedge z^*\big)$ and $(z^*, \bar{u}-c]$.

Subproof. By condition (C2) in Theorem 2.3.1, it suffices to show that $v(x)+\rho(x) > J_p(x)$ for $x \in \big(\eta^*, (\eta^*+c)\wedge z^*\big) \cup (z^*, \bar{u}-c]$. This is proved by slightly modifying the proof of Observation (3) in the proof of Proposition 2.3.1 (Step 3 in Appendix 2.B.3). See footnotes 46 and 47 for details. (Notice that $r_p^{min} = r_p^{max} = r_p$ now under Assumption 2.3.2.) □

Observation (2). $G^*(x) = F_{U-c}(x)$ for $x \in [\underline{u}-c, \eta^*]$.

Subproof. Since $G^* \preceq_{MPS} F_{U-c}$, we must have $G^*(\underline{u}-c) = F_{U-c}(\underline{u}-c) = 0$ (by assumption, F_{U-c} is atomless). Thus it suffices to show the result assuming $\eta^* > \underline{u}-c$.

By condition (C3) in Theorem 2.3.1, we must have $\int v(x)dG^*(x) = \int v(x)dF_{U-c}(x)$. By two steps of integration by parts (similar to the procedure in Case 1), this implies that

$$\int_{(\underline{u}-c,\bar{u}-c]} \Big(\int_{\underline{u}-c}^{x}[G^*(t)-F_{U-c}(t)]dt\Big)dv'(x) = 0$$

and we thus have $\int_{\underline{u}-c}^{x}[G^*(t) - F_{U-c}(t)]dt = 0$ for $dv' - a.e.\, x$, where v' denotes the right-continuous (almost everywhere) derivative of convex function v, and dv' is the positive measure induced by it. As $v(x) = J_p(x)$ for $x \in [\underline{u} - c, \eta^*]$, $v'(x) = j_p(x)$ is strictly increasing over $[\underline{u} - c, \eta^*)$. Thus by similar argument as in Case 1, we have $G^*(t) = F_{U-c}(t)$ for all $t \in [\underline{u} - c, \eta^*)$.

Finally, consider the point η^*. Notice $G^*(t) = F_{U-c}(t)\, \forall t \in [\underline{u} - c, \eta^*)$ and F_{U-c} being continuous at η^* imply that $G^*(\eta^*) \geq F_{U-c}(\eta^*)$ (otherwise G^* must jump down at η^*, violating its monotonicity). Suppose $G^*(\eta^*) > F_{U-c}(\eta^*)$. Then by right-continuity $G^*(x) > F_{U-c}(x)$ over a right-neighborhood of η^*. This implies that $\int_{\underline{u}-c}^{x}[G^*(t) - F_{U-c}(t)]dt > 0$ for some $x > \eta^*$, violating $G^* \preceq_{MPS} F_{U-c}$. Thus $G^*(\eta^*) = F_{U-c}(\eta^*)$. □

Observation (3). If $z^* > \eta^* + c$, then $G^*(x) = F_U(x)$ for $x \in [\eta^* + c, z^*)$.

Subproof. Condition (C3) in Theorem 2.3.1 requires that $\int \rho(x)dG^*(x) = \int \rho(x)dF_U(x)$. This implies that

$$\begin{aligned} 0 &= \int \rho(x)dG^*(x) - \int \rho(x)dF_U(x) \\ &= \rho(x)(G^*(x) - F_U(x))\Big|_{\underline{u}-c}^{\bar{u}} - \int_{\underline{u}-c}^{\bar{u}} [G^*(x) - F_U(x)]d\rho(x) \\ &= -\int_{\underline{u}-c}^{\bar{u}} [G^*(x) - F_U(x)]d\rho(x) \end{aligned}$$

where the integration by parts formula for the 2nd equality is valid because ρ is continuous by construction. Since $G^*(x) \geq F_U(x)$ by FOSD and ρ is increasing, we must have $G^*(x) = F_U(x)$ for $d\rho - a.e.\, x$ in $[\underline{u} - c, \bar{u}]$. By footnote 45, we know ρ is strictly increasing over $[\eta^* + c, z^*]$. Thus $G^*(x) = F_U(x)$ for all $x \in [\eta^* + c, z^*)$ (using right-continuity of CDF). □

Together with the requirement that $supp\{G^*\} \subset [\underline{u} - c, \bar{u} - c]$ (since $G^* \preceq_{MPS} F_{U-c}$), the above observations pin down a unique formula for G^*, which is the same as G_{η^*}. Finally, notice that two upper-censorship signals with different thresholds necessarily lead to different G. Thus the optimal upper-censorship signal is unique under Assumption 2.3.2. *Q.E.D.*

Part 2: uniqueness of the optimal joint distribution of (U, Z)

Proof. I first need a simple observation:

Observation (4). Under any signal, we have: (i) $\mathbb{E}[U - c - U \wedge Z|Z] = 0$ $a.s.$; (ii) $\mathbb{P}(U \wedge Z = Z|Z) > 0$ $a.s.$

Subproof. By definition of Z, we have $\mathbb{E}[U - c - U \wedge Z|S] = \mathbb{E}[(U - Z)_+ - c|S] = 0$. This implies (i) since Z is measurable w.r.t. S. Suppose (ii) does not hold. Then with positive probability $\mathbb{P}(U \wedge Z = U|Z) = 1$ and thus $\mathbb{E}[U - c - U \wedge Z|Z] = -c$, which contradicts with (i). □

Given any signal inducing G_{η^*}, I show two claims below.

Claim 1: On the event $\{U - c \leq \eta^*\}$, $Z = U - c$ almost surely (i.e., $\mathbb{P}(U - c \leq \eta^*, Z \neq U - c) = 0$).

Notice under any signal inducing G_{η^*}, for any $t \leq \eta^*$ we have:

$$
\begin{aligned}
&\mathbb{E}[(U - c - t)\mathbb{1}\{U - c \leq t\}] \\
=&\mathbb{E}[(U \wedge Z - t)\mathbb{1}\{U \wedge Z \leq t\}] \\
=&\mathbb{E}[(U \wedge Z - t)\big(\mathbb{1}\{Z \leq t\} + \mathbb{1}\{Z > t \geq U\}\big)] \\
\geq&\mathbb{E}[(U \wedge Z - t)\mathbb{1}\{Z \leq t\}] + \mathbb{E}[(U - c - t)\mathbb{1}\{Z > t \geq U\}] \\
=&\mathbb{E}[(U - c - t)\mathbb{1}\{Z \leq t\}] + \mathbb{E}[(U - c - t)\mathbb{1}\{Z > t \geq U\}] \\
=&\mathbb{E}[(U - c - t)\mathbb{1}\{U \wedge Z \leq t\}]
\end{aligned}
$$

where the first equality holds because $G_{\eta^*}(x) = F_{U-c}(x)$ for $x \leq \eta^*$; the inequality in the fourth row holds because $U \wedge Z = U > U - c$ when $Z > U$; the equality in the 5th row holds because of result (i) in Observation (4) above. This result further implies:

$$
\begin{aligned}
0 =&\mathbb{E}[(U - c - t)\big(\mathbb{1}\{U - c \leq t\} - \mathbb{1}\{U \wedge Z \leq t\}\big)] \\
=&\mathbb{E}[(U - c - t)\big(\mathbb{1}\{U - c \leq t; U \wedge Z > t\}\big)] - \mathbb{E}[(U - c - t)\big(\mathbb{1}\{U - c > t; U \wedge Z \leq t\}\big)]
\end{aligned}
$$

Notice that in the second line, the first term is non-positive and the second term is non-negative. Thus they must both be zero. This implies that $\mathbb{P}(U - c < t; U \wedge Z > t) = \mathbb{P}(U - c > t; U \wedge Z \leq t) = 0$. Since this holds for arbitrary $t \leq \eta^*$, we have for any $t_1 < t_2 \leq \eta^*$, $\mathbb{P}(U - c \in (t_1, t_2), U \wedge Z \notin [t_1, t_2]) = 0$. Since U is continuous, we have a stronger result that $\mathbb{P}(U - c \in [t_1, t_2], U \wedge Z \notin [t_1, t_2]) = 0$.

Now, notice that whenever $U - c \leq \eta^*$ and $U - c \neq U \wedge Z$, we can always find $t_1, t_2 \in \mathbb{Q}$

such that $t_1 < t_2 \leq \eta^*$, $U - c \in [t_1, t_2]$ and $U \wedge Z \notin [t_1, t_2]$. Thus

$$\mathbb{P}(U - c \leq \eta^*; U - c \neq U \wedge Z) \leq \sum_{t_1, t_2 \in \mathbb{Q}:\, t_1 < t_2 \leq \eta^*} \mathbb{P}(U - c \in [t_1, t_2]; U \wedge Z \notin [t_1, t_2])$$

The conclusion above implies the RHS equals to 0. Thus $\mathbb{P}(U - c \leq \eta^*; Z \neq U - c) = \mathbb{P}(U - c \leq \eta^*; U - c \neq U \wedge Z) = 0$.

Claim 2: On the event $\{U - c > \eta^*\}$, $Z = z(\eta^*)\ a.s.$ (i.e., $\mathbb{P}(U - c > \eta^*; Z \neq z(\eta^*)) = 0$.)

First notice that Observation (4)(ii) implies that the support of Z is a subset of the support of $U \wedge Z$.[49] Thus we have $supp\{Z\} \subset [\underline{u} - c, \eta^*] \cup [(\eta^* + c) \wedge z(\eta^*), z(\eta^*)]$. Also recall that in the proof of Claim 1, I have shown $\mathbb{P}(U - c > t; U \wedge Z \leq t) = 0$, $\forall t \leq \eta^*$. With $t = \eta^*$, this implies $\mathbb{P}(U - c > \eta^*; Z \leq \eta^*) = 0$. So, on the event $\{U - c > \eta^*\}$ we have $Z \in [(\eta^* + c) \wedge z(\eta^*), z(\eta^*)]\ a.s.$ Therefore, when $\eta^* + c \geq z(\eta^*)$, Claim 2 has been proved; when $\eta^* + c < z(\eta^*)$, it suffices to show $\mathbb{P}\big(Z \in \big[\eta^* + c, z(\eta^*)\big)\big) = 0$.

Now, assume $\eta^* + c < z(\eta^*)$. We have:

$$\begin{aligned} &\mathbb{E}[(U \wedge Z)\mathbb{1}\{U \geq \eta^* + c\}] = \mathbb{E}[(U \wedge Z)\mathbb{1}\{U \wedge Z \geq \eta^* + c\}] \\ &= \mathbb{E}\big[\big(U \wedge z(\eta^*)\big)\mathbb{1}\{U \wedge z(\eta^*) \geq \eta^* + c\}\big] = \mathbb{E}\big[\big(U \wedge z(\eta^*)\big)\mathbb{1}\{U \geq \eta^* + c\}\big] \end{aligned}$$

where the first equality holds because almost surely $U \geq \eta^* + c \Rightarrow Z \geq \eta^* + c$ as shown above (the event $\{U = \eta^* + c\}$ has zero probability and thus can be ignored); the second equality holds because G_{η^*} agrees with the distribution of $U \wedge z(\eta^*)$ over $[\eta^* + c, +\infty)$. Since $Z \leq z(\eta^*)\ a.s.$, $U \wedge Z \leq U \wedge z(\eta^*)\ a.s.$ and thus the above result implies $U \wedge Z = U \wedge z(\eta^*)\ a.s.$ on the event $\{U - c > \eta^*\}$.

Notice that Claim 1 implies almost surely $Z > \eta^* \Rightarrow U - c > \eta^*$. Thus we have $Z \in \big[\eta^* + c, z(\eta^*)\big) \Rightarrow U \wedge Z = U \wedge z(\eta^*)\ a.s.$ Also notice when $Z < z(\eta^*)$, $U \wedge Z = U \wedge z(\eta^*)$ only if $U < Z$. This further implies that almost surely $Z \in \big[\eta^* + c, z(\eta^*)\big) \Rightarrow U < Z$. Since Obervation (4)(ii) implies $\mathbb{P}(Z \leq U | Z) > 0\ a.s.$, we must have $\mathbb{P}\big(Z \in \big[\eta^* + c, z(\eta^*)\big)\big) = 0$. This completes the proof for Claim 2.

The two claims above pin down a unique distribution of Z conditional on U. Thus the optimal joint distribution of (U, Z) is unique. *Q.E.D.*

[49] Otherwise, there exists set B s.t. $\mathbb{P}(Z \in B) > 0$ but $\mathbb{P}(U \wedge Z \in B) = 0$. This implies $\mathbb{E}[\mathbb{P}(U \wedge Z = Z|Z)|Z \in B] = \mathbb{P}(U \wedge Z = Z|Z \in B) = \mathbb{P}(U \wedge Z = Z; Z \in B)/\mathbb{P}(Z \in B) = 0$, which contradicts with Observation (4)(ii).

2.B.6 Proof for Proposition 2.3.3

Proof. I show the existence of optimal (p^*, η^*) when $J(\cdot)$ is log-concave below. The other part of the proposition is evident from the discussion in the main text. W.l.g., I assume $J(\bar{u}-c) > 0$. If this does not hold, profit is always 0 for any $p > 0$ and thus every policy is trivially optimal. Notice that by continuity of J, $J(\bar{u}-c) > 0$ implies that the probability of sale is strictly positive for some $p > 0$ under fully revealing signal. Thus the seller's maximal expected profit is strictly positive.

Taking the formula of G_η (equation (2.23)) into the optimization in the proposition, we get:

$$\max_{\substack{p \geq 0 \\ \eta \in [\underline{u}-c, \bar{u}-c]}} \Bigg\{ p \Big[\int_{\underline{u}-c}^{\eta} J(x-p) dF_{U-c}(x) + \int_{(\eta+c)\wedge z(\eta)}^{z(\eta)} J(x-p) dF_U(x) \\ + J\big(z(\eta)-p\big)\Big(1 - F_U\big(\max\{\eta+c, z(\eta)\}\big)\Big) \Big] \Bigg\}$$

Because F_U is continuous, the integrals are continuous in their limits. Due to the continuity of $z(\cdot)$ (Lemma 2.B.1(b)) and $J(\cdot)$, the objective function is continuous in (p, η). Thus, by the Maximum Theorem, the solution exists if we can show that any sufficiently large p is suboptimal (so that the feasible set can be shrunk to be compact).

To show this, notice that the seller's profit is always bounded by $\Pi(p) := pJ(\bar{u}-c-p)$ given any p. Since J is log-concave, $\log(J(\bar{u}-c-p))$ is concave in p and thus $\frac{\partial \log(J(\bar{u}-c-p))}{\partial p}$ is decreasing in p. Because J is a CDF, $\frac{\partial \log(J(\bar{u}-c-p))}{\partial p}$ is strictly negative for some p. Thus there exists $\epsilon > 0$ such that $\frac{\partial \log(J(\bar{u}-c-p))}{\partial p} < -\epsilon$ when p is large enough.

Now, notice that $\frac{\partial \log(\Pi(p))}{\partial p} = 1/p + \frac{\partial \log(J(\bar{u}-c-p))}{\partial p}$. The above result then implies that $\frac{\partial \log(\Pi(p))}{\partial p} < -\epsilon/2$ for p large enough. This implies that $\lim_{p\to+\infty} \log(\Pi(p)) = -\infty$ and thus $\lim_{p\to+\infty} \Pi(p) = 0$. Therefore, any p that is sufficiently large is suboptimal.

Q.E.D.

2.C Proofs for Section 2.4

2.C.1 Proof for Lemma 2.4.1

Proof. Notice:

$$\int \left[\left(x - \mathbb{E}_{U\sim\phi}[U-c]\right)_+ - c\right]\phi(dx) \geq \int \left[x - \mathbb{E}_{U\sim\phi}[U-c] - c\right]\phi(dx)$$
$$= \int (x-c)\phi(dx) - \mathbb{E}_{U\sim\phi}[U-c] = 0 = \int [(x-z_\phi)_+ - c]\phi(dx)$$

where the last equality holds by the definition of z_ϕ. This implies that $\mathbb{E}_{U\sim\phi}[U-c] \leq z_\phi$. Moreover, notice that the inequality above holds as equality if and only if $x - \mathbb{E}_{U\sim\phi}[U-c] \geq 0$ for ϕ-a.e. x. Thus $\mathbb{E}_{U\sim\phi}[U-c] = z_\phi$ if and only if $\inf(supp\{\phi\}) \geq \mathbb{E}_{U\sim\phi}[U-c]$. *Q.E.D.*

2.C.2 Proof for Proposition 2.4.2

The proposition can be proved by applying Theorem 1 in Dworczak & Martini (2019) to the optimization (2.12) – (2.13). Because it is analogous to (but simpler than) the proofs for the search goods case, the details are omitted. An alternative proof, which treats the proposition as a special case of the results for search goods, is also available upon request.

2.C.3 Proof for Proposition 2.4.3

The proof requires a sequence of lemmas that are also useful in some other proofs later. I first show a technical one:

Lemma 2.C.1. *A continuously differentiable function $\Upsilon(\cdot)$ is strictly convex over $[a,t]$ and strictly concave over $[t,b]$. Pick any x, y, w, x', $y' \in [a,b]$ such that $x < y \leq w$, $x' < y'$, $x' \leq x$, $y' \leq y$ and $(x',y') \neq (x,y)$. Then we have $\frac{\Upsilon(y)-\Upsilon(x)}{y-x} \leq \Upsilon'(w) \Rightarrow \frac{\Upsilon(y')-\Upsilon(x')}{y'-x'} < \frac{\Upsilon(y)-\Upsilon(x)}{y-x}$.*

Proof. Assume x, y, w, x' and y' are picked as in the lemma and $\frac{\Upsilon(y)-\Upsilon(x)}{y-x} \leq \Upsilon'(w)$. Notice by the strict concavity of Υ over $[t,b]$, $\frac{\Upsilon(y)-\Upsilon(x)}{y-x} \leq \Upsilon'(w) \Rightarrow x < t$. Moreover, if $y \leq t$, then the result is directly implied by the strict convexity of Υ over $[a,t]$. Thus we only need to consider the case where $x < t < y$.

For any $s_1 < s_2$, let $\overline{\Upsilon'}(s_1,s_2)$ denote the average slope of Υ over $[s_1,s_2]$, i.e., $\overline{\Upsilon'}(s_1,s_2) := \frac{\Upsilon(s_2)-\Upsilon(s_1)}{s_2-s_1}$. Then we have $\overline{\Upsilon'}(x,y) \leq \Upsilon'(w)$ and want to show $\overline{\Upsilon'}(x',y') < \overline{\Upsilon'}(x,y)$.

First, we can show $x' < x \Rightarrow \overline{\Upsilon'}(x', x) < \overline{\Upsilon'}(x, y)$. To see this, notice that by the strict concavity of Υ over $[t, b]$, we have $\overline{\Upsilon'}(t, y) > \Upsilon'(w)$. Thus $\overline{\Upsilon'}(x, y) \leq \Upsilon'(w) \Rightarrow \overline{\Upsilon'}(x, t) < \overline{\Upsilon'}(x, y)$ since $\overline{\Upsilon'}(x, y)$ is a weighted average of $\overline{\Upsilon'}(x, t)$ and $\overline{\Upsilon'}(t, y)$. By the convexity of Υ over $[a, t]$, we have $x' < x \Rightarrow \overline{\Upsilon'}(x', x) \leq \overline{\Upsilon'}(x, t)$. These together imply $x' < x \Rightarrow \overline{\Upsilon'}(x', x) < \overline{\Upsilon'}(x, y)$.

If $y' \leq x$, then the above result directly imply $\overline{\Upsilon'}(x', y') < \overline{\Upsilon'}(x, y)$ since $\overline{\Upsilon'}(x', y') \leq \overline{\Upsilon'}(x', x)$ in this case by the convexity of Υ over $[a, t]$. Thus it remains to consider the case where $y' > x$.

When $y' > x$, we can prove $y' < y \Rightarrow \overline{\Upsilon'}(x, y') < \overline{\Upsilon'}(x, y)$. Since $\overline{\Upsilon'}(x, \cdot)$ is increasing over $(x, t]$ by the convexity of Υ over that region, it suffices to prove this when $y > y' \geq t$. In this case, by the strict concavity of Υ over $[t, b]$, we have $\overline{\Upsilon'}(y', y) > \Upsilon'(w)$. Thus our assumption $\overline{\Upsilon'}(x, y) \leq \Upsilon'(w)$ implies $\overline{\Upsilon'}(x, y') < \overline{\Upsilon'}(x, y)$ since $\overline{\Upsilon'}(x, y)$ is a weighted average of $\overline{\Upsilon'}(x, y')$ and $\overline{\Upsilon'}(y', y)$. This concludes the proof for $y' < y \Rightarrow \overline{\Upsilon'}(x, y') < \overline{\Upsilon'}(x, y)$. Together with the earlier conclusion $x' < x \Rightarrow \overline{\Upsilon'}(x', x) < \overline{\Upsilon'}(x, y)$, this further implies $\overline{\Upsilon'}(x', y') < \overline{\Upsilon'}(x, y)$.

Q.E.D.

Next, I show an important property of the function Γ.

Lemma 2.C.2. *Under Assumption 2.3.2, $\Gamma(\cdot)$ single-crosses zero from below over $[\underline{u} - c, \bar{u} - c)$, i.e., $\Gamma(\eta_1) \geq (>) 0 \Rightarrow \Gamma(\eta_2) \geq (>) 0$ for any $\eta_2 > \eta_1$ in $[\underline{u} - c, \bar{u} - c)$.*

Proof. I prove the contrapositive for this: for any $\eta_1, \eta_2 \in [\underline{u} - c, \bar{u} - c)$ s.t. $\eta_2 > \eta_1$, I show $\Gamma(\eta_2) \leq (<) 0 \Rightarrow \Gamma(\eta_1) \leq (<) 0$. Notice that if $z(\eta_1) < r_p$, the strict convexity of J_p over $[\underline{u} - c, r_p]$ would directly imply $\Gamma(\eta_1) < 0$. Thus it suffices to assume $z(\eta_1) \geq r_p$.

Notice the definition of Γ implies

$$\Gamma(\eta_1) - \Gamma(\eta_2) = \underbrace{\frac{J_p\big((\eta_1 + c) \wedge z(\eta_1)\big) - J_p(\eta_1)}{(\eta_1 + c) \wedge z(\eta_1)}}_{=A} - \underbrace{\frac{J_p\big((\eta_2 + c) \wedge z(\eta_2)\big) - J_p(\eta_2)}{(\eta_2 + c) \wedge z(\eta_2) - \eta_2}}_{=B} + \Big(j_p(z(\eta_2)) - j_p(z(\eta_1))\Big)$$

Since $z(\eta_1) \geq r_p$, we have $\eta_2 > \eta_1 \Rightarrow z(\eta_2) \geq z(\eta_1) \Rightarrow j_p(z(\eta_2)) \leq j_p(z(\eta_1))$. Thus, it now suffices to show that $\Gamma(\eta_2) \leq 0 \Rightarrow A \leq B$, which then implies $\Gamma(\eta_1) \leq \Gamma(\eta_2)$. This is directly implied by Lemma 2.C.1 with $\Upsilon = J_p$, $a = \underline{u} - c$, $b = \bar{u} - c$, $t = r_p$, $x = \eta_2$, $y = (\eta_2 + c) \wedge z(\eta_2)$, $w = z(\eta_2)$, $x' = \eta_1$ and $y' = (\eta_1 + c) \wedge z(\eta_1)$. *Q.E.D.*

With Lemma 2.C.2, the following lemma is almost immediate, which is key to the proof of

Proposition 2.4.3(b).

Lemma 2.C.3. *Assume Assumption 2.3.2 holds. Let η_S^* be the (unique) optimal upper-censorship signal threshold of the non-pricing seller for search goods. For any $x \in [\underline{u} - c, \bar{u} - c)$:*

(a) If $\Gamma(x) < 0$, then $\eta_S^* > x$.

(b) If $\Gamma(x) > 0$, then $\eta_S^* \leq x$, where the inequality is strict unless $x = \underline{u} - c$.

(c) If $\Gamma(x) = 0$, then $\eta_S^* = x$.

Proof. If $r_p = \bar{u} - c$, then J_p is strictly convex over $[\underline{u} - c, \bar{u} - c]$. Thus $\eta^* = \bar{u} - c$ and $\Gamma(x) < 0$ for all $x \in [\underline{u} - c, \bar{u} - c)$. So the results trivially hold. If $r_p < \bar{u} - c$, η^* (as the unique optimal threshold) must satisfy the condition in Proposition 2.3.1(b). For result (a), notice $\Gamma(x) < 0$ implies $\Gamma(x') < 0$ for all $x' \leq x$ by Lemma 2.C.2, so $\eta^* > x$. For result (b), notice $\Gamma(x) > 0$ implies $\Gamma(x') > 0$ for all $x' \geq x$ by Lemma 2.C.2, so either $\eta^* < x$ or $\eta^* = \underline{u} - c \leq x$. For result (c), notice that under Assumption 2.3.2, $\Gamma(x) = 0$ also guarantees $x \in [\eta_0, r_p]$. Indeed, if $x \in (r_p, \bar{u} - c)$, then the strict concavity of Γ over $[r_p, \bar{u} - c]$ would imply $\Gamma(x) > 0$; if $x \in [\underline{u} - c, \eta_0)$ and thus $z(x) < r_p$, then the strict convexity of Γ over $[\underline{u} - c, r_p]$ would imply $\Gamma(x) < 0$. Given this observation, the desired result is directly implied by Proposition 2.3.1(b). (Notice $r_p^{min} = r_p^{max} = r_p$ here under Assumption 2.3.2.) *Q.E.D.*

Now, we are ready to prove Proposition 2.4.3.

Proof for Proposition 2.4.3. When $c \geq \mu(\eta_E^*) - \eta_E^*$, Lemma 2.4.1 implies $z(\eta_E^*) = \mu(\eta_E^*) \leq \eta_E^* + c$. Thus under the EG seller's optimal signal, $U \wedge Z \equiv \mathbb{E}[U - c|S]$ and their common distribution solves the EG seller's problem (2.12) – (2.13). Since the SG seller's problem (2.2) – (2.3) is more constrained, the same distribution also solves it. Thus the same signal is also optimal for the SG seller and we have $\eta_S^* = \eta_E^*$. Moreover, because $z(\eta_E^*) = \mu(\eta_E^*)$, the consumer's search decision will be the same in equilibrium regardless of the product's type; because $z(\eta_E^*) \leq \eta_E^* + c$, we have $U_0 < z(\eta_E^*) \Rightarrow U_0 < \eta_E^* + c$ and thus no one will search without purchase even when the product is a search good. Thus the equilibrium outcomes of the two types of goods are the same. This proves part (a).

For part (b), notice that under Assumption 2.4.1 we have $\Gamma^E(\eta_E^*) = 0 \Rightarrow \eta_E^* \in (\underline{u} - c, \bar{u} - c)$. The results are then directly implied by Lemma 2.C.3.

Q.E.D.

2.C.4 Proof for Properties of Figure 2.4

When Assumption 2.4.1 holds and $\eta_S^* < \eta_E^*$, I prove the properties of Figure 2.4 below. I first show a lemma that is also useful somewhere else:

Lemma 2.C.4. *Assume Assumption 2.4.1 holds and $\mu(\eta_E^*) > \eta_E^* + c$. Then we have $z(\eta_S^*) > \mu(\eta_E^*)$.*

Proof. Suppose $z(\eta_S^*) \le \mu(\eta_E^*)$. Notice by Lemma 2.4.1, we have $\mu(\eta_E^*) > \eta_E^* + c \Rightarrow z(\eta_E^*) > \mu(\eta_E^*)$. Thus $z(\eta_S^*) \le \mu(\eta_E^*)$ implies $\eta_S^* < \eta_E^*$.

Notice Proposition 2.4.2 implies $\Gamma^E(\eta_E^*) = 0$, i.e., $\frac{J_p(\mu(\eta_E^*)) - J_p(\eta_E^*)}{\mu(\eta_E^*) - \eta_E^*} = j_p(\mu(\eta_E^*))$. Invoking Lemma 2.C.1 with $\Upsilon = J_p$, $[a,b] = [\underline{u} - c, \bar{u} - c]$, $t = r_p$, $(x,y) = \big(\eta_E^*, \mu(\eta_E^*)\big)$, $w = \mu(\eta_E^*)$, $(x', y') = \big(\eta_S^*, (\eta_S^* + c) \wedge z(\eta_S^*)\big)$, we then have $\frac{J_p\left((\eta_S^*+c)\wedge z(\eta_S^*)\right) - J_p(\eta_S^*)}{(\eta_S^*+c)\wedge z(\eta_S^*) - \eta_S^*} < \frac{J_p(\mu(\eta_E^*)) - J_p(\eta_E^*)}{\mu(\eta_E^*) - \eta_E^*} = j_p(\mu(\eta_E^*))$. Notice by Proposition 2.3.1(b), we must have $z(\eta_S^*) \ge r_p$ and $\frac{J_p\left((\eta_S^*+c)\wedge z(\eta_S^*)\right) - J_p(\eta_S^*)}{(\eta_S^*+c)\wedge z(\eta_S^*) - \eta_S^*} \ge j_p(z(\eta_S^*))$ (i.e., $\Gamma(\eta_S^*) \ge 0$). These together imply $j_p(\mu(\eta_E^*)) > j_p(z(\eta_S^*))$ but $\mu(\eta_E^*) \ge z(\eta_S^*) \ge r_p$, which contradicts with the strict concavity of J_p over $[r_p, \bar{u} - c]$. Thus we must have $z(\eta_S^*) > \mu(\eta_E^*)$.

Q.E.D.

Notice when Assumption 2.4.1 holds and $\eta_S^* < \eta_E^*$, Proposition 2.4.3 implies that we must have $\mu(\eta_E^*) > \eta_E^* + c$. The above lemma then implies $z(\eta_S^*) > \mu(\eta_E^*)$. Thus we have $z(\eta_S^*) > \mu(\eta_E^*) > \eta_E^* + c$ as is indicated in the figure.

Now, I turn to show the monotonicity properties of red curve, which represents the consumer's surplus change given $U_0 + p$ when the product changes from an EG to an SG. For this, we need a way to compute the consumer's surplus. The following lemma serves this role:

Lemma 2.C.5. *Under any pre-search signal, let G and H respectively denote the distributions of $U \wedge Z$ and $\mathbb{E}[U - c|S]$. Then given any outside option value u_0 and product price p, the consumer's expected utility is $\mathbb{E}_{X\sim G}[\max\{X, u_0 + p\}] - p$ if the product is an SG and is $\mathbb{E}_{X\sim H}[\max\{X, u_0 + p\}] - p$ if the product is an EG.*

Proof. When the product is an EG, the consumer will search and buy the product if and only if $\mathbb{E}[U - c|S] - p \ge u_0$. Thus her expected utility is obviously $\mathbb{E}\big[\max\{\mathbb{E}[U - c|S] - p, u_0\}\big] = \mathbb{E}\big[\max\{\mathbb{E}[U - c|S], u_0 + p\}\big] - p$. When the product is an SG, several recent papers (Kleinberg et al., 2016; Armstrong, 2017; Choi et al., 2018) have shown that the consumer surplus can be computed in a similar way with $\mathbb{E}[U-c|S]$ replaced by $U \wedge Z$. In particular, Corollary 1 in Choi et al. (2018) implies the desired result.

Q.E.D.

Recall that G_η is defined as the distribution of $U \wedge Z$ under the upper-censorship signal with threshold η; similarly, I define H_η as the distribution of $\mathbb{E}[U-c|S]$ under such a signal. Then Lemma 2.C.5 implies that given $U_0 = u_0$, the consumer's surplus change when the product changes from an EG to an SG will be:

$$\int \max\{u_0+p, x\}[G_{\eta_S^*}(dx) - H_{\eta_E^*}(dx)] = \int \max\{u_0+p-x, 0\}[G_{\eta_S^*}(dx) - H_{\eta_E^*}(dx)]$$
$$= \int_{\underline{u}-c}^{u_0+p} (u_0+p-x)[G_{\eta_S^*}(dx) - H_{\eta_E^*}(dx)] = \int_{\underline{u}-c}^{u_0+p} (G_{\eta_S^*}(x) - H_{\eta_E^*}(x))dx \tag{2.24}$$

where the first equality holds because $\int x[G_{\eta_S^*}(dx) - H_{\eta_E^*}(dx)] = 0$ by the mean-preserving property; the second equality holds even if $u_0 + p < \underline{u} - c$ because both $G_{\eta_S^*}$ and $H_{\eta_E^*}$ support within $[\underline{u}-c, \bar{u}-c]$; the third equality holds by integration-by-parts. The derivative of the last expression above w.r.t. u_0+p is $G_{\eta_S^*}(u_0+p) - H_{\eta_E^*}(u_0+p)$.

Recall that the formula of G_η is provided in Lemma 2.B.2. The formula of H_η is easily seen to be:

$$H_\eta(x) = \begin{cases} F_{U-c}(x) & x < \eta \\ F_{U-c}(\eta) & \eta \le x < \mu(\eta) \\ 1 & \mu(\eta) \le x \end{cases} \tag{2.25}$$

When $\eta_S^* < \eta_E^* < \eta_E^* + c < \mu(\eta_E^*) < z(\eta_S^*)$, the relative patterns of $G_{\eta_S^*}$ and $H_{\eta_S^*}$ are illustrated in Figure 2.10. In particular, we can see $G_{\eta_S^*}(x) - H_{\eta_E^*}(x)$ is strictly negative over $(\eta_S^*, \eta_E^* + c) \cup (\mu(\eta_E^*), z(\eta_S^*))$, is strictly positive over $(\eta_E^* + c, \mu(\eta_E^*))$, and is zero elsewhere. Thus we have the monotonicity of the surplus change as a function of $u_0 + p$ in each of these intervals as is depicted in Figure 2.4.

2.C.5 Proof for Proposition 2.4.4

Given any (F_U, c, r_p) satisfying Assumption 2.4.2 and $\kappa > 0$, I define $\delta := \mu(r_p) - r_p - c$ and $A := \min_{\eta \in [\underline{u}-c, r_p]} \int_{\eta+c}^{\eta+c+\delta} (\eta + c + \delta - x) f_U(x)dx$, where the minimum is achievable by the

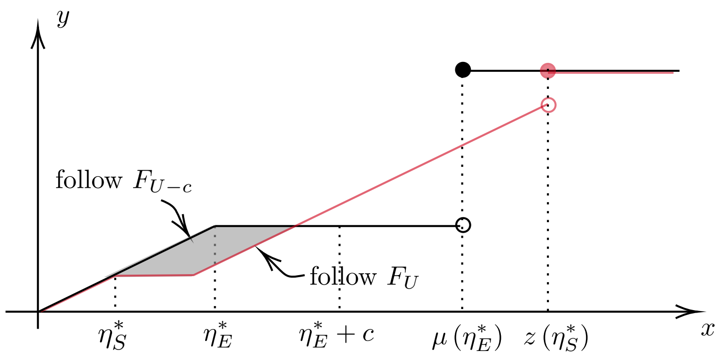

Figure 2.10: Comparison between $G_{\eta_S^*}$ and $H_{\eta_E^*}$. The black curve represents $H_{\eta_E^*}$ and the red curve represents $G_{\eta_S^*}$. For clarity, I have picked F_U to be a uniform distribution.

Weierstrass Theorem. By condition (2) in Assumption 2.4.2, we have $\delta > 0$. Since f_{U-c} has full support on $[\underline{u} - c, \bar{u} - c]$, this further implies $A > 0$. Also define $M_f := \sup_{x \in [\underline{u}, \bar{u}]} f_U(x)$. It is finite since f_U is log-concave by Assumption 2.4.2.

Now, I choose ν to be the largest strictly positive number satisfying:

$$\frac{\kappa c}{2(\bar{u} - \underline{u})}\left(A - [1 - F_U(r_p)]\nu - \frac{1}{2}M_f \nu^2\right) - [1 - F_U(r_p)]\nu \geq 0 \tag{2.26}$$

Notice that the LHS is continuous and decreasing in ν, and is strictly positive when $\nu = 0$. Thus such a ν exists.

Now, to prove Proposition 2.4.4, it suffices to show the following result:

Proposition 2.C.1. *Given (F_u, c, r_p) satisfying Assumption 2.4.2 and $\kappa > 0$, let ν be chosen as above. Assume J_p satisfies Assumption 2.3.2 with its mode being r_p and (i) $j(r_p + \nu) < j(\eta_0)$; (ii) $j(\eta_0) > \kappa \frac{J_p(\bar{u}-c) - J_p(\underline{u}-c)}{\bar{u} - \underline{u}}$. Then we have $cs^S < cs^E$.*

Proof. I first note that under the current assumptions, Assumption 2.4.1 holds.[50] By Proposition 2.4.2, this implies that the equilibrium outcome with experience goods is unique and $\Gamma^E(\eta_E^*) = 0$. Due to the strict convexity-concavity of J_p, $\Gamma^E(\eta_E^*) = 0$ further implies $\eta_E^* < r_p$ and $\mu(\eta_E^*) > r_p$.

Moreover, under the strictly unimodality of J_p and condition (3) in Assumption 2.4.2, Propositions 2.3.1 and 2.3.2 imply that the equilibrium outcome with search goods is also unique with $\eta_S^* \in [\eta_0, r_p]$ and $\Gamma(\eta_S^*) = 0$. To see this, notice condition (3) in Assumption 2.4.1 states that

[50]The only non-trivial part for checking this is to show $\Gamma^E(\underline{u} - c) < 0$. For this, notice that $z(\underline{u} - c) < r_p$ by condition (3) in Assumption 2.4.1. Since Lemma 2.4.1 implies $\mu(\underline{u} - c) \leq z(\underline{u} - c)$, we then have $\mu(\underline{u} - c) < r_p$. The strict convexity of J_p over $[\underline{u} - c, r_p]$ then implies $\Gamma^E(\underline{u} - c) < 0$.

$r_p < \bar{u} - c$ and $z(\underline{u} - c) < r_p$. The former of these implies that η_S^* must satisfy the conditions in Proposition 2.3.1(b); the latter of these, together with the continuity of $z(\cdot)$ (see Lemma 2.B.1), implies $\eta_0 > \underline{u} - c$ and thus rules out the possibility of $\eta_S^* = \underline{u} - c$.

I now show the following properties regarding η_E^* and η_S^*:

Observation (1). We have $\eta_0 \leq \eta_S^* < \eta_E^* < \eta_E^* + c < \mu(\eta_E^*) < z(\eta_S^*) \leq \mu(\eta_E^*) + \nu$ and $F_U(\eta_E^* + c) - F_U(\eta_S^* + c) \geq \big(A - [1 - F_U(r_p)]\nu - \frac{1}{2}M_f\nu^2\big)/c$.

Subproof. The proof consists of several parts.

- Part 1: $\mu(\eta_E^*) - (\eta_E^* + c) > \delta$.

By definition, $\delta = \mu(r_p) - (r_p + c)$, which is > 0 by Assumption 2.4.2. Since $\eta_E^* < r_p$ as mentioned earlier, it suffices to show $\mu(\eta) - \eta = \mathbb{E}[U - c - \eta | U - c - \eta \geq 0]$ is decreasing in η. Notice the log-concavity of f_U implies that $f_U(x+(c+\eta))$ is log-supermodular in $\big(x, -(c+\eta)\big)$.[51] Thus $f_{U-c-\eta}(x) = f_U(x + (c + \eta))$ decreases in the likelihood-ratio order when $c + \eta$ increases. This implies that we indeed have $\mathbb{E}[U - c - \eta | U - c - \eta \geq 0]$ decreasing in η.[52]

- Part 2: $z(\eta_S^*) > \eta_S^* + c$.

By Proposition 2.3.1, we have $\eta_S^* \leq r_p$. The similar proof as in Part 1 then implies $\mu(\eta_S^*) - (\eta_S^* + c) > \delta$. Since $z(\eta_S^*) \geq \mu(\eta_S^*)$ by Lemma 2.4.1, we have $z(\eta_S^*) > \eta_S^* + c$.

- Part 3: $0 < z(\eta_S^*) - \mu(\eta_E^*) < \nu$.

Given the conclusion in Part 1 and that Assumption 2.4.1 holds here as I mentioned earlier, the result $z(\eta_S^*) - \mu(\eta_E^*) > 0$ is directly implied by Lemma 2.C.4.

Now, suppose $z(\eta_S^*) - \mu(\eta_E^*) \geq \nu$. Since $\Gamma^E(\eta_E^*) = 0 \Rightarrow \mu(\eta_E^*) > r_p$ due to the strict convexity of J_p over $[\underline{u} - c, r_p]$, we will then have $z(\eta_S^*) > r_p + \nu$. The condition $j(r_p + \nu) < j(\eta_0)$ together with the strict quasi-concavity of $j(\cdot)$ then implies that $j\big(z(\eta_S^*)\big) < j(x)$ for all $x \in \big[\eta_0, z(\eta_S^*)\big)$. Since $\eta_S^* \geq \eta_0$ as is mentioned earlier, we then must have $\Gamma(\eta_S^*) = \frac{J_p((\eta_S^*+c)\wedge z(\eta_S^*)) - J_p(\eta_S^*)}{(\eta_S^*+c)\wedge z(\eta_S^*) - \eta_S^*} - j(z(\eta_S^*)) > 0$, which contradicts with my earlier conclusion of $\Gamma(\eta_S^*) = 0$. Therefore, we must have $z(\eta_S^*) - \mu(\eta_E^*) < \nu$.

- Part 4: $F_U(\eta_E^* + c) - F_U(\eta_S^* + c) \geq \big(A - [1 - F_U(r_p)]\nu - \frac{1}{2}M_f\nu^2\big)/c$.

By the definitions of $z(\cdot)$ and $\mu(\cdot)$, we have:

$$\int_{\eta_S^*+c}^{\bar{u}} \big[\big(x - z(\eta_S^*)\big)_+ - c\big] F_u(dx) = 0; \quad \int_{\eta_E^*+c}^{\bar{u}} [x - \mu(\eta_E^*) - c] F_u(dx) = 0$$

[51] See, for example, Lemma 2.6.2(b) in Topkis (1998). Apply it to $\log\big(f_U(\cdot)\big)$.

[52] See, e.g., Theorem 1.4.6 in Müller & Stoyan (2002).

Taking difference of these two equations, we get:

$$\int_{\eta_S^*+c}^{\eta_E^*+c} cF_U(dx) = \int_{\eta_S^*+c}^{\bar{u}} \big(x - z(\eta_S^*)\big)_+ F_u(dx) - \int_{\eta_E^*+c}^{\bar{u}} [x - \mu(\eta_E^*)]F_u(dx) \tag{2.27}$$

The LHS above equals to $c[F_U(\eta_E^* + c) - F_U(\eta_S^* + c)]$. Notice $\int_{\eta_S^*+c}^{\bar{u}} \big(x - z(\eta_S^*)\big)_+ F_u(dx) = \int_{z(\eta_S^*)}^{\bar{u}} [x - z(\eta_S^*)]F_u(dx)$ because $z(\eta_S^*) > \eta_S^* + c$ as shown in Part 2. Thus the RHS equals to

$$\int_{z(\eta_S^*)}^{\bar{u}} [\mu(\eta_E^*) - z(\eta_S^*)]F_U(dx) + \int_{\mu(\eta_E^*)}^{z(\eta_S^*)} [\mu(\eta_E^*) - x]F_U(dx) + \int_{\eta_E^*+c}^{\mu(\eta_E^*)} [\mu(\eta_E^*) - x]F_U(dx)$$

Because $z(\eta_S^*) \geq r_p$ and $z(\eta_S^*) - \mu(\eta_E^*) < \nu$, the first term above is greater than $-[1 - F_U(r_p)]\nu$. For the second term, we have $\int_{\mu(\eta_E^*)}^{z(\eta_S^*)} [\mu(\eta_E^*) - x]F_U(dx) \geq \int_{\mu(\eta_E^*)}^{z(\eta_S^*)} [\mu(\eta_E^*) - x]M_f dx \geq \int_{\mu(\eta_E^*)}^{\mu(\eta_E^*)+v} [\mu(\eta_E^*) - x]M_f dx = -M_f\nu^2/2$, where the second inequality is again due to $z(\eta_S^*) - \mu(\eta_E^*) < \nu$. For the third term, we have $\int_{\eta_E^*+c}^{\mu(\eta_E^*)} [\mu(\eta_E^*) - x]F_U(dx) \geq \int_{\eta_E^*+c}^{\eta_E^*+c+\delta} [\eta_E^* + c + \delta - x]F_U(dx) \geq A$, where the first inequality is because $\mu(\eta_E^*) - (\eta_E^* + c) > \delta$ as proved in Part 1, and the second inequality is due to the definition of A and that $\eta_E^* \in [\underline{u} - c, r_p]$. These together imply that the RHS of equation (2.27) is greater than $A - \big(1 - F_U(r_p)\big)\nu - \frac{1}{2}M_f\nu^2$. The desired inequality is thus implied by equation (2.27).

- Part 5: $\eta_E^* > \eta_S^* \geq \eta_0$.

Since $\nu > 0$ satisfies condition (2.26), we must have $A - \big(1 - F_U(r_p)\big)\nu - \frac{1}{2}M_f\nu^2 > 0$. Thus $\eta_E^* > \eta_S^*$ is directly implied by the result of Part 4. Moreover, $\eta_S^* \geq \eta_0$ since $\eta_S^* \in [\eta_0, r_p]$, as I have mentioned in the main proof. □

Recall that I have defined G_η and H_η as the distributions of $U \wedge Z$ and $\mathbb{E}[U - c|S]$ under the upper-censorship signal with threshold η respectively (see expressions (2.23) and (2.25)). In Appendix 2.C.4, I have shown that given $U_0 = u_0$, a consumer's surplus change when the product changes from an EG to an SG will be $\int_{\underline{u}-c}^{u_0+p}(G_{\eta_S^*}(x) - H_{\eta_E^*}(x))dx$, where the patterns of $G_{\eta_S^*}$ and $H_{\eta_E^*}$ are illustrated in Figure 2.10 when $\eta_S^* < \eta_E^* < \eta_E^* + c < \mu(\eta_E^*) < z(\eta_S^*)$. This implies $cs^S - cs^E = \int[\int_{\underline{u}-c}^{y}(G_{\eta_S^*}(x) - H_{\eta_E^*}(x))dx]j_p(y)dy$. To show this is negative, I prove two more observations below.

Observation (2). $\int_{\eta_E^*}^{\eta_E^*+c}[\int_{\underline{u}-c}^{y}(H_{\eta_E^*}(x)-G_{\eta_S^*}(x))dx]j_p(y)dy \geq \frac{1}{2}j_p(\eta_0)c\big(A-[1-F_U(r_p)]\nu-\frac{1}{2}M_f\nu^2\big)$.

Subproof. First, one can show $\int_{\underline{u}-c}^{\eta_E^*+c}(H_{\eta_E^*}(x) - G_{\eta_S^*}(x))dx = [F_U(\eta_E^*+c) - F_U(\eta_S^*+c)]c$. To see this, notice that the integral represents the space of the gray area in Figure 2.10. The area's height is $F_U(\eta_E^*+c)-F_U(\eta_S^*+c)$ at any x, and its width is c at any y. Thus the integral equals to $[F_U(\eta_E^*+c)-F_U(\eta_S^*+c)]c$.[53]

Next, notice that over the region $[\eta_E^*, \eta_E^*+c]$, we have $H_{\eta_E^*}(x)-G_{\eta_S^*}(x) \leq F_U(\eta_E^*+c)-F_U(\eta_S^*+c)$. This implies that for any $y \in [\eta_E^*, \eta_E^*+c]$, we have $\int_y^{\eta_E^*+c}(H_{\eta_E^*}(x)-G_{\eta_S^*}(x))dx \leq [F_U(\eta_E^*+c)-F_U(\eta_S^*+c)](\eta_E^*+c-y)$. Together with the conclusion $\int_{\underline{u}-c}^{\eta_E^*+c}(H_{\eta_E^*}(x)-G_{\eta_S^*}(x))dx = [F_U(\eta_E^*+c)-F_U(\eta_S^*+c)]c$ above, this implies that for any $y \in [\eta_E^*, \eta_E^*+c]$, $\int_{\underline{u}-c}^{y}(H_{\eta_E^*}(x)-G_{\eta_S^*}(x))dx \geq [F_U(\eta_E^*+c)-F_U(\eta_S^*+c)](y-\eta_E^*)$. This further implies $\int_{\eta_E^*}^{\eta_E^*+c}[\int_{\underline{u}-c}^{y}(H_{\eta_E^*}(x)-G_{\eta_S^*}(x))dx]j_p(y)dy \geq \int_{\eta_E^*}^{\eta_E^*+c}[F_U(\eta_E^*+c)-F_U(\eta_S^*+c)](y-\eta_E^*)j_p(y)dy$.

Now, it suffices to show the RHS of the last inequality is greater than $\frac{1}{2}j_p(\eta_0)c\big(A-[1-F_U(r_p)]\nu-\frac{1}{2}M_f\nu^2\big)$. Since $\eta_0<\eta_E^*<r_p$, we have $j_p(\eta_E^*)>j_p(\eta_0)$. Moreover, because $\Gamma^E(\eta_E^*)=0$, we must have $j_p(\mu(\eta_E^*)) \geq j_p(\eta_E^*)$. By the strict quasi-concavity of j_p, these together imply $j_p(y)>j_p(\eta_0)$ for all $y \in [\eta_E^*, \mu^*(\eta_E^*)]$. Thus we have $\int_{\eta_E^*}^{\eta_E^*+c}[F_U(\eta_E^*+c)-F_U(\eta_S^*+c)](y-\eta_E^*)j_p(y)dy \geq [F_U(\eta_E^*+c)-F_U(\eta_S^*+c)]j_p(\eta_0)\int_{\eta_E^*}^{\eta_E^*+c}(y-\eta_E^*)dy = \frac{1}{2}j_p(\eta_0)[F_U(\eta_E^*+c)-F_U(\eta_S^*+c)]c^2 \geq \frac{1}{2}j_p(\eta_0)c\big(A-[1-F_U(r_p)]\nu-\frac{1}{2}M_f\nu^2\big)$, where the last inequality is implied by Observation (1). □

Observation (3). $\int_{\underline{u}-c}^{y}(G_{\eta_S^*}(x)-H_{\eta_E^*}(x))dx \leq [1-F_U(r_p)]\nu$ for any $y \in [\underline{u}-c, \bar{u}-c]$.

Subproof. As is illustrated in Figure 2.4, $\int_{\underline{u}-c}^{y}(G_{\eta_S^*}(x)-H_{\eta_E^*}(x))dx$ achieves its maximum at $y=\mu(\eta_E^*)$. It thus suffices to show $\int_{\underline{u}-c}^{\mu(\eta_E^*)}(G_{\eta_S^*}(x)-H_{\eta_E^*}(x))dx \leq [1-F_U(r_p)]\nu$. By the mean-preserving condition, we have $\int_{\underline{u}-c}^{z(\eta_S^*)}(G_{\eta_S^*}(x)-H_{\eta_E^*}(x))dx = 0$ (notice $G_{\eta_S^*}(x)=H_{\eta_E^*}(x)=1$ for $x>z(\eta_S^*)$). Since $\mu(\eta_E^*)>r_p$, we have $G_{\eta_S^*}(x)-H_{\eta_E^*}(x)>F_U(r_p)-1$ for all $x \geq \mu(\eta_E^*)$. Since $z(\eta_S^*)-\mu(\eta_E^*) \leq \nu$ by Observation (1), this implies $\int_{\mu(\eta_E^*)}^{z(\eta_S^*)}(G_{\eta_S^*}(x)-H_{\eta_E^*}(x))dx \geq (F_U(r_p)-1)\nu$. Together with the condition $\int_{\underline{u}-c}^{z(\eta_S^*)}(G_{\eta_S^*}(x)-H_{\eta_E^*}(x))dx=0$, this further implies $\int_{\underline{u}-c}^{y}(G_{\eta_S^*}(x)-H_{\eta_E^*}(x))dx \leq [1-F_U(r_p)]\nu$. □

Observation (2), together with condition $j_p(\eta_0)>\kappa\frac{J_p(\bar{u}-c)-J_p(\underline{u}-c)}{\bar{u}-\underline{u}}$, imply $\int_{\eta_E^*}^{\eta_E^*+c}[\int_{\underline{u}-c}^{y}(H_{\eta_E^*}(x)-G_{\eta_S^*}(x))dx]j_p(y)dy > \frac{1}{2}\kappa\frac{J_p(\bar{u}-c)-J_p(\underline{u}-c)}{\bar{u}-\underline{u}}c\big(A-[1-F_U(r_p)]\nu-\frac{1}{2}M_f\nu^2\big)$. Observation (3) implies $\int_{[\underline{u}-c,\bar{u}-c]\setminus(\eta_E^*,\eta_E^*+c)}[\int_{\underline{u}-c}^{y}(G_{\eta_S^*}(x)-H_{\eta_E^*}(x))dx]j_p(y)dy \leq [1-F_U(r_p)]\nu[J_p(\bar{u}-c)-J_p(\underline{u}-c)]$.

[53] Notice this argument does not require the left and right sides of the area to be linear. We only need them to be parallel, which is true because $H_{\eta_E^*}$ follows F_{U-c} over $[\eta_S^*, \eta_E^*]$ and $G_{\eta_S^*}$ follows F_U over $[\eta_S^*+c, \eta_E^*+c]$.

Combining these inequalities, we know $\int[\int_{\underline{u}-c}^{y}(G_{\eta_S^*}(x) - H_{\eta_E^*}(x))dx]j_p(y)dy$ is less than

$$\Big[[1 - F_U(r_p)]\nu - \frac{\kappa c}{2(\bar{u} - \underline{u})}\big(A - [1 - F_U(r_p)]\nu - \frac{1}{2}M_f\nu^2\big)\Big][J_p(\bar{u} - c) - J_p(\underline{u} - c)]$$

which is strictly negative as ν satisfies condition (2.26) by construction. Thus we indeed have $cs^S - cs^E < 0$.

Q.E.D.

2.D Proofs for Section 2.5

2.D.1 Proof for Proposition 2.5.1

Proof. Assume $J_{p_2}^2$ dominates $J_{p_1}^1$ in likelihood ratio order (abbr. $J_{p_1}^1 \preceq_{LRD} J_{p_2}^2$). This means that for any $x < y$, we have $j_{p_1}^1(x)j_{p_2}^2(y) \geq j_{p_1}^1(y)j_{p_2}^2(x)$. Let $r_{p_1}^1$ and $r_{p_2}^2$ denote the modes of $j_{p_1}^1$ and $j_{p_2}^2$ over $[\underline{u} - c, \bar{u} - c]$ respectively. Let (η_0^1, Γ_1) and (η_0^2, Γ_2) be the pair of (η_0, Γ) defined in Section 2.3.3 given distributions $J_{p_1}^1$ and $J_{p_2}^2$ respectively. Notice $J_{p_1}^1 \preceq_{LRD} J_{p_2}^2$ implies $r_{p_1}^1 \leq r_{p_2}^2$. By the definition of η_0, This further implies $\eta_0^1 \leq \eta_0^2$ because $z(\cdot)$ is increasing.

Let η_1^* and η_2^* denote the thresholds of the optimal upper-censorship signals given $J_{p_1}^1$ and $J_{p_2}^2$ respectively. Under Assumption 2.3.2, they are fully characterized by conditions in Proposition 2.3.1 (given $J_{p_1}^1$ and $J_{p_2}^2$ respectively). Consider following cases:

Case 1: $r_{p_1}^1 = \bar{u} - c$.

In this case, $r_{p_2}^2$ also equals to $\bar{u} - c$ since $r_{p_2}^2 \geq r_{p_1}^1$. Thus $\eta_2^* = \bar{u} - c \geq \eta_1^*$.

Case 2: $r_{p_1}^1 < \bar{u} - c$.

According to Proposition 2.3.1(b), in this case we must have either (i) $\Gamma_1(\eta_1^*) = 0$ or (ii) $\Gamma_1(\eta_1^*) \geq 0$ and $\eta_1^* = \underline{u} - c$. If (ii) holds, then $\eta_1^* \leq \eta_2^*$ trivially. Thus I assume w.l.g. that (i) holds. I then have the following observation:

Observation. $\Gamma_1(\eta_1^*) = 0 \Rightarrow \Gamma_2(\eta_1^*) \leq 0$.

Subproof. Notice if $j^1_{p_1}(z(\eta_1^*))$ and $j^2_{p_2}(z(\eta_1^*))$ are non-zero, we have:

$$\begin{aligned}\Gamma_1(\eta_1^*)=0 \Leftrightarrow &\int_{\eta_1^*}^{(\eta_1^*+c)\wedge z(\eta_1^*)} \big[j^1_{p_1}(t)-j^1_{p_1}(z(\eta_1^*))\big]dt=0\\ \Leftrightarrow &\int_{\eta_1^*}^{(\eta_1^*+c)\wedge z(\eta_1^*)} \big[\frac{j^1_{p_1}(t)}{j^1_{p_1}(z(\eta_1^*))}-1\big]dt=0\\ \Rightarrow &\int_{\eta_1^*}^{(\eta_1^*+c)\wedge z(\eta_1^*)} \big[\frac{j^2_{p_2}(t)}{j^2_{p_2}(z(\eta_1^*))}-1\big]dt\le 0\\ \Leftrightarrow &\Big(\int_{\eta_1^*}^{(\eta_1^*+c)\wedge z(\eta_1^*)} \big[j^2_{p_2}(t)-j^2_{p_2}(z(\eta_1^*))\big]dt\Big)\le 0 \qquad \Leftrightarrow \Gamma_2(\eta_1^*)\le 0\end{aligned}$$

where the third row uses $J^1_{p_1}\preceq_{LRD} J^2_{p_2}\Rightarrow \frac{j^1_{p_1}(t)}{j^1_{p_1}(z(\eta_1^*))}\ge\frac{j^2_{p_2}(t)}{j^2_{p_2}(z(\eta_1^*))}$ for $t\le z(\eta_1^*)$.

Now, it suffices to consider either $j^1_{p_1}(z(\eta_1^*))$ or $j^2_{p_2}(z(\eta_1^*))$ being zero. Let I_1 and I_2 denote the supports of j_1 and j_2 within $[\underline{u}-c,\bar{u}-c]$ respectively. Notice the unimodality assumption implies that these supports are intervals.

- Suppose $j^1_{p_1}(z(\eta_1^*))=0$. Then $\Gamma_1(\eta_1^*)=0$ implies that $j^1_{p_1}\equiv 0$ on $[\eta_1^*,(\eta_1^*+c)\wedge z(\eta_1^*)]$. Notice that since $\eta_1^*\le r^1_{p_1}$, this interval must be to the left of I_1. Since $J^1_{p_1}\preceq_{LRD} J^2_{p_2}$ implies $\inf I_1\le\inf I_2$, we must have $j^2_{p_2}\equiv 0$ on $[\eta_1^*,(\eta_1^*+c)\wedge z(\eta_1^*)]$ too. This further implies $\Gamma_2(\eta_1^*)\le 0$.
- Suppose $j^2_{p_2}(z(\eta_1^*))=0$. If $z(\eta_1^*)$ is to the left of I_2, then $j^2_{p_2}(x)=0$ for all $x\le z(\eta_1^*)$ and thus $\Gamma_2(\eta_1^*)=0$. If $z(\eta_1^*)$ is to the right of I_2, then because $J^1_{p_1}\preceq_{LRD} J^2_{p_2}$ implies $\sup I_1\le\sup I_2$, we must also have $j^1_{p_1}(z(\eta_1^*))=0$. This implies $\Gamma_2(\eta_1^*)\le 0$ as has been shown above.

□

By the observation, we have $\Gamma_2(\eta_1^*)\le 0$. Lemma 2.C.3 in Appendix 2.C.3 then implies $\eta_2^*\ge\eta_1^*$.

Q.E.D.

2.D.2 Proof for Proposition 2.5.2

Proof. Let F_ϵ and f_ϵ denote the CDF and PDF of ϵ respectively. For any p, let J_p^W be the conditional distribution of U_0+p conditioning on W. Then, $J_p^W(x)=F_\epsilon(x-(p+W))$. Notice the log-concavity of f_ϵ implies that $f_\epsilon(x-(p+W))$ is log-supermodular in $(x,p+W)$,[54] and

[54]See, for example, Topkis (1998) Lemma 2.6.2(b). Apply it to $\log\big(f_\epsilon(\cdot)\big)$.

thus $J_p^W(x) = F_\epsilon(x-(p+W))$ increases in the likelihood-ratio order when $p+W$ increases. Also notice f_ϵ being strictly log-concave implies that J_p^W satisfies Assumption 2.3.2. The optimal η^* is thus unique and all optimal signals are outcome-equivalent. The desired result for a non-pricing seller is directly implied by Proposition 2.5.1.

For a pricing seller, I prove the result with two steps:

- **Step 1.** Given any p and realized W, the seller's problem is just the non-pricing seller's optimization (2.2) – (2.3) with $J_p(x)$ replaced by $J_p^W(x) = F_\epsilon(x-(p+W))$. Notice this optimization depends on (p, W) only through $\ell := p + W$. Given any ℓ, let $Q^*(\ell)$ denote the optimal value and let $\eta^*(\ell)$ denote the threshold of the optimal upper-censorship signal. Then, the argument above for a non-pricing seller implies $\eta^*(\ell)$ is a singleton and increases in ℓ. Moreover, $Q^*(\ell)$ decreases in ℓ since the objective function decreases in ℓ.
- **Step 2.** Given $Q^*(\cdot)$, the optimization over p can be written as:

$$\max_p\{pQ^*(p+W)\} \Longleftrightarrow \max_\ell\{(\ell-W)Q^*(\ell)\}$$

  Since $Q^*(\ell)$ is decreasing in ℓ, the objective function has increasing differences in (ℓ, W). Thus the optimal ℓ increases (in the strong set order) in W. The desired result is then implied by $\eta^*(\ell)$ being increasing in ℓ.[55]

Q.E.D.

2.D.3 Proof for Corollary 2.5.1

Proof. Without discrimination based on W, the setting is the same as in Section 2.3. Since log-concavity is preserved under convolution, condition (A1) implies that U_0 (and thus U_0+p) admits a log-concave density. The strict log-concavity of f_ϵ and condition (A2) further imply that U_0 has full support over $\mathbb{R}$ and has a single mode. Together with the log-concavity, these imply that the density of U_0+p satisfies Assumption 2.3.2. Thus, without discrimination, there is a unique optimal upper-censorship signal for the seller and all optimal signals are outcome equivalent. Let η^*_{nd} denote the threshold of this optimal upper-censorship signal.

With discrimination, as is shown in the proof of Proposition 2.5.2, for any W there is a unique optimal upper-censorship signal and all optimal signals are outcome-equivalent. Let

[55] Notice that the log-concavity of f_ϵ implies log-concavity of $F_\epsilon(\cdot - W)$ given any realized W. Thus the existence of optimal solution is guaranteed by Proposition 2.3.3.

$\eta^*(W)$ denote the threshold of the optimal upper-censorship signal.

By Proposition 2.5.2, $\eta^*(W)$ is (weakly) increasing in W. Let $w^* = \inf\{W : \eta^*(W) \geq \eta^*_{nd}\}$. Then we have $W > w^* \Rightarrow \eta^*(W) \geq \eta^*_{nd}$ and $W < w^* \Rightarrow \eta^*(W) < \eta^*_{nd}$. This implies that for $W > w^*$, the pre-search signal is more informative with discrimination; for $W < w^*$, the pre-search signal is less informative with discrimination. Thus we have the conclusion in the corollary. *Q.E.D.*

2.D.4 Proof for Proposition 2.5.3

Proof. Given any search cost c, let $z(\eta; c)$ equal to the $z(\eta)$ defined in Section 2.3.3; let $\Gamma(\eta; c)$ equal to the $\Gamma(\eta)$ defined in equation (2.9); let $r_p(c)$ denote the maximum point of j_p over $[\underline{u} - c, \bar{u} - c]$. Let $t^* \in [-\infty, \bar{u}]$ be the maximum point of j_p over $(-\infty, \bar{u}]$. Notice that $r_p(c) = t^*$ if and only if $t^* \in [\underline{u} - c, \bar{u} - c]$.

Pick any search costs c_1 and c_2. Let η_1 and η_2 denote the thresholds of optimal upper-censorship signals given these two search costs respectively, which are unique under condition (A1) (by Proposition 2.3.2) and are characterized by Proposition 2.3.1. Let $z_1 := z(\eta_1; c_1)$ and $z_2 := z(\eta_2; c_2)$. We have the following observation:

Observation (1). If $\eta_1 < \eta_2$ and $\eta_1 + c_1 < \eta_2 + c_2$, then $z_1 \geq z_2$.

Subproof. Since $\underline{u} \leq \eta_i + c_i \leq \bar{u}$ $(i = 1, 2)$, condition $\eta_1 + c_1 < \eta_2 + c_2$ further implies $\underline{u} < \eta_2 + c_2$ and $\eta_1 + c_1 < \bar{u}$. Consider two cases:

- Case 1: $\eta_2 + c_2 = \bar{u}$.

  In this case, full disclosure is optimal under search cost c_2 (since $\eta_2 = \bar{u} - c_2$) but is not optimal under search cost c_1 (since $\eta_1 < \bar{u} - c_1$). Thus J_p is convex over $[\underline{u} - c_2, \bar{u} - c_2]$ but is not convex over $[\underline{u} - c_1, \bar{u} - c_1]$. This implies that $\bar{u} - c_2 \leq t^* < \bar{u} - c_1$.

  Now, suppose $z_1 < z_2$. Since $z_2 \leq \bar{u} - c_2$ (see Lemma 2.B.1(a)), we have:

  $$\underline{u} - c_1 \leq z_1 < z_2 \leq \bar{u} - c_2 \leq t^* < \bar{u} - c_1$$

  Notice $\underline{u} - c_1 < t^* < \bar{u} - c_1$ implies $r_p(c_1) = t^* < \bar{u} - c_1$ and thus η_1 is characterized by the condition in Proposition 2.3.1(b). Then, by Lemma 2.B.3, we must have $z_1 \geq r_p(c_1) = t^*$. This contradicts with the fact $z_1 < t^*$ derived above. Thus $z_1 \geq z_2$.

- Case 2: $\eta_2 + c_2 < \bar{u}$.

In this case, full disclosure is suboptimal given both search costs c_1 and c_2, and thus J_p is not convex over either $[\underline{u} - c_1, \bar{u} - c_1]$ or $[\underline{u} - c_2, \bar{u} - c_2]$. This implies that the optimal signals are characterized by the condition in Proposition 2.3.1(b) under both search costs. Thus $\Gamma(\eta_i; c_i) \geq 0$ for $i = 1, 2$. Moreover, since $\eta_2 + c_2 > \underline{u}$ as is shown earlier, we further have $\Gamma(\eta_2; c_2) = 0$.

Then, it can be shown that:

$$j_p(z_1) \leq \frac{J_p\big((\eta_1 + c_1) \wedge z_1\big) - J_p(\eta_1)}{(\eta_1 + c_1) \wedge z_1 - \eta_1} \leq \frac{J_p\big((\eta_2 + c_2) \wedge z_2\big) - J_p(\eta_2)}{(\eta_2 + c_2) \wedge z_2 - \eta_2} = j_p(z_2)$$

The first inequality is equivalent to $\Gamma(\eta_1; c_1) \geq 0$. The last equality is equivalent to $\Gamma(\eta_2; c_2) = 0$. Given the last equality, the second inequality is implied by Lemma 2.C.1 in Appendix 2.C.3 with $\Upsilon = J_p$, $a = \min\{\underline{u} - c_1, \underline{u} - c_2\}$, $b = \bar{u}$, $t = t^*$, $x = \eta_2$, $y = (\eta_2 + c_2) \wedge z_2$, $w = z_2$, $x' = \eta_1$ and $y' = (\eta_1 + c_1) \wedge z_1$.

By Lemma 2.B.3, we know that $z_1 \geq r_p(c_1)$ and $z_2 \geq r_p(c_2)$. As is mentioned above, J_p is not convex over either $[\underline{u} - c_1, \bar{u} - c_1]$ or $[\underline{u} - c_2, \bar{u} - c_2]$. These together imply that z_1 and z_2 lie in the region where J_p is strictly concave and j_p is strictly decreasing. Thus the fact $j_p(z_1) \leq j_p(z_2)$ derived above implies $z_1 \geq z_2$.

□

Now, I show part (a) and part (b) of the proposition in sequence.

Part (a): $c_1 < c_2 \Rightarrow \eta_1 + c_1 \leq \eta_2 + c_2$.[56]

Suppose $c_1 < c_2$ but $\eta_1 + c_1 > \eta_2 + c_2$. Then we must have $\eta_1 > \eta_2$. According to Observation (1) above, conditions $\eta_1 > \eta_2$ and $\eta_1 + c_1 > \eta_2 + c_2$ imply $z_1 \leq z_2$. (When using the observation, interchange the positions of (c_1, η_1, z_1) and (c_2, η_2, z_2).)

However, because $\eta_1 + c_1 > \eta_2 + c_2$, the posterior belief on U after learning $U - c_1 \geq \eta_1$ is superior to that after learning $U - c_2 \geq \eta_2$ (in terms of first-order stochastic dominance). Together with the assumption that $c_1 < c_2$, it is easy to see that $z_1 > z_2$. This contradicts with the result $z_1 \leq z_2$ above. Thus $c_1 < c_2$ must imply $\eta_1 + c_1 \leq \eta_2 + c_2$.

Part (b): Given any realized U_0, the consumer surplus is decreasing in c.

Assume $c_1 < c_2$. I first show $z_1 \geq z_2$ by considering the following two cases.

- Case 1: $\eta_1 < \eta_2$.

[56] Recall that η_i is a threshold on $U - c_i$ according to Definition 2.3.1. Thus $\eta_i + c_i$ is the corresponding threshold on U and the signal is more informative on U if $\eta_i + c_i$ is larger.

In this case, $c_1 < c_2$ further implies $\eta_1 + c_1 < \eta_2 + c_2$. Then by Observation (1) above, we have $z_1 \geq z_2$.

- Case 2: $\eta_1 \geq \eta_2$.

  Suppose $z_1 < z_2$. We have:

$$
\begin{aligned}
0 =& \mathbb{E}[(U - z_1)_+ - c_1 | U - c_1 \geq \eta_1] \\
\geq& \mathbb{E}[\big((U - c_1) + c_1 - z_1\big)_+ - c_1 | U - c_1 \geq \eta_2] \\
\geq& \mathbb{E}[\big((U - c_1) + c_2 - z_1\big)_+ - c_2 | U - c_1 \geq \eta_2] \\
\geq& \mathbb{E}[\big((U - c_2) + c_2 - z_1\big)_+ - c_2 | U - c_2 \geq \eta_2] \\
=& \mathbb{E}[(U - z_1)_+ - c_2 | U - c_2 \geq \eta_2] \\
\geq& \mathbb{E}[(U - z_2)_+ - c_2 | U - c_2 \geq \eta_2] = 0
\end{aligned}
$$

  The first equality and the last equality hold by the definition of z_1 and z_2. The first inequality holds because $\eta_1 \geq \eta_2$. The second inequality holds because $c_1 < c_2$ and the term $\big((U - c_1) + x - z_1\big)_+ - x$ is (weakly) decreasing in x. To show the third inequality holds, notice under assumption (A2) of the proposition, $c_1 < c_2$ implies that $U - c_1$ dominates $U - c_2$ in the *hazard rate order*.[57] This further implies $U - c_1 \,|\, \{U - c_1 \geq \eta_2\}$ first-order stochastic dominates $U - c_2 \,|\, \{U - c_2 \geq \eta_2\}$ and thus the inequality holds.[58] The last inequality holds since $z_1 < z_2$ as is supposed.

  The above result implies $\mathbb{E}[(U - z_1)_+ - c_2 | U - c_2 \geq \eta_2] = 0$ and thus $z_1 = z_2$. This contradicts with what I supposed (i.e., $z_1 < z_2$). Thus we must have $z_1 \geq z_2$.

Now, for $i = 1, 2$, let Z_i denote the Pandora's index for the realized posterior belief given search cost c_i and the corresponding equilibrium signal. Then we have:

$$
Z_i = \begin{cases} U - c_i & \text{if } U \leq \eta_i + c_i \text{ (full disclosure region)} \\ z_i & \text{if } U > \eta_i + c_i \text{ (pooling region)} \end{cases}
$$

Recall that I have shown $\eta_1 + c_1 \leq \eta_2 + c_2$ (part (a) of the proposition) and $z_1 \geq z_2$ above. These

[57] See Section 1.3 in Müller & Stoyan (2002) for the definition and properties of hazard rate order. Formally, assumption (A2) implies that F_U has increasing hazard rate and thus for any x s.t. $F_U(x + c_2) < 1$, we have $\frac{f_U(x+c_1)}{1-F_U(x+c_1)} \leq \frac{f_U(x+c_2)}{1-F_U(x+c_2)}$, which is equivalent to $\frac{f_{U-c_1}(x)}{1-F_{U-c_1}(x)} \leq \frac{f_{U-c_2}(x)}{1-F_{U-c_2}(x)}$. This implies that $U - c_1$ dominates $U - c_2$ in the hazard rate order by Theorem 1.3.3 in Müller & Stoyan (2002).

[58] This follows from the discussion in Müller & Stoyan (2002) right above their Definition 1.3.2.

imply that $Z_1 \geq Z_2$. (When $U \leq \eta_1 + c_1$, $Z_1 = U - c_1 > U - c_2 = Z_2$; when $\eta_1 + c_1 < U \leq \eta_2 + c_2$, $Z_1 = z_1 \geq z_2 \geq \eta_2 \geq U - c_2 = Z_2$; when $U > \eta_2 + c_2$, $Z_1 = z_1 \geq z_2 = Z_2$.) Thus we have $U \wedge Z_1 \geq U \wedge Z_2$.

When $c = c_i$, as I discussed in Section 2.5.1, the consumer's expected surplus would be $U_0 + \mathbb{E}[\max\{U \wedge Z_i - p - U_0, 0\}]$ given any realized U_0.[59] Thus $U \wedge Z_1 \geq U \wedge Z_2$ implies that the consumer surplus is higher when $c = c_1$ compared to that when $c = c_2$. This concludes the proof for part (b). *Q.E.D.*

2.E Proofs for Section 2.6

2.E.1 Proof for Proposition 2.6.1

Proof. By the definition of SG-equilibrium, it suffices to show:

$$(H_N^{ex}, p_N^{ex}) \in \underset{(G_i, p_i)}{\arg\max}\{p_i \int \left[H_N^{ex}(x - p_i + p_N^{ex})\right]^{N-1} dG_i(x)\} \tag{2.28}$$

$$\text{s.t. } G_i \text{ is a feasible distribution for } U_i \wedge Z_i$$

Since (H_N^{ex}, p_N^{ex}) is an EG-equilibrium, we know:

$$(H_N^{ex}, p_N^{ex}) \in \underset{(H_i, p_i)}{\arg\max}\{p_i \int \left[H_N^{ex}(x - p_i + p_N^{ex})\right]^{N-1} dH_i(x)\} \tag{2.29}$$

$$\text{s.t. } H_i \preceq_{MPS} F_{U-c}$$

Crucially, Lemma 2.3.1 implies that the optimization in (2.28) is more constrained than that in (2.29). Thus, to show condition (2.28) holds, it suffices to check that H_N^{ex} is indeed a feasible distribution for $U_i \wedge Z_i$ when N is large enough. For this purpose, I need the following property of H_N^{ex}:

Observation (1). For any $\epsilon > 0$, there exists $N_\epsilon < \infty$ such that whenever $N \geq N_\epsilon$, we have $H_N^{ex}(x) = F_{U-c}(x)$ for all $x \leq \bar{u} - c - \epsilon$.

Subproof. This observation is directly implied by Proposition 2 and footnote 17 in Hwang et al. (2019) (June 5th version). When referring to Hwang et al. (2019), one need to notice that the

[59] See, for example, Corollary 1 in Choi et al. (2018).

consumer's value for a product in their paper is corresponding to the net-match-utility $U_i - c$ in my paper, so their range for product value $[\underline{v}, \bar{v}]$ is corresponding to $[\underline{u} - c, \bar{u} - c]$ in my model. □

According to the observation, we can find N^* such that $N \geq N^* \Rightarrow H_N^{ex}(x) = F_{U-c}(x)$, $\forall x \leq \bar{u} - 2c$. Now, fix any $N \geq N^*$. Let $(\mathcal{S}_i, \pi_i)$ be a particular signal under which $\mathbb{E}[U_i - c|S_i] \sim H_N^{ex}$. Since $H_N^{ex}(x) = F_{U-c}(x)$ for all $x \leq \bar{u} - 2c$, this signal fully reveals $U_i - c$ less than $\bar{u} - 2c$, and thus $U_i - c \leq \bar{u} - 2c \Rightarrow Z_i = U_i - c < U_i$.[60] Since $Z_i \leq \bar{u} - c$ under any signal, we also have $U_i - c > \bar{u} - 2c \Rightarrow Z_i \leq U_i$. Combining these facts, we always have $Z_i \leq U_i$ under $(\mathcal{S}_i, \pi_i)$. This further implies:

$$0 = \mathbb{E}[(U_i - Z_i)_+ - c|S_i] = \mathbb{E}[U_i - Z_i - c|S_i] = \mathbb{E}[U_i - c|S_i] - Z_i$$

where the first equality holds by the definition of Z_i and the second equality holds because $Z_i \leq U_i$. Thus $\mathbb{E}[U_i - c|S_i] = Z_i = U_i \wedge Z_i$ under $(\mathcal{S}_i, \pi_i)$. Therefore, under $(\mathcal{S}_i, \pi_i)$ we also have $U_i \wedge Z_i \sim H_N^{ex}$. This shows that H_N^{ex} is indeed a feasible distribution for $U_i \wedge Z_i$. *Q.E.D.*

2.E.2 Proof for Corollary 2.6.1

Proof. This is directly implied by Proposition 2.6.1 and Observation (1) in the proof of Proposition 2.6.1 above. *Q.E.D.*

2.F Other Proofs

Proof for Proposition 2.A.1. I first enlarge the model to accommodate search subsidy and price discount. Let $c^\dagger$ denote the objective search cost (without deducting any search subsidy) and let $U^\dagger$ denote the product's (uncertain) consumption utility. Let b and d denote the search subsidy and price discount respectively. Let p denote the original product price without any discount.

Given any search subsidy b and price discount d, let $c := c^\dagger - b$ and $U := U^\dagger + d$. Then the consumer's behavior is characterized in the same way as in the baseline model. In particular, the consumer would search if $Z \geq U_0 + p$ and would purchase if $U \wedge Z \geq U_0 + p$, where Z is

[60] A formal proof for this is similar to that for Claim 1 in part 2 of the proof for Proposition 2.3.2.

defined identically as in Section 2.3 (with respect to the U and c define here).[61] I still use G to denote the distribution of $U \wedge Z$ under any signal and use J to denote the distribution of U_0.

Under strategy A, we have $(b, d) = (y, 0)$ and the seller's maximal expected profit given any original price p is:

$$\max\{p\mathbb{P}(\text{purchase}) - y\mathbb{P}(\text{search})\} \tag{2.30}$$

$$\leq \max\{(p - y)\mathbb{P}(\text{purchase})\} \tag{2.31}$$

$$= \begin{array}{l} \max\limits_{G}\{(p - y)\displaystyle\int J(x - p)dG(x)\} \\ \text{s.t. } G \preceq_{MPS} F_{U-c}\,(= F_{U^\dagger - c^\dagger + y});\; G \preceq_{FOD} F_U\,(= F_{U^\dagger}) \end{array} \tag{2.32}$$

where the maximizations in the first two lines are over all pre-search signals. The inequality holds because $\mathbb{P}(\text{purchase}) \leq \mathbb{P}(\text{search})$; the equality holds because under the unimodality assumption of J, the seller's optimization over purchase probability (given any (p, b, d)) is fully characterized by the Relaxed Problem.

Under strategy B, we have $(b, d) = (0, y)$ and the seller's maximal expected profit given p is:

$$\max\{(p - y)\mathbb{P}(\text{purchase})\} \tag{2.33}$$

$$= \begin{array}{l} \max\limits_{G}\{(p - y)\displaystyle\int J(x - p)dG(x)\} \\ \text{s.t. } G \preceq_{MPS} F_{U-c}\,(= F_{U^\dagger - c^\dagger + y});\; G \preceq_{FOD} F_U\,(= F_{U^\dagger + y}) \end{array} \tag{2.34}$$

Comparing optimization (2.34) with optimization (2.32), one can see that they are the same except that the second constraint in (2.34) is less restrictive. Thus the maximal profit under strategy B is larger than that under strategy A. Also notice that the inequality in the line of (2.31) holds strictly when $\mathbb{P}(\text{purchase}) > 0$. Thus strategy A is strictly dominated as long as there would be some search without purchase under it.

Q.E.D.

[61] Notice all U, c and Z here depend on the underlying (b, d). I suppress this dependence to ease the notations.

Chapter 3

Optimal Disclosure Regulation for Entrepreneur Financing with Ex-Post Moral Hazard

3.1 Introduction

What is the optimal level of disclosure regulation for entrepreneur public financing? This question has become increasingly important in recent years, as many countries experienced significant regulatory reforms and various new entrepreneur financing channels were introduced (e.g., security-based crowdfunding).[1] An important feature of these new financing channels is that they impose much lighter disclosure requirements than traditional approaches (e.g., IPO).[2] Given the conventional view that rigorous disclosure standard is important for investor protection and efficient capital allocation, it is not surprising that this looseness of disclosure regulation has brought lots of controversy.

The main argument supporting light disclosure regulation in practice is that disclosure can be costly for firms and lighter regulation makes public financing more affordable to them. This particularly applies to small businesses and startups, who have little financial resource and expertise to comply sophisticated disclosure requirements. However, is lower cost the only

[1]See Hornuf & Schwienbacher (2017), for example, for a detailed introduction to regulation reforms in several countries on equity crowdfunding.

[2]For example, the JOBS Act in the US passed in 2012 introduced several new financing channels for small businesses, all of which feature looser disclosure requirement than that in traditional IPO. In particular, Regulation Crowdfunding bears the least disclosure burden among them.

reason for looser regulation, or is there any intrinsic benefit from it? Different answers to this question have completely different policy implications. If lower cost is the only reason for lighter regulation, the government should focus on lowering disclosing cost with better technology and institutional design, and gradually enhance disclosure standard as its cost gets lower. If there are intrinsic benefits from less disclosure, on the other hand, the regulator may want to deliberately avoid full disclosure even if it can be achieved in a costless way.

This paper contributes to the above discussion with a novel story supporting the view that less information disclosure has its intrinsic benefit. Specifically, I study socially optimal disclosure in a simple model of public equity financing with consideration of the entrepreneur's post-financing moral hazard problem. While the moral hazard problem causes efficiency loss under full disclosure, this paper shows partial disclosure can help to mitigate it. As a result, a properly designed partial disclosure rule is shown to be socially optimal, although no exogenous disclosure cost has been assumed.

While this implication about the optimal disclosure may be surprising, the moral hazard problem I consider is quite stylized in entrepreneurial finance. After getting funded, the entrepreneur can choose to honestly develop the project or to divert the fund for other uses.[3] Although numerous studies have proposed different mechanisms to solve this problem, most of them took mechanism design approaches and focused on how to use better monitoring or contracting devices to provide post-financing incentives. In contrast, this paper takes an information design approach and, to my best knowledge, is the first to show how a properly designed disclosure rule in the financing campaign can help to alleviate it.

The possibility that less information can improve social welfare has been broadly explored in the literature (e.g., Morris & Shin, 2002; Angeletos & Pavan, 2007; Amador & Weill, 2010). In particular, a growing literature surveyed below has proposed various stories on why restricted information releasing in financial market can be beneficial. Although this paper also derives sub-optimality for full information, the mechanism explored is new to the literature. While existing studies have focused on frictions in the financial market, this paper first shows how an agency problem within the firms makes partial information socially optimal in financing

[3]This type of moral hazard can be particularly important in some recently legalized public financing channels for small businesses, like equity crowdfunding. Besides standard features of small businesses, a particular reason for this is that investors investing via these channels typically only contribute a small fund to each project, which endows them little monitoring incentive. For example, on Wefunder, the largest Regulation Crowdfunding platform in the US, an investor can invest as little as 100 dollars in a project.

campaigns.

In the basic model, I consider two types of entrepreneurs with low value and high value projects respectively. Both types of projects are assumed to have positive net value and thus are socially optimal to develop. Due to the moral hazard problem above, however, the low types will not be able to get financed when types are fully disclosed under certain conditions. Intuitively, for an entrepreneur to behave after financing, she must receive enough incentive rent from developing the project. This is not possible for the low type projects when their net values, although positive, are lower than the incentive rent needed.

To see how partial disclosure can help to mitigate this efficiency loss, notice from investors' perspective, partial disclosure essentially pools some low type entrepreneurs with high types. When this is done properly, we can reduce the financing costs (in terms of shares sold) of pooled low type entrepreneurs such that they receive enough incentive rents to carry out their projects. Although this induces higher financing costs for the high types, enough incentive can still be kept for them. In the end, some low type projects together with all the high type projects would be developed, which generates higher social surplus than the full disclosure case.

With Bayesian persuasion tools, the optimal rule is derived explicitly. Notice although pooling is beneficial, certain degree of disclosure is necessary when the high type projects are not abundant. Intuitively, if too many low types are pooled with high types, their financing costs will not be low enough to get them incentivized. As a result, the low types would choose to run away once funded and the entire market would have frozen at the beginning.

Like most Bayesian persuasion papers, I have assumed exogenous type distribution in my basic model. However, since the disclosure rule affects different types' expected payoffs, it changes the relative incentive for entrepreneurs to acquire a high value project versus to acquire a low value one. Thus, disclosure regulation may change the type distribution in turn. Especially, adverse selection could arise with naively designed partial disclosure. Indeed, a standard opposing argument against loose disclosure regulation is that project qualities would deteriorate when low value projects are not well distinguished from high value ones.[4] To deal with this, I extend the basic model to endogenize project types and consider the designer's problem as choosing a disclosure rule and a target type distribution simultaneously, where the type distribution must be incentivized properly. It is shown that the socially optimal disclosure is still partial disclosure

[4]The fact that poor disclosure leads to inefficient pre-financing investments by firms is well understood in the literature (e.g., Fishman & Hagerty, 1989).

under certain assumptions.

Interestingly, with endogenous type distribution, the optimal disclosure rule may use one more signal realization to facilitate better information revealing compared to the case where the designer can fix the target type distribution without incentivizing for it. In particular, the optimal rule can simultaneously involve a realization fully identifying a low type and a realization fully identifying a high type, which never happens in the basic model. Intuitively, with endogenous types, certain degree of disclosure is needed not only to provide sufficient post-financing incentives for pooled low types, but also to seize adverse selection and maintain desired proportion of high types in the market.

In Appendix 3.A, I further extend the model to allow high type entrepreneurs to privately disclose their types with a cost. I show that any optimal disclosure rule should not induce the entrepreneurs to use private disclosures. It turns out that the optimal rule is still partial disclosure and is similar to that in the main text, but fewer low types would be pooled unless the private disclosing cost is sufficiently high. Interestingly, the costliness of private disclosure is beneficial here, as it prevents unraveling (Milgrom, 1981) and thus makes partial disclosure implementable.

Although the paper just considers a simple moral hazard problem, the intuition explored may still be useful when one replaces it with some general contracting or mechanism design problem following the financing campaign. If the firms' post-financing incentive constraints depend on their financing costs, then the paper suggests partial disclosure in the financing campaigns may help to relax some types' incentive constraints via transferring potential surplus among different types. In this way, outcomes infeasible under full disclosure may be achieved.

In Appendix 3.B, I also adapt the basic model to debt financing and apply it to regulation for banking system disclosure. Alvarez & Barlevy (2015) illustrates with an example that given post-financing moral hazard, full disclosure can outperform non-disclosure even if the bank financing market does not freeze without disclosure, which is in contrast with Goldstein & Leitner (2018). However, their paper does not consider partial disclosure rules. By using a concavification graph, I intuitively show that under certain conditions, although full disclosure dominates non-disclosure, partial disclosure is typically optimal.

Related literature – This paper directly relates to the literature on optimal disclosure in financial market. Although inefficiency caused by asymmetric information has been well recognized since Akerlof (1970), several strands of literature have shown full disclosure is not guaranteed optimal in financial market. One strand exploits the Hirshleifer effect in Hirshleifer (1971) (e.g., Andolfatto et al., 2014; Dang et al., 2017; Monnet & Quintin, 2017b; Goldstein & Leitner, 2018). In these papers, restricted information disclosure helps to maintain risk sharing opportunity either between the security issuer and investors, or among investors in the secondary market, which thus improves efficiency. Another strand considers crowding-out effect of public information disclosure (e.g., Gao & Liang, 2013; Colombo et al., 2014). Basically, better public information may reduce investors' private incentives to gather information and thus decrease aggregate information available to decision makers.[5] Yet another strand of literature considers the possibility that more information disclosure can actually exacerbate information asymmetry when investors are heterogeneous in their ability of interpreting information (Pagano & Volpin, 2012; Monnet & Quintin, 2017a).

This paper complements the above literature with a new story on why partial disclosure can improve welfare over full disclosure. While the studies above focus on frictions in the financial market, this paper focuses on an agency issue within the firm. Indeed, the post-financing moral hazard problem is critical in my story, absent which full disclosure would be optimal. This interplay between disclosure in financing campaigns and corporate governance also relates the paper to a broad literature on corporate finance (e.g., Zingales, 1995; Shleifer & Wolfenzon, 2002).

Several recent papers also highlight the role of moral hazard in information design problems (Rodina, 2016; Georgiadis & Szentes, 2018; Boleslavsky & Kim, 2018). Besides apparent differences in the questions studied, these papers focus on an agency problem before Bayesian persuasion on the outcome of the agent's work, while the moral hazard problem considered in my model is after Bayesian persuasion and they are related only through the entrepreneur's financing cost (shares sold). Thus the mechanism explored in my model is completely different from theirs. This also easily distinguishes this paper from the literature on contract design with information revealing (e.g., Fuchs, 2007; Fu & Trigilia, 2019).[6]

[5]See Goldstein & Yang (2017) for a nice review on this issue.

[6]Another paper, Alexander & Isakin (2015), does consider post-financing moral hazard after disclosure by a credit rating agent, but the incentive problem in their model favors full disclosure. The credit rater chooses partial disclosure in their paper because of an ad hoc assumption that she wants to send as many good ratings as possible.

From practical perspective, the paper relates to a surging literature on crowdfunding.[7] In particular, several recent papers study how properly designed crowdfunding mechanisms help to solve the entrepreneur's post-financing moral hazard problem (Chang, 2016; Strausz, 2017; Chemla & Tinn, 2018). However, none of them have considered the information design problem studied in this paper.

From methodological perspective, the paper relates to the Bayesian persuasion literature (Aumann et al., 1995; Rayo & Segal, 2010; Kamenica & Gentzkow, 2011). Especially, the constrained concavification method I use in Section 4 is from Boleslavsky & Kim (2018). Rosar (2017) also uses a similar Lagrangian concavification method in a constrained Bayesian persuasion problem.

The paper is organized as follows: Section 2 introduces the moral hazard problem and motivates partial disclosure; Section 3 and Section 4 study optimal disclosure with exogenous and endogenous type distributions respectively; Section 5 provides some practical discussions; Section 6 concludes. In addition to the main text, Appendix 3.A (online) extends the model to allow costly private disclosure and Appendix 3.B (online) extends the basic model to debt financing with an application to banking system disclosure.

3.2 Moral Hazard and Efficiency Loss under Full Disclosure

In this section, I introduce the economic agents and the post-financing moral hazard problem. I then show how moral hazard causes efficiency loss under full disclosure and intuitively illustrate the motivation for partial disclosure. For simplicity, I assume all agents in the economy are risk neutral and the discounting factor is normalized to 1.

3.2.1 Entrepreneurs and Moral Hazard Problem

Consider a continuum of entrepreneurs, each of whom has one project to develop. There are two types of projects (thus two types of entrepreneurs) denoted by $v \in \{v_L, v_H\}$ ($v_H > v_L$), where v indicates the project's expected value in date 1 if developed. Each project requires investment K in date 0. Assume $v_L > K$, so both types of projects are socially optimal to develop.[8]

[7] For recent studies on security-based crowdfunding, see Hornuf & Schwienbacher (2018), Brown & Davies (2018), Signori & Vismara (2018), and Walthoff-Borm et al. (2018).

[8] Notice low type projects are not bad projects since they have positive net expected values. Actually, they can be highly valuable projects with low probability of success ex-ante. Many boldly innovative projects can be

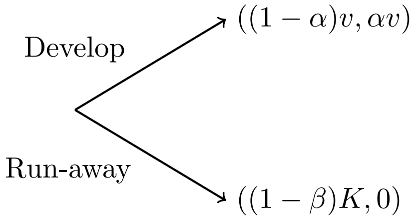

Figure 3.1: Entrepreneur's Moral Hazard Problem

Assume each entrepreneur needs to finance his entire investment K through equity-based public financing (e.g. equity crowdfunding). Specifically, the entrepreneur needs to offer a share of the firm's equity to a representative investor in exchange for funding. Let $\alpha \in [0, 1]$ denote the share of equity she offers. If the investor accepts it, he contributes investment K (in cash) to the firm and receives α share of the firm's future cash flow; otherwise, the project remains not financed and the entrepreneur gets outside payoff 0.

Once a project is financed, the entrepreneur can decide whether to honestly develop it. If not, she can "run away" with the funded money K by paying a cost βK, where $\beta \in [0, 1]$.[9] This causes a moral hazard problem and the parameter β governs how serious it is. With larger β, the cost of running away is higher and the moral hazard problem is thus less severe. In particular, if $\beta = 1$, the entrepreneur gets nothing by running away and there is no moral hazard problem.

The entrepreneur's post-financing decision problem is summarized in Figure 3.1, where the first payoff in each bracket is for the entrepreneur and the second payoff is for the investor. One implicit assumption here is that the entrepreneur cannot divert her project's value, so the expected payoffs to the entrepreneur and the investor are just $(1 - \alpha)v$ and αv respectively after the project is developed. This assumption is reasonable because the project's value is typically only captured when the project is completed and finally goes public with IPO or gets acquired by other firms, which are easily observable events. Thus it is hard to divert a project's value after development. In contrast, diverting cash raised in the financing campaign is much easier and technically feasible for the entrepreneur.

By the entrepreneur's moral hazard problem, a project is developed if and only if the fol-

in this category.

[9] One can think of this literally as running-away with βK being the cost of hiding. Alternatively, one can think of it as a situation where the entrepreneur pretends to carry out the project with investment βK and diverts the remaining money for personal purposes.

lowing incentive constraint holds[10]:

$$(1-\alpha)v \geq (1-\beta)K \tag{3.1}$$

For the constraint to be satisfied, we need v to be large enough relative to K and α to be small. This implies that the entrepreneur must hold enough share of the firm to be incentivized to develop the project.

3.2.2 Investors

There is a representative investor representing a potentially large group of small homogeneous investors. By risk-neutrality,[11] he accepts the entrepreneur's offer if and only if the expected payoff from investing exceeds his current investment, which means the following condition holds[12]:

$$K \leq \mathbb{E}[\alpha v \mathbb{1}\{a(v,\alpha) = \text{Develop}\}|\mathcal{I}] \tag{3.2}$$

where $\mathbb{1}$ is the indicator function, $a(v,\alpha)$ denotes the entrepreneur's post-financing choice and $\mathcal{I}$ denotes the investor's information when making his decision.

3.2.3 Full Disclosure and Efficiency Loss

Now, suppose the project types are fully disclosed to the investor. Then, condition (3.2) implies that the investor would invest in a type v project if and only if $\alpha v \geq K$ and the incentive constraint (3.1) for a type v entrepreneur is satisfied. Formally, a type v project is invested and carried out if and only if:

$$\exists \alpha \in [0,1] \text{ s.t. } (1-\alpha)v \geq (1-\beta)K \text{ and } \alpha v \geq K$$

which is equivalent to $v \geq (2-\beta)K$.

[10] I assume the entrepreneur will develop the project when being indifferent. This guarantees the existence of a minimal posterior belief that can sustain each class of equilibria discussed later.

[11] Risk neutrality of the investor is actually not needed for results in this paper, because any risk induced by partial disclosure is idiosyncratic and a risk averse investor can simply hold a large portfolio to diversify it away. This is particularly applicable to the crowdfunding context, where the minimal investment required for each project is typically small and thus each investor can build a large portfolio with moderate amount of money.

[12] For the same reason as stated in footnote 10, I assume the investor would invest when being indifferent.

Now, assume

$$v_H > (2-\beta)K > v_L > K \tag{3.3}$$

Then, with full disclosure, only the high type projects are financed and developed, which causes an efficiency loss since both types of projects have positive net value. Intuitively, incentivizing an entrepreneur to develop the project requires an incentive rent at least $(1-\beta)K$. However, the total surplus of a low type project is $v_L - K < (1-\beta)K$. Thus when projects are fairly priced with full information, it is not possible to provide enough incentive to the low type entrepreneurs.

To improve efficiency, we have to increase the low type entrepreneurs' incentive rents through lowering their financing costs (in terms of α). It turns out that this can be achieved with partial disclosure under certain condition. By pooling some low type projects with high types, the expected project values conceived by the investor can be high enough for him to accept a lower α, which in turn provides sufficient incentives for the low type entrepreneurs to behave in the moral hazard problem. As to be shown later, this is possible if the following condition holds:

$$\Big(1-\frac{K}{v_H}\Big)v_L > (1-\beta)K \tag{3.4}$$

Notice if a project is believed to be of high type, K/v_H will be the smallest possible share sold for financing capital K (since we need $\alpha v_H \geq K$). Thus the above condition is saying if a low type project is evaluated as high type, the financing cost (in terms of α) could be low enough to satisfy the entrepreneur's incentive constraint. We need the condition holds as strict inequality because it is impossible to make a Bayesian investor believe a low type project is high type with probability 1. Thus a room is needed here. If this condition is violated, there would be no way to get the low type projects developed and partial disclosure cannot be helpful.

Throughout the paper, conditions (3.3) and (3.4) are the maintaining assumptions imposed on the fundamental parameters.

Assumption 3.2.1. Parameters K, v_L, v_H and $\beta \in [0,1]$ satisfy conditions (3.3) and (3.4).

A numerical example satisfying the assumption is: $K = 1$, $v_H = 2$, $v_L = 1.6$ with $\beta \in (0.2, 0.4)$.

3.3 Basic model: Optimal Disclosure with Exogenous Types

In this section, I introduce a simple model of equity-based public financing with disclosure regulation and design the optimal disclosure rule given exogenous entrepreneur type distribution. As discussed in the introduction, the exogeneity of type distribution is a strong assumption, which I will relax in Section 4. I first focus on this simplified case for two reasons: (1) the problem is much easier to solve and can already illustrate important intuitions about the optimal disclosure; (2) the solution derived will be needed for the endogenous types case later.

3.3.1 Financing Campaign and Designer's Problem

As stated in Section 2.1, there are two types' of entrepreneurs seeking for public equity financing. Let p_H denote the proportion of high types, which is exogenous in this section. Before a financing campaign, the regulator can examine the entrepreneur's project and learn its type v. Then, she sends a signal on v to the (representative) investor according to a pre-determined disclosure rule. After the signal is publicly observed, the entrepreneur makes an equity offer α to the investor, who then decides whether to accept it. If a project is successfully financed, the moral hazard problem specified in Section 2 follows.

In accordance with the Bayesian persuasion literature (e.g., Kamenica & Gentzkow, 2011), a regulator's disclosure rule is defined as:

Definition 3.3.1. A disclosure rule $\mathcal{D}$ consists of a finite signal realization space $\mathcal{S}$ and a family of probability distributions $\{\Gamma(\cdot|v_L), \Gamma(\cdot|v_H)\}$ on $\mathcal{S}$.

Essentially, a disclosure rule defines the conditional distribution of signal realization given each project type v. The disclosure rule is assumed to be publicly announced at the very beginning, so the Bayesian inference from each signal realization is commonly understood by all agents.

Given a disclosure rule, the game's timeline is summarized as following:

–1. An entrepreneur learns her type $v \in \{v_L, v_H\}$, with probability p_H for v_H.

–2. A signal $s \in \mathcal{S}$ about v is realized according to the pre-announced disclosure rule and publicly observed.

–3. The entrepreneur makes equity offer $\alpha \in [0, 1]$ in exchange for investment K.

–4. The investor decides whether to invest in the project according to condition (3.2), where the information set $\mathcal{I} = \{s, \alpha\}$.

–5. If the campaign succeeds, the entrepreneur chooses action $a \in \{$Develop, Run-away$\}$ by solving the decision problem in Figure 3.1 and the final payoffs are correspondingly realized. If the campaign fails, the entrepreneur gets payoff 0 and the investor keeps his fund K.

Notice Stages 3 – 4 formally form a signaling game and the signal s disclosed in Stage 2 decides the initial belief[13] of that signaling game.

Designer's problem The designer's problem is to find a disclosure rule that maximizes ex-ante total social surplus. Notice in terms of deciding the social outcome, only the posterior belief distribution induced by the disclosure rule (given p_H) matters. Let π_H denote the posterior belief on $v = v_H$ after a signal realizes. Then by Kamenica & Gentzkow (2011), given prior p_H, a disclosure rule is equivalent to a distribution μ of posterior π_H which satisfies Bayesian feasibility condition $\mathbb{E}_\mu[\pi_H] = p_H$. Thus the designer's problem can be equivalently stated as finding a posterior distribution μ with $\mathbb{E}_\mu[\pi_H] = p_H$ that maximizes ex-ante social welfare. This allows one to solve the problem using the concavification method (Aumann et al., 1995; Kamenica & Gentzkow, 2011).

One concern here is that in the signaling game starting from Stage 3, the offer α can potentially serve as another signal conveying information about the project's type. If we have a separating equilibrium in this signaling game (two types offer different α), then the outcome will not depend on the initial belief induced by s, which makes the information design problem irrelevant. As we will see in Section 3.2, however, except for trivial cases where no project gets financed, only pooling equilibria exist in this signaling game. The intuition is simple. If a low type entrepreneur is identified, she cannot be financed and gets payoff 0. Since being financed is always better than not (at least one can run away with the money), a low type would always want to mimic the high type to raise some fund if possible.[14]

As the offer α can convey no additional information, an alternative way to model the financing campaign is to let the investor move first and make an offer. Letting the less informed side

[13] I save the word "prior belief" to refer to p_H.

[14] This intuition and result remain valid even if one allows the entrepreneurs to offer more sophisticated financing contracts.

move first can save us from discussing multiple equilibria typically associated with a signaling game and is often employed in the literature. However, in this paper I stick to the natural timing above and formally show that the offer cannot convey additional information. I will specify equilibrium selection rules in the next subsection to pin down a unique equilibrium.

3.3.2 The Equilibrium

Given any belief π_H induced by the signal realization in Stage 2, I solve here the (weak sequential) equilibria for Stages 3 – 5. Since things are trivial when $\pi_H \in \{0,1\}$, with which the outcome has been discussed in Section 2.3, I focus on the case where $\pi_H \in (0,1)$. All proofs are provided in Appendix 3.C.1 (online).

First, as mentioned above, we have the following observation:

Observation 3.3.1. *Any equilibrium in which some projects are financed is a pure strategy pooling equilibrium.*

Notice by Observation 3.3.1, if some projects are financed, all projects must be financed with the same offer. Thus there may be three classes of equilibria:

- Class 1: Both types of projects are financed and developed.
- Class 2: Both types are financed, but only high type projects are developed.[15]
- Class 3: No project is financed.

To ease notations, define constants:

$$A := \frac{K}{v_H - (1-\beta)K}; \quad B := \left[\frac{K}{v_L - (1-\beta)K} - 1\right]\frac{v_L}{v_H - v_L}$$

Notice $A, B \in (0,1)$ by Assumption 3.2.1 and either of them can be bigger than the other. The following lemma characterizes equilibrium outcomes.

Lemma 3.3.1. *Class 1 equilibria exist if and only if $\pi_H \geq B$; Class 2 equilibria exist if and only if $\pi_H \geq A$; only Class 3 equilibria exist if $\pi_H < min\{A, B\}$. Moreover, there can be multiple equilibrium α within each class.*[16]

By Lemma 3.3.1, there are two kinds of equilibrium multiplicity in the model. One is the co-existence of several equilibrium classes with some initial belief π_H. This multiplicity is

[15]Notice it cannot be the case where low types are develop while high types are not since $v_H > v_L$.

[16]In the proof, I show that all these equilibria survive the intuitive criterion in Cho & Kreps (1987).

largely fundamental and comes from a kind of self-fulfilling belief, which is similar to that in a self-fulfilling debt crisis (e.g., Cole & Kehoe (1996)). Intuitively, if the investor believes a low type entrepreneur will run away, he would require a higher α for compensation, which in turn makes the low types indeed choose to run away. For this multiplicity, I assume the designer is free to choose the most efficient equilibria. Notice Class 1 equilibria are more efficient than Class 2 equilibria, which are yet more efficient than Class 3 equilibria, so I impose the following selection rule:

Equilibrium Selection Rule 3.3.1. Select Class 1 equilibria if they exist; otherwise, select Class 2 equilibria if they exist. If no Class 1 or Class 2 equilibrium exists, select Class 3 equilibria.

Another multiplicity is about multiple equilibria within each class, which is due to the arbitrariness of off-path belief in weak sequential equilibria. For all results in the current section, this kind of multiplicity is innocuous because we only care about social efficiency and all equilibria within the same class produce the same social surplus. The specific equilibrium result matters only in deciding the surplus split between the entrepreneur and the investor. However, I will endogenize type distribution in Section 4, in which entrepreneurs' perceived expected payoffs before financing would be important. Thus I need to pin down a unique equilibrium from each class that specifies the entrepreneur's share of surplus. As shown in the proof of Lemma 3.3.1, intuitive criterion does not help here, so I adopt the following ad-hoc selection rule:

Equilibrium Selection Rule 3.3.2. Within each class of equilibria, select the one with lowest α. (Note: This is not needed for results in this section.)

The main motivation for this rule is that in many entrepreneur public financing campaigns (e.g., equity crowdfunding), the entrepreneur makes a "take or leave" offer to a large crowd of investors. Thus it is reasonable to give full bargaining power to the entrepreneur. By selecting the equilibrium with lowest α, I leave the investor just indifferent and give full surplus to the entrepreneur. This is also consistent with models assuming the investor moves first to bid a fairly valued offer, as is often seen in the literature.

That being said, it is actually not hard to extend analysis to scenarios with other well-behaved equilibrium selection rules. However, if the investor obtains part of the surplus, in the endogenous types model studied in Section 4, acquiring a high value project (v.s. a low value one) by the entrepreneur would have positive externality to the investor, which adds another source of inefficiency. My selection rule here avoids this complication to allow us to solely focus on the friction caused by moral hazard.

Combining the two equilibrium selection rules with Lemma 1, we have:

Proposition 3.3.1. *When $\pi_H \geq B$, both types are financed and developed with equity offer $\alpha = \frac{K}{\pi_H v_H + (1-\pi_H)v_L}$; when $B > \pi_H \geq A$, both types are financed with $\alpha = \frac{K}{\pi_H v_H}$, but only the high type projects are developed; when $\pi_H < \min\{A, B\}$, no project is financed. (Notice the second case may happen only when $B > A$.)*

3.3.3 Optimal Disclosure Rule

For a project with type v and posterior belief π_H induced by the regulator's disclosure, Proposition 3.3.1 implies that the ex-post social surplus from this project is:

$$
\begin{aligned}
w(v, \pi_H) := & \big[(v_H - K)\mathbb{1}\{v = v_H\} + (v_L - K)\mathbb{1}\{v = v_L\}\big]\mathbb{1}\{\pi_H \geq B\} \\
& + \big[(v_H - K)\mathbb{1}\{v = v_H\} - \beta K \mathbb{1}\{v = v_L\}\big]\mathbb{1}\{B > \pi_H \geq A\}
\end{aligned}
$$

The first line is corresponding to the case of $\pi_H \geq B$, where both types of projects are developed; the second line is corresponding to the case of $B > \pi_H \geq A$, where a high type project is developed while a low type project's entrepreneur runs away with cost βK. Notice when $\pi_H < \min\{A, B\}$, no project is financed and thus social surplus is 0.

In order to use concavification method to design the optimal disclosure rule, we need to derive an indirect social welfare function on the posterior π_H, taking expectation of which delivers the ex-ante social welfare. This is done by taking expectation of the ex-post welfare function conditional on π_H. Specifically, define indirect social welfare function as:

$$
\begin{aligned}
W(\pi_H) := & \mathbb{E}[w(v, \pi_H)|\pi_H] \\
= & [(v_H - K)\pi_H + (v_L - K)(1 - \pi_H)]\mathbb{1}\{\pi_H \geq B\} \\
& + [(v_H - K)\pi_H - \beta K(1 - \pi_H)]\mathbb{1}\{B > \pi_H \geq A\}
\end{aligned}
\tag{3.5}
$$

where the second equality comes from the fact that $\mathbb{P}[v = v_H|\pi_H] = \pi_H$ and $\mathbb{P}[v = v_L|\pi_H] =$

$1 - \pi_H$ since π_H is just the posterior probability for $v = v_H$. Then the ex-ante total social welfare is obtained as $\mathbb{E}[W(\pi_H)]$, in which the distribution of π_H is induced by the disclosure rule given prior p_H. Following Kamenica & Gentzkow (2011), the designer's problem is then formulated as:

$$V_0(p_H) := \max_{\mu \in \Delta([0,1])} \mathbb{E}_\mu[W(\pi_H)] \tag{3.6}$$

$$\text{s.t. } \mathbb{E}_\mu(\pi_H) = p_H \tag{3.7}$$

where $\Delta([0,1])$ denotes the set of all distributions over $[0,1]$ (in which π_H takes values) and $V_0(\cdot)$ denotes the value function of this optimization problem.

The solution can be easily derived by concavification and is provided in the following proposition:

Proposition 3.3.2. *Given exogenous high type proportion p_H, we have:*

(i) If $p_H < B$, the optimal distribution of posteriors supports on $\{0, B\}$ with $\mu(B) = p_H/B$ and $\mu(0) = 1 - p_H/B$.

(ii) If $p_H \geq B$, any distribution of posteriors supporting on a subset of $[B, 1]$ with $\mathbb{E}_\mu[\pi_H] = p_H$ is optimal.

Moreover, the value function is:

$$V_0(p_H) = \begin{cases} (v_H - K)p_H + (v_L - K)\frac{(1-B)p_H}{B} & \text{if } p_H < B \\ (v_H - K)p_H + (v_L - K)(1 - p_H) & \text{if } p_H \geq B \end{cases}$$

Proof Sketch. The concavification graphs for the case $B > A$ and the case $B \leq A$ are provided in Figure 3.2. The blue curve plots function $W(\cdot)$ and the red dashed curve shows its concavification. See Appendix 3.C.1 for details. *Q.E.D.*

To see the intuition behind this optimal disclosure rule, recall that B is the smallest posterior belief on $v = v_H$ that can sustain Class 1 equilibria, where the low type projects are developed. Therefore, ideally we want as many low types as possible to receive posterior π_H higher than B. When $p_H < B$, we do not have enough high type projects in the sense that it is not possible to pool all low types with the high types while keeping π_H higher than B, so certain degree of

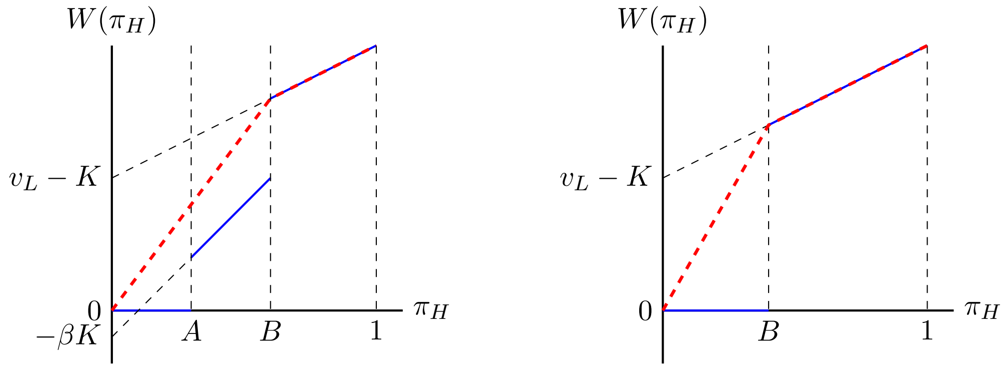

(a) Case 1: $B > A$

(b) Case 2: $B \leq A$

Figure 3.2: Concavification for $W(\cdot)$ (with exogenous p_H)

disclosure is necessary. In this case, the proposition shows that the optimal signal is binary, in which one realization pools some low types with high types to just induce $\pi_H = B$ and the other fully reveals a low type. In this way, maximum amount of low types can be developed. When $p_H \geq B$, on the other hand, the prior is already high enough to sustain Class 1 equilibria, so all projects can be developed without any disclosure. In this case, it is never desirable to disclose a project as low type (or with posterior $\pi_H < B$) since that will just prevent the project from developing.

A desirable feature of this optimal disclosure rule is that there will be no running away by low type entrepreneurs. This is true even if $\beta = 0$ (with Assumption 3.2.1 still being held), in which case there is no direct social loss from running away.[17] Intuitively, there is no benefit to get a project financed when we know the entrepreneur is not going to develop it, but there is a cost with it, which is making the investor reluctant to finance other pooled projects without larger compensation. Therefore, if low types are going to run away with a signal realization (with $\pi_H \in [A, B)$), it is better to reveal types for some of them such that the financing costs become low enough for the rest of them to develop their projects.

A direct implication of Proposition 3.3.2 is:

Corollary 3.3.1. *Full disclosure is not optimal with exogenous type distribution (under Assumption 3.2.1).*

For later use, given any exogenous p_H, I specify a particular optimal disclosure rule in the form of that in Definition 3.3.1:

[17]When $\beta = 0$, the segment of $W(\cdot)$ corresponding to $[A, B)$ interval in Figure 3.2a is still lower than the concavification function and thus the result does not change.

Definition 3.3.2. Given any p_H, define disclosure rule $\mathcal{D}^u(p_H)$ as follows:

- If $p_H < B$, $\mathcal{D}^u(p_H)$ has two signal realizations generating posterior beliefs $\pi_H = 0$ and B respectively.
- If $p_H > B$, $\mathcal{D}^u(p_H)$ has two signal realizations generating posterior beliefs $\pi_H = B$ and 1 respectively.
- If $p_H = B$, $\mathcal{D}^u(p_H)$ has only one signal realization with posterior $\pi_H = B$.

For reason that will be clear later, I will call $\mathcal{D}^u(p_H)$ the unconstrained optimal rule given p_H. This special disclosure rule is going to be useful in characterizing the range of potentially optimal target p_H when entrepreneur types are endogenous in Section 4.

3.3.4 Social Benefit from Increasing p_H

It is interesting to compare the marginal social benefit from increasing p_H (upgrading some low types to high types) under the optimal disclosure with that under full disclosure or in the first best. In the first best, all projects are developed, so the marginal benefit of increasing p_H is the difference between a high type's value and a low type's value, which is $MB^{fb} = v_H - v_L$. With full disclosure, only high type projects are developed, so the marginal benefit of increasing p_H is just the net value of a high type project, which is $MB^{fd} = v_H - K$. Under the optimal disclosure rule in Proposition 3.3.2, the marginal benefit of increasing p_H is measured by the derivative of the value function:

$$V_0'(p_H) = \begin{cases} (v_H - K) + (v_L - K)\frac{1-B}{B} & \text{if } p_H < B \\ v_H - v_L & \text{if } p_H > B \end{cases}$$

It is easy to see:

$$\text{when } p_H < B,\ V_0'(p_H) > MB^{fd} > MB^{fb}$$
$$\text{when } p_H > B,\ MB^{fd} > V_0'(p_H) = MB^{fb}$$

Intuitively, when high type projects are scarce ($p_H < B$), increasing p_H under the optimal disclosure rule not only increases the number of high type projects, which are all to be developed, but also helps to get more low types developed through pooling. In other words, with optimal disclosure the high types are valuable not only in their own values, but also in mitigating the efficiency loss caused by low types' post-financing moral hazard problem. Therefore, the marginal benefit exceeds a high type's own net value and is thus higher than that under full disclosure or in the first best. On the other hand, when high types are abundant ($p_H > B$), all low types can already be developed under optimal disclosure, so the trade-off with raising p_H is the same as that in the first best. Therefore, the marginal benefit under optimal disclosure just equals to MB^{fb} when $p_H > B$.

3.4 Optimal Disclosure with Endogenous Types

One important limitation of the basic model above is that the type distribution is treated as fixed. This can be problematic because when designing the disclosure rule, we are changing financing costs faced by the two types of entrepreneurs, which then affects their ex-ante entrepreneuring incentives. For instance, suppose we start with full disclosure and observe the proportion of high types being $p_H < B$. After imposing the disclosure rule in Proposition 3.3.2, the financing cost for high types will increase and that for low types will decrease, which lowers the incentive of becoming a high type relative to becoming a low type. As a result, the proportion of high types may drop below p_H. This deterioration of quality due to asymmetric information has been well recognized since Akerlof (1970). After the investors realize this, the signal realization intended to induce $\pi_H = B$ will only induce a posterior belief lower than B, which then makes the market freeze with all projects unfunded. This destroys the entire design.

To deal with this issue, I extend the basic model to endogenize project types and design the optimal disclosure rule while taking the change in type distribution into consideration. All proofs in this section are provided in Appendix 3.C.2 (online).

3.4.1 Pre-Financing Upgrading and Designer's Problem

Consider a unit measure of entrepreneurs indexed by x with uniform distribution over $[0, 1]$, each of whom is initially endowed with a low type project. Entrepreneur x can choose to upgrade her project to high type by paying an upgrading cost $\Psi(x)$, where $\Psi(\cdot)$ is an (weakly) increasing function on $[0, 1]$.[18] Assume the regulator's disclosure rule is publicly announced before the upgrading decisions. Then an entrepreneur would upgrade to high type only if her upgrading

[18] Notice we can induce any probability distribution for upgrading costs by just letting Ψ be the generalized inverse cdf of that distribution.

cost is weakly less than the difference between two types' expected payoffs under the disclosure rule. The high type proportion p_H is therefore endogenously decided by the disclosure rule now.

Throughout this section, I assume:

Assumption 3.4.1. (1) $\Psi(0) = 0$; (2) $\Psi(1) > v_H - K$; (3) $\Psi(\cdot)$ is continuous.

Given (3), conditions (1) and (2) guarantee that we would always have both high type projects and low type projects in the end (i.e., $p_H \in (0, 1)$). They are not essential but help to avoid trivial cases and ease the discussion. Condition (3) guarantees each $p_H \in (0, 1)$ can be induced by just one value of expected upgrading benefit, which is equal to $\Psi(p_H)$ (i.e., the marginal entrepreneur must be indifferent). This also simplifies the analysis.

Notice this setup is quite general in the sense that $\Psi(\cdot)$ can be any increasing function that satisfies Assumption 3.4.1. Thus the framework can (approximately) capture any positive dependence of high type proportion on the entrepreneurs' perceived benefit from upgrading. The only requirement is that the designer knows this dependence when designing the disclosure rule.

The model's full timeline is given by replacing Stage 1 in the basic model with the entrepreneur's upgrading problem:

–1. Each entrepreneur decides whether to upgrade to high type given her upgrading cost $\Psi(x)$. The project type v is then decided correspondingly.

(Stages 2 – 5 remain the same as in the basic model.)

Designer's problem Let $\mathcal{D}$ denote a disclosure rule (as defined in Definition 3.3.1) and p_H still denote the proportion of high type projects. The designer's problem is to choose a pair of $(p_H, \mathcal{D})$ such that the expected social welfare is maximized and p_H is indeed an equilibrium high type proportion given $\mathcal{D}$.

It is instructive to first consider optimal disclosure without moral hazard ($\beta = 1$) in this setup. Since all projects are to be developed in that case, the efficient level of p_H satisfies $\Psi(p_H) = v_H - v_L$, which means the marginal cost of upgrading equals to its marginal social benefit. Notice the required upgrading incentive $v_H - v_L$ is provided to the entrepreneurs only under full disclosure (since both types get developed), so full disclosure would be the unique optimal disclosure rule without moral hazard.

	$\pi_H \geq B$	$B > \pi_H \geq A$	$\pi_H < \min\{A, B\}$
$v = v_H$	$(1 - \frac{K}{\pi_H v_H + (1-\pi_H) v_L}) v_H$	$(1 - \frac{K}{\pi_H v_H}) v_H$	0
$v = v_L$	$(1 - \frac{K}{\pi_H v_H + (1-\pi_H) v_L}) v_L$	$(1 - \beta) K$	0

Table 3.1: Ex-post Project Payoff $u(\pi_H; v)$

3.4.2 Equilibrium Condition for p_H

To solve the designer's problem, I first derive the equilibrium condition for p_H. For language simplicity, I will use the term "project payoff" to refer to an entrepreneur's payoff without deducting her upgrading cost.

Let $u(\pi_H; v)$ denote the project payoff for a type v entrepreneur when posterior belief π_H is realized. According to Proposition 3.3.1, values for $u(\pi_H; v)$ are as listed in Table 3.1.[19] Given any p_H, the disclosure rule decides the joint distribution over π_H and v. Thus for any pair of $(p_H, \mathcal{D})$, the expected project payoffs for a high type and a low type are respectively:

$$u_H(\mathcal{D}; p_H) := \mathbb{E}_{(p_H, \mathcal{D})}[u(\pi_H; v_H) | v = v_H]; \quad u_L(\mathcal{D}; p_H) := \mathbb{E}_{(p_H, \mathcal{D})}[u(\pi_H; v_L) | v = v_L]$$

where the subscripts to the expectation signs indicate the probability measure is induced by $(p_H, \mathcal{D})$. Notice the expected benefit from upgrading perceived by each entrepreneur is just the difference between these two terms. Thus the equilibrium condition for a proportion p_H of entrepreneurs to upgrade in Stage 1 is then:

$$u_H(\mathcal{D}; p_H) - u_L(\mathcal{D}; p_H) = \Psi(p_H) \tag{3.8}$$

which guarantees the marginal entrepreneur is just indifferent between upgrading and not upgrading, taking aggregate p_H as given. When and only when this condition holds, p_H is an equilibrium proportion of high types and I will say the disclosure rule $\mathcal{D}$ incentivizes p_H.[20]

Notice in general, given a disclosure rule $\mathcal{D}$, there may be multiple p_H that satisfy the

[19] Proposition 3.3.1 still applies since the signaling game starting from Stage 3 is the same as that in the basic model.

[20] The necessity of this condition is guaranteed by Assumption 3.4.1. To see this, notice: first, continuity of $\Psi(\cdot)$ guarantees the marginal entrepreneur must be indifferent when $p_H \in (0, 1)$; second, condition (2) in the assumption implies equilibrium $p_H < 1$ since upgrading benefit cannot exceeds a high type's net value; third, $p_H = 0$ requires the upgrading benefit to be weakly less than $\Psi(0) = 0$, which trivially implies them to be equal since upgrading benefit is non-negative.

condition and thus can be incentivized. When this multiplicity exists, I assume the designer has freedom to pick her target p_H. Actually, one can show that if a disclosure rule can incentivize multiple p_H, the designer will always pick the largest one in optimum.[21]

To use constrained Bayesian persuasion tools later, I need to rewrite the LHS of (3.8) as an expectation of some function on the posterior belief π_H. Define

$$h(\pi_H; p_H) := \begin{cases} \left(1 - \frac{K}{\pi_H v_H + (1-\pi_H) v_L}\right)\left(\frac{\pi_H v_H}{p_H} - \frac{(1-\pi_H) v_L}{1-p_H}\right) & \text{if } \pi_H \geq B \\ \left(1 - \frac{K}{\pi_H v_H}\right)\frac{\pi_H v_H}{p_H} - \frac{1-\pi_H}{1-p_H}(1-\beta)K & \text{if } B > \pi_H \geq A \\ 0 & \text{otherwise} \end{cases} \tag{3.9}$$

Then, condition (3.8) is equivalent to

$$\mathbb{E}_{(p_H, \mathcal{D})}[h(\pi_H; p_H)] = \Psi(p_H) \tag{3.10}$$

(The derivation is provided in Appendix 3.C.2.)

3.4.3 Solution to Designer's Problem

The designer's problem can be solved in two steps. First, given any target p_H, I derive the optimal disclosure rule with condition (3.10) satisfied using the constrained concavification method (Boleslavsky & Kim, 2018). Second, I compute backwardly to characterize the optimal target p_H.

Constrained Bayesian persuasion targeting on p_H

Notice the indirect welfare function on π_H (taking expectation of which gives total social surplus without deducting upgrading costs) is the same as that in the basic model, so the designer's

[21] Also notice the LHS of (3.8) is typically not continuous in p_H due to selection among multiple equilibria. Thus it is possible that there is no equilibrium p_H given some disclosure rules. Such rules are not considered potentially optimal in my analysis.

constrained Bayesian persuasion problem given target p_H is:

$$V(p_H) := \max_{\mu \in \Delta([0,1])} \mathbb{E}_\mu[W(\pi_H)] \tag{3.11}$$

$$\text{s.t. } \mathbb{E}_\mu(\pi_H) = p_H \tag{3.12}$$

$$\mathbb{E}_\mu[h(\pi_H; p_H)] = \Psi(p_H) \tag{3.13}$$

where $W(\cdot)$ is as defined in (3.5). Compared to the problem in Section 3, the key difference here is the inclusion of incentive constraint (3.13), which is just condition (3.10) with subscript $(p_H, \mathcal{D})$ replaced by the induced posterior belief distribution μ. This constraint guarantees the derived disclosure rule can indeed incentivize target p_H. Notice without constraint (3.13), the problem just becomes optimization (3.6) – (3.7) and the unconstrained optimal rule $\mathcal{D}^u(p_H)$ defined in Definition 3.3.2 (with its induced posterior distribution) would be optimal.

Before solving the problem, I first characterize a range for potentially optimal target p_H, which eases the discussion when applying the constrained concavification method. Define:

$$\rho(p_H) := u_H(\mathcal{D}^u(p_H); p_H) - u_L(\mathcal{D}^u(p_H); p_H) \tag{3.14}$$

Then, $\rho(p_H)$ is the difference between two types' expected project payoffs (thus the expected benefit from upgrading) under $\mathcal{D}^u(p_H)$ given that the equilibrium high type proportion is indeed p_H. Thus if target p_H satisfies $\rho(p_H) < \Psi(p_H)$, to incentivize it we must deviate from $\mathcal{D}^u(p_H)$ to induce higher upgrading incentive. If $\rho(p_H) > \Psi(p_H)$, exactly the opposite needs to be done. If $\rho(p_H) = \Psi(p_H)$, the unconstrained optimal rule $\mathcal{D}^u(p_H)$ can just incentivize p_H and is thus also optimal for the constrained problem.

The range for potentially optimal target p_H is provided in the following lemma:

Lemma 3.4.1. *Define constants:*

$$p_H^0 := \sup\{p \in [0,1] : \rho(p) \geq \Psi(p)\} \tag{3.15}$$

$$p_H^1 := \sup\{p \in [0,1] : \Psi(p) \leq v_H - K\} \tag{3.16}$$

Then, the optimal target p_H is in interval $[p_H^0, p_H^1]$. Moreover, we have $1 > p_H^1 > p_H^0 > 0$ with $\rho(p_H^0) = \Psi(p_H^0)$ and $\Psi(p_H^1) = v_H - K$.

Most importantly, the lemma says any optimal $p_H \geq p_H^0$. By the definition of p_H^0, this implies that for any potentially optimal p_H, we never need to induce lower upgrading incentive than that under $\mathcal{D}^u(p_H)$, which suggests that constraint (3.13) is restrictive only in one direction for any potentially optimal p_H.[22] This effectively simplifies the discussion needed when applying the constrained concavification method. The lemma also provides an upper bound p_H^1 for the optimal target p_H. This upper bound is somewhat trivial because p_H^1 is the overall highest p_H one can ever incentivize, since upgrading benefit never exceeds the high type net value $v_H - K$.

It is useful to provide an explicit expression for function $\rho(\cdot)$, because pieces of it appear in the main proposition below.

$$\rho(p_H) = \begin{cases} \rho_L(p_H) := y_1(B) - \frac{(1-B)p_H}{B(1-p_H)} y_2(B) & \text{if } p_H < B \\ \rho_R(p_H) := \frac{B(1-p_H)}{(1-B)p_H} y_1(B) - y_2(B) + \frac{p_H - B}{(1-B)p_H}(v_H - K) & \text{if } p_H \geq B \end{cases} \tag{3.17}$$

where $y_1(B) := (1 - \frac{K}{Bv_H + (1-B)v_L})v_H$ and $y_2(B) := (1 - \frac{K}{Bv_H+(1-B)v_L})v_L$. (See Appendix 3.C.2 for the derivation and more properties of this function.)

Now, by Lemma 3.4.1, I only need to solve problem (3.11) – (3.13) for $p_H \in [p_H^0, p_H^1]$. The result is as follows:

Proposition 3.4.1. *For target $p_H \in [p_H^0, p_H^1]$, the solution to designer's problem (3.11) – (3.13) supports on a subset of $\{0, B, 1\}$ with:*

$$\mu(0) = \frac{\Psi(p_H) - \rho_R(p_H)}{v_H - K - \rho_L(p_H)} \cdot \frac{(1-B)p_H}{B}; \quad \mu(B) = \frac{v_H - K - \Psi(p_H)}{v_H - K - \rho_L(p_H)} \cdot \frac{p_H}{B};$$
$$\mu(1) = \frac{\Psi(p_H) - \rho_L(p_H)}{v_H - K - \rho_L(p_H)} \cdot p_H$$

where $\rho_L(\cdot)$ and $\rho_R(\cdot)$ are as defined in (3.17). Moreover, the value function is: $V(p_H) = (v_H - K)p_H + (v_L - K)(1-B)\mu(B)$.

It is easy to check that when target $p_H = p_H^0$, the optimal rule just becomes $\mathcal{D}^u(p_H^0)$.[23]

[22] Boleslavsky & Kim (2018) also proves a similar result that any optimal target agent effort in their model is weakly higher than the effort readily incentivized by the unconstrained optimal rule. However, their result's logic is much simpler because the designer in their model does not care about the agent's cost in making effort. In contrast, the designer in my model takes the entrepreneurs' upgrading costs into consideration, so the result is much less trivial. A rough intuition here is: because of pooling, one's upgrading always exerts positive externality on other entrepreneurs under any partial disclosure rule, so it is never desirable to induce a low p_H when a higher one can be incentivized without distorting from the unconstrained optimal rule.

[23] To see this, notice either $\Psi(p_H^0) = \rho_L(p_H^0)$ or $\Psi(p_H^0) = \rho_R(p_H^0)$ (or both) depending on whether $p_H^0 \leq B$ or $p_H^0 \geq B$ (or both).

When target p_H satisfies $\Psi(p_H) = v_H - K$, we have $\mu(B) = 0$ and thus the optimal rule is full disclosure. Except for these two cases, the optimal μ has full support on $\{0, B, 1\}$ and three signal realizations would be needed accordingly.[24] For later references, I will call these three signal realizations as low signal (with $\pi_H = 0$), pooling signal (with $\pi_H = B$) and high signal (with $\pi_H = 1$) respectively.

The potential use of a third signal realization when target $p_H > p_H^0$ is an important feature of the optimal disclosure with endogenous types. The realization(s) additional to those in $\mathcal{D}^u(p_H)$ helps to enhance the entrepreneurs' upgrading incentives, so that the target p_H can be incentivized. For instance, when $p_H \in (p_H^0, B)$, while $\mathcal{D}^u(p_H)$ only sends low and pooling signals, the optimal disclosure here also involves the high signal to encourage upgrading through lowering high types' expected financing costs; when $p_H > \max\{p_H^0, B\}$, while $\mathcal{D}^u(p_H)$ only sends pooling and high signals, the optimal rule also uses the low signal to spur upgrading by leaving some low type projects not financed. Overall, the need to incentivize for target $p_H > p_H^0$ requires more informative disclosure than $\mathcal{D}^u(p_H)$. Notice full disclosure is not used, however, unless one wants to induce a highest possible p_H that satisfies $\Psi(p_H) = v_H - K$.

It is also instructive to understand Proposition 3.4.1 as a result about what posterior π_H should not be induced under optimal disclosure. First, as in Proposition 3.3.2, we should never induce $\pi_H \in (0, B)$, since those posteriors cannot get the pooled low types developed and thus do not help to alleviate the moral hazard problem. Second, any $\pi_H \in (B, 1)$ should not be induced either, even if target $p_H \in (B, 1)$, which is in contrast with the result in Proposition 3.3.2. Intuitively, when target $p_H > p_H^0$,[25] we need to provide higher upgrading incentive than that under $\mathcal{D}^u(p_H)$. This makes it valuable to be able to enhance upgrading incentive without preventing pooled low types from developing. Thus, if some $\pi_H \in (B, 1)$ was initially induced, we should split it into two posteriors B and 1 with more disclosure, which then provides higher upgrading incentive while keeping all projects receiving those posteriors developed. As a conclusion, any information that does not hinder low types from developing should be disclosed.

[24] In general, Boleslavsky & Kim (2018) shows with two states (types) and one incentive constraint, optimality can be achieved with no more than three signal realizations. This cardinality bound is tight in this case.

[25] When $p_H = p_H^0$, the implication is trivial since p_H^0 is just induced by $\mathcal{D}^u(p_H^0)$ and by design, $\mathcal{D}^u(p_H^0)$ provides the highest upgrading incentive among all optimal rules in the basic model given exogenous $p_H = p_H^0$, which does not induce any $\pi_H \in (B, 1)$.

Optimization over p_H

Proposition 3.4.1 derives the value function $V(\cdot)$, which gives the maximal social surplus when targeting on $p_H \in [p_H^0, p_H^1]$ without deducting the upgrading costs. Computing backwards, the designer's optimization problem over p_H is:

$$\max_{p_H \in [p_H^0, p_H^1]} \{V(p_H) - \int_0^{p_H} \Psi(x)dx\} \tag{3.18}$$

where the integration computes total upgrading cost by integrating over the individual entrepreneurs' upgrading costs. Since $\Psi(\cdot)$ is assumed to be continuous, it is easy to see that the objective function is continuous. Thus the solution exists by Weierstrass Theorem. Solving this problem then delivers the optimal target p_H. Formally,

Proposition 3.4.2. *A pair of $(p_H, \mathcal{D})$ solves the designer's problem if and only if p_H solves optimization (3.18) and the posterior belief distribution induced by $(p_H, \mathcal{D})$ is as given in Proposition 3.4.1.*

To better understand the trade-off underlying optimization (3.18), consider the derivative of its objective function (assuming differentiability):

$$[v_H - K - \Psi(p_H)] + (v_L - K)(1 - B)\frac{d\mu_B}{dp_H}$$

where $\mu_B := \mu(B)$ in Proposition 3.4.1. By FOC, this derivative should be non-positive at any optimal p_H (a careful examination of the function shows this holds trivially when $p_H = p_H^1$). Notice $v_H - K - \Psi(p_H) \geq 0$ for any $p_H \leq p_H^1$. Thus, at any optimal p_H, we must have $\frac{d\mu_B}{dp_H} \leq 0$. This highlights the key trade-off in the optimization around any optimum: by marginally increasing p_H, one obtains more high type projects, but would have fewer low type projects being developed (whose mass is $(1 - B)\mu_B$).

Why increasing target p_H may require reducing the number of low types to be developed? There are two possible forces behind it. To see them intuitively, we can consider increasing target p_H marginally by dp_H and hypothetically decompose the change in the optimal posterior distribution into two steps. First, we increase $\mu(1)$ by dp_H so that all the new high types are incorporated into the high signal. If initial $\mu(0) = 0$, we have to source these upgradings by reducing a mass dp_H of low types currently receiving the pooling signal. This acts as a direct force to reduce the number of low types being developed. I call it *replacement effect*. If initial $\mu(0) > 0$, on the other hand, we can source the upgradings from those who are receiving the low signal. In this case, $\mu(B)$ is not changed and the replacement effect is zero. After this first step, the posterior distribution gets consistent with the new p_H, but upgrading incentive is typically not just right for incentivizing it. Thus in the second step, we need to further change the posterior distribution in a mean preserving way to achieve the right upgrading incentive. I call the impact of this on the number of pooled low types *incentivizing effect*. If the upgrading incentive after the first step is inadequate, we must provide higher incentive by supplying more information (i.e., making a mean preserving spread on the posterior distribution). Then the incentivizing effect would be negative.[26]

Proposition 3.4.2 characterizes the optimal $(p_H, \mathcal{D})$ in two steps. Given the structure of the optimal posterior support in Proposition 3.4.1, it is also possible to characterize it with a single optimization problem:

Corollary 3.4.1. *A pair of* $(p_H, \mathcal{D})$ *solves the designer's problem if and only if the posterior belief distribution* μ *induced by it supports on a subset of* $\{0, B, 1\}$ *and* $(p_H, \mu(B))$ *solves:*

$$\max_{(p_H, \mu_B)} \{(v_H - K)p_H + (v_L - K)(1 - B)\mu_B - \int_0^{p_H} \Psi(x)dx\} \tag{3.19}$$

$$\begin{aligned} s.t.\ & 1 - p_H - (1 - B)\mu_B \geq 0 \\ & \mu_B \geq 0 \\ & p_H - B\mu_B \geq 0 \\ & \Big[1 - \frac{K}{Bv_H + (1 - B)v_L}\Big]\Big[\frac{Bv_H}{p_H} - \frac{(1 - B)v_L}{1 - p_H}\Big]\mu_B \\ & + \frac{v_H - K}{p_H}(p_H - B\mu_B) = \Psi(p_H) \end{aligned}$$

This characterization does not require solving p_H^0 and is useful in proving several features of the optimal rule below.

3.4.4 Structure of Optimal Disclosure

Proposition 3.4.2 shows the optimal posterior distribution supports on a subset of $\{0, B, 1\}$, which still leaves the structure of the optimal disclosure rule indeterminate. Especially, full

[26] Notice if the optimal p_H satisfies $\Psi(p_H) < v_H - K$ (which is true under a very weak condition as to be shown later), the FOC implies that one of these effects must be strictly negative around the optimum.

disclosure (μ supporting on $\{0, 1\}$) has not been ruled out. In this section, I provide further characterizations on the structure of the optimal rule. In particular, I show: (i) Full disclosure is not optimal under a weak condition on $\Psi(\cdot)$; (ii) Under a slightly stronger condition, each of $\{0, B\}$, $\{B, 1\}$ and $\{0, B, 1\}$ can be the support of the optimal posterior distribution for some set of parameters. Moreover, a comparative statics result with respect to the moral hazard parameter β is also provided. To simplify the analysis, I assume:

Assumption 3.4.2. $\Psi(\cdot)$ is continuously differentiable.[27]

Then, we have the following conclusion:

Proposition 3.4.3. *If $\Psi'(p_H^1) > 0$, then the optimal p_H is strictly less than p_H^1 and full disclosure is not optimal.*

The proposition states that as long as the individual upgrading cost $\Psi(\cdot)$ is increasing in first order at p_H^1, full disclosure would be suboptimal. Intuitively, when p_H increases to p_H^1, the marginal benefit from having more high types developed net of upgrading costs vanishes (since $\Psi(p_H)$ approaches to the high type net value $v_H - K$), while the condition $\Psi'(p_H^1) > 0$ guarantees the loss from reducing pooling that is needed to incentivize higher p_H (negative incentivizing effect) remains first order. Thus it is not optimal to induce p_H^1, which is the only p_H requiring full disclosure to incentivize given condition $\Psi'(p_H^1) > 0$. Therefore, full disclosure is not optimal. Also notice the proposition's condition is sufficient but not necessary. Indeed, even if $\Psi'(p_H^1) = 0$, it is still a special case for the optimal p_H to require full disclosure to incentivize.

Proposition 3.4.3 implies that even with endogenous types and adverse selection concerns, it is still a rare case for full disclosure to be optimal. As long as the designer understands how the type distribution endogenously changes in upgrading incentive, she is able to design a proper partial disclosure rule to alleviate welfare loss caused by moral hazard while sustaining the target type distribution.

To introduce further characterizations on the optimal rule, fix a set of parameters (K, v_L, v_H) and define

$$\underline{\beta} := 1 - (\frac{1}{K} - \frac{1}{v_H})v_L; \ \bar{\beta} := 2 - \frac{v_L}{K}$$

[27] This guarantees the objective function in optimization (3.18) is continuously differentiable, so standard analytical tools are applicable.

Then, β (together with the other parameters) satisfies Assumption 3.2.1 if and only if $\beta \in (\underline{\beta}, \bar{\beta}) \cap [0,1]$. To ease discussion and focus on the most interesting case, I assume $\underline{\beta} \geq 0$ from now on.[28]

For some intuition on how β affects the optimal disclosure, notice when $\beta \to \bar{\beta}$, the moral hazard problem gets less serious and the supporting point $B \to 0$; when $\beta \to \underline{\beta}$, the problem gets more serious and $B \to 1$. It is easy to see (by Proposition 3.4.1) that in both cases, the optimal disclosure rule converges to full disclosure. However, in the former case, the social outcome converges to the first best with p_H satisfying $\Psi(p_H) = v_H - v_L$, which is not surprising as the moral hazard friction vanishes; while in the latter case, the social outcome converges to a situation where only high types are developed and $p_H = p_H^1$ (supposing $\Psi'(p_H^1) > 0$). For later references, I denote the first best p_H as p^{fb} and notice $p_H^1 > p^{fb}$.

Since B and p_H^0 depend on β by their definitions, I denote them as $B(\beta)$ and $p_H^0(\beta)$ respectively to highlight the dependences. One characterization for the optimal disclosure is as follows:

Proposition 3.4.4. *If there exists $\delta > 0$ s.t. $\Psi'(p_H) \geq \delta \, \forall p_H$, then:*

(a) There exists $a > 0$ s.t. for all $\beta \in (\bar{\beta} - a, \bar{\beta})$, we have: the optimal p_H is unique and equals to $p_H^0(\beta)$; the optimal μ supports on $\{B(\beta), 1\}$.

(b) There exists $b > 0$ s.t. for all $\beta \in (\underline{\beta}, \underline{\beta} + b)$, we have: the optimal p_H is unique and equals to $p_H^0(\beta)$; the optimal μ supports on $\{0, B(\beta)\}$.

Under the condition on $\Psi(\cdot)$, the proposition says when the moral hazard problem is sufficiently mild or sufficiently serious, the optimal disclosure rule coincides with the unconstrained optimal rule (given the optimal target p_H) and just uses two signal realizations. In particular, when the problem is mild enough, the optimal disclosure only sends pooling and high signals, so all low types are pooled and developed; when the problem is serious enough, the optimal disclosure only sends pooling and low signals, so all high types are pooled while some low types are left not financed.

A simple result implied by Proposition 3.4.4 is:

Corollary 3.4.2. *Under the condition of Proposition 3.4.4, we have:*

(a) There exists $a > 0$ s.t. the optimal p_H is unique and increases in β when $\beta \in (\bar{\beta} - a, \bar{\beta})$.

[28] For example, $(K, v_H, v_L) = (1, 2, 1.6)$ satisfies Assumption 3.2.1 with $(\underline{\beta}, \bar{\beta}) = (0.2, 0.4)$.

(b) There exists $b > 0$ s.t. the optimal p_H is unique and decreases in β when $\beta \in (\underline{\beta}, \underline{\beta} + b)$.

Corollary 3.4.2 implies that (under its condition) the optimal p_H eventually increases to p^{fb} as $\beta \to \bar{\beta}$ and eventually increases to p_H^1 as $\beta \to \underline{\beta}$. (More comparative statics results are provided in Proposition 3.4.6 below.)

While Proposition 3.4.4 shows (under its condition) only two signal realizations are needed for extreme values of β, the next proposition shows that the optimal disclosure may indeed require a third signal realization for some intermediate values of β. Define $\nu(p) := (1 - \frac{K}{pv_H + (1-p)v_L})(v_H - v_L)$. Then, we have:

Proposition 3.4.5. *Assume the condition in Proposition 3.4.3 holds. Suppose $\Psi(p) = \nu(p) \Rightarrow \Psi'(p) < H(p)$, where $H(\cdot)$ is a positive function provided in the proof. Then there exists a non-degenerate interval I s.t. for all $\beta \in I$, the optimal p_H is strictly bigger than $p_H^0(\beta)$ and the optimal μ has full support on $\{0, B(\beta), 1\}$.*

The proposition is easiest to understand when the result in Proposition 3.4.4 holds (although it does not require so). Roughly speaking, the technical condition in the proposition guarantees that the optimal μ never solely supports on the singleton $\{B(\beta)\}$. Thus when its support transits from $\{0, B(\beta)\}$ to $\{B(\beta), 1\}$ as β increases, it must pass the full support case for some interval of β by a continuity result shown in the proof (Lemma 3.C.2 in Appendix 3.C.2). For these intermediate values of β, the optimal disclosure simultaneously involves a low signal and a high signal, which never happens in the basic model with exogenous p_H.

Figure 3.3 provides a numerical example and plots the optimal target p_H and posterior distribution support for different values of β. As suggested by Proposition 3.4.4 and Proposition 3.4.5, the optimal posterior distribution supports on $\{0, B(\beta)\}$ or $\{B(\beta), 1\}$ when β is sufficiently small or large respectively, and has full support $\{0, B(\beta), 1\}$ for intermediate values of β.

Moreover, two additional features of the graph deserve highlighting. First, the optimal disclosure rule does not use the low signal (i.e., $0 \notin supp\{\mu\}$) if and only if β is higher than a threshold ($\beta = 0.376$). Second, the optimal p_H changes in a U-shape when β evolves, with its minimum achieved at that threshold. These comparative statics patterns turn out to hold more generally, which are explored in the next proposition.

For simplicity, I only focus on the case where the optimal p_H is always unique, which tends to hold when $\Psi(\cdot)$ increases fast or is sufficiently convex (not very concave) in the relevant

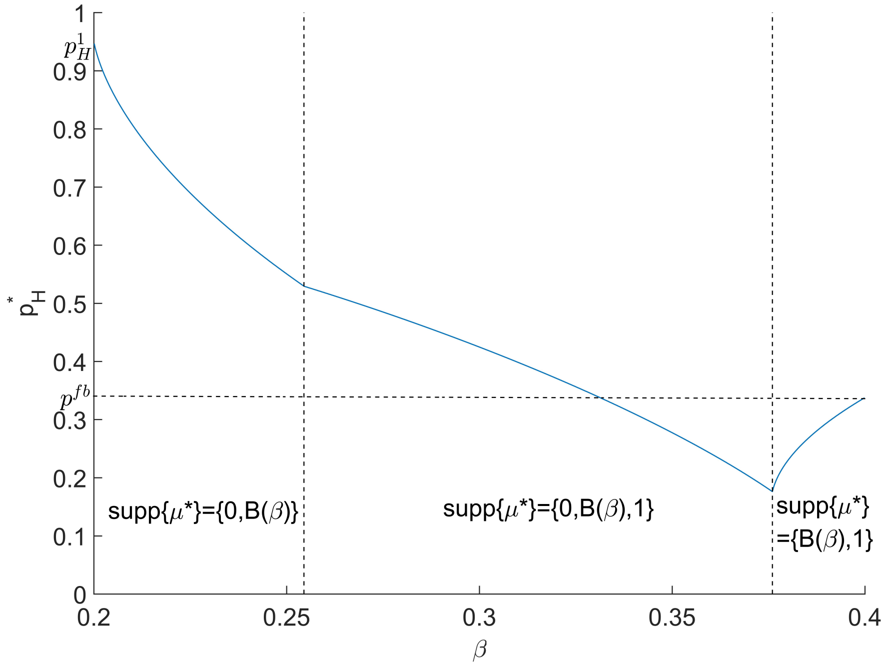

Figure 3.3: p_H^* and $supp\{\mu^*\}$ for different β (with $v_H = 2$, $v_L = 1.6$, $K = 1$ and $\Psi(x)$ has distribution $Normal(0.5, 0.3)$ truncated for the positive part).

region. For any β, let $p_H^*(\beta)$ and $\mu^*(\cdot;\beta)$ denote the corresponding optimal p_H and posterior distribution $\mu(\cdot)$. We have:

Proposition 3.4.6. *Assume the condition in Proposition 3.4.3 holds. Suppose for all* $\beta \in (\underline{\beta}, \bar{\beta})$*, the optimal* p_H *is unique and the objective function in optimization (3.18) is piecewise monotone on* $[0, p_H^1]$*.*[29] *Then, there exists* $\beta_c \in (\underline{\beta}, \bar{\beta}) \cup \{+\infty\}$ *s.t.*

(a) $\mu^*(\cdot;\beta)$ *supports on* $\{B(\beta), 1\}$ *or* $\{B(\beta)\}$ *if and only if* $\beta \geq \beta_c$*. (Singleton* $\{B(\beta)\}$ *may be the support only when* $\beta = \beta_c$*.)*

(b) $p_H^*(\cdot)$ *is decreasing when* $\beta \leq \beta_c$ *and strictly increasing when* $\beta \geq \beta_c$*.*

Moreover, β_c *can be defined as* $\inf\{\beta \in (\underline{\beta}, \bar{\beta}) : p_H^*(\beta) = p_H^0(\beta) \geq B(\beta)\}$*.*

To interpret the results, recall that higher β implies less severe moral hazard problem. Part (a) says there is a threshold β_c such that the optimal disclosure does not use the low signal if and only if $\beta \geq \beta_c$ (i.e., moral hazard is milder than a threshold), in which case all projects can be developed. Part (b) says that the optimal p_H as a function of β is globally U-shaped, which achieves its minimum at β_c (if $\beta_c < +\infty$). Thus as moral hazard problem gets milder from the

[29]Piecewise monotonicity is a rather weak condition that rules out functions oscillating violently. See the proof for its definition.

most serious scenario, the optimal target p_H first decreases from p_H^1 to its global minimum and then increases to p^{fb} (as seen in Figure 3.3). Notice if $\beta_c = +\infty$, the set of $\beta \geq \beta_c$ is empty, so the low signal is always used and the optimal p_H is globally decreasing in β. However, I note that it is more common for β_c to be finite, which is the case, for example, under the condition of Proposition 3.4.4.

Notice the fact that all projects are developed when $\beta \geq \beta_c$ does not imply first best welfare is achieved then, since part (b) implies $p_H^*(\beta)$ would be too little compared to p^{fb} when $\beta \in (\beta_c, \bar{\beta})$. Intuitively, given all projects are developed when $\beta \in (\beta_c, \bar{\beta})$, certain degree of pooling needed to overcome the moral hazard problem necessarily makes upgrading incentive lower than that in the first best, so distortion in p_H compared to the first best is inevitable. This is in contrast to the exogenous types case, where first best welfare can be achieved as long as β is large enough such that $B(\beta)$ is smaller than the exogenous p_H.

To gain better intuition behind the U-shape relation between the optimal p_H and β, we need to consider the two ranges separately. When $\beta \geq \beta_c$, as β increases, the incentive rent needed for low types in the moral hazard problem decreases, so higher upgrading incentive can be provided while keeping all projects developed. Therefore, the optimal p_H, which equals to $p_H^0(\beta)$ in this case, naturally increases. When $\beta < \beta_c$, the intuition is more subtle. First notice in this case, the optimal posterior distribution has $\mu(0) > 0$ by part (a), so the replacement effect around $p_H^*(\beta)$ is zero, which implies that the key cost of increasing target p_H around $p_H^*(\beta)$ comes from the negative incentivizing effect. After β increases, the incentive rent paid to each pooled low type entrepreneur under any optimal disclosure decreases (since B decreases). Thus to increase upgrading incentive by the same amount, more low types must be expelled from pooling and become undeveloped now. This implies that with higher β, the incentivizing effect around the original optimal p_H becomes more negative. Therefore, the trade-off aforementioned in Section 4.3.2 favors lower p_H after β increases.

Finally, I note that even if the optimal p_H is not globally unique, the proposition's results still hold for any subinterval of β where it is unique. For example, if uniqueness holds when $\beta \in (\beta_1, \beta_2)$, then we can find β_c such that $p_H^*(\cdot)$ is decreasing on $(\beta_1, \beta_c]$ and increasing on $[\beta_c, \beta_2)$, where β_c can be defined in the same manner as in the proposition with $(\underline{\beta}, \bar{\beta})$ replaced by (β_1, β_2).

3.5 Discussion

In this section, I provide more discussions regarding implementation of the optimal disclosure rule characterized above from a practical perspective.

3.5.1 Negative Value Projects

When interpreting this paper's results, it is important to notice that both types of projects in the model are socially optimal to carry out. The low type projects are problematic only because they cannot get financed and developed on their own given the moral hazard problem. In this case, the paper shows that it is efficiency-improving to get some low types financed and developed through pooling them with high types in the financing campaigns. In practice, however, some projects are of zero or negative net values (e.g., fraudulences). For those projects, there is certainly no benefit to finance them and the regulator should just eliminate them from the market by disclosing their values whenever possible. The partial disclosure rule designed in this paper should only be applied to the remaining projects with positive net values.

3.5.2 Implementing Optimal Disclosure with Random Inspection

The Bayesian persuasion technique used in this paper is powerful in characterizing the optimal information structure in an abstract (probabilistic) way. However, one may wonder how it can be implemented in practice, especially how the signals leading to different posterior beliefs can be generated in the entrepreneur financing context. For sure, the real world information environment is far more complicated than my parsimonious model. But in the model's stylized two-types setup, a simple random inspection approach can be used to implement the optimal rule (p_H^*, μ^*) in Proposition 3.4.2.

Specifically, suppose there are N projects in the market seeking for financing. The regulator can randomly select $N \cdot \max\{\frac{\mu^*(0)}{1-p_H^*}, \frac{\mu^*(1)}{p_H^*}\}$ of them to investigate.[30] This may involve sending experts to examine the projects or requiring the entrepreneurs to provide supplemental materials to justify their project values. When N is large, in equilibrium, proportion $1 - p_H^*$ of these selected projects would be low types and the rest of them would be high types. Then, a simple calculation shows that after the investigation, at least $\mu^*(0)N$ projects can be deemed as low

[30] It can be shown that this is the minimal number of projects that we need to investigate to implement the information structure.

types and at least $\mu^*(1)N$ projects can be deemed as high types. With these results, we can disclose exactly $\mu^*(0)N$ low types and $\mu^*(1)N$ high types to the public. For the rest of projects, the regulator can just keep silent and let them go for financing directly. This then nicely induces the desired posterior distribution μ^*. The key here is that we need to maintain the posterior belief π_H just high enough for those with null disclosure, so that they can get financed and developed.

This kind of random sampling is widely used by governments to examine and control the quality of physical products (e.g., foods). The discussion above suggests it may also be useful for financial products, although the motivation here is different. In usual contexts, people use random sampling instead of censoring mainly to lower cost. But in our context, it is desired also for inducing partial disclosure.

3.5.3 Cost of Disclosure

It is hard to incorporate disclosure cost in the current framework. There are two reasons: first, it is not clear how to measure the costs for different disclosure rules, which is more of an empirical question; second, the concavification method can only deal with a particular kind of disclosure cost (written as expectation of a function on the posterior beliefs), which may not be reasonable in the current context.

However, it is safe to say that the partial disclosure rule designed in this paper should cost less than full disclosure. As can be seen in the random inspection approach, only some of the projects need to be investigated for partial disclosure. Thus given the presence of disclosure cost, the designer would be more willing to use partial disclosure and probably even less information should be disclosed.

3.5.4 Information Gathering by Investors

One concern about implementing the optimal disclosure rule in practice is that the investors may be able to collect information themselves, which makes it hard to pool different types in the end. However, I argue that in many recently legalized public financing channels for entrepreneurs, the individual investors' incentive and ability in collecting information are very limited.

Taking equity crowdfunding for example, the crowd investors typically do not have expertise in the business field, which limits their ability in collecting and interpreting information. Moreover, in each crowdfunding campaign, the total number of investors is often large and each of them only contributes a small amount of fund. As the total capital raised for a project is usually publicly observable during the funding process, each investor has strong incentive to free ride on others' information and has little incentive to collect information himself. Given these frictions, it is unlikely for voluntary information gathering by investors to generate substantial information on project values. Therefore, the regulator's due-diligence and disclosure would be critical in shaping the information environment, which makes it possible for her to implement pooling results with partial disclosure.

3.5.5 Private Disclosure by Entrepreneurs

Another concern on implementing partial disclosure rules is that the entrepreneurs may have access to private signaling devices. For example, an entrepreneur can hire professionals to certify her business plan or invite potential investors to try the firm's prototype products. If the high types are going to use these signaling devices, we may not be able to pool them with the low types in the end. On this issue, two things can be said.

First, private signals can be too costly to send for small businesses, especially startups. Many small firms have little amount of cash when they go for financing, which makes it hard for them to afford the certification needed for credible signals. This is in remarkable contrast with big firms in IPO, who can hire accounting companies to audit financial statements and hire investment banks to certify their value through underwriting. These means are typically unaffordable for small firms. Moreover, many innovative startups are reluctant to release detailed information about their technology and business pattern to avoid imitation by competitors. Actually, one important praise for the current Regulation Crowdfunding in the United States is that it significantly reduces the disclosure burden of small firms and thus lowers their financing costs. This suggests that, at least within some relevant range, disclosure regulation is binding for small businesses in the sense that they do not voluntarily disclose more than required.

Second, in some cases, sending private signals is costly but not infeasible. When this is true, we can explicitly take it into consideration when designing the optimal disclosure rule. This is studied in Appendix 3.A. It is shown that optimal disclosure would not induce the high types to use private signals and a properly designed partial disclosure rule still improves efficiency over full disclosure.

3.5.6 Application to Banking System Disclosure

While the paper has been focused on entrepreneur public equity financing, its basic intuition is more widely applicable. In Appendix 3.B, I extend the basic model to debt financing and apply it to optimal disclosure regulation for banking system. Alvarez & Barlevy (2015) illustrates with an example that given post-financing moral hazard, full disclosure for the banks can outperform non-disclosure even if the bank financing market does not freeze without disclosure. However, their paper does not consider partial disclosure rules. By using a concavification graph, I intuitively show that under certain conditions, although full disclosure dominates non-disclosure, partial disclosure is typically optimal.

3.6 Conclusion

This paper has studied optimal disclosure regulation for entrepreneur public financing while taking a post-financing moral hazard problem into consideration. It is shown that partial disclosure can help to alleviate efficiency loss caused by the moral hazard problem and thus improve efficiency over full disclosure, even when full disclosure would be costless. This remains true after allowing entrepreneurs' type distribution to be endogenously influenced by the disclosure rule.

With constrained Bayesian persuasion tools, the optimal disclosure rule with endogenous type distribution is fully characterized. Most notably, three features of the optimal disclosure with endogenous types are highlighted. First, it is never optimal to induce a pooling posterior that cannot get the pooled low types developed. In particular, this rules out any post-financing running away by the low types. Second, with the need to incentivize target type distribution, any information that does not prevent pooled low types from developing should be disclosed. Third, to incentivize the optimal target type distribution, more information may need to be disclosed compared to the case where the designer can fix target type distribution without incentivizing for it. Especially, optimal disclosure may simultaneously involve a signal realization that fully identifies a low type and a realization that fully identifies a high type, which never happens with exogenous type distribution.

While the paper has focused on a simple model of entrepreneur equity financing, I believe the main intuition that partial disclosure can help to (partly) solve post-financing moral hazard problem is more generally applicable. In particular, two extensions of the paper may deserve further explorations. One is to apply the model to other financing environments with similar policy concerns. This is pursued as a preliminary attempt in Appendix 3.B, where I adapt the basic model to debt financing and apply it to banking system disclosure regulation. Another possible extension is to replace the moral hazard problem in the model with a more complex mechanism or contract design problem, in which the entrepreneurs' incentive constraints depend on their financing costs and thus can be affected by the disclosure rule in financing campaigns. In practice, such models can be useful to venture capitalists, for example, who may need to reveal information on a bunch of startups they hold in subsequent financing rounds with post-financing incentive concerns. I leave this kind of exercises for future studies.

Appendix

3.A Private Disclosure by High Types

In this section, I extend the analysis by allowing high type entrepreneurs to privately disclose their types to investors after the public signal is realized. A crucial assumption is that this private disclosure is costly, so the high types will not use it if posterior belief π_H induced by the public signal is already above certain threshold. Under this assumption, it is still feasible for the designer to pool some low types with high types without unraveling.[31]

Specifically, assume all high type entrepreneurs have access to a costly certification service that can perfectly reveal their types. Let ξ denote the certification cost and assume $\xi < v_H - K$ (otherwise, the certification is never used). Then, given a posterior π_H induced by the public signal, a high type entrepreneur will not pursue certification if ξ exceeds the benefit from raising the investor's belief on $v = v_H$ from π_H to 1, i.e., if π_H satisfies:

$$\Big(1 - \frac{K}{\pi_H v_H + (1-\pi_H)v_L}\Big)v_H \geq v_H - K - \xi \quad \text{for } \pi_H \geq B \tag{3.20}$$

$$\Big(1 - \frac{K}{\pi_H v_H}\Big)v_H \geq v_H - K - \xi \quad \text{for } B > \pi_H \geq A \tag{3.21}$$

To guarantee the existence of solution to designer's problem, I assume high types will not use

[31] All proofs are put into Appendix 3.C.3.

the certification service when being indifferent.

In this environment, it is actually without loss of generality to focus on public disclosure rules that do not induce high type entrepreneurs to use the costly certification service. With exogenous p_H (basic model), this is obvious since the designer can induce any (Bayesian feasible) information structure for the economy by just using public disclosure. If a high type is going to use the private certification, the designer can simply induce posterior $\pi_H = 1$ for her with public signal and achieve the same social outcome except for saving the entrepreneur's certification cost. With endogenous p_H, however, this is not obvious because the potential certification cost affects entrepreneurs' upgrading incentives and thus influences p_H. Therefore, I formally prove the following lemma:

Lemma 3.A.1. *Any disclosure rule that causes some high type entrepreneurs to use the certification service is not optimal.*

By Lemma 3.A.1, it suffices to consider disclosure rules which only induce posterior beliefs that satisfy condition (3.20) or (3.21). This implies that, to find the optimal disclosure rule, we just need to add an additional constraint on the range of posterior belief π_H when solving the (constrained) Bayesian persuasion problem. Notice condition (3.20) is equivalent to:

$$\pi_H \geq \max\left\{1 - \frac{\xi v_H}{(K+\xi)(v_H - v_L)}, B\right\} \tag{3.22}$$

If $1 - \frac{\xi v_H}{(K+\xi)(v_H - v_L)} \leq B$, then $\pi_H = B$ satisfies the condition and thus high type entrepreneurs will not use the certification when $\pi_H \geq B$ is realized for them. In this case, we do not need to change the optimal rule found in the main text, since the certification is not used under it. Therefore, I only focus on the case where:

$$B < 1 - \frac{\xi v_H}{(K+\xi)(v_H - v_L)} \tag{3.23}$$

In this case, conditions (3.20) and (3.21) lead to the following observation:

Observation 3.A.1. *Given (3.23), costly certification is not used if and only if posterior belief* $\pi_H \in \{0\} \cup [\underline{\pi}, 1]$*, where* $\underline{\pi} = 1 - \frac{\xi v_H}{(K+\xi)(v_H - v_L)}$.

By Lemma 3.A.1 and Observation 3.A.1, the access to private disclosure just imposes an additional constraint that $\pi_H \in \{0\} \cup [\underline{\pi}, 1]$ under any optimal disclosure rule. This can be easily handled with the concavification methods used in the main text by ignoring the interval $(0, \underline{\pi})$ for π_H or, equivalently, by assuming the indirect welfare function $W(\cdot)$ takes negative infinity over that posterior range. In the next two subsections, I characterize the optimal disclosure rules with exogenous p_H and endogenous p_H respectively.

3.A.1 Exogenous p_H

The following proposition is analogous to Proposition 3.3.2.

Proposition 3.A.1. *Given exogenous high type proportion p_H,*

(i) If $p_H < \underline{\pi}$, the optimal distribution of posteriors supports on $\{0, \underline{\pi}\}$ with $\mu(\underline{\pi}) = p_H/\underline{\pi}$ and $\mu(0) = 1 - p_H/\underline{\pi}$.

(ii) If $p_H \geq \underline{\pi}$, any distribution of posteriors supporting on a subset of $[\underline{\pi}, 1]$ with $\mathbb{E}_\mu[\pi_H] = p_H$ is optimal.

It is easy to see as long as $\xi > 0$, $\underline{\pi}$ is smaller than 1 and the optimal disclosure rule is partial disclosure. Intuitively, as long as the certification is not free, the designer can still pool some low types with high types without invoking a high type's self-certificating. This intuition and the optimality of partial disclosure remain valid even if some high type entrepreneurs can get certification with no cost. In that case, these costless certifications are always used, but we can still pool certain amount of low types with the rest of high types whose certifications are costly. The existence of private signaling cost prevents unraveling from happening when the disclosure rule is properly designed.

3.A.2 Endogenous p_H

With endogenous p_H, I first derive the optimal disclosure rule given any target p_H. By Lemma 3.A.1 and Observation 3.A.1, we simply need to add the additional constraint on posterior belief to problem (3.11) – (3.13), which becomes

$$V(p_H; \xi) := \max_{\mu \in \Delta([0,1])} \mathbb{E}_\mu[W(\pi_H)] \tag{3.24}$$

$$\text{s.t. } \mathbb{E}_\mu(\pi_H) = p_H \tag{3.25}$$

$$\mathbb{E}_\mu[h(\pi_H; p_H)] = \Psi(p_H) \tag{3.26}$$

$$\mu\big((0, \underline{\pi})\big) = 0 \tag{3.27}$$

Notice the access to costly certification does not affect functions $W(\cdot)$ and $h(\cdot;p_H)$ as long as (3.27) holds so that the certification is never used.

To solve this problem, one can just follow the same procedures as those for optimization (3.11) – (3.13). Because the proofs are tedious and basically the same as their counterparts in Section 4, I omit them and directly present the main results.

Define function $\rho(\cdot)$ similar to that in (3.17) but with B replaced by $\underline{\pi}$:

$$\rho(p_H;\xi) := \begin{cases} \rho_L(p_H;\xi) := y_1(\underline{\pi}) - \frac{(1-\underline{\pi})p_H}{\underline{\pi}(1-p_H)} y_2(\underline{\pi}) & \text{if } p_H < \underline{\pi} \\ \rho_R(p_H;\xi) := \frac{\underline{\pi}(1-p_H)}{(1-\underline{\pi})p_H} y_1(\underline{\pi}) - y_2(\underline{\pi}) + \frac{p_H-\underline{\pi}}{(1-\underline{\pi})p_H}(v_H - K) & \text{if } p_H \geq \underline{\pi} \end{cases} \tag{3.28}$$

where $y_1(\underline{\pi}) := (1 - \frac{K}{\underline{\pi}v_H+(1-\underline{\pi})v_L})v_H$ and $y_2(\underline{\pi}) := (1 - \frac{K}{\underline{\pi}v_H+(1-\underline{\pi})v_L})v_L$. The following result is analogous to Lemma 3.4.1 and Proposition 3.4.1:

Proposition 3.A.2. *Given $\xi > 0$, define:*

$$p_H^0 := \sup\{p \in [0,1] : \Psi(p) \leq \rho(p;\xi)\};\ p_H^1 := \sup\{p \in [0,1] : \Psi(p) \leq v_H - K\}$$

Then, the optimal target $p_H \in [p_H^0, p_H^1]$. Moreover, for any target $p_H \in [p_H^0, p_H^1]$, the solution to designer's problem (3.24) – (3.27) supports on a subset of $\{0, \underline{\pi}, 1\}$ with:

$$\mu(0) = \frac{\Psi(p_H) - \rho_R(p_H;\xi)}{v_H - K - \rho_L(p_H;\xi)} \cdot \frac{(1-\underline{\pi})p_H}{\underline{\pi}}; \quad \mu(\underline{\pi}) = \frac{v_H - K - \Psi(p_H)}{v_H - K - \rho_L(p_H;\xi)} \cdot \frac{p_H}{\underline{\pi}};$$
$$\mu(1) = \frac{\Psi(p_H) - \rho_L(p_H;\xi)}{v_H - K - \rho_L(p_H;\xi)} \cdot p_H$$

Thus the optimal disclosure rule given any potentially optimal target p_H is just similar to that in Proposition 3.4.1, but with B replaced by $\underline{\pi}$.

Finally, I extend Proposition 3.4.3 to the current setup:

Proposition 3.A.3. *With Assumptions 1 – 3 (in the main text) and $\xi > 0$, if $\Psi'(p_H^1) > 0$, then full disclosure is not optimal.*

In conclusion, although high types' access to costly certification limits the designer's pooling ability as $\underline{\pi} > B$ under condition (3.23), a properly designed partial disclosure rule can still dominate full disclosure. Again, costliness of private disclosure prevents unraveling from happening and thus allows the designer to implement certain degree of partial disclosure.

3.B Debt Financing and Application to Banking System

In this section, I extend the basic model to debt financing and apply it to (socially) optimal disclosure for banking system.

In Alvarez & Barlevy (2015), the authors give an example where banks are facing post-financing moral hazard problem. They show that full disclosure improves welfare over non-disclosure even if the market does not freeze without disclosure, which is in contrast with Goldstein & Leitner (2018). The main intuition is that without disclosure, the low types can get financed but will misbehave afterwards, which causes a social loss. Imposing full disclosure can prevent these low types from getting financed and thus save this loss. However, their paper does not consider partial disclosure. A key lesson from my model is that partial disclosure can help to incentivize some of the low types in the moral hazard problem to behave and thus improves efficiency over full disclosure. Therefore, as to be shown, a properly designed partial disclosure can dominate full disclosure in social welfare.

I first adapt the setup in Section 3 to debt financing and motivate partial disclosure under certain assumptions. Then, I solve the model's equilibria and derive the optimal disclosure rule. Throughout this section, I assume every agent is risk neutral and the discounting factor is normalized to 1. All assumptions on fundamental parameters are summarized in Assumption 3.B.1. All proofs are put into Appendix 3.C.4.

3.B.1 Model Setup

Agents and the moral hazard problem

There is a bank with two possible initial equity levels in period 0, denoted as $e \in \{e_L, e_H\}$ ($e_H > e_L$).[32] I refer to e as the bank's type and assume the probability of high type (e_H) is $p_H \in (0, 1)$. The bank has a project that delivers random payoff $z \sim F$ in period 1 if developed, where F is the cdf for z. A unit investment ($K = 1$) is required for the project to be developed and the bank must finance it through issuing debt in period 0 (equity is assumed to be illiquid in period 0). Assume $\mathbb{E}(z) > 1$, so it is socially optimal to develop the project.

There is a representative investor, who forms belief about the bank's type based on public information available in the market. The investor can either offer the bank a standard debt

[32] As usual, one can interpret the model as having a continuum of banks.

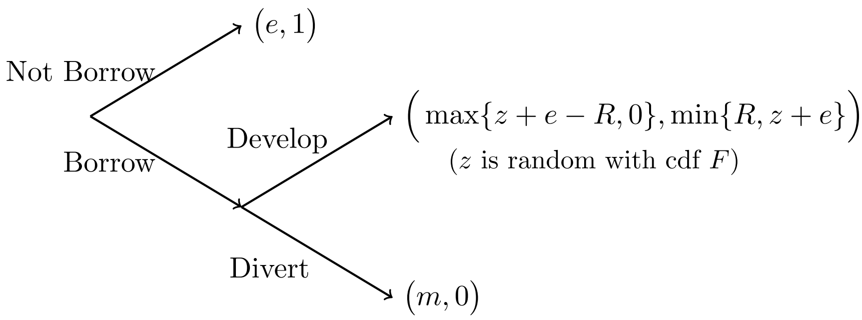

Figure 3.4: Bank's decision problem after being offered with gross interest rate R in date 0. (First payoff is for the bank; second payoff is for the investor.)

contract to borrow for its investment, or just save his money in a risk free asset which pays zero interest rate.

Suppose an offer is made to the bank with R being the gross interest rate. Then, the bank can decide whether to borrow with that offer. If not, the bank just keeps its current equity without developing the project and the investor saves his fund in the risk-free asset; if yes, after getting financed, the bank can decide whether to honestly develop the project, or to divert the fund for other uses (e.g., engaging in high risk adventures). For simplicity, assume that if the bank diverts the fund, it will get final payoff m and the investor will get payoff 0. The fact that these payoffs do not depend on e or R is certainly not without loss of generality, but can ease the discussion significantly. If the bank develops the project, payoffs are realized according to the standard debt contract.

The bank's decision problem is as summarized in Figure 3.4, where the first elements in the parentheses are payoffs for the bank and the second elements are payoffs for the investor. (Notice z is random for both agents, while e is known by the bank but not by the investor.) There are several implicit assumptions underlying these payoffs. First, the bank has limited liability; second, the bank's equity is assumed to be liquid in period 1, so it can be used to repay the debt; third, the new debt holder investing in period 0 is junior to any existing debt holders. These assumptions together imply that the bank's payoff is always non-negative and the investor's payoff is limited by $z + e$ after investing.

Bank's incentive

To keep things simple and interesting, I assume:

$$e_L < m < \min\{e_H, e_L + 1\} \tag{3.29}$$

This leads to three features for the model. First, the low type bank always chooses to borrow once it gets an offer, since it can at least divert the fund and $m > e_L$. Second, the high type never chooses to divert the fund since $m < e_H$, so it either develops the project or chooses not to borrow. These reasonably reflect general firm behaviors in reality. Firms with little equities are often eager to borrow at whatever interest rate and have severe moral hazard problem, while those in solid financial status are less willing to issue high-yield debt and less likely to exploit the debt holders. Third, there is a social loss from diverting the fund, as $m < e_L + 1$.

Define $S(e, R) := \mathbb{E}[\max\{z + e - R, 0\}|e]$. Then, by the decision problem in Figure 3.4 and the discussions above, we can see that a high type bank borrows if and only if $S(e_H, R) \geq e_H$ and a low type bank develops the project if and only if $S(e_L, R) \geq m$.[33] Because $S(e, R)$ is decreasing in R, we can conveniently express these observations as the following:

Observation 3.B.1. *Let $R_1, R_2 > 0$ be such that $S(e_L, R_1) = m$ and $S(e_H, R_2) = e_H$. Then, R_1 and R_2 are uniquely defined. Moreover, when the investor offers with gross interest rate R, we have:*

- *A low type bank will borrow and develop its project if $R \leq R_1$; otherwise, it will accept the offer but divert the fund.*
- *A high type bank will borrow and develop its project if $R \leq R_2$; otherwise, it will not borrow.*

To focus on the most interesting case and compare to Alvarez & Barlevy (2015), I shall assume $R_2 > R_1$. Notice otherwise, the low type can never divert the fund in equilibrium.[34]

[33]When being indifferent, the bank is assumed to behave in favor of social efficiency.

[34]Suppose $R_1 \geq R_2$. For the low type to divert the fund, we must have $R > R_1 \geq R_2$. But with this R, only the low type borrows. Thus the investor would not lend anticipating the low type to divert the fund.

Investor's incentive

For the investor, let $B(e, R) := \mathbb{E}[\min\{R, z + e\}|e]$. Then, the investor would like to offer the bank with interest rate R if and only if:

$$\mathbb{E}\big[B(e, R)\mathbb{1}\{a(e, R) = \text{Borrow and Develop}\} + \mathbb{1}\{a(e, R) = \text{Not Borrow}\}|\mathcal{I}\big] \geq 1 \tag{3.30}$$

where $a(e, R)$ denotes the bank's choice in Figure 3.4 and $\mathcal{I}$ denotes the investor's information when making the decision. Notice the investor is not optimizing over R, as he potentially represents a large group of investors. The exact value of R decides how the investors would split surplus with the bank, which depends on the detailed market structure. For example, if there are many investors who competitively bid their offers, then all surplus goes to the bank and equilibrium R should satisfy (3.30) with equality. For my purpose here, however, exact value of R is not needed and it suffices to know that any equilibrium R satisfies the inequality and if such an R exists, the bank will receive an offer.

Full disclosure and efficiency loss

Under full disclosure, the investor knows the bank's type. With Observation 3.B.1, condition (3.30) is simplified to:

$$(B(e_H, R) - 1)\mathbb{1}\{R \leq R_2\} \geq 0 \text{ when the bank is of high type} \tag{3.31}$$

$$B(e_L, R)\mathbb{1}\{R \leq R_1\} - 1 \geq 0 \text{ when the bank is of low type} \tag{3.32}$$

Assume $B(e_L, R_1) < 1 < B(e_H, R_2)$. Then, one can see $B(e_L, R)\mathbb{1}\{R \leq R_1\} < 1$ for all R (notice $B(e, \cdot)$ is increasing in R). Thus the low types can not get financed, which causes an efficiency loss compared to the first best. To mitigate this loss, we may use partial disclosure to reduce the low type's financing cost, so as to provide enough incentive for it to behave in the moral hazard problem. This turns out to require the following condition to hold:

$$B(e_H, R_1) > 1 \tag{3.33}$$

which is reminiscent of condition (3.4).[35]

The following assumption summarizes all assumptions I made on the parameters and distribution of z.

Assumption 3.B.1. e_L, e_H, m and distribution of z satisfies:

(a) $\mathbb{E}[z] > 1$;

(b) $e_L < m < \min\{e_H, e_L + 1\}$;

(c) $B(e_L, R_1) < 1 < B(e_H, R_1)$;

(d) $R_2 > R_1$.

(Recall R_1 and R_2 are defined by $S(e_L, R_1) = m$ and $S(e_H, R_2) = e_H$.)

According to discussions above, (c) is the most important assumption we need. The other assumptions are only for simplification or for keeping the question interesting. I will show that under Assumption 3.B.1, partial disclosure can improve welfare over full disclosure. An example that satisfies the assumption is: $e_L = 0.34$, $e_H = 1.5$, $m = 0.9$ and $z \sim Uniform[0, 3]$.

Timeline and the designer's problem

At the beginning of period 0, the bank's type is realized and is learned by the regulator (e.g., through stress tests). Then, the regulator sends a signal on it to the market according to a pre-announced disclosure rule. A disclosure rule is defined as:

Definition 3.B.1. A disclosure rule consists of a finite signal realization space $\mathcal{S}$ and a family of distributions $\{\Gamma(\cdot|e_L), \Gamma(\cdot|e_H)\}$ on $\mathcal{S}$.

After the signal is realized, the investor makes a debt offer to the bank with gross interest rate R satisfying condition (3.30) if such an R exists. Then, the bank decides whether to accept the offer and if yes, whether to develop the project.

Given a disclosure rule, the game's timeline is summarized as the following:

–1. The bank learns its type $e \in \{e_L, e_H\}$ with probability p_H for e_H.

–2. A signal realization s on e is disclosed according to the pre-announced disclosure rule.

–3. The investor makes a debt offer with gross interest rate R that satisfies condition (3.30) if such an R exists, where the information set is $\mathcal{I} = \{s\}$. If no such R exists, the bank remains not financed and the game ends.

[35] Intuitively, this guarantees the low type's financing cost can get sufficiently low when it is recognized as a high type by the investor.

–4. When offered with R, the bank solves the decision problem in Figure 3.4. Payoffs then realize accordingly.

Compared to the setup of financing game in Section 3, a difference here is that I assume the investor is to make the offer. This saves us from dealing with the signaling game and simplifies the analysis. However, one can also let the bank make the offer and follow the procedures in Section 3 to show only pooling equilibria exist in the signaling game. Thus the equilibrium outcomes of these two setups are going to be identical.

The designer's problem is to design a disclosure rule to maximize the total social surplus. Notice for any signal realization s, the investor forms a belief on the bank's type, which then decides the equilibrium outcome in Stages 3 – 4. Let π_H denote the belief on $e = e_H$ induced by the signal. Given any π_H, the (subgame perfect) equilibrium is solved in the next subsection.

3.B.2 Equilibrium given π_H

Just like the financing game in Section 3, there can be three classes of equilibria:

- Class 1: $R \leq R_1$, so both types borrow and develop the project.
- Class 2: $R_1 < R \leq R_2$, so both types borrow but the low type diverts the fund.
- Class 3: no offer is made by the investor.

  (Notice the investor never offers $R > R_2$, since only the low type would accept it and the fund will be diverted.)

Define constants:

$$C_1 := \frac{1 - B(e_L, R_1)}{B(e_H, R_1) - B(e_L, R_1)}; \quad C_2 := \frac{1}{B(e_H, R_2)}$$

By Assumption 3.B.1, it is easy to see $C_1, C_2 \in (0, 1)$.[36] Then, the basic result is the following:

Lemma 3.B.1. *Class 1 equilibria exist if $\pi_H \geq C_1$; Class 2 equilibria exist if $\pi_H \geq C_2$; only Class 3 equilibria exist if $\pi_H < \min\{C_1, C_2\}$.*

When multiple classes of equilibria exist, I select the one that is most efficient. Notice by Assumption 3.B.1, $\mathbb{E}(z) > 1$ and $m < e_L + 1$, so each developed project delivers a positive (expected) social surplus and diverting causes a social loss. Therefore, the selection rule is: select Class 1 equilibria if they exist; otherwise, select Class 2 equilibria if they exist; otherwise,

[36] Notice it is not sure whether $C_1 > C_2$ or $C_1 \leq C_2$. Both cases can happen in general.

select Class 3 equilibria. This selection rule allows us to focus on the informational problem here, instead of equilibrium multiplicity in debt financing.[37]

According to the selection rule, the equilibrium outcome is:

Proposition 3.B.1. *When $\pi_H \geq C_1$, both types develop the project; when $C_2 \leq \pi_H < C_1$, the high type develops the project and the low type diverts the fund; when $\pi_H < \min\{C_1, C_2\}$, no bank is financed.*

3.B.3 Optimal Disclosure

To use Bayesian persuasion tools, I first derive the indirect social welfare function. Let b denote the expected net surplus from each developed project, i.e., $b := \mathbb{E}[z] - 1$. Let ℓ denote the loss from a low type's diverting, i.e., $\ell := 1 + e_L - m$. Then, according to the equilibrium outcomes in Proposition 3.B.1, the social surplus given bank type e and posterior belief π_H is:

$$\begin{aligned} w(e, \pi_H) :=& b\big[\mathbb{1}\{\pi_H \geq C_1\} + \mathbb{1}\{C_1 > \pi_H \geq C_2\}\mathbb{1}\{e = e_H\}\big] \\ &- \ell\mathbb{1}\{C_1 > \pi_H \geq C_2\}\mathbb{1}\{e = e_L\} \end{aligned}$$

Taking expectation conditional on π_H leads to the indirect social welfare function:

$$W(\pi_H) := \mathbb{E}[w(e, \pi_H)|\pi_H] = \begin{cases} b & \text{if } \pi_H \geq C_1 \\ b\pi_H - \ell(1 - \pi_H) & \text{if } C_1 > \pi_H \geq C_2 \\ 0 & \text{otherwise} \end{cases}$$

Then, the Bayesian persuasion problem of the designer is:

$$\begin{aligned} &\max_{\mu \in \Delta([0,1])} \mathbb{E}_\mu[W(\pi_H)] \\ &\text{s.t. } \mathbb{E}_\mu(\pi_H) = p_H \end{aligned} \tag{3.34}$$

The following proposition provides the solution to it.

Proposition 3.B.2. *If $p_H < C_1$, the optimal distribution of posteriors supports on $\{0, C_1\}$, with $\mu(C_1) = p_H/C_1$ and $\mu(0) = 1 - p_H/C_1$; if $p_H \geq C_1$, any distribution of posteriors supporting*

[37] For my purpose here, I don't need to specify which equilibrium to select within a class.

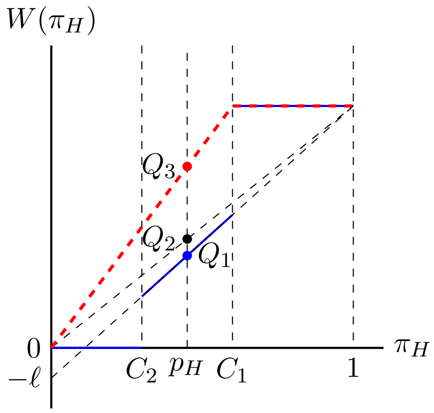

Figure 3.5: Comparison for Disclosure Rules

Note: The blue curve is function $W(\pi_H)$; the red dotted curve is the concavification boundary.

on a subset of $[C_1, 1]$ *with* $\mathbb{E}_\mu[\pi_H] = p_H$ *is optimal.*

In particular, the proposition implies full disclosure is never optimal (under Assumption 3.B.1). In the next subsection, I compare this result with that in Alvarez & Barlevy (2015).

3.B.4 Discussion

Alvarez & Barlevy (2015) proposes an example, in which without any disclosure, the market does not freeze but a low type diverts the fund once getting financed. In my model, this is corresponding to the case of $C_2 \leq p_H < C_1$, in which Class 2 equilibria realize without disclosure. Similar to the result in Alvarez & Barlevy (2015), my model also predicts that full disclosure improves welfare over non-disclosure in this case. The novel implication of my model, however, is that partial disclosure further dominates full disclosure when properly designed. These insights are intuitively illustrated with Figure 3.5.

Specifically, Figure 3.5 plots the indirect welfare function $W(\cdot)$ together with its concavification for the case $C_2 \leq p_H < C_1$. Without any disclosure, the social welfare is given by $W(\cdot)$ evaluated at p_H, which is represented by point Q_1. With full disclosure, the social welfare is the convex combination of $W(0)$ and $W(1)$ with the weight on $W(1)$ being p_H, which is represented by point Q_2. Under the optimal disclosure, the social welfare is the concavification curve evaluated at p_H, which is given by point Q_3. It is easy to see although Q_2 is better than Q_1, Q_3 strictly dominates both of them.

Intuitively, full disclosure outperforms non-disclosure in that it prevents low type banks from ever diverting the fund. However, when doing so, it necessarily leaves all low types not financed. In contrast, the optimal disclosure rule in Proposition 3.B.2 is able to make some low types develop their projects while keeping the rest of them not financed. In the end, there is still no loss from diverting, but we can enjoy higher surplus from more projects being developed.

3.C Proofs

3.C.1 Proofs in Section 3

Proof for Observation 3.3.1

Proof. The proof has 3 steps:

- Step 1: Either all on-path offers are financed or all on-path offers are not financed.

  Suppose in an equilibrium there are two offers α_1 and α_2, where α_1 is financed but α_2 is not. The entrepreneur playing α_2 currently gets payoff 0, so she can profitably deviate to α_1 to be financed and receive positive payoff (at least she can run away with the fund). Thus we cannot be in equilibrium.

- Step 2: If some projects are financed, there is no on-path offer that is only played by the low type.

  Otherwise, the low type is recognized with that offer and cannot get financed due to the discussion in Section 2.3. This violates the result in step 1.

- Step 3: If some projects are financed, there is only one on-path offer.

  Due to step 2, all on-path offers are used by the high type with positive probability. Suppose there are two on-path offers α_1 and α_2 ($\alpha_2 > \alpha_1$). By step 1, both offers are financed. This implies the high type is going to develop her project with both offers. (Otherwise, both types will divert with at least one of the offers and that offer cannot be financed in equilibrium then.) Since $\alpha_2 > \alpha_1 \Rightarrow (1-\alpha_2)v_H < (1-\alpha_1)v_H$, the high type strictly prefers α_1 to α_2, so they cannot both be on equilibrium path.

Q.E.D.

Proof for Lemma 3.3.1

Proof. Consider the three classes of equilibria one by one:

- Class 1: Both types of projects are financed and developed.

  By Observation 3.3.1, this kind of equilibrium must be pure strategy pooling equilibrium.

Let α denote the equilibrium offer. We must have that incentive condition (3.1) holds for both types and condition (3.2) holds so that the investor would like to accept the offer. Specifically, we need:

$$(1-\alpha)v_H \geq (1-\beta)K \tag{3.35}$$

$$(1-\alpha)v_L \geq (1-\beta)K \tag{3.36}$$

$$\alpha[\pi_H v_H + (1-\pi_H)v_L] \geq K \tag{3.37}$$

where the third condition is from condition (3.2) with the expectation computed given initial belief π_H (since we are considering pooling equilibria). Since $v_H > v_L$, the first condition is redundant. Combining (3.36) and (3.37) we obtain:

$$1-(1-\beta)\frac{K}{v_L} \geq \alpha \geq \frac{K}{\pi_H v_H + (1-\pi_H)v_L} \tag{3.38}$$

There exists an α satisfying the condition if and only if the upper bound is higher than the lower bound, which is equivalent to

$$\pi_H \geq \left[\frac{K}{v_L-(1-\beta)K} - 1\right]\frac{v_L}{v_H - v_L} = B \tag{3.39}$$

Notice Assumption 3.2.1 guarantees the RHS is in $(0,1)$.

The above analysis shows the necessity of condition (3.39). For sufficiency, notice when condition (3.39) is satisfied, for any α^* that satisfies (3.38), we have a weak sequential equilibrium: both types offer share α^*; the investor invests if and only if the offer is α^* (with on-path belief $\mathbb{P}(v = v_H) = \pi_H$ and off-path belief, e.g., $\mathbb{P}(v = v_H) = 0$); the entrepreneur develops her project when (3.35) or (3.36) holds for the two types respectively. As α^* satisfies both types' incentive constraints, both types are developed in this equilibrium. To check this is indeed an equilibrium, notice α^* satisfying condition (3.37) guarantees the investor will not deviate from investing. Moreover, the off-path belief guarantees any offer other than α^* will not be financed and thus the entrepreneurs will not deviate to other offers.

(These equilibria satisfy the intuitive criterion (Cho & Kreps, 1987). For any $\alpha > \alpha^*$, no type wants to deviate to it since it just increases the entrepreneur's financing cost even if she still gets financed; for any $\alpha < \alpha^*$ satisfying $\alpha v_H \geq K$, both types would be better off when deviating to it if they can be recognized as a high type and thus get financed with less equity sold; for any $\alpha < \alpha^*$ satisfying $\alpha v_H < K$, no type wants to deviate to it as it is never accepted by the investor.)

- Class 2: Both types are financed, but only high type projects are developed.

  By Observation 3.3.1 again, we only need to consider pure strategy pooling equilibria. Let α be the equilibrium offer. We have conditions:

$$(1-\alpha)v_H \geq (1-\beta)K > (1-\alpha)v_L \tag{3.40}$$

$$\alpha v_H \pi_H \geq K \tag{3.41}$$

  where the first condition guarantees the high type will develop the project and the low type will run away after getting funded; the second condition guarantees the investor does invest despite anticipating a low type entrepreneur to run away, which is derived from condition (3.2) with on-path belief π_H.

  There exists an α that satisfies both conditions if and only if:

$$\pi_H \geq \frac{K}{v_H - (1-\beta)K} = A \tag{3.42}$$

  Notice the RHS is always between 0 and 1 since $v_H \geq (2-\beta)K$ by Assumption 3.2.1.

  When (3.42) holds, for any α^* that satisfies (3.40) and (3.41), we have the following equilibrium: both types offer α^* in the campaign; the investor invests in the project if and only if the offer is α^* (with on-path belief the same as initial belief π_H and off-path belief chosen as $\mathbb{P}(v = v_H) = 0$); after getting funded, the entrepreneur's choice is governed by the incentive condition (3.1). Since α^* satisfies (3.40), only the high type project is developed in this equilibrium. To check this is indeed an equilibrium, notice condition (3.41) guarantees the investor will accept offer α^*; the off-path belief guarantees any other offer is not accepted and thus the entrepreneurs will not deviate from α^*.

  (Any such equilibrium also survives intuitive criterion. The argument is exactly the same as that for Class 1 equilibria.)

- Class 3: No project is financed.

  This class of equilibria always exists. For example, we can let both types offer $\alpha = 0$ and set all off-equilibrium beliefs to $\mathbb{P}[v = v_H] = 0$. This kind of equilibrium also survives intuitive criterion since both types are receiving the lowest possible payoffs for them.

Q.E.D.

Proof for Proposition 3.3.1

Proof. The only thing requiring clarification here is the equilibrium α in the first two cases. In the case of $\pi_H \geq B$, it comes from the lower bound in condition (3.38) for Class 1 equilibria. In the case of $B > \pi_H \geq A$, by conditions (3.40) and (3.41), we need an α from a Class 2 equilibrium to satisfy:

$$\alpha > 1 - (1-\beta)\frac{K}{v_L} \text{ and } \alpha \geq \frac{K}{\pi_H v_H}$$

Notice when $\pi_H < B$, we have (3.39) violated, so

$$1 - (1-\beta)\frac{K}{v_L} < \frac{K}{\pi_H v_H + (1-\pi_H)v_L} < \frac{K}{\pi_H v_H}$$

Therefore, the smallest α selected from Class 2 equilibria is $\frac{K}{\pi_H v_H}$. *Q.E.D.*

Proof for Proposition 3.3.2

Proof. • Case 1: $B > A$

After some simplifications, the indirect social welfare function becomes:

$$W(\pi_H) = \begin{cases} (v_H - v_L)\pi_H + v_L - K & \text{if } \pi_H \geq B \\ \big(v_H - (1-\beta)K\big)\pi_H - \beta K & \text{if } B > \pi_H \geq A \\ 0 & \text{if } \pi_H < A \end{cases}$$

The indirect welfare function together with its concavification is drawn in Figure 3.2a (main text). It is easy to see the second piece of $W(\cdot)$ is below the segment joining $(0,0)$ and $(B, W(B))$. The red dashed curve then represents the concavification function, whose graph consists of two segments: $(0,0) \to (B, W(B))$ and $(B, W(B)) \to (1, W(1))$.

Therefore, when $p_H < B$, the optimal μ supports on $\{0, B\}$. By Bayesian feasibility, $\mu(0) \times 0 + \mu(B)B = p_H \Rightarrow \mu(B) = p_H/B$. When $p_H \geq B$, any μ supporting on $[B, 1]$ satisfying Bayesian feasibility condition is optimal, since the second piece of the concavification function coincides with $W(\cdot)$.

- Case 2: $B \leq A$

  In this case, we simply do not have the second piece of $W(\cdot)$, so the concavification graph is as given in Figure 3.2b. We have the same expression for the concavification as in Case 1 above, except that now there is no segment corresponding to interval $[A, B)$. Thus the same result as in Case 1 is obtained.

Under the optimal rule, when $p_H \geq B$, all projects are developed, so the social welfare is $(v_H - K)p_H + (v_L - K)(1 - p_H)$; when $p_H < B$, all high types and those low types receiving posterior $\pi_H = B$ are developed, so the social welfare is $(v_H - K)p_H + (v_L - K)\frac{(1-B)p_H}{B}$. (Notice the measure of low types receiving $\pi_H = B$ is $(1 - B)\mu(B) = \frac{(1-B)p_H}{B}$.) Combining these gives the expression for the value function. *Q.E.D.*

3.C.2 Proofs in Section 3.4

Derivation of the incentive condition (3.10)

We want to show $u_H(\mathcal{D}; p_H) - u_L(\mathcal{D}; p_H) = \mathbb{E}_{(p_H,\mathcal{D})}[h(\pi_H; p_H)]$, where

$$h(\pi_H; p_H) := \begin{cases} \left(1 - \frac{K}{\pi_H v_H + (1-\pi_H)v_L}\right)\left(\frac{\pi_H v_H}{p_H} - \frac{(1-\pi_H)v_L}{1-p_H}\right) & \text{if } \pi_H \geq B \\ (1 - \frac{K}{\pi_H v_H})\frac{\pi_H v_H}{p_H} - \frac{1-\pi_H}{1-p_H}(1 - \beta)K & \text{if } B > \pi_H \geq A \\ 0 & \text{otherwise} \end{cases}$$

Notice:

$$\begin{aligned}
u_H(\mathcal{D};p_H) =& \mathbb{E}_{(p_H,\mathcal{D})}\Big[(1-\frac{K}{\pi_H v_H+(1-\pi_H)v_L})v_H \mathbb{1}\{\pi_H \geq B\} \\
& + (1-\frac{K}{\pi_H v_H})v_H \mathbb{1}\{B > \pi_H \geq A\}\Big| v = v_H\Big] \\
=& \mathbb{E}_{(p_H,\mathcal{D})}\Big[(1-\frac{K}{\pi_H v_H+(1-\pi_H)v_L})v_H \mathbb{1}\{\pi_H \geq B\}\mathbb{1}\{v = v_H\} \\
& + (1-\frac{K}{\pi_H v_H})v_H \mathbb{1}\{B > \pi_H \geq A\}\mathbb{1}\{v = v_H\}\Big]\Big/\mathbb{P}_{(p_H,\mathcal{D})}[v = v_H] \\
=& \frac{1}{p_H}\mathbb{E}_{(p_H,\mathcal{D})}\Bigg[\mathbb{E}_{(p_H,\mathcal{D})}\Big[(1-\frac{K}{\pi_H v_H+(1-\pi_H)v_L})v_H \mathbb{1}\{\pi_H \geq B\}\mathbb{1}\{v = v_H\} \\
& + (1-\frac{K}{\pi_H v_H})v_H \mathbb{1}\{B > \pi_H \geq A\}\mathbb{1}\{v = v_H\}\Big|\pi_H\Big]\Bigg] \\
=& \frac{1}{p_H}\mathbb{E}_{(p_H,\mathcal{D})}\Big[(1-\frac{K}{\pi_H v_H+(1-\pi_H)v_L})\pi_H v_H \mathbb{1}\{\pi_H \geq B\} \\
& + (1-\frac{K}{\pi_H v_H})\pi_H v_H \mathbb{1}\{B > \pi_H \geq A\}\Big]
\end{aligned}$$

The first equality is by definition; the second equality is by elementary definition of conditional expectation; the third equality is by Law of Iterated Expectation; the last equality is by the fact that $\mathbb{E}(\mathbb{1}\{v = v_H\}|\pi_H) = \mathbb{P}(v = v_H|\pi_H) = \pi_H$, since π_H itself is the posterior belief on $v = v_H$.

Similarly, we have:

$$\begin{aligned}
u_L(\mathcal{D};p_H) =& \mathbb{E}_{(p_H,\mathcal{D})}\Big[(1-\frac{K}{\pi_H v_H+(1-\pi_H)v_L})v_L \mathbb{1}\{\pi_H \geq B\} \\
& + (1-\beta)K\mathbb{1}\{B > \pi_H \geq A\}\Big| v = v_L\Big] \\
=& \mathbb{E}_{(p_H,\mathcal{D})}\Big[(1-\frac{K}{\pi_H v_H+(1-\pi_H)v_L})v_L \mathbb{1}\{\pi_H \geq B\}\mathbb{1}\{v = v_L\} \\
& + (1-\beta)K\mathbb{1}\{B > \pi_H \geq A\}\mathbb{1}\{v = v_L\}\Big]\Big/\mathbb{P}_{(p_H,\mathcal{D})}[v = v_L] \\
=& \frac{1}{1-p_H}\mathbb{E}_{(p_H,\mathcal{D})}\Bigg[\mathbb{E}_{(p_H,\mathcal{D})}\Big[(1-\frac{K}{\pi_H v_H+(1-\pi_H)v_L})v_L \mathbb{1}\{\pi_H \geq B\}\mathbb{1}\{v = v_L\} \\
& + (1-\beta)K\mathbb{1}\{B > \pi_H \geq A\}\mathbb{1}\{v = v_L\}\Big|\pi_H\Big]\Bigg] \\
=& \frac{1}{1-p_H}\mathbb{E}_{(p_H,\mathcal{D})}\Big[(1-\frac{K}{\pi_H v_H+(1-\pi_H)v_L})(1-\pi_H)v_L \mathbb{1}\{\pi_H \geq B\} \\
& + (1-\beta)K(1-\pi_H)\mathbb{1}\{B > \pi_H \geq A\}\Big]
\end{aligned}$$

In the last step, I use $\mathbb{E}(\mathbb{1}\{v = v_L\}|\pi_H) = \mathbb{P}(v = v_L|\pi_H) = 1 - \pi_H$.

We obtain the result by taking difference and rearranging terms according to $\pi_H \geq B$ or $B > \pi_H \geq A$.

Characterization for function $\rho(\cdot)$ in equation (3.17)

The following observation characterizes function $\rho(\cdot)$ as defined in equation (3.14).

Observation 3.C.1.

$$\rho(p_H) = \begin{cases} \rho_L(p_H) := y_1(B) - \frac{(1-B)p_H}{B(1-p_H)} y_2(B) & \text{if } p_H < B \\ \rho_R(p_H) := \frac{B(1-p_H)}{(1-B)p_H} y_1(B) - y_2(B) + \frac{p_H - B}{(1-B)p_H}(v_H - K) & \text{if } p_H \geq B \end{cases}$$

where $y_1(B) := (1 - \frac{K}{Bv_H + (1-B)v_L})v_H$ and $y_2(B) := (1 - \frac{K}{Bv_H+(1-B)v_L})v_L$.
(For later references, also notice $y_1(B) = (1-\beta)\frac{K}{v_L} v_H$ and $y_2(B) = (1-\beta)K$ by definition of B.)

Moreover, $\rho(\cdot)$ is continuous, decreasing on $[0, B]$ and increasing on $[B, 1]$, with $\rho(0) = y_1(B)$, $\rho(B) = y_1(B) - y_2(B)$ and $\rho(1) = v_H - K - y_2(B)$.

Proof. **Part 1:** $p_H < B$

Under $\mathcal{D}^u(p_H)$, all high projects are developed with $\alpha^* = \frac{K}{Bv_H+(1-B)v_L}$, so

$$u_H(\mathcal{D}^u(p_H); p_H) = (1 - \frac{K}{Bv_H + (1-B)v_L})v_H = y_1(B)$$

A low type entrepreneur gets positive project payoff if and only if she receives posterior $\pi_H = B$, in which case her project payoff is $(1 - \frac{K}{Bv_H+(1-B)v_L})v_L = y_2(B)$. We know the mass of low types getting this posterior is $(1-B)\mu(B) = (1-B)\frac{p_H}{B}$ and the total mass of low types is $1 - p_H$, so the probability for a low type to have posterior B is $\frac{(1-B)p_H}{B(1-p_H)}$. Therefore,

$$u_L(\mathcal{D}^u(p_H); p_H) = \frac{(1-B)p_H}{B(1-p_H)} y_2(B)$$

Taking difference, we obtain the expression of $\rho_L(\cdot)$. It is obvious to see the function is decreasing.

Part 2: $p_H \geq B$

In this case, $\mathcal{D}^u(p_H)$ induces posteriors in $\{B, 1\}$ by definition, so all low types get posterior B

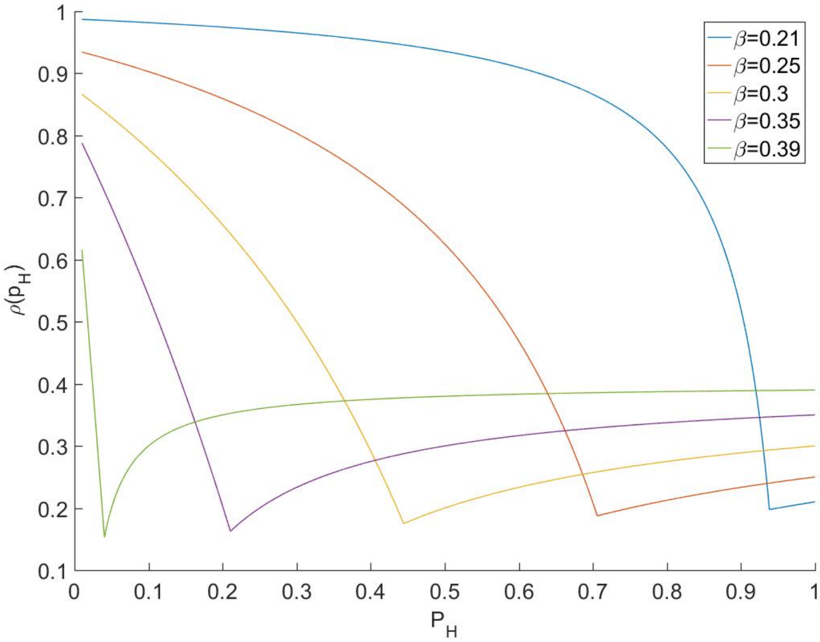

Figure 3.6: $\rho(\cdot)$ with different β ($v_H = 2$, $v_L = 1.6$ and $K = 1$).

under $\mathcal{D}^u(p_H)$ and thus

$$u_L(\mathcal{D}^u(p_H); p_H) = y_2(B)$$

The mass of high types getting posterior B is $\mu(B)B = \frac{(1-p_H)B}{1-B}$ (notice $\mu(B)B + 1 - \mu(B) = p_H$ by Bayesian feasibility) and the total mass of high types is p_H. When a high type gets posterior 1, she can get all the surplus from the project, which is $v_H - K$. Therefore,

$$u_H(\mathcal{D}^u(p_H); p_H) = \frac{B(1-p_H)}{(1-B)p_H} y_1(B) + \Big[1 - \frac{B(1-p_H)}{(1-B)p_H}\Big](v_H - K)$$

Taking difference, we obtain the expression of $\rho_R(\cdot)$.

Notice $y_1(B) = (1 - \frac{K}{Bv_H + (1-B)v_L})v_H \leq (v_H - K)$ and the weight $\frac{B(1-p_H)}{(1-B)p_H}$ is decreasing in p_H, so $u_H(\mathcal{D}^u(p_H); p_H)$ is increasing in p_H when $p_H \geq B$. Thus, $\rho_R(\cdot)$ is increasing.

Finally, the special values are straightforward to compute. For continuity, just notice the expressions of the two pieces of $\rho(\cdot)$ coincide at $p_H = B$. *Q.E.D.*

Figure (3.6) plots the $\rho(\cdot)$ function for different moral hazard parameters β. This "two wings" shape and how it shifts when β changes are important for some proofs in Section 4.4.

Proof for Lemma 3.4.1:

Proof. By conditions (1) and (2) in Assumption 3.4.1, $\Psi(0) = 0 < \rho(0)$ and $\Psi(1) > v_H - K > \rho(1)$. By the continuity of $\Psi(\cdot)$ and $\rho(\cdot)$, we know p_H^0 is in $(0, 1)$ with $\Psi(p_H^0) = \rho(p_H^0)$. Similarly, we have $p_H^1 \in (0, 1)$ and $\Psi(p_H^1) = v_H - K$. Because $\Psi(\cdot)$ is increasing and it is easy to see $\rho(\cdot) < v_H - K$ by Observation 3.C.1, we have $p_H^1 > p_H^0$.

For the range of potentially optimal p_H, first notice it is not possible to incentivize $p_H > p_H^1$ since the highest possible incentive of upgrading is $v_H - K$, which is achieved only under full disclosure. Thus the optimal $p_H \le p_H^1$.

To show any $p_H < p_H^0$ is suboptimal, notice $V_0(p_H) \ge V(p_H)$ for any p_H since the optimization problem with $V_0(\cdot)$ does not involve the incentive constraint (3.13). Therefore,

$$V(p_H) - \int_0^{p_H} \Psi(x)dx \le V_0(p_H) - \int_0^{p_H} \Psi(x)dx \quad \forall p_H$$

Denote the RHS as $\mathcal{V}_0(p_H)$. By Proposition 3.3.2, $V_0(\cdot)$ is continuous on $[0, 1]$ and differentiable on $(0, B) \cup (B, 1)$ and so is $\mathcal{V}_0(p_H)$. Taking derivative we have:

$$\mathcal{V}_0'(p_H) = \begin{cases} (v_H - K) + (v_L - K)\frac{1-B}{B} - \Psi(p_H) & \text{for } p_H \in (0, B) \\ (v_H - v_L) - \Psi(p_H) & \text{for } p_H \in (B, 1) \end{cases}$$

If $p_H^0 < B$, it is easy to see $\mathcal{V}_0'(p_H) > 0$ for all $p_H \le p_H^0$, since $\Psi(p_H) \le v_H - K$ for all $p_H \le p_H^1$ implies the first piece of $\mathcal{V}_0'(\cdot)$ is positive for $p_H \le p_H^0 < p_H^1$.

If $p_H^0 \ge B$, for any $p_H \le p_H^0$ we have $\Psi(p_H) \le \Psi(p_H^0) = \rho(p_H^0) \le \rho(1) = v_H - K - y_2(B) = v_H - (2 - \beta)K < v_H - v_L$, where the second inequality is for $\rho(\cdot)$ is increasing on $[B, 1]$ and the last inequality is by condition (3.3). Thus $\mathcal{V}_0'(\cdot) > 0$ on both $(0, B)$ and $(B, p_H^0]$.

Therefore, in either case we have $\mathcal{V}_0(\cdot)$ is strictly increasing on $[0, p_H^0]$. Thus,

$$\begin{aligned} V(p_H) - \int_0^{p_H} \Psi(x)dx &\le \mathcal{V}_0(p_H) \\ &< \mathcal{V}_0(p_H^0) \\ &= V(p_H^0) - \int_0^{p_H^0} \Psi(x)dx \quad \forall p_H \le p_H^0 \end{aligned}$$

where the last equality holds because $\rho(p_H^0) = \Psi(p_H^0)$ implies the incentive constraint (3.13) automatically holds under the unconstrained optimal rule $\mathcal{D}^u(p_H^0)$ and thus the optimal values are equal for the two problems: $V(p_H^0) = V_0(p_H^0)$.

In conclusion, any $p_H < p_H^0$ is suboptimal. *Q.E.D.*

Proof for Proposition 3.4.1

I first introduce the following theorem by Boleslavsky & Kim (2018) (adapted to my problem):

Theorem 3.C.1. *(Boleslavsky and Kim (2018) Proposition 3.3) A distribution μ over posteriors (with finite support) is a solution to the optimization problem (3.11) – (3.13) if and only if it satisfies the constraints (3.12) and (3.13), and there exist coefficients λ, a and b such that:*

$$\mathcal{L}(\pi_H; \lambda) := W(\pi_H) + \lambda h(\pi_H) \leq a + b\pi_H \quad \forall \pi_H \in [0, 1]$$

which holds as equality for all π_H with $\mu(\pi_H) > 0$.

This theorem provides a practical method to solve the constrained Bayesian persuasion problem. First, one plots the Lagrangian function $\mathcal{L}$ with some multiplier λ and its concavification. Second, one can find an affine function (a line) supporting the concavification function from above at $\pi_H = p_H$. Then, the points at which the affine function equals to $\mathcal{L}$ would be the candidate supporting points for the optimal posterior distribution. Finally, the optimal μ over these supporting points can be computed by making the two constraints hold.

One difficulty here is that to start the procedure, we need to properly guess a range for λ, which decides the shape of Lagrangian function. This can be tricky and in general may involve discussion for several cases. Moreover, it is not clear from the theorem whether we may get multiple solutions from multiple supporting lines. If yes, then we need to consider all possible λ values to find all the solutions. Fortunately, this is not the case. Specifically, I propose the following result:

Theorem 3.C.2. *Suppose there are two sets of coefficients (a_1, b_1, λ_1) and (a_2, b_2, λ_2) with which we find solutions μ_1 and μ_2 respectively satisfying the conditions in Theorem 3.C.1. Then, for any $\pi_H \in supp\{\mu_2\}$, we have $W(\pi_H) + \lambda_1 h(\pi_H) = a_1 + b_1\pi_H$. (Thus π_H can also be found as a candidate supporting point with the first set of coefficients (a_1, b_1, λ_1).)*

Proof. First, because both μ_1 and μ_2 are optimal solutions, we have

$$\mathbb{E}_{\mu_k}[W(\pi_H)] = V(p_H)$$
$$\mathbb{E}_{\mu_k}[h(\pi_H)] = \Psi(p_H)$$
$$\mathbb{E}_{\mu_k}[\pi_H] = p_H$$

for $k \in \{1, 2\}$.

Since μ_1 satisfies the conditions in Theorem 3.C.1 with (a_1, b_1, λ_1), all supporting points of μ_1 satisfy the inequality in Theorem 3.C.1 with equality and thus we have:

$$\mathbb{E}_{\mu_1}[W(\pi_H)] + \lambda\mathbb{E}_{\mu_1}[h(\pi_H)] = a_1 + b_1\mathbb{E}_{\mu_1}[\pi_H]$$

Therefore,

$$V(p_H) + \lambda_1\Psi(p_H) = a_1 + b_1 p_H$$

Suppose the conclusion does not hold, there exists $\pi_H^* \in supp\{\mu_2\}$ s.t. $W(\pi_H^*) + \lambda_1 h(\pi_H^*) < a_1 + b_1\pi_H^*$. Moreover, for any other $\pi_H \in supp\{\mu_2\}$, we have $W(\pi_H) + \lambda_1 h(\pi_H) \leq a_1 + b_1\pi_H$. Therefore,

$$\begin{aligned}
&\sum_{\pi_H \in supp\{\mu_2\}} \mu_2(\pi_H)[W(\pi_H) + \lambda_1 h(\pi_H)] < a_1 + b_1 \sum_{\pi_H \in supp\{\mu_2\}} \mu_2(\pi_H)\pi_H \\
\Rightarrow &\mathbb{E}_{\mu_2}[W(\pi_H)] + \lambda_1\mathbb{E}_{\mu_2}[h(\pi_H)] < a_1 + b_1\mathbb{E}_{\mu_2}[\pi_H] \\
\Rightarrow &V(p_H) + \lambda_1\Psi(p_H) < a_1 + b_1 p_H
\end{aligned}$$

which leads to a contradiction. *Q.E.D.*

This theorem implies that it suffices to find only one set of coefficients that satisfies the conditions in Theorem 3.C.1 to identify all possible supporting points of the optimal posterior distribution(s). Therefore, once a solution is found, there is no need to consider other possible values for λ, which can significantly simplify the discussion. This simple result can also be useful for other applications using constrained Bayesian persuasion. It is also straightforward to extend the theorem to higher dimensional cases with more types and incentive constraints.

Now, it is ready to prove Proposition 3.4.1.

Proof for Proposition 3.4.1. To ease notation, I suppress the dependence of $h(\cdot; p_H)$ on p_H.

Case 1: $A < B$ Define function $\mathcal{L}(\pi_H) := W(\pi_H) + \lambda h(\pi_H)$ and assume $\lambda > 0$. Then the shape of function $\mathcal{L}(\cdot)$ is as shown in Figure 3.7a (solid blue curve), although we do not know whether the three points corresponding to $\pi_H = 0$, B and 1 are necessarily on the same line for now. ($\mathcal{L}(B)$ may even be negative for some λ.) Two important features of $\mathcal{L}(\cdot)$ captured in the figure are proved in the following lemma:

Lemma. (i) The graph of $\mathcal{L}(\cdot)$ on $[A, B)$ is below the segment between $(0, \mathcal{L}(0))$ and $(B, \mathcal{L}(B))$; (ii) $\mathcal{L}(\cdot)$ is strictly convex on $[B, 1]$.

Subproof. For part (i), notice the graph of $W(\cdot)$ on $[A, B)$ is below the segment between $(0, W(0))$ and $(B, W(B))$ by the proof of Proposition 3.3.2. Since $\lambda > 0$, it then suffices to show the same holds for function $h(\cdot)$.

By the definition of $h(\cdot)$, when $B > \pi_H \geq A$ we have

$$h(\pi_H) = (1 - \frac{K}{\pi_H v_H})\frac{\pi_H v_H}{p_H} - \frac{1-\pi_H}{1-p_H}(1-\beta)K = \frac{\pi_H v_H}{p_H} - \frac{K}{p_H} - \frac{1-\pi_H}{1-p_H}(1-\beta)K$$

Notice this piece of $h(\cdot)$ is linear in π_H, so if we define the last expression on the RHS as $\hat{h}(\pi_H)$ for all $\pi_H \in [0, B]$, it then suffices to show $\hat{h}(B) < h(B)$ and $\hat{h}(0) < h(0)$.

It is easy to see:

$$h(B) = \frac{B}{p_H} y_1(B) - \frac{1-B}{1-p_H} y_2(B) = \frac{B}{p_H} y_1(B) - \frac{1-B}{1-p_H}(1-\beta)K$$

(See Observation 3.C.1 for the expressions of $y_1(B)$ and $y_2(B)$, and the fact that $y_2(B) = (1-\beta)K$.) Thus we have:

$$h(B) - \hat{h}(B) = \frac{B}{p_H} y_1(B) - \frac{B}{p_H}(1 - \frac{K}{Bv_H})v_H > 0$$

since $y_1(B) = \big(1 - \frac{K}{Bv_H + (1-B)v_L}\big)v_H > (1 - \frac{K}{Bv_H})v_H$. Moreover, $\hat{h}(0) = -\frac{K}{p_H} - \frac{1}{1-p_H}(1-\beta)K < 0 = h(0)$. This completes the proof for part (i).

For part (ii), notice $W(\cdot)$ is affine on $[B, 1]$, so it suffices to show $h(\cdot)$ is convex on the range.

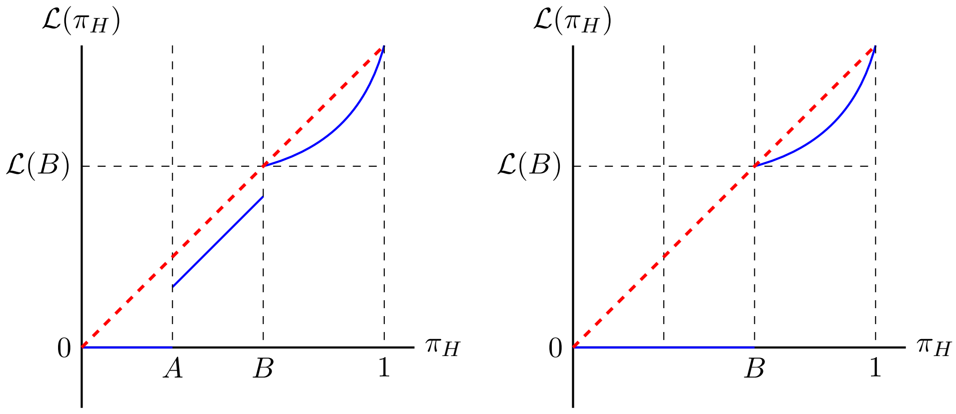

(a) Case 1: $B > A$

(b) Case 2: $B \leq A$

Figure 3.7: Lagrangian concavification for the constrained problem

One can check:

$$h''(\pi_H) = \frac{2Kv_Hv_L(v_H - v_L)}{(1-p_H)p_H[v_L + (v_H - v_L)\pi_H]^3} > 0 \quad \forall \pi_H > B$$

Thus $h(\cdot)$ is strictly convex on $[B, 1]$. $\square$

The lemma together with the facts that $\mathcal{L}(1) > 0$ and $\mathcal{L} = 0$ on $[0, A)$ implies the candidate supporting points for the optimal μ are $\{0, B, 1\}$.

Now, for any optimal μ, by the two constraints we must have: $\mu(B)B + \mu(1) = p_H$; $\mu(B)h(B) + \mu(1)h(1) = \Psi(p_H)$. Solving the system gives:

$$\mu(B) = \frac{h(1)p_H - \Psi(p_H)}{Bh(1) - h(B)} = \frac{v_H - K - \Psi(p_H)}{[v_H - K - \rho_L(p_H)]\frac{B}{p_H}}$$
$$\mu(1) = \frac{B\Psi(p_H) - p_Hh(B)}{Bh(1) - h(B)} = \frac{B[\Psi(p_H) - \rho_L(p_H)]}{[v_H - K - \rho_L(p_H)]\frac{B}{p_H}}$$

and

$$\mu(0) = 1 - \mu(B) - \mu(1) = \frac{(1-B)[\Psi(p_H) - \rho_R(p_H)]}{[v_H - K - \rho_L(p_H)]\frac{B}{p_H}}$$

(To check the final expressions, notice $h(1) = \frac{v_H - K}{p_H}$ and $h(B) = \frac{B}{p_H}y_1(B) - \frac{1-B}{1-p_H}y_2(B)$.)

To verify this is indeed a solution, we need to check all three probabilities are positive. First, notice by the results in Observation 3.C.1, $\rho_L(p_H) \leq \rho_L(1) < v_H - K$, so the denominators are all positive. Moreover, $p_H \leq p_H^1$ implies $\Psi(p_H) \leq v_H - K$, so we have $\mu(B) \geq 0$.

To check $\mu(0)$ and $\mu(1)$ are positive, consider two scenarios:

- Scenario (i): $p_H \leq B$

  Since $p_H \geq p_H^0$, we have $\Psi(p_H) \geq \rho(p_H) = \rho_L(p_H)$ and thus $\mu(1) \geq 0$.

  Now, since $\mu(B) \geq 0$, $\mu(1) \geq 0$ and the probabilities sum up to 1, by Bayesian feasibility condition we must have $\mu(0) \geq 0$ as $p_H \leq B$.

- Scenario (ii): $p_H > B$

  Since $p_H \geq p_H^0$, we have $\Psi(p_H) \geq \rho(p_H) = \rho_R(p_H)$ and thus $\mu(0) \geq 0$.

  Now, since $\mu(0) \geq 0$, $\mu(B) \geq 0$ and the probabilities sum up to 1, by Bayesian feasibility condition we must have $\mu(1) \geq 0$ as $p_H > B$.

Finally, what is remaining to show is that there is indeed a $\lambda > 0$ that makes the three points on the graph of $\mathcal{L}(\cdot)$ corresponding to $\pi_H = 0, B$ and 1 on the same line. Those points are on the same line if and only if:

$$\frac{W(B) + \lambda h(B)}{B} = \frac{W(1) + \lambda h(1)}{1} \iff \lambda = \frac{W(B) - BW(1)}{Bh(1) - h(B)}$$

The denominator is equal to $[v_H - K - \rho_L(p_H)]\frac{B}{p_H} > 0$ as shown above. For the numerator, notice $W(B) - BW(1) = (v_L - K)(1 - B) > 0$. Thus such λ is indeed positive.

By Theorem 3.C.1, we can conclude the μ we found is indeed a solution. The concavification for $\mathcal{L}(\cdot)$ is depicted by the red dashed line in Figure 3.7a. By Theorem 3.C.2, the set $\{0, B, 1\}$ includes all possible supporting points in any optimal solution and thus the solution given above is the unique solution. (Thus we do not need to consider $\lambda \leq 0$.)

Case 2: $A \geq B$ The proof for this case is exactly the same as that in Case 1, except that we do not need to consider the interval $[A, B)$ now. Thus exactly the same solution is obtained. The concavification graph is as shown in Figure 3.7b.

To derive the value function $V(\cdot)$, notice under the optimal μ, all high types are developed and the low types receiving posterior belief B are developed. Therefore, the value function is

$$V(p_H) = (v_H - K)p_H + (v_L - K)(1 - B)\mu(B)$$

where $(1 - B)\mu(B)$ is the measure of low types who receive posterior B. *Q.E.D.*

Proof for Corollary 3.4.1

Proof. By Proposition 3.4.1, any optimal posterior belief distribution μ supports on a subset of $\{0, B, 1\}$. Thus $(p_H, \mathcal{D})$ is optimal if and only if the induced posterior distribution μ supports on a subset of $\{0, B, 1\}$ and $(p_H, \mu(0), \mu(B), \mu(1))$ maximizes expected net social surplus while satisfying Bayesian feasibility and incentive constraint for p_H.

Notice given $(p_H, \mu(B))$, $\mu(0)$ and $\mu(1)$ are pinned down by Bayesian feasibility and probabilities summing up to one as $\mu(0) = 1 - p_H - (1 - B)\mu(B)$ and $\mu(1) = p_H - B\mu(B)$. Thus it suffices to consider the control variables as $(p_H, \mu(B))$.

It is easy to see that the first three constraints in optimization (3.19) are equivalent to that μ (with $\mu(B) = \mu_B$) is a valid probability distribution and is Bayesian feasible given p_H; the last constraint is the incentive constraint (3.13) evaluated with μ supporting on (a subset of) $\{0, B, 1\}$, which then guarantees p_H is incentivized given $\mu(B) = \mu_B$. For the objective function, notice projects are developed with posterior believes above B, so all high types and mass $\mu_B(1-B)$ of low types would be developed. Thus, the objective function indeed measures expected net social surplus.

Overall, optimization (3.19) captures the designer's problem after restricting the posterior distribution support according to Proposition 3.4.1. *Q.E.D.*

Proof for Proposition 3.4.3

Proof. Suppose p_H^1 is optimal. We have FOC for optimization (3.18):

$$V'(p_H^1) - \Psi(p_H^1) \geq 0$$

By the expression of $V(\cdot)$, we have:

$$V'(p_H^1) = (v_H - K) + (v_L - K)\frac{1-B}{B}R(p_H^1) + (v_L - K)\frac{1-B}{B}p_H^1 R'(p_H^1)$$

where $R(p_H) := \frac{v_H - K - \Psi(p_H)}{v_H - K - \rho_L(p_H)}$. Since $\Psi(p_H^1) = v_H - K$ and thus $R(p_H^1) = 0$, the FOC becomes:

$$(v_L - K)\frac{1-B}{B}p_H^1 R'(p_H^1) \geq 0 \Leftrightarrow R'(p_H^1) \geq 0$$

However, it is easy to see

$$R'(p_H^1) = \frac{-\Psi'(p_H^1)}{v_H - K - \rho_L(p_H^1)} < 0$$

since $\rho_L(\cdot) < v_H - K$ and $\Psi'(p_H^1) > 0$ by assumption. This gives a contradiction. Thus optimal p_H is less than p_H^1.

Notice under the condition $\Psi'(p_H^1) > 0$, p_H^1 is the unique p_H that satisfies $\Psi(p_H^1) = v_H - K$. As the optimal $p_H < p_H^1$, the upgrading incentive required by any optimal p_H is strictly smaller than $v_H - K$, so full disclosure is not optimal. *Q.E.D.*

Proof for Proposition 3.4.4

The proof uses the characterization of the optimal p_H with optimization (3.18), as stated in Proposition 3.4.2. Before going to the main proof, it is useful to provide an equivalent expression for the value function $V(\cdot)$ in Proposition 3.4.1. Specifically, it is easy to check:

$$V(p_H) = \begin{cases} (v_H - K)p_H + (v_L - K)\frac{(1-B)p_H}{B}Q(p_H) & \text{if } p_H^0 \le p_H < B \\ (v_H - K)p_H + (v_L - K)(1 - p_H)Q(p_H) & \text{if } B \le p_H \le p_H^1 \end{cases} \tag{3.43}$$

where $Q(p_H) = \frac{v_H - K - \Psi(p_H)}{v_H - K - \rho(p_H)}$. Notice when $B > p_H^1$, only the first case matters and when $B \le p_H^0$, only the second case matters.

By Observation 3.C.1, we know $\rho(\cdot)$ is nonnegative and is decreasing on $[0, B]$ and increasing on $[B, 1]$, with $\rho(0) = y_1(B) < v_H - K$ and $\rho(1) = v_H - K - y_2(B) < v_H - v_L$. Moreover, function $\rho(\cdot)$ varies with β in a particular pattern as illustrated in Figure 3.6. These facts are frequently used below. In particular, Lemma (i) and Lemma (ii) below just formalize some patterns in Figure 3.6. Readers convinced by the figure may skip their proofs.

Proof for Proposition 3.4.4. Notice by definition, B is strictly decreasing in β with $\lim_{\beta \downarrow \underline{\beta}} B(\beta) = 1$ and $\lim_{\beta \uparrow \bar{\beta}} B(\beta) = 0$, so $\beta \to \bar{\beta}$ is equivalent to $B \to 0$ and $\beta \to \underline{\beta}$ is equivalent to $B \to 1$. To ease notations, however, I suppress the dependences of B, $\rho(\cdot)$ and $V(\cdot)$ on β, but one should keep in mind that B changes in β as specified above.

Part (a): Since $\Psi(0) = 0$, $\Psi(p_H^1) = v_H - K$ and $\Psi(\cdot)$ is continuous, there exists $\hat{p} \in (0, p_H^1)$ s.t. $\Psi(\hat{p}) = \frac{v_H - v_L}{2}$. Since we are thinking about β being large enough and correspondingly B being close to 0, it is w.l.g. to assume $B < \hat{p}$.

First, we need a lemma:

Lemma (i). When β is large enough, $p_H^0(\beta) > \hat{p}$. Moreover, as $\beta \to \bar{\beta}$: (1) $\rho(\cdot)$ uniformly converges to $v_H - v_L$ on $[\hat{p}, p_H^1]$; (2) $\rho'(\cdot) \geq 0$ and uniformly converges to 0 on $[\hat{p}, p_H^1]$.

Subproof. For any $p_H \geq \hat{p}$, we have:

$$\rho(p_H) = \frac{B(1-p_H)}{(1-B)p_H} y_1(B) - y_2(B) + \frac{p_H - B}{(1-B)p_H}(v_H - K)$$

As $\beta \to \bar{\beta}$, we have $B \to 0$, $y_1(B) \to (1 - \frac{K}{v_L})v_H$ and $y_2(B) \to v_L - K$, so $\rho(p_H) \to v_H - v_L$ as $\beta \to \bar{\beta}$ for all $p_H \in [\hat{p}, p_H^1]$. Since $\Psi(\hat{p}) = \frac{v_H - v_L}{2} < v_H - v_L$, when β is large enough we have $\Psi(\hat{p}) < \rho(\hat{p})$ and therefore $p_H^0(\beta) > \hat{p}$. Moreover, because $\rho(\cdot)$ is increasing but always less than $v_H - v_L$ when $p_H > B$, its uniform convergence on $[\hat{p}, p_H^1]$ is implied by the convergence at point $p_H = \hat{p}$.

For the uniform convergence of $\rho'(\cdot)$ on $[\hat{p}, p_H^1]$, notice when $p_H > B$

$$\rho'(p_H) = \frac{B}{(1-B)p_H^2}(v_H - K - y_1(B))$$

which is decreasing and pointwisely converges to 0 from above as $B \to 0$. Thus the uniform convergence is implied by the convergence at point $p_H = \hat{p}$. □

By the lemma, we assume β is large enough such that $p_H^0(\beta) > \hat{p} > B$. Moreover, from the uniform convergences of $\rho(\cdot)$ and $\rho'(\cdot)$, we have:

$$\lim_{\beta \uparrow \bar{\beta}} \rho(p_H^0(\beta)) = v_H - v_L \tag{3.44}$$

$$\lim_{\beta \uparrow \bar{\beta}} \left(\sup_{p_H \in [\hat{p}, p_H^1]} \{\rho'(p_H)\} \right) = 0 \tag{3.45}$$

To show statement (a), it suffices to show when β is large enough, $V'(p_H) - \Psi(p_H) < 0$, $\forall p_H > p_H^0(\beta)$, so that the objective in optimization (3.18) is strictly decreasing on the feasible set, which implies $p_H^0(\beta)$ is uniquely optimal. Since $B < p_H^0(\beta)$ when β is large enough, this then implies that the optimal μ supports on $\{B(\beta), 1\}$ by the discussion right below Proposition 3.4.1.

From expression (3.43), we have for $p_H \geq p_H^0(\beta)$:

$$V'(p_H) - \Psi(p_H) = \underbrace{(v_H - K) - (v_L - K)Q(p_H) - \Psi(p_H)}_{\text{term 1}} + \underbrace{(v_L - K)(1 - p_H)Q'(p_H)}_{\text{term 2}}$$

(recall we assume $p_H^0(\beta) > \hat{p} > B$ w.l.g.)

For term 1, notice for any $p_H \geq p_H^0(\beta)$, we have:

$$\begin{aligned}
&(v_H - K) - (v_L - K)Q(p_H) - \Psi(p_H) \\
=&(v_H - K) - (v_L - K)\frac{v_H - K - \Psi(p_H)}{v_H - K - \rho(p_H)} - \Psi(p_H) \\
=&(v_H - K) - \frac{[v_H - v_L - \rho(p_H)]\Psi(p_H) + (v_L - K)(v_H - K)}{v_H - K - \rho(p_H)} \\
\leq&(v_H - K) - \frac{(v_L - K)(v_H - K)}{v_H - K - \rho(p_H)} \\
\leq&(v_H - K) - \frac{(v_L - K)(v_H - K)}{v_H - K - \rho(p_H^0(\beta))}
\end{aligned}$$

where the second last inequality is because $\rho(p_H) < v_H - v_L$ for all $p_H > B$ and the last inequality is because $\rho(\cdot)$ is increasing when $p_H > B$. Notice by (3.44), the last line converges to 0 as $\beta \to \bar{\beta}$. Thus,

$$\limsup_{\beta \uparrow \bar{\beta}} \sup_{p_H \in [p_H^0(\beta), p_H^1]} \{(v_H - K) - (v_L - K)Q(p_H) - \Psi(p_H)\} \leq 0 \tag{3.46}$$

For term 2, notice for any $p_H \in [p_H^0(\beta), p_H^1]$, we have

$$\begin{aligned}
Q'(p_H) &= -\frac{\Psi'(p_H)}{v_H - K - \rho(p_H)} + \frac{[v_H - K - \Psi(p_H)]\rho'(p_H)}{[v_H - K - \rho(p_H)]^2} \\
&\leq -\frac{\Psi'(p_H)}{v_H - K - \rho(p_H^0(\beta))} + \frac{[v_H - K - \Psi(p_H)]\rho'(p_H)}{[v_H - K - \rho(p_H)]^2} \\
&\leq -\frac{\delta}{v_H - K - \rho(p_H^0(\beta))} + \frac{v_H - K}{(v_L - K)^2} \sup_{p_H \in [\hat{p}, p_H^1]} \{\rho'(p_H)\}
\end{aligned}$$

The first inequality is because $\rho(\cdot)$ is increasing on $[B, 1]$. The second inequality is because: $\Psi'(\cdot) \leq \delta$ by assumption; $\Psi(\cdot) \geq 0$; and $\rho(p_H) < v_H - v_L$ when $p_H > B$ for any β. By (3.44) and (3.45), the last line converges to $-\frac{\delta}{v_L-K}<0$ as $\beta\to\bar{\beta}$. Thus,

$$\limsup_{\beta\uparrow\bar{\beta}}\sup_{p_H\in[p_H^0(\beta),p_H^1]}\{(v_L-K)(1-p_H)Q'(p_H)\}\le-\delta(1-p_H^1)\tag{3.47}$$

Combining (3.46) and (3.47), we have

$$\limsup_{\beta\uparrow\bar{\beta}}\sup_{p_H\in[p_H^0(\beta),p_H^1]}[V'(p_H)-\Psi(p_H)]\le-\delta(1-p_H^1)<0$$

Thus there exists $a>0$ such that

$$\sup_{p_H\in[p_H^0(\beta),p_H^1]}[V'(p_H)-\Psi(p_H)]\le-\frac{\delta(1-p_H^1)}{2},\ \forall\beta\in(\bar{\beta}-a,\bar{\beta})$$

Therefore, $V'(p_H)-\Psi(p_H)<0\,\forall p_H>p_H^0(\beta)$ and thus $p_H^0(\beta)$ is (uniquely) optimal when $\beta>\bar{\beta}-a$.

Part (b): Since $\Psi(0)=0$, $\Psi(p_H^1)=v_H-K$ and $\Psi(\cdot)$ is continuous, there exists $\hat{p}\in(0,p_H^1)$ s.t. $\Psi(\hat{p})=\frac{v_H-K}{2}$. Since we are thinking about β being small enough and thus B being close to 1, it is w.l.g. to assume $B>p_H^1$. First, we need two lemmas:

Lemma (ii). When β is small enough, $p_H^0(\beta)>\hat{p}$ and $\rho(p_H)$ uniformly converges to v_H-K on $[0,p_H^1]$ as $\beta\to\underline{\beta}$.

Subproof. When $p_H\le p_H^1<B$, we have:

$$\rho(p_H)=y_1(B)-\frac{(1-B)p_H}{B(1-p_H)}y_2(B)$$

As $\beta\to\underline{\beta}$, we have $B\to 1$ and $y_1(B)\to v_H-K$, so $\rho(p_H)\to v_H-K$ pointwisely on $[0,p_H^1]$. Notice $\rho(\cdot)$ is decreasing and is always smaller than v_H-K when $p_H<B$, so the uniform convergence of $\rho(\cdot)$ is implied by convergence at point $p_H=p_H^1$.

Moreover, since $\Psi(\hat{p})=(v_H-K)/2<v_H-K$, we have $\Psi(\hat{p})<\rho(\hat{p})$ and thus $p_H^0(\beta)>\hat{p}$ when β is small enough. □

By the lemma, we assume β is small enough such that $B>p_H^1>p_H^0(\beta)>\hat{p}$ without loss of generality. Moreover, the lemma implies:

$$\lim_{\beta \downarrow \underline{\beta}} \Psi(p_H^0(\beta)) = \lim_{\beta \downarrow \underline{\beta}} \rho(p_H^0(\beta)) = v_H - K \tag{3.48}$$

Another lemma we need is:

Lemma (iii). There exists a constant $c > 0$ s.t.

$$c \leq \liminf_{\beta \downarrow \underline{\beta}} \inf_{p_H \in [p_H^0(\beta), p_H^1]} [\frac{1-B}{v_H - K - \rho(p_H)}]$$

Subproof. Since $\rho(\cdot)$ is decreasing on $[0, B]$, we have $\frac{1-B}{v_H - K - \rho(p_H^1)} \leq \frac{1-B}{v_H - K - \rho(p_H)}$ for $p_H \leq p_H^1$ $(< B)$. Thus,

$$\inf_{p_H \in [p_H^0(\beta), p_H^1]} [\frac{1-B}{v_H - K - \rho(p_H)}] \geq \frac{1-B}{v_H - K - \rho(p_H^1)} = \frac{1}{\frac{v_H - K - y_1(B)}{1-B} + \frac{p_H^1}{B(1-p_H^1)} y_2(B)}$$

Notice

$$\lim_{\beta \downarrow \underline{\beta}} \frac{v_H - K - y_1(B)}{1-B} = \lim_{B \to 1} \frac{K(v_H - v_L)v_H}{(Bv_H + (1-B)v_L)^2} = \frac{K(v_H - v_L)}{v_H}$$

where I used the L'Hospital's rule. Moreover, as $\beta \to \underline{\beta}$, $B \to 1$ and $y_2(B) \to (1 - \frac{K}{v_H})v_L$. Thus,

$$0 < \frac{1}{\frac{K(v_H - v_L)}{v_H} + \frac{p_H^1}{(1-p_H^1)}(1 - \frac{K}{v_H})v_L} \leq \liminf_{\beta \downarrow \underline{\beta}} \inf_{p_H \in [p_H^0(\beta), p_H^1]} [\frac{1-B}{v_H - K - \rho(p_H)}]$$

□

To show statement (b), it suffices to show when β is small enough, $V'(p_H) - \Psi(p_H) < 0, \forall p_H > p_H^0(\beta)$, which implies $p_H^0(\beta)$ is uniquely optimal. Since $B > p_H^0(\beta)$ when β is small enough, the optimal μ then supports on $\{0, B(\beta)\}$ by the discussion right below Proposition 3.4.1.

Recall that w.l.g., we assume β is small enough such that $B > p_H^1 > p_H^0(\beta) > \hat{p}$. By (3.43), we have for $p_H \in [p_H^0(\beta), p_H^1]$:

$$\begin{aligned}
& V'(p_H) - \Psi(p_H) \\
=& v_H - K - \Psi(p_H) + (v_L - K)\frac{1-B}{B}Q(p_H) + (v_L - K)\frac{1-B}{B}p_H Q'(p_H) \\
=& v_H - K - \Psi(p_H) + (v_L - K)\frac{1-B}{B}Q(p_H) - (v_L - K)\frac{(1-B)p_H \Psi'(p_H)}{B[v_H - K - \rho(p_H)]} \\
& + \frac{(1-B)[v_H - K - \Psi(p_H)]\rho'(p_H)}{B[v_H - K - \rho(p_H)]^2}(v_L - K)p_H \\
\le& v_H - K - \Psi(p_H^0(\beta)) + (v_L - K)\frac{1-B}{B}Q(p_H) - (v_L - K)\frac{(1-B)p_H \Psi'(p_H)}{B[v_H - K - \rho(p_H)]} \\
\le& v_H - K - \Psi(p_H^0(\beta)) + (v_L - K)\frac{1-B}{B} - (v_L - K)\frac{(1-B)p_H \Psi'(p_H)}{B[v_H - K - \rho(p_H)]}
\end{aligned}$$

The 3rd line is by computing $Q'(\cdot)$. The 4th line is because $\rho'(p_H) < 0$ when $p_H < B$ and $\Psi(\cdot)$ is increasing. The 5th line is because $Q(\cdot)$ is decreasing on $[0, B]$ and $Q(p_H^0) = 1$ since $\rho(p_H^0) = \Psi(p_H^0)$.

Notice for any $p_H \in [p_H^0(\beta), p_H^1]$, the last term

$$\begin{aligned}
-(v_L - K)\frac{(1-B)p_H \Psi'(p_H)}{B[v_H - K - \rho(p_H)]} &\le -(v_L - K)\frac{(1-B)p_H \delta}{B[v_H - K - \rho(p_H)]} \\
&\le -\delta(v_L - K)\hat{p} \inf_{p_H \in [p_H^0(\beta), p_H^1]} \frac{1-B}{v_H - K - \rho(p_H)}
\end{aligned}$$

where I used $\Psi'(\cdot) \ge \delta$ (be assumption), $p_H^0(\beta) > \hat{p}$ and $B < 1$. Thus,

$$\begin{aligned}
\sup_{p_H \in [p_H^0(\beta), p_H^1]} [V'(p_H) - \Psi(p_H)] \le & v_H - K - \Psi(p_H^0(\beta)) \\
& + (v_L - K)\frac{1-B}{B} \\
& - \delta(v_L - K)\hat{p} \inf_{p_H \in [p_H^0(\beta), p_H^1]} \left[\frac{1-B}{v_H - K - \rho(p_H)}\right]
\end{aligned}$$

When $\beta \to \underline{\beta}$, the first line of the RHS goes to 0 by (3.48); the second line goes to 0 as $B \to 1$. Therefore,

$$\begin{aligned}
& \limsup_{\beta \downarrow \underline{\beta}} \sup_{p_H \in [p_H^0(\beta), p_H^1]} [V'(p_H) - \Psi(p_H)] \\
\le & -\delta(v_L - K)\hat{p} \liminf_{\beta \downarrow \underline{\beta}} \inf_{p_H \in [p_H^0(\beta), p_H^1]} \left[\frac{1-B}{v_H - K - \rho(p_H)}\right] \le -c\delta(v_L - K)\hat{p} < 0
\end{aligned}$$

where c is as found in Lemma (iii). Thus there exists $b > 0$ such that

$$\sup_{p_H \in [p_H^0(\beta), p_H^1]} [V'(p_H) - \Psi(p_H)] \leq -\frac{c\delta(v_L - K)\hat{p}}{2} < 0, \ \forall \beta \in (\underline{\beta}, \underline{\beta} + b)$$

and thus $p_H^0(\beta)$ is the (unique) optimal solution when $\beta \in (\underline{\beta}, \underline{\beta} + b)$. *Q.E.D.*

Proof for Corollary 3.4.2

First, the following observation is useful both here and later on. Although the proof is somewhat tedious, its intuition fully comes from the fact that as β increases, the "left wing" of function $\rho(\cdot, \beta)$ shifts downwards and the "right wing" of it shifts upwards. (See Figure 3.6).

Lemma 3.C.1. *If β_0 satisfies $p_H^0(\beta_0) \geq B(\beta_0)$, then $p_H^0(\cdot)$ is strictly increasing on $[\beta_0, \bar{\beta})$; If β_0 satisfies $p_H^0(\beta_0) \leq B(\beta_0)$, then $p_H^0(\cdot)$ is strictly decreasing on $(\underline{\beta}, \beta_0]$.*

Proof. For the first part, pick any β_0 such that $p_H^0(\beta_0) \geq B(\beta_0)$ and pick any $\beta_2 > \beta_1 > \beta_0$. Since $B(\beta)$ decreases in β, $B(\beta_2) < B(\beta_1) < B(\beta_0) \leq p_H^0(\beta_0)$. Thus as long as we can show $p_H^0(\beta_1) > p_H^0(\beta_0)$, the fact that $p_H^0(\beta_1) > B(\beta_1)$ further implies $p_H^0(\beta_2) > p_H^0(\beta_1)$ in the same way. Therefore, it suffices to show $p_H^0(\beta_1) > p_H^0(\beta_0)$.

It is easy to check $p_H^0(\beta_0) \geq B(\beta_0) \Rightarrow \rho(p_H^0(\beta_0), \beta_1) > \rho(p_H^0(\beta_0); \beta_0)$. (Intuitively, the "right wing" of function $\rho(\cdot; \beta)$ shifts up when β increases, as shown in Figure 3.6.) Since $\rho(p_H^0(\beta_0); \beta_0) = \Psi(p_H^0(\beta_0))$, we know $\rho(p_H^0(\beta_0), \beta_1) > \Psi(p_H^0(\beta_0))$ and thus $p_H^0(\beta_1) > p_H^0(\beta_0)$ by definition of p_H^0.

For the second part, pick any β_0 such that $p_H^0(\beta_0) \leq B(\beta_0)$ and pick any $\beta_2 < \beta_1 < \beta_0$. Then, $B(\beta_2) > B(\beta_1) > B(\beta_0) \geq p_H^0(\beta_0)$. We can show $p_H^0(\beta_1) \leq B(\beta_1)$. Supposing not, then $p_H^0(\beta_1) > B(\beta_1) > B(\beta_0)$ and we have $\Psi(p_H^0(\beta_1)) = \rho(p_H^0(\beta_1), \beta_1) < \rho(p_H^0(\beta_1); \beta_0)$ (since the "right wing" of function $\rho(\cdot; \beta)$ shifts up when β increases). This implies $p_H^0(\beta_0) > p_H^0(\beta_1) > B(\beta_1)$ – contradiction!

Now, as $p_H^0(\beta_0) \leq B(\beta_0) \Rightarrow p_H^0(\beta_1) \leq B(\beta_1)$, it suffices to show $p_H^0(\beta_1) > p_H^0(\beta_0)$. Then, $p_H^0(\beta_2) > p_H^0(\beta_1)$ follows in the same way.

It is easy to check $\rho(p_H^0(\beta_0); \beta_0) < \rho(p_H^0(\beta_0); \beta_1)$. (Intuitively, the "left wing" of function $\rho(\cdot)$ shifts down when β increases.) Therefore, $\rho(p_H^0(\beta_0), \beta_1) > \Psi(p_H^0(\beta_0))$ and thus $p_H^0(\beta_1) > p_H^0(\beta_0)$. *Q.E.D.*

Proof for Corollary 3.4.2. By Proposition 3.4.4, the optimal $p_H = p_H^0(\beta)$ when β is large enough or small enough. Also notice when β is small enough, $B(\beta)$ is close to 1 and thus larger than $p_H^0(\beta)$; when β is large enough, $B(\beta)$ is close to 0 and thus smaller than $p_H^0(\beta)$, which is bounded away from zero (see Lemma (i) in the proof for Proposition 3.4.4). Thus the result follows from Lemma 3.C.1. *Q.E.D.*

Proof for Proposition 3.4.5

Before proving the proposition, I first introduce an important lemma that is useful both here and later. Let V^* denote the optimal value of the designer's problem (solved by optimization (3.18) or (3.19)). We have:

Lemma 3.C.2. *V^* is continuous in the moral hazard parameter β. Moreover, the set of optimal $(p_H, \mu(B))$ is compact and upper hemi-continuous in β.*

Proof. We use the characterization of (3.19). Denote the objective function in (3.19) as $\Omega(p_H, \mu_B; B)$. Since there is a homeomorphism between β and B, it is equivalent to consider the continuous dependences on B for $B \in (0, 1)$.

First, the following lemma shows we can relax the incentive constraint of optimization (3.19) in one direction.

Lemma. Optimization (3.19) is equivalent to the following optimization:

$$V^* = \max_{p_H, \mu_B} \{(v_H - K)p_H + (v_L - K)(1 - B)\mu_B - \int_0^{p_H} \Psi(x)dx\} \tag{3.49}$$

$$\text{s.t. } 1 - p_H - (1 - B)\mu_B \geq 0 \tag{3.50}$$

$$\mu_B \geq 0 \tag{3.51}$$

$$p_H - B\mu_B \geq 0 \tag{3.52}$$

$$\left[1 - \frac{K}{Bv_H + (1 - B)v_L}\right]\left[\frac{Bv_H}{p_H} - \frac{(1 - B)v_L}{1 - p_H}\right]\mu_B + \frac{v_H - K}{p_H}(p_H - B\mu_B) \geq \Psi(p_H) \tag{3.53}$$

Subproof. It suffices to show (3.53) is always binding in optimum. Supposing not, let (p_H^*, μ_B^*) be a solution with (3.53) being slack. Notice $p_H^* \leq p_H^1 < 1$ by Lemma 3.4.1. Consider two cases:

Case 1: (3.50) is slack at (p_H^*, μ_B^*)

Notice the LHS of (3.53) measures an entrepreneur's upgrading incentive, which is easily shown to be $\leq v_H - K$, so (3.53) being slack at (p_H^*, μ_B^*) implies $\frac{\partial \Omega}{\partial p_H}|_{(p_H^*, \mu_B^*)} = v_H - K - \Psi(p_H^*) > 0$. Since the functions in all of the constraints are continuous, we can increase p_H a little bit (while keeping μ_B constant) without violating any constraint. This improves the objective function's value and thus leads to a contradiction.

Case 2: (3.50) is binding at (p_H^*, μ_B^*)

We can increase p_H and decrease μ_B in a particular way such that (3.50) keeps binding. Specifically, let $\hat{\mu}_B(p_H) := \frac{1-p_H}{1-B}$. Then, initially we are at $(p_H^*, \hat{\mu}_B(p_H^*))$, where all constraints are satisfied. Notice: (1) By design, $(p_H, \hat{\mu}_B(p_H))$ satisfies (3.50) and (3.51) for all p_H; (2) $(p_H^*, \hat{\mu}_B(p_H^*))$ satisfying (3.52) implies $(p_H, \hat{\mu}_B(p_H))$ satisfies the condition for all $p_H > p_H^*$; (3) (3.53) being slack initially together with its continuity implies we can increase p_H a little bit from p_H^* while keeping the constraint satisfied for $(p_H, \hat{\mu}_B(p_H))$. These together implies $(p_H, \hat{\mu}_B(p_H))$ is feasible for any p_H higher than but close enough to p_H^*.

Now, it is easy to compute:

$$\frac{d\Omega(p_H, \hat{\mu}_B(p_H); B)}{dp_H} = v_H - v_L - \Psi(p_H)$$

Notice (3.53) being slack and (3.50) being binding at the beginning implies $\Psi(p_H^*) < (v_H - K) - (v_L - K) = v_H - v_L$. (Intuitively, when (3.50) is binding, all projects are developed, so a low type's expected project payoff is $> v_L - K$ and a high type's expected project payoff is $< v_H - K$.) Therefore, $\frac{d\Omega(p_H, \hat{\mu}_B(p_H); B)}{dp_H}|_{p_H = p_H^*} > 0$, which means marginally increasing p_H from p_H^* (while keeping $\mu_B = \hat{\mu}_B(p_H)$) increases the objective function's value. This leads to a contradiction since such deviation is feasible as shown above. □

By the lemma, it suffices to focus on optimization (3.49) – (3.53). To avoid worrying about $p_H \to 0$ or 1, we can further restrict the range of p_H. By Lemma 3.4.1, we can w.l.g. restrict $p_H \leq p_H^1 < 1$. Moreover, it is easy to see the minimum value of function $\rho(\cdot)$ given B is $y_1(B) - y_2(B) > (1 - \frac{K}{v_L})(v_H - v_L) > 0$. Since $\Psi(0) = 0$, there exists $a > 0$ s.t. $p_H^0(\beta) \geq a$ for any β. Thus by Lemma 3.4.1, we can w.l.g. restrict $p_H \in [a, p_H^1]$.

Now, for any B, let $\mathcal{G}(B)$ denote the feasible set (after restricting $p_H \in [a, p_H^1]$). It is easy to see $\mathcal{G}(B)$ is a subset of $[0,1]^2$ and thus bounded. It is also closed because all the functions in the constraints are continuous. Therefore, $\mathcal{G}$ is compact-valued. Since the objective function is continuous, by Maximum Theorem, it suffices to show $\mathcal{G}(B)$ is a continuous correspondence on $B \in (0,1)$.

(1) Upper hemi-continuity of $\mathcal{G}$

Pick any $\hat{B} \in (0,1)$ and any sequence $\{B_i\}$ such that $\lim B_i = \hat{B}$. For each i, arbitrarily pick $(p_H^i, \mu_B^i) \in \mathcal{G}(B_i)$. Then, it suffices to show $\{(p_H^i, \mu_B^i)\}$ has a subsequence that converges to some point in $\mathcal{G}(\hat{B})$ (see, for example, Proposition 9.8 in Sundaram (1996)). Since $\{(p_H^i, \mu_B^i)\} \subset [0,1]^2$, it has a converging subsequence, say $\{(p_H^{i_k}, \mu_B^{i_k})\}_{k=1}^{\infty}$. Let $(\hat{p}_H, \hat{\mu}_B)$ denotes its limit. Then $(p_H^{i_k}, \mu_B^{i_k}, B_{i_k})$ satisfies the constraints (3.50) – (3.53) for all k implies $(\hat{p}_H, \hat{\mu}_B, \hat{B})$ also satisfies them by the continuity of functions in the constraints. Therefore, $(\hat{p}_H, \hat{\mu}_B) \in \mathcal{G}(\hat{B})$.

(2) Lower hemi-continuity of $\mathcal{G}$

For any $\hat{B} \in (0,1)$, pick any $(\hat{p}_H, \hat{\mu}_B) \in \mathcal{G}(\hat{B})$. We want to show for any $\epsilon > 0$, there exists $\delta > 0$ s.t. when $B \in \mathcal{B}_\delta(\hat{B})$, there exists $(p_H, \mu_B) \in \mathcal{G}(B) \cap \mathcal{B}_\epsilon((\hat{p}_H, \hat{\mu}_B))$.[38] Consider two cases:

Case 1: $\hat{\mu}_B = 0$

In this case, holding $\mu_B = 0$, any $p_H \in [a, p_H^1]$ satisfies all the constraints for all B since they don't depend on B. Thus the claim holds as $(\hat{p}_H, \hat{\mu}_B)$ remains feasible for all B.

Case 2: $\hat{\mu}_B > 0$

We can rewrite the constraints (3.50), (3.52) and (3.53) into:

$$\mu_B \leq \frac{1 - p_H}{1 - B}$$
$$\mu_B \leq \frac{p_H}{B}$$
$$\mu_B \leq \frac{v_H - K - \Psi(p_H)}{\frac{B}{p_H}[v_H - K - y_1(B)] + \frac{1-B}{1-p_H} y_2(B)}$$

where $y_1(B)$ and $y_2(B)$ are as given in Observation 3.C.1. Notice all the RHS's are continuous in B, so when B deviates from $\hat{B}$ but keeps close enough to it, holding $p_H = \hat{p}_H$, the changes in the RHS's are all within ϵ. Thus, there exists $\mu_B \in \mathcal{B}_\epsilon(\hat{\mu}_B)$ s.t. $(\mu_B, \hat{p}_H)$ satisfies all the constraints. Therefore, there exists $\delta > 0$ s.t. for all $B \in \mathcal{B}_\delta(\hat{B})$ there exists μ_B satisfying $(\hat{p}_H, \mu_B) \in \mathcal{G}(B)$ and $(\hat{p}_H, \mu_B) \in \mathcal{B}_\epsilon((\hat{p}_H, \hat{\mu}_B))$.

Therefore, $\mathcal{G}$ is a continuous correspondence on B and thus the lemma is true by Maximum

[38] $\mathcal{B}_\epsilon(x)$ denotes ϵ open ball around x.

Theorem. *Q.E.D.*

Now, we can go to prove Proposition 3.4.5.

Proof for Proposition 3.4.5. The function $H(\cdot)$ is defined as:

$$H(p) := \frac{[v_H - K - y_1(p) + y_2(p)]^2}{(1-p)(v_L - K)} + \frac{v_H - K - y_1(p) + y_2(p)}{p} - \frac{y_2(p)}{p(1-p)} \tag{3.54}$$

where $y_1(\cdot)$ and $y_2(\cdot)$ are as defined in Observation 3.C.1. As $y_1(p) - y_2(p) < v_H - v_L$ for any p, one can verify $H(p) > \frac{v_H - K - y_1(p)}{p(1-p)} > 0$.

In the following proof, to highlight the functions $\rho(\cdot)$ and $\rho_L(\cdot)$ depend on β, I will denote them as $\rho(\cdot;\beta)$ and $\rho_L(\cdot;\beta)$ respectively. Let $p_H^*(\beta)$ denote the smallest optimal p_H given any β (which exists because the solution set is compact as shown in Lemma 3.C.2). The following lemma illustrates the implication needed from the proposition's condition:

Lemma (i). If β_0 satisfies $\Psi(B(\beta_0)) = \nu(B(\beta_0))$, then under the proposition's condition we have $p_H^*(\beta_0) > B(\beta_0)$.

Subproof. To ease notation, let $B_0 := B(\beta_0)$. By the definition of $\nu(\cdot)$, it is easy to check $\Psi(B_0) = \nu(B_0) \Rightarrow \Psi(B_0) = \rho(B_0;\beta_0) = \rho_L(B_0;\beta_0)$, so $p_H^0(\beta_0) \geq B_0$. If $p_H^0(\beta_0) > B_0$, the result trivially holds. If $p_H^0(\beta_0) = B_0$, then one can check:

$$V'(B_0;\beta_0) - \Psi(B_0) = \frac{(1-B_0)(v_L - K)}{v_H - K - y_1(B_0) + y_2(B_0)}[H(B_0) - \Psi'(B_0)]$$

(To check this, use the facts: (i) $\rho_L(B_0;\beta_0) = \Psi(B_0)$; (ii) $\rho_L(B_0;\beta_0) = y_1(B_0) - y_2(B_0)$; (iii) $\rho_L'(B_0;\beta_0) = -\frac{y_2(B_0)}{(1-B)B}$.) Notice the LHS is just the derivative at $p_H = B_0$ of the objective function in optimization (3.18) given that $B = B_0$, so the condition $H(B_0) > \Psi'(B_0)$ implies the FOC for optimization (3.18) does not hold at B_0. Thus $p_H^*(\beta_0) > B_0$. □

Define $\beta^* = \inf\{\beta \in (\underline{\beta}, \bar{\beta}) : p_H^0(\beta) \geq B(\beta)\}$. It's easy to see $\beta^* < \bar{\beta}$, since as $\beta \to \bar{\beta}$, we have $B(\beta) \to 0$ but $p_H^0(\beta)$ is bounded by a positive number from below (see Lemma (i) in the proof of Proposition 3.4.4). It's also easy to see $\beta^* > \underline{\beta}$, since as $\beta \to \underline{\beta}$, $B(\beta) \to 1$ but $p_H^0(\beta) < p_H^1 < 1$.

We have:

Lemma (ii). $p_H^0(\beta^*) \geq B(\beta^*)$.

Subproof. Suppose $p_H^0(\beta^*) < B(\beta^*)$ and pick $a \in (p_H^0(\beta^*), B(\beta^*))$. Define

$$\tau(\beta) := \max_{p_H \in [a,1]} \{\rho(p_H; \beta) - \Psi(p_H)\}$$

(max is achieved by continuity and compactness of $[a, 1]$). Then, $p_H^0(\beta^*) < a$ implies $\tau(\beta^*) < 0$, since $\Psi(\cdot)$ must be above $\rho(\cdot; \beta)$ after $p_H^0(\beta)$.

Since $\rho(\cdot; \cdot)$ and $\Psi(\cdot)$ are continuous, the Maximum Theorem implies $\tau(\cdot)$ is continuous. Since $B(\beta)$ is also continuous in β, there exists $\epsilon > 0$ s.t. for all $\beta \in (\beta^*, \beta^* + \epsilon)$ we have $\tau(\beta) < 0$ and $B(\beta) > a$, which implies $p_H^0(\beta) \leq a < B(\beta)$. This contradicts with the definition of β^*, so $p_H^0(\beta^*) \geq B(\beta^*)$.

□

If $p_H^0(\beta^*) > B(\beta^*)$, then $p_H^*(\beta^*) > B(\beta^*)$ trivially. If $p_H^0(\beta^*) = B(\beta^*)$, then $\Psi(\cdot)$ intersects $\rho(\cdot; \beta^*)$ at $B(\beta^*)$, which implies $\Psi(B(\beta^*)) = \nu(B(\beta^*))$. Then, by Lemma (i), we also have $p_H^*(\beta^*) > B(\beta^*)$. Therefore, $p_H^*(\beta^*) > B(\beta^*)$ anyways.

Let $z := [p_H^*(\beta^*) + B(\beta^*)]/2$. Since $B(\beta)$ is continuously decreasing in β, when β is smaller than but close enough to β^*, we have $B(\beta) < z$. Moreover, by the definition of β^*, we have $p_H^0(\beta) < B(\beta)\ \forall \beta < \beta^*$. Combining these, there exists $\epsilon_1 > 0$ s.t. $p_H^0(\beta) < z$ for all $\beta \in (\beta^* - \epsilon_1, \beta^*)$.

Now, suppose for any $\epsilon \in (0, \epsilon_1)$, there exists $\beta \in (\beta^* - \epsilon, \beta^*)$ s.t. $p_H^0(\beta)$ is optimal given β. Then, we can pick a sequence $\{\beta^i\}_{i=1}^{\infty} \subset (\beta^* - \epsilon_1, \beta^*)$ s.t. $\lim \beta^i = \beta^*$ and $p_H^0(\beta^i)$ is optimal given β^i for all i. But because $p_H^0(\beta) < z$ for all $\beta \in (\beta^* - \epsilon_1, \beta^*)$, we know $p_H^*(\beta^*) - p_H^0(\beta^i) \geq p_H^*(\beta^*) - z > 0$ for all i. Since $p_H^*(\beta^*)$ is the smallest optimal p_H given β^*, this implies that the sequence $\{p_H^0(\beta^i)\}$ is bounded away from the optimal set of p_H given β^*, which violates the optimal set of p_H being upper hemi-continuous.

Therefore, there exists $\epsilon > 0$ s.t. for any $\beta \in (\beta^* - \epsilon, \beta^*)$, $p_H^0(\beta)$ is not optimal and thus $p_H^*(\beta) > p_H^0(\beta)$. Since Proposition 3.4.3 has ruled out the optimality of full disclosure, the optimal μ then has full support on $\{0, B(\beta), 1\}$ for all such β according to the discussion right below Proposition 3.4.1. *Q.E.D.*

Proof for Proposition 3.4.6

First, I explain the technical condition of piecewise monotonicity. Formally, I say a continuous function f is piecewise monotone on $[a, b]$ if there exists a finite partition $a = t_0 < t_1 < ... < t_n = b$ s.t. f is monotone on each subinterval $[t_i, t_i + 1]$ for $0 \leq i \leq n-1$. Indeed, this is a rather weak condition and almost all continuous functions one commonly sees satisfy it on any bounded domain.

The condition is useful for the following (almost obvious) observation: if continuous function f is piecewise monotone on $[a, b]$, then for any local maximum of it, say x^*, there exists $\epsilon > 0$ s.t. f is increasing on $(x^* - \epsilon, x^*] \cap [a, b]$ and decreasing on $[x^*, x^* + \epsilon) \cap [a, b]$. (For an example where this conclusion fails when f is C^1 but not piecewise monotone, consider $f(x) = x^4(\sin(\frac{1}{x}) - 1)$ with $f(0) := 0$ around local maximum $x = 0$.)

Now, I go to the main proof:

Proof. The central part is the following lemma, which shows a local version of decreasing differences property for the objective function in optimization (3.18).

Lemma (i). Let $\Upsilon(p_H, \beta)$ denote the objective function in optimization (3.18) given β. Then, for any $\beta_0 \in (\underline{\beta}, \bar{\beta})$, there exist β_l, β_u, p_1 and p_2 s.t. $(p_H^*(\beta_0), \beta_0) \in (p_1, p_2) \times (\beta_l, \beta_h)$ and $\frac{\partial^2 \Upsilon}{\partial p_H \partial \beta} < 0$ on $(p_1, p_2) \times (\beta_l, \beta_h)$.

Subproof. Define function

$$\begin{aligned}\gamma(p_H, \beta) :=& \Big[(1-\beta)Kp_H - [(2-\beta)K - v_L](1-p_H)\Big][v_H - K - \Psi(p_H)] \\ &+ [(2-\beta)K - v_L](1-p_H)p_H\Psi'(p_H) + (1-p_H)p_H^2(v_L - K)\Psi'(p_H)\end{aligned}$$

Then, one can check (maybe with a computer) that: $sign(\frac{\partial^2 \Upsilon}{\partial p_H \partial \beta}) = -sign(\gamma(p_H, \beta))$.

Now, for any $\beta_0 \in (\underline{\beta}, \bar{\beta})$, by FOC we have:

$$v_H - K - \Psi(p_H^*(\beta_0)) + \frac{1-B}{B}(v_L - K)\frac{\partial[p_H R(p_H; \beta)]}{\partial p_H}|_{(p_H^*(\beta_0), \beta_0)} \leq 0$$

Since $v_H - K - \Psi(p_H^*(\beta_0)) > 0$ by Proposition 3.4.3,[39] we have $\frac{\partial[p_H R(p_H;\beta)]}{\partial p_H}|_{(p_H^*(\beta_0),\beta_0)} < 0$.

One can check (maybe with a computer) that: $sign(\frac{\partial[p_H R(p_H;\beta)]}{\partial p_H}) = -sign\big(\gamma(p_H, \beta) + p_H(1-$

[39]Recall a p_H with $v_H - K - \Psi(p_H) = 0$ is induced only under full disclosure.

$p_H)(v_L - K)[\Psi(p_H) - (v_H - K)]\big)$. Thus,

$$\frac{\partial[p_H R(p_H;\beta)]}{\partial p_H}|_{(p_H^*(\beta_0),\beta_0)} < 0 \Rightarrow \gamma(p_H^*(\beta_0),\beta_0) > 0 \Rightarrow \frac{\partial^2\Upsilon}{\partial p_H \partial\beta}|_{(p_H^*(\beta_0),\beta_0)} < 0$$

where I used $v_H - K - \Psi(p_H^*(\beta_0)) > 0$ again in the first step. Finally, the result holds by the continuity of the cross partial derivative under Assumption 3. □

As mentioned in the proposition, define $\beta_c := \inf\{\beta \in (\underline{\beta},\bar{\beta}) : p_H^*(\beta) = p_H^0(\beta) \geq B(\beta)\}$. Notice when β is close to $\underline{\beta}$, $B(\beta) > p_H^1 > p_H^0(\beta)$, so $\beta_c > \underline{\beta}$. Thus $\beta_c \in (\underline{\beta},\bar{\beta})$ when the set is non-empty and $\beta_c = +\infty$ otherwise.

Part (a): First notice $p_H^*(\beta) = p_H^0(\beta) \geq B(\beta)$ is equivalent to that $\mu^*(;\beta)$ supports on $\{B(\beta),1\}$ or $\{B(\beta)\}$, so the "only if" part is trivial. For the "if" part, if $\beta_c = +\infty$, the conclusion is also trivial, so it suffices to assume $\beta_c < \bar{\beta}$.

By continuity proved in Lemma 3.C.2, we know when $\beta = \beta_c$, the optimal $\mu(0) = 0$ and thus the optimal μ supports on $\{B(\beta)\}$ or $\{B(\beta),1\}$, which implies $p_H^*(\beta_c) = p_H^0(\beta_c) \geq B(\beta_c)$. Notice $p_H^0(\beta_c) \geq B(\beta_c)$ further implies $p_H^0(\beta)$ is increasing in β when $\beta \geq \beta_c$ (Lemma 3.C.1), so $p_H^0(\beta) > B(\beta)$ for all $\beta > \beta_c$ (recall B decreases in β). Therefore, it suffices to show $p_H^*(\beta) = p_H^0(\beta)$ for all $\beta > \beta_c$.

Now, suppose this is not true (i.e., $p_H^*(\beta) > p_H^0(\beta)$ for some $\beta > \beta_c$). Define $\beta^* = \inf\{\beta \in [\beta_c,\bar{\beta}) : p_H^*(\beta) > p_H^0(\beta)\}$. Then, $\beta^* < \bar{\beta}$.

If $\beta^* = \beta_c$, we have $p_H^*(\beta^*) = p_H^0(\beta^*)$. If $\beta^* > \beta_c$, then for any $\beta \in (\beta_c,\beta^*)$, we have $p_H^*(\beta) = p_H^0(\beta)$. By continuity of $p_H^*(\cdot)$, we have $p_H^*(\beta^*) = \lim_{\beta\uparrow\beta^*}(p_H^*(\beta)) = \lim_{\beta\uparrow\beta^*}(p_H^0(\beta)) \leq p_H^0(\beta^*)$ (where the last inequality is by Lemma 3.C.1). This again implies $p_H^*(\beta^*) = p_H^0(\beta^*)$. Thus in either case, $p_H^*(\beta^*) = p_H^0(\beta^*)$.

However, we have the following lemma:

Lemma (ii). If β_0 satisfies $p_H^*(\beta_0) = p_H^0(\beta_0) \geq B(\beta_0)$, then there exists $\epsilon > 0$ s.t. $p_H^*(\beta) = p_H^0(\beta)$ for all $\beta \in (\beta_0, \beta_0 + \epsilon)$.

Subproof. Let β_0 satisfies the lemma's condition. By Lemma (i), there exist β_l, β_u, p_1 and p_2 s.t. $(p_H^*(\beta_0),\beta_0) \in (p_1,p_2)\times(\beta_l,\beta_h)$ and $\frac{\partial^2\Upsilon}{\partial p_H\partial\beta} < 0$ on $(p_1,p_2)\times(\beta_l,\beta_h)$. Since $p_H^*(\beta_0)$ is the unique optimal p_H given β_0, the technical condition that $\Upsilon(\cdot;\beta)$ is piecewise monotone implies $\Upsilon(\cdot;\beta_0)$ is decreasing on $[p_H^*(\beta_0), p_H^*(\beta_0)+\delta)$ for some $\delta > 0$. Thus we can w.l.g. assume p_2 is small enough s.t. $\Upsilon(\cdot;\beta_0)$ is decreasing on $[p_H^*(\beta_0), p_2)$. Then, the decreasing differences condition $\frac{\partial^2\Upsilon}{\partial p_H \partial \beta} < 0$ implies $\Upsilon(\cdot;\beta)$ is decreasing on $[p_H^*(\beta_0), p_2)$ for all $\beta \in [\beta_0, \beta_h)$.

Since $p_H^0(\beta_0) \geq B(\beta_0)$, $p_H^0(\beta)$ is increasing in β on $[\beta_0, \bar{\beta})$ (Lemma 3.C.1). Thus, $p_H^0(\beta) \geq p_H^0(\beta_0) = p_H^*(\beta_0)$ for any $\beta > \beta_0$. By the local monotonicity of $\Upsilon(\cdot;\beta)$, we know for any $\beta \in (\beta_0, \beta_h)$, it cannot be the case that $p_H^0(\beta) < p_H^*(\beta) < p_2$. Therefore, either $p_H^*(\beta) = p_H^0(\beta)$ or $p_H^*(\beta) \geq p_2$ (or both). However, by the continuity of $p_H^*(\cdot)$, as $\beta \downarrow \beta_0$, $p_H^*(\beta) \to p_H^*(\beta_0) < p_2$. Thus, when β is close to β_0 enough, we must have $p_H^*(\beta) = p_H^0(\beta)$, which is the lemma's conclusion. □

Now, since $p_H^*(\beta^*) = p_H^0(\beta^*) \geq B(\beta^*)$, the lemma implies that there exists $\epsilon > 0$ s.t. $p_H^*(\beta) = p_H^0(\beta)$ for all $\beta \in (\beta^*, \beta^* + \epsilon)$. This contradicts with the definition of β^*, so the result in part (a) holds. (As a side note, $p_H^0(\beta) > B(\beta)$ for all $\beta > \beta_c$ implies that the optimal μ may support on $\{B\}$ only when $\beta = \beta_c$.)

Part (b): First notice for $\beta \in [\beta_c, \bar{\beta})$, part (a) implies $p_H^*(\beta) = p_H^0(\beta)$, so the increasing result is implied by the fact that $p_H^0(\cdot)$ is strictly increasing on $[\beta_c, \bar{\beta})$ as mentioned above. For $p_H^*(\cdot)$ being decreasing on $(\underline{\beta}, \beta_c)$, we show the following result:

Lemma (iii). For any $\beta_0 \in (\underline{\beta}, \beta_c)$, there exists $\epsilon > 0$ s.t. $p_H^*(\beta) \geq p_H^*(\beta_0)$ for all $\beta \in (\beta_0 - \epsilon, \beta_0)$.

Subproof. We consider two cases:

Case 1: $p_H^0(\beta_0) \leq B(\beta_0)$.

By Lemma 3.C.1, we know $p_H^0(\cdot)$ is decreasing on $(\underline{\beta}, \beta_0]$, so if $p_H^*(\beta_0) = p_H^0(\beta_0)$, then the result is trivial since for all $\beta < \beta_0$, we have $p_H^*(\beta) \geq p_H^0(\beta) > p_H^0(\beta_0)$.

If $p_H^*(\beta_0) > p_H^0(\beta_0)$, then there exists $\epsilon > 0$ s.t. $\Upsilon(p_H;\beta_0)$ is increasing on $(p_H^*(\beta_0) - \epsilon, p_H^*(\beta_0)]$ by the piecewise monotonicity. Moreover, by Lemma (i), there exist β_l, β_u, p_1 and p_2 s.t. $(p_H^*(\beta_0), \beta_0) \in (p_1, p_2) \times (\beta_l, \beta_h)$ and $\frac{\partial^2\Upsilon}{\partial p_H \partial \beta} < 0$ on $(p_1, p_2) \times (\beta_l, \beta_h)$. We can assume p_1 is large enough s.t. $p_1 > p_H^*(\beta_0) - \epsilon$. Then, the decreasing differences condition implies $\Upsilon(p_H;\beta)$ is increasing on $(p_1, p_H^*(\beta_0)]$ for all $\beta \in (\beta_l, \beta_0]$. Therefore, for all $\beta \in (\beta_l, \beta_0]$, either $p_H^*(\beta) \leq p_1$ or $p_H^*(\beta) \geq p_H^*(\beta_0)$ (notice if $p_H^*(\beta) < p_H^*(\beta_0)$, $p_H^*(\beta_0)$ is feasible). However, by continuity $p_H^*(\beta) \to p_H^*(\beta_0)$ as $\beta \uparrow \beta_0$, so when β is close enough to β_0, we must have $p_H^*(\beta) > p_1$. Therefore, $p_H^*(\beta) \geq p_H^*(\beta_0)$ when $\beta < \beta_0$ but close enough to it, which is the lemma's conclusion.

Case 2: $p_H^0(\beta_0) > B(\beta_0)$.

In this case, part (a) implies $p_H^*(\beta_0) > p_H^0(\beta_0)$, otherwise $\mu^*(;\beta_0)$ would support on $\{B(\beta_0), 1\}$ or $\{B(\beta_0)\}$. Then, the argument is exactly the same as the second paragraph in Case 1. $\square$

Finally, the proof is completed by the following simple fact:

Lemma (iv). If a continuous function $f : (a, b) \to \mathbb{R}$ satisfies for any $x_0 \in (a, b)$, there exists $\epsilon > 0$ s.t. $f(x) \geq f(x_0)$ for all $x \in (x_0 - \epsilon, x_0)$, then f is decreasing on (a, b).

Subproof. Supposing not, there exist x_1 and x_2 s.t. $a < x_1 < x_2 < b$ and $f(x_2) > f(x_1)$. By continuity of f, $\text{argmax}_{x\in[x_1,x_2]}\{f(x)\}$ is compact. Let x_m be its minimum element. Then, $x_m > x_1$ since $f(x_1) < f(x_2)$. This implies that $f(x_m) > f(x)$ for all $x \in (x_1, x_m)$, which contradicts the lemma's assumption. $\square$

Combining Lemma (iii) and Lemma (iv), we conclude $p_H^*(\cdot)$ is decreasing on $(\underline{\beta}, \beta_c)$ *Q.E.D.*

3.C.3 Proofs in Appendix 3.A

Proof for Lemma 3.A.1

Proof. It suffices to prove this for the endogenous p_H case.

First notice condition (3.20) is equivalent to:

$$\pi_H \geq \max\{1 - \frac{\xi v_H}{(K+\xi)(v_H - v_L)}, B\}$$

Define the RHS as π^*. Then, the high types will not use costly certification when $\pi_H \geq \pi^*$.

Now, let $(p_H^a, \mathcal{D})$ be a pair of target p_H and disclosure rule that induces some high type entrepreneurs to use the certification service. Let S be the set of signal realizations of $\mathcal{D}$ and Γ be its family of conditional distribution (so $\mathcal{D} = \{S, \Gamma\}$). Then, there must be some signal realizations in S inducing $\pi_H < \pi^*$ given p_H^a. Let $T \subset S$ be the subset including all these kind of signal realizations.

Now, suppose the designer can fix p_H at p_H^a, but design a new disclosure rule $\mathcal{D}'$, which is the same as $\mathcal{D}$ except that whenever signals in T are realized under $\mathcal{D}$, a full disclosure signal is realized under $\mathcal{D}'$. Specifically, $\mathcal{D}' = \{S', \Gamma'\}$ s.t.

$$S' = (S \setminus T) \cup \{s_0, s_1\}$$
$$\Gamma'(s|v) = \Gamma(s|v) \; \forall s \in S \setminus T, \; \forall v$$
$$\Gamma'(s_0|v_L) = \Gamma(T|v_L), \; \Gamma'(s_1|v_L) = 0, \; \Gamma'(s_0|v_H) = 0, \; \Gamma'(s_1|v_H) = \Gamma(T|v_H)$$

where $\Gamma(T|v)$ denotes the sum of probabilities of signal realizations in T given type v under rule $\mathcal{D}$.

Notice under $\mathcal{D}$, any signal realization in T causes a high type to be developed (at least the entrepreneur can use certification) and causes a low type not to be developed (due to $\pi_H < B$ or high types identify themselves by certification). Then, it's obvious that the pair $(p_H^a, \mathcal{D}')$ dominates the pair $(p_H^a, \mathcal{D})$ in welfare because all projects that are developed under $\mathcal{D}$ are still developed under $\mathcal{D}'$ (given p_H^a), while the high types don't need to pay the certification costs. Moreover, all low types used to receive $s \in T$ under $\mathcal{D}$ get payoff 0 under $\mathcal{D}'$ (by receiving s_0) and all high types used to receive $s \in T$ under $\mathcal{D}$ get payoff $v_H - K$ under $\mathcal{D}'$ (by receiving s_1). Thus the upgrading incentive under $(p_H^a, \mathcal{D}')$ is higher than that under $(p_H^a, \mathcal{D})$ (strictly so since high types don't need to pay certification costs under $\mathcal{D}'$). Although p_H^a is no longer incentivized by $\mathcal{D}'$, we have the following lemma:

Lemma (i). There exists $p_H^b > p_H^a$ s.t. p_H^b is incentivized by $\mathcal{D}'$.

Subproof. Since $\mathcal{D}'$ does not induce any non-zero posterior $\pi_H < \pi^*$ given p_H^a, so is true for $\mathcal{D}'$ given any $p_H \geq p_H^a$. Also since $\pi^* \geq B$, we have for all $p_H \geq p_H^a$:

$$\begin{aligned} u_H(\mathcal{D}'; p_H) =& \mathbb{E}_{(p_H, \mathcal{D}')}\Big[(1 - \frac{K}{\pi_H v_H + (1-\pi_H)v_L})v_H \mathbb{1}\{\pi_H > 0\}\Big] \\ =& \sum_{s \in \hat{S}'} \Big[(1 - \frac{K}{\pi(s;p_H)v_H + (1-\pi(s;p_H))v_L})v_H \Gamma'(s|v_H)\Big] \\ u_L(\mathcal{D}'; p_H) =& \mathbb{E}_{(p_H, \mathcal{D}')}\Big[(1 - \frac{K}{\pi_H v_H + (1-\pi_H)v_L})v_L \mathbb{1}\{\pi_H > 0\}\Big] \\ =& \sum_{s \in \hat{S}'} \Big[(1 - \frac{K}{\pi(s;p_H)v_H + (1-\pi(s;p_H))v_L})v_L \Gamma'(s|v_L)\Big] \end{aligned}$$

where $\hat{S}' := \{s \in S' : \Gamma'(s|v_H) > 0\}$ is the set of signal realizations in S' that induce non-zero posterior π_H (which then must be $\geq \pi^*$ given $p_H \geq p_H^a$); and $\pi(s; p_H) := \frac{\Gamma'(s|v_H)p_H}{\Gamma'(s|v_H)p_H + \Gamma'(s|v_L)(1-p_H)}$

is the posterior π_H induced by $s \in \mathcal{D}'$ given p_H. It's easy to see $\pi(\cdot; p_H)$ is continuous in p_H and thus $u_H(\mathcal{D}'; p_H) - u_L(\mathcal{D}'; p_H)$ is continuous in p_H.

As shown above, upgrading incentive under $(p_H^a, \mathcal{D}')$ is strictly higher than that under $(p_H^a, \mathcal{D})$, so we have:

$$u_H(\mathcal{D}'; p_H^a) - u_L(\mathcal{D}'; p_H^a) > u_H(\mathcal{D}; p_H^a) - u_L(\mathcal{D}; p_H^a) = \Psi(p_H^a)$$

Since $u_H(\mathcal{D}'; p_H^a) - u_L(\mathcal{D}'; p_H^a) \leq v_H - K < \Psi(1)$, by continuity of $[u_H(\mathcal{D}'; p_H) - u_L(\mathcal{D}'; p_H)]$ and $\Psi(p_H)$ in p_H, there exists $p_H^b > p_H^a$ s.t.

$$u_H(\mathcal{D}'; p_H^b) - u_L(\mathcal{D}'; p_H^b) = \Psi(p_H^b)$$

which means p_H^b is incentivized by $\mathcal{D}'$. □

Now, it suffices to show $(p_H^b, \mathcal{D}')$ (weakly) dominates $(p_H^a, \mathcal{D}')$ in welfare, which then dominates $(p_H^a, \mathcal{D})$. Let W_1 denote the net social surplus under $(p_H^a, \mathcal{D}')$ and W_2 denote that under $(p_H^b, \mathcal{D}')$. Then,

$$\begin{aligned} W_2 - W_1 &= (p_H^b - p_H^a)[(v_H - K) - (v_L - K)\Gamma'(\hat{S}'|v_L)] - \int_{p_H^a}^{p_H^b} \Psi(x)dx \\ &= \int_{p_H^a}^{p_H^b} [(v_H - K) - (v_L - K)\Gamma'(\hat{S}'|v_L) - \Psi(x)]dx \end{aligned}$$

where $\Gamma'(\hat{S}'|v_L)$ is the probability for a low type to get $s \in \hat{S}'$ under $\mathcal{D}'$ and thus get financed and developed given $p_H \geq p_H^a$. Notice the expected project payoff for high type entrepreneurs is always smaller than $v_H - K$ and the expected project payoff for low type entrepreneurs under $(p_H^b, \mathcal{D}')$ is higher than $(1-\beta)K\Gamma'(\hat{S}'|v_L) \geq (v_L - K)\Gamma'(\hat{S}'|v_L)$. Therefore,

$$\begin{aligned} (v_H - K) - (v_L - K)\Gamma'(\hat{S}'|v_L) &\geq u_H(\mathcal{D}'; p_H^b) - u_L(\mathcal{D}'; p_H^b) \\ &= \Psi(p_H^b) \geq \Psi(p_H) \; \forall p_H \in [p_H^a, p_H^b] \end{aligned}$$

so the integrand above is positive and thus $W_2 - W_1 \geq 0$. *Q.E.D.*

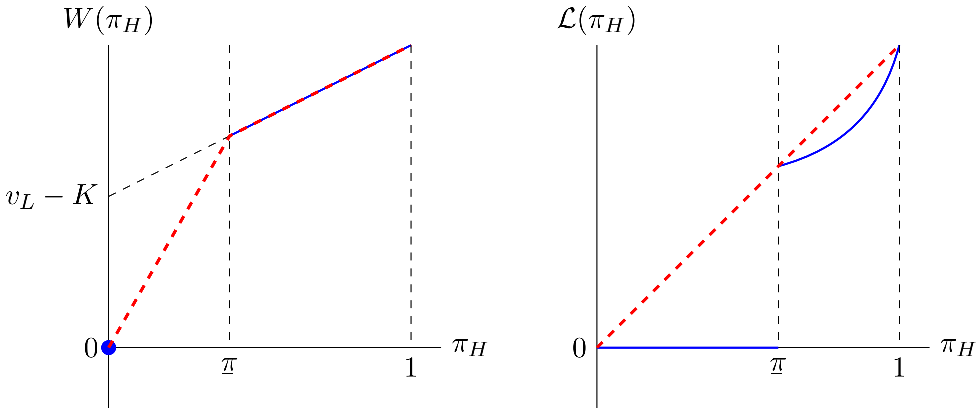

(a) Exogenous p_H

(b) Endogenous p_H

Figure 3.8: Cancavification with costly certification

Proof for Observation 3.A.1

Proof. The range includes $\pi_H = 0$ since it can only be induced for low types. For $\pi_H \geq B$, notice condition (3.23) says $\underline{\pi} > B$, so (3.22) is equivalent to $\pi_H \geq \underline{\pi}$. Therefore, certification is not used when $\pi_H \in [\underline{\pi}, 1]$ and is used when $\pi_H \in [B, \underline{\pi}]$. For $\pi_H \in (0, B)$, we have $(1 - \frac{K}{\pi_H v_H})v_H \leq (1 - \frac{K}{\pi_H v_H + (1-\pi_H)v_L})v_H \leq (1 - \frac{K}{Bv_H + (1-B)v_L})v_H < (1 - \frac{K}{\underline{\pi} v_H + (1-\underline{\pi})v_L})v_H = v_H - K - \xi$, where the last inequality is because $B < \underline{\pi}$ by (3.23). Thus certification is used with $\pi_H \in (0, B)$. *Q.E.D.*

Proof for Proposition 3.A.1

Proof. The proof is basically the same as that for Proposition 3.3.2 with B replaced by $\underline{\pi}$. The graph for concavification is in Figure 3.8a. The solid blue curve is for $W(\cdot)$ on domain $\{0\} \cup [\underline{\pi}, 1]$ and the dotted red curve is the concavification. *Q.E.D.*

Proofs for Proposition 3.A.2 and Proposition 3.A.3

The proofs are largely the same as those in Section 3.C.2, with B replaced by $\underline{\pi}$ and function $\rho(\cdot)$ redefined correspondingly. The concavification graph for the constrained Bayesian persuasion problem is presented in Figure 3.8b.

3.C.4 Proofs in Appendix 3.B

Proof for Observation 3.B.1

Proof. We mainly need to show R_1 and R_2 are well defined. Notice it is easy to see $S(e, R)$ is continuous and decreasing in R, which is strictly decreasing on $\{R : \mathbb{P}(z > R - e) > 0\}$. Moreover, notice:

(1) $S(e_L, 0) = \mathbb{E}[\max\{z + e_L, 0\}] \geq \max\{\mathbb{E}(z + e_L), 0\} > e_L + 1 > m$;

(2) $S(e_H, 0) = \mathbb{E}[\max\{z + e_H, 0\}] \geq \max\{\mathbb{E}(z + e_H), 0\} > e_H + 1 > e_H$;

(3) $\lim_{R\to\infty} S(e, R) = 0$ for any e.

Thus by continuity, there exist $R_1, R_2 > 0$ s.t. $S(e_L, R_1) = m$ and $S(e_H, R_2) = e_H$.

For uniqueness, take R_1 for example. Suppose there exists $R_1' > R_1$ s.t. $S(e_L, R_1') = S(e_L, R_1) = m$. Then, $S(e_L, \cdot)$ is constant on $[R_1, R_1']$, which implies $P(z > R - e_L) = 0 \Rightarrow S(e_L, R) = 0$ for $R \in (R_1, R_1')$. This contradicts with $S(e_L, R_1') = m$. Similarly, R_2 is also unique.

Finally, by monotonicity of $S(e, R)$ in R, we have $S(e_H, R) \geq e_H \Leftrightarrow R \leq R_2$ and $S(e_L, R) \geq m \Leftrightarrow R \leq R_1$. Thus we have the bank's behaviors as given in the observation. *Q.E.D.*

Proof for Lemma 3.B.1

Proof.

- Class 1: $R \leq R_1$

  When $R \leq R_1$, both types accept the offer and develop the project. Then, investor's condition (3.30) implies this kind of equilibria exist if and only if:

$$\exists R \leq R_1, \text{ s.t. } \pi_H B(e_H, R) + (1 - \pi_H) B(e_L, R) \geq 1$$

  Since $B(e, R)$ is increasing in R, this is equivalent to:

$$\pi_H B(e_H, R_1) + (1 - \pi_H) B(e_L, R_1) \geq 1$$

  By definition of C_1, this is equivalent to $\pi_H \geq C_1$.

- Class 2: $R_1 < R \leq R_2$

  In this case, both types accept the debt offer, but only the high type carries out the project while the low type diverts. Thus the investor's condition (3.30) is satisfied for this case if

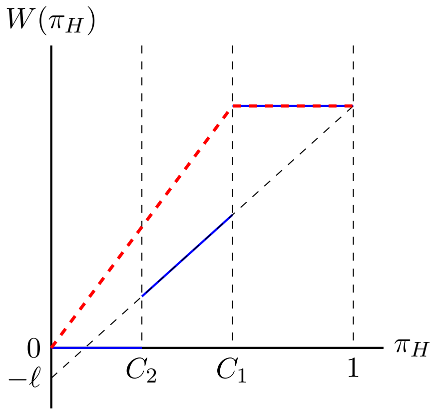

Figure 3.9: Concavification for $W(\pi_H)$ in (3.34)

Note: Blue curve is function $W(\pi_H)$; Red dotted curve is the concavification boundary.

and only if:

$$\exists R \in (R_1, R_2], \text{ s.t. } \pi_H B(e_H, R) \geq 1$$

Since $B(e, R)$ is increasing in R, this is equivalent to $\pi_H \geq \frac{1}{B(e_H, R_2)} = C_2$.

- Class 3: no debt offer is made by the investor

  This happens when $\pi_H < \min\{C_1, C_2\}$, so no R can make condition (3.30) satisfied. (Notice the investor never offers $R > R_2$ since only the low type accepts it, who is going to divert with it.)

Q.E.D.

Proof for Proposition 3.B.2

Proof. **Case 1:** $C_1 \geq C_2$

The indirect welfare function together with its concavification is drawn in Figure 3.9. The red dashed line is the concavification of $W(\cdot)$. It consists of two segments: $(0, 0) \rightarrow (C_1, W(C_1))$ and $(C_1, W(C_1)) \rightarrow (1, W(1))$.

Thus, when $p_H < C_1$, the optimal μ supports on $\{0, C_1\}$. By Bayesian feasibility, $\mu(0) \times 0 + \mu(C_1)C_1 = p_H \Rightarrow \mu(C_1) = p_H/C_1$. When $p_H \geq C_1$, any μ supporting on $[C_1, 1]$ satisfying Bayesian feasibility condition is optimal.

Case 2: $C_1 < C_2$

In this case, we simply don't have the second piece for indirect welfare function, so the concavification is the same as that in Case 1, expect that there is no segment corresponding to interval $[C_2, C_1]$. Therefore, the result is the same as in Case 1. *Q.E.D.*

References

Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of economic Literature*, *54*(2), 442–92.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, *84*(3), 488–500.

Alexander, D., & Isakin, M. (2015). Credit ratings, credit crunches, and the pricing of collateralized debt obligations. *Working Paper*.

Altman, E. (1999). *Constrained markov decision processes: stochastic modeling.* Routledge.

Alvarez, F., & Barlevy, G. (2015). Mandatory disclosure and financial contagion. *NBER Working Paper*.

Amador, M., & Weill, P.-O. (2010). Learning from prices: Public communication and welfare. *Journal of Political Economy*, *118*(5), 866–907.

Anderson, S. P., & Renault, R. (2006). Advertising content. *American Economic Review*, *96*(1), 93–113.

Anderson, S. P., & Renault, R. (2013). The advertising mix for a search good. *Management Science*, *59*(1), 69–83.

Andolfatto, D., Berentsen, A., & Waller, C. (2014). Optimal disclosure policy and undue diligence. *Journal of Economic Theory*, *149*, 128–152.

Angeletos, G.-M., & Pavan, A. (2007). Efficient use of information and social value of information. *Econometrica*, *75*(4), 1103–1142.

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association*, *15*(5), 989–1024.

Ashour, S. K., & Abdel-hameed, M. A. (2010). Approximate skew normal distribution. *Journal of Advanced Research*, *1*(4), 341–350.

Au, P. H. (2018). Competition in designing Pandora's boxes. *Available at SSRN 3141387*.

Aumann, R. J., Maschler, M., & Stearns, R. E. (1995). *Repeated games with incomplete information*. MIT press.

Bagnoli, M., & Bergstrom, T. (2005). Log-concave probability and its applications. *Economic theory*, *26*(2), 445–469.

Bahar, G., Smorodinsky, R., & Tennenholtz, M. (2015). Economic recommendation systems. *arXiv preprint arXiv:1507.07191*.

Bahar, G., Smorodinsky, R., & Tennenholtz, M. (2021). Recommendation systems and self motivated users. *arXiv preprint arXiv:1807.01732*.

Ball, I. (2019). Dynamic information provision: Rewarding the past and guiding the future. *Available at SSRN 3103127*.

Bar-Isaac, H., Caruana, G., & Cuñat, V. (2010). Information gathering and marketing. *Journal of Economics & Management Strategy*, *19*(2), 375–401.

Bar-Isaac, H., Caruana, G., & Cuñat, V. (2012). Search, design, and market structure. *American Economic Review*, *102*(2), 1140–60.

Best, J., & Quigley, D. (2020). Persuasion for the long run. *Available at SSRN 2908115*.

Beutler, F. J., & Ross, K. W. (1985). Optimal policies for controlled markov chains with a constraint. *Journal of mathematical analysis and applications*, *112*(1), 236–252.

Beutler, F. J., & Ross, K. W. (1986). Time-average optimal constrained semi-markov decision processes. *Advances in Applied Probability*, *18*(2), 341–359.

Board, S., & Lu, J. (2018). Competitive information disclosure in search markets. *Journal of Political Economy*, *126*(5), 1965–2010.

Boleslavsky, R., Cotton, C. S., & Gurnani, H. (2017). Demonstrations and price competition in new product release. *Management Science*, *63*(6), 2016–2026.

Boleslavsky, R., & Kim, K. (2018). Bayesian persuasion and moral hazard. *Working paper*.

Brown, D. C., & Davies, S. (2018). Financing efficiency of securities-based crowdfunding. *Working Paper*.

Chang, J.-W. (2016). The economics of crowdfunding. *Working Paper*.

Che, Y.-K., & Hörner, J. (2018). Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, *133*(2), 871–925.

Chemla, G., & Tinn, K. (2018). Learning through crowdfunding. *Working Paper*.

Chen, B., Frazier, P., & Kempe, D. (2018). Incentivizing exploration by heterogeneous users. In *Conference on learning theory* (pp. 798–818).

Cho, I.-K., & Kreps, D. M. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics*, *102*(2), 179–221.

Choi, M., Dai, A. Y., & Kim, K. (2018). Consumer search and price competition. *Econometrica*, *86*(4), 1257–1281.

Choi, M., Kim, K., & Pease, M. (2019). Optimal information design for search goods. In (Vol. 109, pp. 550–56).

Cole, H. L., & Kehoe, T. J. (1996). A self-fulfilling model of mexico's 1994–1995 debt crisis. *Journal of International Economics*, *41*(3-4), 309–330.

Colombo, L., Femminis, G., & Pavan, A. (2014). Information acquisition and welfare. *The Review of Economic Studies*, *81*(4), 1438–1483.

Dang, T. V., Gorton, G., Holmström, B., & Ordonez, G. (2017). Banks as secret keepers. *American Economic Review*, *107*(4), 1005–29.

Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory*, *3*(2), 156–168.

Dogan, M., & Hu, J. (2018). Consumer search and optimal information. *Available at SSRN 3308157*.

Durrett, R. (2019). *Probability: theory and examples* (Vol. 49). Cambridge university press.

Dworczak, P., & Kolotilin, A. (2019). The persuasion duality. *arXiv preprint arXiv:1910.11392*.

Dworczak, P., & Martini, G. (2019). The simple economics of optimal persuasion. *Journal of Political Economy*, *127*(5), 1993–2048.

Ely, J. C. (2017). Beeps. *American Economic Review*, *107*(1), 31–53.

Ely, J. C., & Szydlowski, M. (2020). Moving the goalposts. *Journal of Political Economy*, *128*(2), 468–506.

Feinberg, E. A., & Piunovskiy, A. B. (2000). Multiple objective nonatomic markov decision processes with total reward criteria. *Journal of mathematical analysis and applications*, *247*(1), 45–66.

Fishman, M. J., & Hagerty, K. M. (1989). Disclosure decisions by firms and the competition for price efficiency. *The Journal of Finance*, *44*(3), 633–646.

Folland, G. B. (1999). *Real analysis: modern techniques and their applications* (Vol. 40). John Wiley & Sons.

Fu, S., & Trigilia, G. (2019). Voluntary disclosure, moral hazard and default risk. *Working Paper*.

Fuchs, W. (2007). Contracting with repeated moral hazard and private evaluations. *American Economic Review*, *97*(4), 1432–1448.

Gao, P., & Liang, P. J. (2013). Informational feedback, adverse selection, and optimal disclosure policy. *Journal of Accounting Research*, *51*(5), 1133–1158.

Georgiadis, G., & Szentes, B. (2018). Optimal monitoring design. *Working Paper*.

Glazer, J., Kremer, I., & Perry, M. (2021). The wisdom of the crowd when acquiring information is costly. *Management Science*, *67*(10), 6443–6456.

Goldstein, I., & Leitner, Y. (2018). Stress tests and information disclosure. *Journal of Economic Theory*, *177*, 34–69.

Goldstein, I., & Yang, L. (2017). Information disclosure in financial markets. *Annual Review of Financial Economics*, *9*, 101–125.

Guo, Y., & Shmaya, E. (2019). The interval structure of optimal disclosure. *Econometrica*, *87*(2), 653–675.

Hinnosaar, T., & Kawai, K. (2020). Robust pricing with refunds. *The RAND Journal of Economics*.

Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. *American Economic Review*, *61*(4), 561–574.

Hörner, J., & Skrzypacz, A. (2017). Learning, experimentation and information design. *Advances in Economics and Econometrics*, *1*, 63–98.

Hornuf, L., & Schwienbacher, A. (2017). Should securities regulation promote equity crowdfunding? *Small Business Economics*, *49*(3), 579–593.

Hornuf, L., & Schwienbacher, A. (2018). Market mechanisms and funding dynamics in equity crowdfunding. *Journal of Corporate Finance*, *50*, 556–574.

Hwang, I., Kim, K., & Boleslavsky, R. (2019). Competitive advertising and pricing. *Working paper*.

Ichihashi, S. (2020). Online privacy and information disclosure by consumers. *American Economic Review*, *110*(2), 569–95.

Immorlica, N., Mao, J., Slivkins, A., & Wu, Z. S. (2019). Bayesian exploration with heterogeneous agents. In *The world wide web conference* (pp. 751–761).

Ivanov, M. (2013). Information revelation in competitive markets. *Economic Theory*, *52*(1), 337–365.

Johnson, J. P., & Myatt, D. P. (2006). On the simple economics of advertising, marketing, and product design. *American Economic Review*, *96*(3), 756–784.

Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, *101*(6), 2590–2615.

Kleinberg, R., Waggoner, B., & Weyl, E. G. (2016). Descending price optimally coordinates search. *arXiv preprint arXiv:1603.07682*.

Klenke, A. (2020). *Probability theory: a comprehensive course.* Springer Cham.

Kolotilin, A. (2018). Optimal information disclosure: A linear programming approach. *Theoretical Economics*, *13*(2), 607–635.

Kolotilin, A., Mylovanov, T., & Zapechelnyuk, A. (2021). Censorship as optimal persuasion. *Available at SSRN 3783291*.

Kolotilin, A., Mylovanov, T., Zapechelnyuk, A., & Li, M. (2017). Persuasion of a privately informed receiver. *Econometrica*, *85*(6), 1949–1964.

Komiyama, J., & Noda, S. (2021). Deviation-based learning. *arXiv preprint arXiv:2109.09816*.

Kremer, I., Mansour, Y., & Perry, M. (2014). Implementing the "wisdom of the crowd". *Journal of Political Economy*, *122*(5), 988–1012.

Küçükgül, C., Özer, Ö., & Wang, S. (2022). Engineering social learning: Information design of time-locked sales campaigns for online platforms. *Management Science*, *68*(7), 4899–4918.

Lewis, T. R., & Sappington, D. E. (1994). Supplying information to facilitate price discrimination. *International Economic Review*, 309–327.

Liang, Y., Sun, P., Tang, R., & Zhang, C. (2020). *Efficient resource allocation contracts to reduce adverse events* (Tech. Rep.). Working Paper.

Lorecchio, C. (2021). Pursuading crowds. *unpublished manuscript*.

Lorecchio, C., & Monte, D. (2021). Dynamic information design under constrained communication rules. *American Economic Journal: Microeconomics*.

Luenberger, D. G. (1997). *Optimization by vector space methods.* John Wiley & Sons.

Mansour, Y., Slivkins, A., & Syrgkanis, V. (2020). Bayesian incentive-compatible bandit exploration. *Operations Research*, *68*(4), 1132–1161.

Mansour, Y., Slivkins, A., Syrgkanis, V., & Wu, Z. S. (2016). Bayesian exploration: Incentivizing exploration in bayesian games. *arXiv preprint arXiv:1602.07570*.

Martin, K., Ryan, C. T., & Stern, M. (2016). The slater conundrum: duality and pricing in infinite-dimensional optimization. *SIAM Journal on Optimization*, *26*(1), 111–138.

Mathevet, L., Pearce, D., & Stacchetti, E. (2019). Reputation and information design. *Unpublished paper, New York University*.

Matyskova, L. (2018). Bayesian persuasion with costly information acquisition. *CERGE-EI Working Paper Series*(614).

Meurer, M., & Stahl II, D. O. (1994). Informative advertising and product match. *International Journal of Industrial Organization*, *12*(1), 1–19.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 380–391.

Monnet, C., & Quintin, E. (2017a). Limited disclosure and hidden orders in asset markets. *Journal of Financial Economics*, *123*(3), 602–616.

Monnet, C., & Quintin, E. (2017b). Rational opacity. *The Review of Financial Studies*, *30*(12), 4317–4348.

Morris, S., & Shin, H. S. (2002). Social value of public information. *American Economic Review*, *92*(5), 1521–1534.

Müller, A., & Stoyan, D. (2002). *Comparison methods for stochastic models and risks* (Vol. 389). Wiley.

Myerson, R. B. (1986). Multistage games with communication. *Econometrica: Journal of the Econometric Society*, 323–358.

Nelson, P. (1970). Information and consumer behavior. *Journal of political economy*, *78*(2), 311–329.

Orlov, D., Skrzypacz, A., & Zryumov, P. (2020). Persuading the principal to wait. *Journal of Political Economy*, *128*(7), 2542–2578.

Pagano, M., & Volpin, P. (2012). Securitization, transparency, and liquidity. *The Review of Financial Studies*, *25*(8), 2417–2453.

Papanastasiou, Y., Bimpikis, K., & Savva, N. (2018). Crowdsourcing exploration. *Management Science*, *64*(4), 1727–1746.

Rayo, L., & Segal, I. (2010). Optimal information disclosure. *Journal of Political Economy*, *118*(5), 949–987.

Renault, J., Solan, E., & Vieille, N. (2017). Optimal dynamic information provision. *Games and Economic Behavior*, *104*, 329–349.

Rodina, D. (2016). Information design and career concerns. *Northwestern University Job Market Paper*.

Rosar, F. (2017). Test design under voluntary participation. *Games and Economic Behavior*, *104*, 632–655.

Shaked, M., & Shanthikumar, J. G. (2007). *Stochastic orders*. Springer.

Shleifer, A., & Wolfenzon, D. (2002). Investor protection and equity markets. *Journal of Financial Economics*, *66*(1), 3–27.

Signori, A., & Vismara, S. (2018). Does success bring success? the post-offering lives of equity-crowdfunded firms. *Journal of Corporate Finance*, *50*, 575–591.

Smith, L., & Tian, J. (2018). Informational inertia. *unpublished mimeo*.

Smolin, A. (2021). Dynamic evaluation design. *American Economic Journal: Microeconomics*, *13*(4), 300–331.

Strausz, R. (2017). A theory of crowdfunding: A mechanism design approach with demand uncertainty and moral hazard. *American Economic Review*, *107*(6), 1430–76.

Sun, M. (2011). Disclosing multiple product attributes. *Journal of Economics & Management Strategy*, *20*(1), 195–224.

Sundaram, R. K. (1996). *A first course in optimization theory*. Cambridge University Press.

Topkis, D. M. (1998). *Supermodularity and complementarity*. Princeton university press.

Walthoff-Borm, X., Schwienbacher, A., & Vanacker, T. (2018). Equity crowdfunding: First resort or last resort? *Journal of Business Venturing*, *33*(4), 513–533.

Wang, C. (2017). Advertising as a search deterrent. *The RAND Journal of Economics*, *48*(4), 949–971.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 641–654.

Whitmeyer, M. (2018). Persuading a consumer to visit. *arXiv e-prints*, arXiv–1802.

Whitmeyer, M. (2020). Persuasion produces the (Diamond) paradox. *arXiv preprint arXiv:2011.13900*.

Zingales, L. (1995). Insider ownership and the decision to go public. *The Review of Economic Studies*, *62*(3), 425–448.