

Essays on Ignorance and Living Well

By  
Shanna K. Slank

A dissertation submitted in partial fulfillment of  
the requirements for the degree of

Doctor of Philosophy  
(Philosophy)

at the  
UNIVERSITY OF WISCONSIN-MADISON  
2019

Date of final oral examination: 8/9/2018

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

Russ Shafer-Landau, Professor, Philosophy (co-Chair)  
Michael Titelbaum, Professor, Philosophy (co-Chair)  
Harry Brighouse, Professor, Philosophy  
Paula Gottlieb, Professor, Philosophy  
Karola Kreitmair, Assistant Professor, Medical History and Bioethics

*To my parents*

## Acknowledgments

My first debt of thanks is to my advisors and co-chairs, Russ Shafer-Landau and Mike Titelbaum. It is a rare kind of luck if you find a kind but challenging, smart and supportive PhD advisor. Mine is a greedy kind of luck because I found two. They have taught me that good philosophy is found in that elusive grey area between the careful and the daring. And they have showed me that doing good philosophy can be just one part of a rich and full life. I am grateful for their guidance and, especially, for their patience.

I would like, secondly, to thank Harry Brighthouse. Harry's generosity and mentoring have been invaluable to me, as has his ability to elicit a laugh in just about any circumstance. He's taught me a lot—about philosophy, about being a philosopher, and very importantly, about being a teacher. I can only hope, in my own career, to emulate his unwavering commitment to his students' well-being.

Third, I would like to thank the UW-Madison Philosophy Department, which has been just about the best PhD department one could hope for. It is a department chock-full of very talented, well-renown philosophers. But these same folks, remarkably, are dedicated to maintaining a philosophical community that is present, grounded, congenial, and supportive. I feel very proud to be a Wisconsin philosopher.

Fourth, thanks to the very fine philosophers with whom I had the pleasure of being colleagues. My cohort: Frank Cabrera, Hadley Cooney, Justin Morton, David O'Brien, and Eric Sampson. And others: Augusta Moore and Ben Schwan. I'm very grateful for our many, many discussions, philosophical and otherwise.

Fifth, I would like to give a special thanks to Hayley Clatterbuck, Andrea Kreuzer, and Gina Schouten—for our conversations and for their endless support. Theirs is impressive company to keep—not just philosophically and professionally, but personally. I am fortunate to have their friendship.

My penultimate thank-you is arguably my greatest: to my parents. All of the things that I am glad that I have done would never have been possible without their generosity, their support, and their love. I could never repay them for these things, of course. My only hope is that I can be a parent to my son as they have been parents to me.

Lastly and most importantly, thanks to my husband, Reuben Stern. He is an impressive philosopher, but a better partner. My most demanding interlocutor, but my fiercest advocate. Ours is a very full life, with its fair share of chaos and challenge. But it is made possible by enormous quantities of laughter. Thanks for your encouragement, and for being such a good dad. I can't wait to see how our next adventure unfolds.

## Abstract

This dissertation consists of three freestanding essays. The first, *Rethinking the Imposter Phenomenon*, argues for the rationality of certain “Imposter” beliefs that are usually thought to be irrational. The second, *On the Practical Insignificance of ‘Ought Implies Can’*, argues that we’re often not in a position to rule out moral requirements on the basis of the principle ‘Ought Implies Can’ because it is usually very hard to acquire the justified belief that we cannot fulfill some candidate obligation. The third, *Authenticity, Imagination, and Transformative Choice*, argues that we can sometimes authentically choose to do something that we’ve never done before because we can sometimes adequately imagine what that thing is like.

**Table of contents**

Chapter 1:	Rethinking the Imposter Phenomenon	1
Chapter 2:	On the Practical Insignificance of ‘Ought Implies Can’	25
Chapter 3:	Authenticity, Imagination, and Transformative Choice	49
Bibliography		83

## Chapter 1: Rethinking the Imposter Phenomenon<sup>†</sup>

**Abstract:** The Imposter Phenomenon—i.e., the phenomenon of feeling like a fraud and like your successes aren’t really yours—is typically construed not just as a crisis of confidence, but as a failure of rationality. On the standard story, “imposters” have bad beliefs about their talents because they dismiss the evidence provided by their successes. Here I suggest that this standard picture could be mistaken, and that these “imposters” may actually be more rational (on average) than non-imposters. Why? Accounting for the non-talent causes of your successes may require you to lower your confidence in your talents, in which case, “imposter” beliefs are rational. I then go on to suggest a second reason to worry about the standard picture: It does not adequately address the possible role that one’s environment has in the production of the phenomenon. To give an example, I hypothesize that environments that host a “culture of genius” can alter our evidential landscape in a way that promotes the Imposter Phenomenon. Finally, I argue that my alternative picture of the Imposter Phenomenon should prompt us to opt for a conception of self-worth that is more deeply tied to virtues like intellectual humility than to relative talent possession.

### 1. Introduction

This paper is about what psychologists call the “Imposter Phenomenon,” but what you may know as “Imposter Syndrome.” Roughly, “IP” is the phenomenon of feeling like a fraud and believing that you do not belong. If you experience IP, then you experience your success as not really *yours*. Though IP is perhaps best known because of its considerable coverage in popular media, it is increasingly being discussed among academic philosophers.<sup>1</sup>

While it is good that IP is being brought to the fore of our collective awareness, most

---

<sup>†</sup> I’m grateful to Harry Brighouse, Gina Schouten, Russ Shafer-Landau, Michael Titelbaum, and audiences at the 2017 APA Pacific meeting and the Munich Center for Mathematical Philosophy’s 2017 “Values in Science” workshop for their helpful criticisms and suggestions. I owe special thanks to Erin Beeghly for her comments and conversation at the APA Pacific and to Reuben Stern for extensive discussion.

<sup>1</sup> My evidence of this is mostly anecdotal, although, see Hawley’s 2016 talk, “Philosophical Reflections on Imposter Syndrome,” and the following 2017 interview from the *What it’s like to be a philosopher* blog for examples: <http://www.whatisitliketobeaphilosopher.com/rebecca-tuvel/>.

<sup>2</sup> There is some reason to think that the labeling, though, is not entirely accurate. For example, see Hawley’s 2016 talk, “Philosophical Reflections on Imposter Syndrome,” and the following 2017 interview from the *What it’s like to be a philosopher* blog for examples: <http://www.whatisitliketobeaphilosopher.com/rebecca-tuvel/>.

popular commentary and psychological research on IP is at least on one count wrong-headed. The standard line is that individuals who experience IP (“IPPs,” let’s say) systematically cognitively err: They underestimate their own talents by ignoring evidence for those talents—namely, the evidence of their successes—and are therefore irrational.<sup>2</sup>

I think that the question of whether IPPs have irrational talent beliefs is more complicated than the standard view suggests. I will argue that there exists a rational route to some IP beliefs and thus, that it is quite possible that many IPPs are rational. Recognizing this possibility is important not only because it forces us to reconsider our perception of IPPs, but also because it leads us to countenance another possibility—namely, that it when it comes to beliefs about talent, it may be those who do *not* experience IP (“non-IPPs”) who are (on average) the irrational ones. That said, while many IPPs might be rational, it is in an importantly qualified sense: This rationality is *non-ideal* insofar as it is achieved against a socially given epistemic background that is rife with misleading evidence and ill-conceived conceptual frameworks. This point leads to the second task of this paper, which is to argue that there are environmental factors that may be contributing to the presence of IP that the standard view has failed to recognize. I’ll specifically entertain the hypothesis that environments with a “culture of genius” promote IP. Lastly, I will consider whether IPPs should ever strive to be more confident when doing so involves a departure from rationality. On my view, IPPs can be rational, but nevertheless less confident than their non-IPP (but possibly irrational) counterparts. Since diminished confidence can be harmful, this may suggest that we should prescribe irrationality. Against this possibility, I will argue that we shouldn’t sacrifice our rationality for increased confidence because there is an alternative

---

<sup>2</sup> There is some reason to worry that labeling those who experience IP as IPPs perniciously essentializes these individuals. I am sympathetic to this point and don’t advocate using this label in a public setting. Here, however, brevity demands its use.

available to us. If we revise our conception of self-worth so that it rests centrally on the possession of virtues like intellectual humility rather than relative talent possession, then we may be able to avoid some of the psychological distress associated with IP.

A disclaimer is in order. Though I can provide plausible hypotheses about the nature of IP beliefs, the way in which certain environmental factors may encourage IP, or how the cultivation of intellectual virtues may stymie psychological distress, these are ultimately empirical questions that cannot be definitively answered from the armchair. What I can do is use the existing empirical literature to argue for a novel and highly plausible explanation of the phenomenon that additionally draws our attention to the problematic assumptions at play in the current proposal on the table.<sup>3</sup> If my argument is successful, then (i) there is strong reason for empirical scientists to further investigate the cognitive mechanisms that fuel IP and (ii) in the meantime, we should reconsider our current practices premised on the epistemic defectiveness of IPPs.

## 2. IP, the standard view

Clance and Imes (1978) first identified IP in a paper describing their therapeutic work with college and professional women. In the course of this work, they noticed that many of their patients expressed insecurity about their school or work performance and experienced feelings of incompetence and fraudulence, despite having been largely successful in their pursuits. These observations prompted Clance and Imes to ask why certain

---

<sup>3</sup> I do not take my explanation of IP to be complete in the sense that it should (or is likely to) explain every instance of IP. For instance, while my explanation assumes that the affective components of IP are causally downstream of the doxastic components of the phenomenon, it is highly plausible that there are cases of IP in which the causal order is reversed, so that IP beliefs are a product of IP feelings and other attitudes. Since my model won't capture such cases, it can't fully explain IP. The *extent* to which my account is explanatory is an empirical matter that we currently cannot settle (since we don't know enough about how IP beliefs and feelings are causally related).



successful individuals but not others develop these tendencies. Since this initial exploration, three different scales have been developed to measure IP in individuals. The 14-item Harvey Scale, the 51-item Perceived Fraudulence Scale, and the 20-item Clance Scale each use a Likert-type scale to determine how much respondents identify with a series of statements.<sup>4</sup> Examples of these statements (taken from the Clance test) include “I sometimes think I obtained my present position or gained my present success because I happened to be in the right place at the right time or knew the right people,” “When I’ve succeeded at something and received recognition for my accomplishments, I have doubts that I can keep repeating that success,” and “I often compare my ability to those around me and think they may be more intelligent than I am.” Though various studies link IP to characteristic clinical symptoms (e.g., anxiety and depression), behaviors (e.g., severe procrastination and discounting of praise), and even personality traits (e.g., neuroticism and perfectionism), psychologists have used the three IP scales to argue that IP is a valid psychological construct, discriminable from these and other negative affective phenomena. As should be clear from the diagnostic tools, IP is not a matter of having it or not. Rather, one can experience some or all of the aspects of IP to a greater or lesser extent. So we should think of IP as a spectrum phenomenon and recognize that many individuals fall somewhere on the spectrum. We can think of those who identify most strongly with the diagnostic statements as the paradigm instances of IPPs.

Clance and Imes’ initial investigation into IP led them to believe that it was a phenomenon especially prevalent among women, but later research has raised doubts about this claim. A handful of small-scale studies have interrogated the claim directly and found

---

<sup>4</sup> See Harvey 1981 and 1985, Kolligian and Sternberg 1991, and Clance 1985 for development of these respective tests. In this paper, I rely exclusively on the Clance test.

that men experience IP at equal rates as women.<sup>5</sup> But the vast majority of the remaining research doesn't focus on whether IP correlates with gender. It instead investigates patterns of IP in, e.g., student populations, among academic faculty, and in members of helping professions (such as doctors and nurses), and has looked for (and revealed) correlations between IP and traits such as status in one's profession and achieving beyond what is typical in one's family. In other words, the empirical (non-anecdotal) evidence that we have on the matter of IP and gender is limited and inconclusive; we simply don't know enough to say whether on average women experience IP more often than men. This point generalizes. Given the dearth of empirical evidence, we simply can't say whether, more generally, IP tends to be especially prevalent among individuals from historically marginalized groups.

There are two points related to this one that are important. The first has to do with the fact that despite the status of our empirical evidence, the belief that women primarily experience IP persists—especially in popular discussions. The way in which this belief currently exists in the public's mind's eye is problematic, so we should be wary of associating IP with women *even if* ultimately the association is empirically vindicated. In public discourse, the belief that women “suffer” IP has been absorbed into the familiar narrative where women are dispositionally unsuited to the pressures of competition and achievement, which explains why they on average do not succeed to the extent that their male counterparts do. The general line is that IP is a condition where women are unable to acknowledge their talents and are too timid to pursue achievement at the risk of failure. But the narrative complicates: Women are not just afraid of failure; they're also afraid of success! They don't believe in themselves! The solution to the problem (it continues) is to find a way to help women correct the pathological aberrations in their thinking so that they can get out of their

---

<sup>5</sup> For studies that show men experiencing IP at equal rates as women, see Lawler 1984, Flewelling 1985, Beard 1990, and Cozzarelli & Major 1990.

own way.<sup>6</sup> To the extent that the standard view facilitates this narrative, it likely has the unfortunate effect of encouraging women who do experience IP to think that their negative experiences are a problem due to them and them alone, which is both harmful and false.<sup>7</sup>

The second point relating to the connection between IP and marginalized populations has to do with whether we should in the first place be concerned with IP. If we have no reason to believe that it is especially prevalent among individuals from historically marginalized groups, we might wonder: Why care? First, though the evidence doesn't tell us anything definitive, there may be other reasons to *suspect* that IP tracks marginalization. For example, it seems plausible that certain features of marginalization make such individuals more vulnerable to evidential distortions of the sort that I will discuss below. So despite the verdict's being out, we can—armed with plausible suspicion—still proceed in inquiring for the usual reasons. But second, our interest in IP need not to depend on IP's being predominantly experienced by members of the familiar marginalized identity groups. It is worthy of our attention in its own right, as a phenomenon that negatively affects a significant number of individuals, and that may be the product of certain contingent but problematic social arrangements and ideologies (as I argue below)—even if these are not, for

---

<sup>6</sup> Sometimes this help comes in the form of reassuring women that many famous people experience IP (Kate Winslet!). Other times it comes in the form of pithy advice:

“Lean in!”

“Be less nice!”

“Steeple instead of slouch!”

“Fake it until you make it!”

<sup>7</sup> In an interview, Clance reports that from the outset that she and Imes were concerned not to have IP be taken as “another ‘defect’ in women or a pathologizing of women” (Clance & O’Toole 1987, 160). In fact, this reason is why they were deliberate in calling this experience a *phenomenon* rather than a syndrome since the latter can connote defect or disease. Nevertheless, they themselves first link the phenomenon to women in particular, which no doubt has influenced how it has been adapted into the public discourse. Simmons 2016 also traces the intellectual origins of IP to earlier well-known 20<sup>th</sup> century psychological research programs that pathologize women—e.g., programs studying “the need for achievement” (posited as the universal source of cultural and economic development) and “the fear of success” (posited to explain why in tests women displayed greater drive to achieve, but in practice were much less successful than men).

example, the standard sexist and racist arrangements and ideologies that often (rightly) worry us.<sup>8</sup>

So: That is “the who” of IP. Let’s now turn to “the what.” We can think of the phenomenon generally as a characteristic set of beliefs, feelings, and behaviors (e.g., the *belief* that your success was obtained through luck or the *feeling* of fear of being found out). Because it will aid our arguments below, here is some additional conceptual structure for the doxastic components of IP. Using the statements from the IP scales as proxies for IPP beliefs, we can categorize these beliefs into three groups: Beliefs about *talent*—e.g., that my success is due to luck or that I am not as talented as my superiors believe. Beliefs about *fit*—e.g., that I will be discovered as a fraud or that I do not belong in my group. And beliefs about *desert*—e.g., that I do not deserve the praise and recognition given to me. For the purposes of this paper, we will focus primarily on beliefs about talent—or, *T-beliefs*. This isn’t to say that T-beliefs are the most theoretically important aspect of IP and certainly isn’t to say that they are the most salient aspect in the experience of IP.<sup>9</sup> The dialectical decision to focus on T-beliefs is largely because the standard view itself tends to focus on T-beliefs (see below) and our aim here is to criticize the standard view. For the purposes of this paper, we will also grant that there is such a thing as *talent*, which, as a cause of our successes, can be meaningfully contrasted with environmental, luck, and effort causes. Though we should be concerned about the validity of talent as a concept and of the aforementioned partition of

---

<sup>8</sup> Thanks to an anonymous reviewer for forcing me to clarify the connection between IP and historically disadvantaged groups. My own suspicion is that the “cultures of genius” that I discuss below are ultimately bound up with sexist and racist (and classist) institutions and ideologies, though in complicated ways. But since untangling these relations is beyond the scope of this paper, the relevant point is that we have independent reason to investigate IP, and therefore needn’t first do this untangling.

<sup>9</sup> For an example of a recent popular discussion of IP that focuses primarily on fit-beliefs, see Amy Olberding’s 2018 essay “The Outsider.” Olberding reports feeling out of place in Academia because of her class background, but interestingly, does not talk about this estrangement as leading to doubts about her ability to succeed in Academia.

causes, for better or worse, these are the terms of ordinary parlance in which the debate is had. So we adopt them here. Again, we should think of this choice as dialectical, motivated on the grounds that it is more fitting given the paper's aspiration to engage with the standard view on its own terms.<sup>10</sup>

Diagnostic tests indicate that IPPs have at least three different kinds of T-beliefs: There are *backward-looking* T-beliefs, or beliefs about the extent to which one's talents were causally responsible for producing past successes. These are expressed in test statements like "At times, I feel my success has been due to some kind of luck" and "I sometimes think I obtained my present position or gained my present success because I happened to be in the right place at the right time or knew the right people." There are *forward-looking* T-beliefs, or beliefs about one's talents being able to produce new successes. The statement "I'm often afraid that I may fail at a new assignment or undertaking even though I generally do well at what I attempt" is one example of this kind of belief. And finally, there are *comparative* T-beliefs, which comes in two varieties: First, those about how much talent one has relative to others (as expressed, e.g., in the statement "I often compare my ability to those around me and think they may be more intelligent than I am"), and second, those about how much talent IPPs believe others perceive them to have (as expressed, e.g., in the statement "I'm afraid people important to me may find out that I'm not as capable as they think I am").<sup>11</sup>

What does the standard narrative have to say about T-beliefs? It starts with the recognition that among successful individuals, some of them have the beliefs, feelings, and

---

<sup>10</sup> We should also recall that the claim here is that IPPs are *non-ideally* rational. Potentially one way in which it is non-ideal is that their beliefs are formed against a background framework that assumes 'talent' as a meaningful concept. But it is worth noting that IPPs and non-IPPs alike have beliefs about how talented they are, so to the extent that this is erroneous, it's not a failure that is particular to IPPs.

<sup>11</sup> These categories are not mutually exclusive, of course. One can have, for instance, a comparative backwards-looking T-belief.

behaviors characteristic of IP. And it tends to assume both that IP beliefs *cause* these feelings and behaviors and, further, that if these beliefs can be prevented or changed, then so too can the negative feelings and behaviors. A similar logic is applied more specifically to T-beliefs: These beliefs can be prevented or changed by understanding and intervening on *their* causes.<sup>12</sup> Here are Clance and Imes to illustrate:

“However, despite their earned degrees, scholastic honors, high achievement on standardized tests, praise and professional recognition from colleagues and respected authorities, these women do not experience an internal sense of success. They consider themselves to be “impostors.” Women who experience the impostor phenomenon maintain a strong belief that they are not intelligent; in fact they are convinced that they have fooled anyone who thinks otherwise. [...] In other words, these women find innumerable means of negating any external evidence that contradicts their belief that they are, in reality, unintelligent.”<sup>13</sup>

So according to the standard view, the negative IP feelings (such as not having an internal sense of success) are due to T-beliefs (“they believe they are not intelligent”), which are in turn due to IPPs discounting (“negating”) evidence of their talent. What evidence? Their successes (“degrees, honors, achievements,” etc.). The standard explanation tends to come with two further assumptions. First, that IPPs are making a *rational* mistake insofar as they are failing to count some of their evidence. And second, that this explanation accounts for how IPPs think differently from non-IPPs about their talents—non-IPPs don’t have a problem because they are able to take appropriate stock of their successes. In the next section, we’ll take aim at this explanation and its accompanying assumptions.

---

<sup>12</sup> In fn. 3 above, I mention the important possibility of IP cases where beliefs are causally downstream from the affective components of the phenomenon, which we are setting aside in our present discussion in order to focus on cases where beliefs are causally upstream of the affective stuff. We can now better see the reason for setting aside this other kind of case—viz., that our target in the paper, the standard view, itself assumes that beliefs are causally upstream of affectivity.

<sup>13</sup> Clance & Imes 1978, 1.

### 3. Reasoning about your talents

Let's begin with backwards-looking T-beliefs by considering a simple example about other kinds of backwards-looking beliefs. Imagine that you woke up in the middle of the night with terrible food poisoning. Naturally you wonder what made you sick and you begin to consider the possible causes. Well, you had two meals yesterday—lunch and dinner—so you come to think that either lunch or dinner poisoned you. Then you hear on the news that the food inspector has uncovered an outbreak of campylobacter bacteria at your favorite lunch food cart. Obviously, when you learn this information, your confidence that it was your lunch that poisoned you goes up. But what now happens to your confidence in the possibility that your dinner poisoned you? It seems that you should become less confident that dinner was poisonous. You don't have to *rule out* the possibility of a poisonous dinner altogether; after all, you could have fallen victim to a double whammy of campylobacter. But you do *adjust* your views about dinner's causal role in your sickness. With the evidence of a poisonous lunch, you come to think it is less likely that dinner was poisonous.<sup>14</sup>

It is likely that our rational backwards-looking T-beliefs should replicate the pattern in our food poisoning case. We begin from the assumption that the causal structure of the kinds of successes we're interested in—e.g., getting into college, getting the prestigious residency, publishing in the top journal—is complex in the following sense: There are a variety of agential powers (including our talents) but also non-agential powers that causally contribute to their production. We have to think about all of these causes when forming rational beliefs about the causal structure of our successes. And getting evidence that various non-talent causes played a role in bringing about those successes requires us to accordingly

---

<sup>14</sup> This example is a favorite of Reuben Stern's. I'm grateful to him for letting me borrow it.

decrease our confidence in our talent's causal efficacy from what it was prior to taking stock of the non-talent causes.

Once we grant that non-talent evidence can play an important role in determining the rationality of our backwards-looking T-beliefs, we must more seriously acknowledge that there are usually considerable non-talent causal forces at work in the production of our various successes. We tend to call the fluky forces *luck*: Had Senior Professor So-and-so not chosen to retire *this* past year, Acclaimed Institution would not have hired in my area of expertise. Had Big Wig CEO not turned out to be a dutiful Lions fan like me (who suffers from the same disappointment and dejection season after season), she likely wouldn't know me from the rest of the Adams at the bottom of the company totem pole. We usually call those forces that are more systematic, whose mechanisms are embedded within our social norms, ideologies, and institutions *environmental* or *structural*: middle class Mom and Dad paying for college, rent, application fees, and test prep classes matters for your prestigious acceptance. Uncle Buck's casual nepotism matters for your stable job. So, too, does the company's affirmative action hiring policy.<sup>15</sup> But regardless of how we ultimately systematize these forces, the uncontroversial point is supposed to be that the story about our agential efficacy in the causal world is much more complicated than our first-personal experience

---

<sup>15</sup> We might not like having to recognize affirmative action ("AA") interventions as evidence of non-talent causes since it'd be unfortunate if interventions designed to correct opportunity discrepancies or increase diversity rationally require recipients to lower their confidence in their talents. But it seems to me that this is so. To return to the analogy above, it's like AA recipients getting evidence that it was lunch that poisoned them, but non-AA recipients getting evidence that it *wasn't* lunch that poisoned them. One might insist that this can't be the rational way to think about AA policies because when operating well, they are merely tie-breakers among equally talented individuals. But if there is widespread belief that the policies aren't being ideally enacted, then individuals benefitting from AA may rationally treat their awards as evidence against their talents anyhow. Further, even *if* the policies are ideally enacted (as tie-breakers), it still isn't clear how this should boost the confidence of rational AA-recipients since this information just tells them definitively that it was something other than their talent that *made the difference* in their being awarded the scholarship, job, award, etc. For a nice discussion of what attitudes one *all-things-considered* should have towards being the beneficiary of affirmative action policies, see Gheaus 2015.



sometimes suggests—or than our psychology and perhaps ideology dispose us to acknowledge. Especially so for our socially-embedded successes, which tend to be subject to greater influence by structural causes. What’s more, if we’re being honest and paying attention, evidence of these non-talent causes is *not* hard to see as we try to explain our successes. This means that if IP backward-looking T-beliefs like “my success was due to luck” or “I’m not that talented” are explained by the fact that IPPs take stock of these non-talent factors (while perhaps others ignore them), and if the same kind of evidential reasoning that is appropriate for thinking about food poisoning is appropriate here (as it seems to be), then it is rational for IPPs to be less confident in their talents.<sup>16</sup>

How do one’s successes evidentially bear on these beliefs? It is true that presence of the effect is evidence of its causes and so a success tells you *something* about your talents. But success also tells you something about all of the other causes that produced it. Upon first discovering your food poisoning, you become more convinced that something was off about lunch *and* more convinced that something was off about dinner. The same goes for success—i.e., you should initially become more confident of *all* of its causes, and thus should become more confident of both your talents and your luck. But since these causes are neither necessary nor sufficient for success, and since the precise probabilities of these causes given our success escapes our cognitive grasp, we lack definitive evidence about *the extent* to which we are talented, about *the extent* to which we are lucky, and about whether we aren’t luckier than talented. Curiosity thus calls for a richer explanation of our success and asks us to seek out further evidence. Enter here our story about non-talent causes and the subsequent decreases in talent confidences—i.e., IPPs may attend to (and acquire evidence

---

<sup>16</sup> Whether they are irrational may also depend on whether I update my beliefs appropriately in other ways—e.g., that I don’t decrease my confidences *too* much, or that I have coherent beliefs. Since my aim here is just to identify *a* rational method of belief-updating that may fuel IP, we needn’t dwell on these other ways in which such agents can fail to be rational.

about) the perspicuous non-talent causes of their success, and subsequently (rationally) decrease their confidence in their talents.

So with respect to backward-looking T-beliefs, one's successes aren't *that* evidentially informative. But having a preponderance of non-talent evidence pushes the balance in favor of decreased confidence in one's talent. A similar story provides a rational vindication for forward-looking T-beliefs and comparative T-beliefs. In the case of forward-looking T-beliefs—i.e., beliefs indicating uncertainty about one's talents producing future successes, the story is parasitic on the one told about the backwards-looking beliefs. If you believe that a lot of non-talent (and more broadly, non-agential) causes have contributed to the production of your successes in the past, then you can reasonably infer by induction that any future successes will be similar in this way. And if you believe that many things beyond just your talent will figure in to the production of future successes, then you have some reason not to be entirely confident that *you* will continue to be successful—that is, that your talents can reliably produce future successes. Furthermore and once again, evidence of past successes does not alter this picture. Given what we know about the causal structure of those successes, it would be inappropriate to use them as inductive evidence for beliefs about *your* ability to produce future successes.<sup>17</sup>

A rational vindication of comparative T-beliefs requires a slightly more complicated story, one that relies on the idea of an evidential asymmetry. Remember that these beliefs come in two varieties—beliefs about what others think of IPPs (e.g., “My superiors probably overestimate my talents”) and beliefs about IPPs' talents relative to their peers (e.g. “I doubt that I'm as competent as my peers”). In the case of my beliefs about what others think of

---

<sup>17</sup> You might reasonably infer that you'll continue to be the beneficiary of the mechanisms of privilege that you believe helped give you your current successes (if you think these mechanisms are stable). But this belief is subtly but importantly different than thinking *you* qua *your abilities* can be successful again.

me, these beliefs may be produced by the fact that the non-talent causes of my successes may be opaque to others but transparent to me. So I come to reason that they have overestimated my talents because they (unlike me) are not aware of all of the non-talent causes of my successes—i.e., they haven't seen as much of my life as I have.

Of course, there may be cases where others are in a better position to assess the causal histories of our accomplishments than we are—e.g., when perfectionist novelists are appalled with their own work despite massive critical acclaim. But this is only sometimes true, and even when true, it is hard to argue that it's *rational* to defer. If I am of the mind that I have strictly more evidence than you about why I got where I got (as many IPPs plausibly are), then I should defer to your judgment only if I regard you as a guru of some sort. And though there are some gurus in life—maybe some of whom are our advisors—they're rare. So it seems that IPPs are at least *often* rationally justified in thinking that they are in a better position than others to causally explain their own successes.

In the case of beliefs about my talents relative to others', I also have better access to the non-talent causes of my own successes than I do to the non-talent causes of others' successes. This, again, is for the simple reason that I'm better positioned to know the facts of my own history than those of others', more so the more they are strangers to me. This asymmetry might lead me to decrease my confidence in my own talents according to the lunch-dinner reasoning from the previous section, but not make similar adjustments in my confidence of others' talents. The result? A negative interpersonal comparison. In the next section, we'll elaborate on this point when we see how various kinds of environments might exacerbate these evidential asymmetries.

#### 4. Imposter environments

The aim so far has been to show how the three types of IP T-beliefs can be rational. This picture relies on the assumption that the successes usually at stake in IP cases are causally complex, produced not just by the agent's talents, but also by a variety of non-talent causes that vary in their causal power. The hypothesis is that IP T-beliefs can be rational when formed on the basis of the evidence of the non-talent causes of one's successes. One obvious lesson from the discussion so far is that the standard view is flawed the extent to which it focuses on the evidence of successes and ignores all other kinds of evidence relevant to forming beliefs about our talents as we try to explain our successes. We now recognize that other kinds of evidence can matter and that the question of IPPs' rationality very much depends on whether they are paying attention to this other evidence. A second flaw of the standard view is that it presents an overly individualistic picture of IP, often ignoring some of the possible ways in which environments can play a role in promoting IP. In this way, the standard view runs afoul of the familiar feminist lesson that some phenomena cannot be adequately appreciated unless we widen the scope of our view, shifting focus from individual psychological mechanisms to the social structures through which those mechanisms operate. In this section, I motivate this call to think more about the social underpinnings of IP by providing one illustration of the way in which environmental factors might be especially conducive to the incidence of IP. The hypothesis is that one kind of proximate environment—viz., one that harbors a “culture of genius”—might promote IP by altering the evidential landscape, thereby enabling these environments to exploit the rational talent-belief forming processes described above.

Given the picture developed in the previous section, we can expect such an exploitation to look as follows. IP targets backwards-looking, forward-looking, and

comparative T-beliefs. I have argued that IP backward-looking and forward-looking T-beliefs can be rational when they are formed on the basis of evidence of non-talent causes of success. From this story, we can infer that an environment that increases the appearance of non-talent causes of one's successes will be one in which people who appropriately take stock of their evidence will have lesser opinions of their talent. Regarding comparative T-beliefs, the earlier claim was that discounting one's own talents but not others' can be rational when there is an evidential asymmetry between the non-talent causes of my successes and the non-talent causes of others' successes. We can expect environments that widen this asymmetry to increase these negative interpersonal assessments. So if a culture of genius both increases the appearance of non-talent causes and promotes an asymmetry in appearance of non-talent causes between self and other, then we should expect it to promote a higher incidence of IP. Does a culture of genius do these things?

The notion of a culture of genius comes from social psychologists Carol Dweck and Mary Murphy, who study the individual cognitive and behavioral effects produced by an organization's theory of intelligence—i.e., the views about intelligence held by most of the members of an organization and expressed in the norms and values that give it structure.<sup>18</sup> In their work, Dweck and Murphy show that whether an organization has an “entity theory” of intelligence (i.e., views intelligence as fixed and innate) or an “incremental theory” (i.e., views intelligence as malleable and teachable) alters the beliefs and behaviors of both its members and those seeking membership. These individuals, for example, display changes in how they present themselves to others, as well as in how they evaluate others—trying to signal as well as seek out in others the traits believed to be evidence for possession of intelligence of the relevant sort. On Dweck and Murphy's framework, the primary distinguishing feature of a

---

<sup>18</sup> Murphy and Dweck 2010.

genius culture is that its background beliefs, norms, and values express a commitment to the view that intelligence is fixed and innate rather than malleable and teachable (i.e., that it holds an entity theory). Such cultures are further distinguished by their acceptance of the idea that there are true or genuine successes in their organization (or their field, more generally) and that successes of this kind require talent, so-conceived.<sup>19</sup>

What makes it plausible that genius culture promotes IP is that its central commitment to intelligence as fixed and innate has consequences for the way in which effort is understood in such environments. One common conception of effort is that it is not an ability, but a tool that we use to develop our abilities. On this picture, effort is roughly democratic—each of us is endowed with roughly the same capacity for it. It is voluntaristic—it can be “switched on” for deployment through sheer force of will. And it is renewable—there is no limit on how hard or for how long we can try (apart from the limits set by choosing to give up). Effort, thus, is *contrasted* with our talents, which have none of these three traits, but are instead the idiosyncratic collection of capacities whose distribution in kind and amount is subject to greater variation across the human population and whose quantities in individuals are roughly finite and unalterable. Because the entity theory of intelligence conceives of intelligence (the relevant talent here) as fixed and innate and as something that cannot *through effort* be much developed, it aligns itself naturally with this dualist conception of effort and talent as two very different sorts of things at work in our agential economy. And by treating effort as something distinct from and opposed to talent,

---

<sup>19</sup> An anonymous reviewer asks: Where, exactly, are the cultures of genius? This is a good question that I can’t satisfactorily answer, but that is worth flagging. Clearly, Dweck and Murphy’s research identifies cultures of genius in various corporate environments. Recent work in Leslie et al. 2015 and Meyer and Leslie 2015 investigate cultures of genius in academia (and find that this culture appears to be especially prevalent in philosophy). Are there more systemic patterns that we can identify and use to predict what gives rise to cultures of genius? This is an important question worthy of future empirical investigation.

the entity theory makes effort out to be a non-talent cause of success. Thus, just like luck and structural forces in our earlier examples, effort operates as evidence against one's talent. Noticing of yourself that the achievement of a particular success required a significant amount of effort on your part is evidence of a piece with learning that your lunch was poisoned: It forces you to revise your confidence in the causal contribution of your talent just as you revise your confidence in dinner's being poisonous.

So if it is true that genius culture's acceptance of an entity theory of intelligence will make effort appear as evidence of a non-talent cause of success within such cultures, then genius culture looks primed to promote both backwards-looking and forwards-looking T-beliefs. This conception of effort may additionally lead to a widening of the kind of evidential asymmetry that gives rise to comparative T-beliefs. Genius cultures may promote *this* effect by way of encouraging the practice of hiding effort—i.e., of attempting to minimize the extent to which you appear to others as needing to exert a lot of effort. This may lead to a situation where agents' own efforts are highly salient to themselves, but where agents are blind to others' efforts. And when effort counts as non-talent evidence, this means that one gets a particular kind of evidence against her talents that she rarely gets against others'.

Why would genius cultures promote effort-hiding? Because when individuals think that legitimate successes and legitimate membership attach to the possession of an innate and fixed brilliance, they have serious incentive to signal the possession of such brilliance. But since exerting effort is antithetical to such brilliance, signaling possession requires hiding effort. In such circumstances, one has special access to her own effort, but this effort looks even more unusual in interpersonal comparisons since it has become more difficult than it

might otherwise be to observe the efforts of others.<sup>20</sup> Indeed, Murphy and Dweck’s research lends empirical support to this hypothesis since exactly what they show is that individuals alter their self-presentations in order to align with what they perceive to be the theory of intelligence dominant in their domain. In genius environments, this means shifting one’s behaviors so as to give the appearance of having natural, raw, and outsized intellectual firepower.

To be clear, the worry about genius culture is not that it looks capable of promoting IP T-beliefs, which we want to avoid because T-beliefs are bad. The lesson of the previous section was that T-beliefs may very well be the right beliefs to have as far as rationality is concerned. And for all we know, they may also turn out to be the kind of beliefs that when held by a good number of folks and when revered instead of ridiculed, make for more amicable and cooperative social arrangements. Indeed, we might find some environments *desirable* were they to make non-talent causes more salient to believers—for instance, were they to make it easier to recognize various privileges as such. The worry about genius culture, then, is more specifically that it makes something out to be a non-talent cause that almost certainly isn’t, thereby creating an environment in which one can justifiably infer false IP T-beliefs.<sup>21</sup> In addition, since, by hypothesis, IPPs are more epistemically sensitive than

---

<sup>20</sup> One might balk here at the suggestion that effort-hiding is prominent in such cultures, thinking perhaps of the “careerist” proclivities prominent in many professional domains (including philosophy). It is important to distinguish between effort and what we might call, by contrast, hard work. *Effort* is the strenuous cognitive exertion put forth when one finds accomplishing a task difficult, whereas *hard work* is the devotion of most of one’s time and attention to a task. To my mind, most things that require effort also demand hard work, but not everything towards which one works hard requires effort. In this way, widespread effort-hiding is consistent with widespread hard-work-signaling—e.g., being very public about one’s long hours, grueling travel schedule, bursting inbox, etc.

<sup>21</sup> The idea that effort is much more a talent than non-talent is one that is made often in discussions about egalitarian justice. For one nice example, see Brighthouse and Swift 2013, where they discuss some of the complicated ways in which effort looks to be a kind of talent, and further, in which talent and effort can both depend crucially on environment (16-18). Why we tend to ignore the fact that effort itself is a kind of talent is difficult to say, though one guess to venture is that recognizing



their non-IPP peers to the non-talent causes of their successes, genius culture disadvantages them in particular.

##### 5. The turn to intellectual humility

The story I have constructed about IP is meant to be emancipatory. IPPs are not necessarily irrational since it is possible that they are often simply integrating their evidence as they should, some of which may be distorted if they find themselves in certain kinds of environments. So we need not think of them—and they need not think of themselves—as defective. Perhaps for the individuals who experience IP, seeing their views in this light could ameliorate some of the negative affective and behavioral aspects of the phenomenon. It is important, however, to be realistic about the limitations of the alternative picture painted here. The suggestion is that there is a way to explain why IPPs are less confident than some of their equally successful peers such that IPPs are not irrational on a certain measure. But there still may be other ways in which IPPs are irrational or mistaken. (We haven't, remember, evaluated IPP desert- and fit-beliefs. We've also acknowledged along the way various global worries, e.g., about the terms in which these beliefs are put or about whether they are, in fact, caused by assessments of evidence rather than by certain affective states.) Nevertheless, our efforts, though modest, are important since they draw attention to our default orientation towards the self-assured and the unsure, which assumes that there must be something doxastically amiss with being less confident (despite having success). They also draw our attention to our tacit sanctioning of a conception of the successful as seamless integrators of their successes, which somehow makes us stubborn to the possibility of the successful as apprehensive interrogators of the background conditions that enable

---

effort as such doesn't easily square with the American ideology of success as available to anyone who is willing to work for it.

their successes.

Painting IP in this light does not settle whether we collectively should do what we can to reduce the incidence of IP, for we must not dismiss the serious negative behavioral and affective aspects of IP—i.e., the pain and anguish that can accompany even a justified lack of confidence in one’s talents. Indeed, we might even have the worry that the revised understanding of IP can make these aspects of IP worse since (at least on one gloss) it amounts to insisting to IPPs that they’re exactly as untalented as they’ve feared all along!<sup>22</sup> If it turns out that the psychological connection between well-being and confidence (read: having high enough talent credences) is so strong that having rational talent beliefs usually threatens our well-being—even once we see such beliefs as rational—should we simply be irrational?

To be clear, the underlying suggestion of the argument is that most of us probably aren’t as talented as some non-IPPs would believe and that some of us are better at realistic self-appraisal. It is *not* that IPPs are indeed less talented than their non-IPP successful peers. Still, we can wonder whether we’d better off living the lives of irrational non-IPPs. On this matter, I contend that we can have our cake and eat it, too—i.e., that we can be happy and rational—if we opt for a different conception of self-worth than the one that gives rise to this dilemma.

The conception of self-worth that seemingly fuels the difficulty in being content with being less confident is one that gives pride of place to talent possession. If you are under the impression that how valuable you are largely depends on how talented you are, and come to believe that you are not *that* talented, then no wonder that it might be small consolation to be told that you at least believe rationally. There’s a lot to be said against this conception of

---

<sup>22</sup> Thanks to Erin Beeghly for first pointing out this way of seeing the revised view.

self-worth (which tends to be prominent in hyper-competitive and “meritocratic” cultures like the U.S.’s), but appraising this conception would take us too far afield. Instead, we can simply point to an alternative that appears to be a promising solution to the present problem.

The alternative conception of self-worth is one rooted in the long-standing tradition of virtue ethics and underpinned by the possession of virtues. On this vision of self-worth, your value is bound up with being virtuous, including—and important for our purposes—possessing epistemic virtues like humility. To the virtuous agent, the usual sorts of socially embedded institutional and professional achievements that provide access to greater positional goods and increased social standing don’t hold a candle to what is genuinely valuable—for instance, epistemic goods such as truth and understanding. To the extent that an agent regulated by this ideal of self-worth does care about his own talents, he does so with intellectual humility. He faces his epistemic limitations honestly, is undefeated by shortcomings, and is hopeful that he can improve—since he believes that his talents are valuable insofar as they help realize these coveted goods and recognizes that these attitudes towards his talents facilitate this pursuit.<sup>23</sup> Were IPPs to embrace this conception of self-worth, they might come to see their T-beliefs in a different, less negative light.<sup>24</sup>

---

<sup>23</sup> Here I follow Alessandra Tanesini’s account of epistemic humility (Tanesini 2016). Her account of epistemic humility is notable because it does not require the humble individual to *ignore* their talents and successes. Very often the advice is given to IPPs to stop caring about their talents and to stop engaging in interpersonal comparisons, which I find has the pernicious implication that non-IPP’s must already be doing these things. But there is no evidence to support the idea that non-IPP’s care less about their talents or make fewer interpersonal comparisons. The advice is bad for the further reasons that (1) it’s very psychologically demanding to withdraw care and refrain from comparisons altogether and (2) there remain some important reasons to engage in these practices.

<sup>24</sup> Though, as an anonymous review points out to me, even this account of self-worth may leave one feeling badly about one’s (lack of) talents given that one may (appropriately) care so much about the realization of certain epistemic goods so as to lament not being able to realize them. Does this put the virtue-based account of self-worth right back on a par with the talent-based one? I think not. If we’re damned to bad feelings any which way, those that allow us to preserve our rationality *and* that reflect recognition of the loss of something genuinely valuable are the better bad feelings.

## 6. Conclusion

I have suggested that the Imposter Phenomenon may not necessarily be what we think it is. It is possible that IPPs have rational beliefs about their talents because of the preponderance of evidence of the non-talent causes of our successes. Moreover, some environments—like those that harbour a culture of genius—may exacerbate the presence of IP by altering the evidential landscape, making some things look like evidence of a non-talent cause.

If my picture of IP turns out to be right, then we must revise our thinking in several ways. First, and obviously, we should revise our thinking about IPPs. But secondly, we might also revise our thinking about non-IPPs. The standard view assumes that these folks have the right attitudes; they have no problem internalizing their successes as their own! But if, as we have been imagining, evidence of non-talent causes is rampant—that is, if the world is constantly presenting us *all* with clues telling us that lots of non-talent stuff is lending us a causal hand—then it seems that we all shouldn't be *so* confident that our talent is such a powerful and central causal player in our successes. Perhaps then those who are so confident are not paying attention to their evidence or are not incorporating it properly into their reasoning.

Finally, we should reevaluate our thinking about how best to alleviate the presence of IP. If the causal mechanisms that lead to IP are not just located in the individual but are also found in her environment, then we should look not just to individual interventions, but environmental ones as well. Since we have seen that, in particular, environments with cultures of genius may aggravate the incidence of IP, we should perhaps begin by directing our attention to interventions that would diminish the existence of such cultures. One way

of doing this might be to reevaluate our thinking about the nature of talents, generally, and intelligence, in particular.

## Chapter 2: On the practical insignificance of “ought implies can”<sup>†</sup>

**Abstract:** The principle “ought implies can” implies that if you cannot do something, then it is not the case that you ought to do it. Can we use this principle in practical deliberation to rule out purported moral obligations on the grounds that we cannot fulfill them? I argue that we typically cannot because it is usually very hard to acquire the justified belief that we cannot fulfill some candidate obligation. This follows from the general fact that it can be very hard for us to discern our own inabilities.

### 1. Introduction

Figuring out what morality requires of *me*, individually, is hard. It is even harder to figure out what justice requires of *us*, collectively. Many philosophers have thought that there is at least one principle that helps make these notoriously difficult problems slightly easier: “Ought implies can” (OIC). With this principle in hand, we can rule out at least some of the many possible moral constraints on action by identifying those that we cannot follow. For example, if we cannot (in the right sense) adhere to a principle of justice that permits inequalities only to the extent that they benefit the least well off, then that principle of justice is not, properly speaking, one that demands anything of us.<sup>25</sup>

In order to use OIC in the service of winnowing down some of our possible moral obligations, we must say exactly what “can” means and implies. To do this, the usual *modus operandi* is to analyze the “can” in terms of some other concept. Some have pursued the idea that we ought do only what we are *able*.<sup>26</sup> Others have thought that we ought do only what is

---

<sup>†</sup> I’m grateful to Harry Brighouse, Dan Hausman, Russ Shafer-Landau, Reuben Stern, and Mike Titelbaum for their helpful comments and discussion.

<sup>25</sup> See Rawls 1971, Wiens 2012.

<sup>26</sup> See, e.g., Schwan 2017 and Vranas 2007.

*feasible*.<sup>27</sup> Still others have argued that we are only required to do those things that are *accessible* to us.<sup>28</sup> The standard next step is then to provide a further analysis of one's preferred intermediary concept: some action is feasible, accessible, etc. just in case such and such conditions obtain. But whatever intermediary concept one chooses, the aim of the analysis is always the same—to provide the contours of the “can.” OIC says that the set of required actions is a subset of the set of all logically possible actions. The name of the game is to say which kinds of considerations set the boundaries of that subset.<sup>29</sup>

I think of this approach as taking the theoretician's stance on OIC. It aims at a general theory of OIC, the point of which is to improve our understanding of the nature of our obligations and perhaps in certain circumstances to help us evaluate whether some person failed to meet their obligations. In this paper, I would like to think about OIC from the perspective of the moral deliberator—i.e., the perspective of someone who is contemplating what is morally required of her as she makes a choice. So rather than asking when OIC lets agents off the moral hook, I would like to ask when agents can justifiably let themselves off the moral hook by judging some possible moral obligations to be in violation of OIC.<sup>30</sup> What use, if any, does OIC have as we try to lead good lives?

In focusing on what I call the *deliberative* aspect of OIC, we focus on what the agent needs to think about when determining what she can or cannot do as she seeks to get a

---

<sup>27</sup> See, e.g., Brennan and Southwood 2007 and Gilabert and Lawford-Smith 2012.

<sup>28</sup> See, e.g., Lawford-Smith 2012.

<sup>29</sup> This has been the central question for those who accept OIC, but also for some of those who do not. (See, e.g., Brouwer 1969, Graham 2011, and Saka 2000.) If I am inclined to reject OIC, but am also inclined, for instance, to believe that I cannot be morally obligated to sprout wings and fly, I have to say how these two beliefs are consistent; doing so usually includes saying something about the “can.”

<sup>30</sup> Here's how we'll think about the relation between requirements and actions. For morally conscientious agents, what is required of us determines to an important extent which actions we attempt. If we have some moral requirement R, since such requirements are understood as action-guiding, R is going to provide instructions as to which actions I must, may, or must not take. In this way, determining my abilities matters for my accepting or rejecting requirements to the extent that such requirements demand that I take certain actions.

better idea of what she should do. This is not entirely independent from the theoretician's problem, since the nature of the "can" in OIC certainly implies constraints on what the agent should consider as she determines what she can do. But when the deliberator looks to *apply* OIC in her moral-practical reasoning, she not only needs to have a view about what the correct interpretation of "can" is, but also must be able to say of herself whether she, on the particular occasion in question, is able to do what is purportedly required of her.<sup>31</sup> That is, she must have beliefs about what she can or cannot do, given her best theory of "can," in order to legitimately use OIC to rule out certain moral directives. More precisely, the agent must have a justified belief that she cannot do the thing in question.

In this paper, I will argue that this is a very tall order for most ordinary deliberators. Once we take careful stock of what an agent has to judge in order to form a justified belief that she cannot do something, we shall see that it is probably very difficult for her to form a justified belief that she cannot do something—at least when it comes to many of the actions that would be required of her by some of the most widely discussed candidates for moral directives. Thus in many cases of moral-practical reasoning, it is hard to be justified in ruling out some moral requirement on the grounds that it violates OIC. The upshot of this is unclear. Should considerations of what we can (and cannot) do play no role as we engage in moral deliberation? Or are there right and wrong ways to incorporate uncertainty about what we cannot do as we settle what is (morally) right to do? I briefly consider these questions in the last section of this paper.

---

<sup>31</sup> Since I am interested in the deliberative question, "What morally ought I do?" I will talk about *moral-practical deliberation*. I will not take a stand here on whether this question is independent of a more general question, "What ought I do?" or whether there are any instances of deliberation in which moral considerations are not germane.



## 2. Preliminaries

Let us set the stage with three preliminary remarks. First, the broader dialogue within which I will situate my discussion is one where what is at stake are the demands of justice. These demands are thought to be a subset of the demands of morality. They may apply to individuals, groups, or institutions, and may call for individual action or collective action. This complicates things a bit since I, the moral deliberator, must consider not only what *I* can do as I try to apply OIC, but also what *we* can do, including, e.g., what kind of social arrangements we can construct. To make things simpler in the present context, I limit consideration to the individual case and often use examples that count straightforwardly as moral requirements but perhaps not as requirements of justice. If it's hard to determine what *I*, the individual, can do, it's plausibly even harder to determine what *we*, the collective, can do.

Second, let us clarify what we mean when we say that OIC is “requirement-blocking.” Political philosophers (among others) make use of the distinction between evaluative principles and action-guiding principles. Principles of the first kind tell us what is just or good, but they do not necessarily offer advice about what to do. Principles of the second kind (*directives* or *requirements*) *do* guide action—i.e., they demand that we perform (or at least pursue) certain actions. It is the latter sort of principle that is thought to be subject to OIC. I remain silent here on the nature of the relationship between the two types of principles.

Finally, while many philosophers have contributed to the discussion about the restrictions on the demands of justice, I focus on a particular exchange between David Estlund and David Wiens as I develop my position. Partly, this is because the two are paradigmatic representations of either side of the debate in this literature. Estlund is the

“idealist” or the “ideal theorist,” and contends that (i) there is value in pure intellectual exploration of the ideal of justice and (ii) the demands of justice are *aspirational* (in the sense that they are not sensitive to considerations of feasibility). Wiens is the “realist” or the “nonideal theorist,” and contends that no exploration of questions of justice is appropriate unless one takes seriously the many practical limitations that we face towards meaningful moral action in the real world. By using their discussion of OIC, we can see how the debate over OIC is often ultimately a debate about how *ideal* or *grounded* morality is—or put differently, how *demanding* morality is.

Anchoring the discussion in their exchange is also nice insofar as it enables us to help ourselves to their combined treatment of “can,” which is both admirably clear and relatively neutral. This latter feature justifies some hope that our case against the practical significance of OIC is not parasitic on any particular analysis of “can”—i.e., that it goes through successfully (perhaps with some finesse) if you swap in your favorite analysis of “can” for the one assumed here.<sup>32</sup> Still, it is worth noting that this generality cannot be demonstrated here (since this would require surveying the many treatments of “can” in the literature). As such, we should think of this paper as an important first chapter in the case against the practical significance of OIC, rather than the whole story.

### 3. Estlund and Wiens on OIC

With these preliminary remarks in place, we can turn to Estlund, who is particularly interested in whether OIC permits our motivational incapacities to limit the demands of

---

<sup>32</sup> One might think that our uncertainty about what the “can” means is itself an obstacle to using OIC justifiably in deliberation since I might not be able to justifiably believe that I cannot do something if I am unsure what it means exactly to say that I cannot do it. I won’t address this possibility here and leave it open whether this consideration is also a problem for the practical import of OIC.

justice. He argues that it does not. Estlund thinks that what we can do is best understood in terms of what we are *able* to do, and his analysis of “can” thus turns on his analysis of “is able.” He defends the following analysis of “is able”:

[OIC<sub>E</sub>] A person is able to (can)  $\phi$  if and only if, were she to try and not give up, she would tend to successfully  $\phi$ .<sup>33</sup>

This analysis, says Estlund, will get the normatively appropriate extension of OIC: it will permit some, but not too many, constraints on what requirements there are. For example, according to OIC<sub>E</sub>, I am not obligated to fly because no matter how hard I try (even if I were never to give up), I would not tend to succeed at flying. But according to Estlund, OIC<sub>E</sub> does not block moral requirements on the basis of motivational deficiencies. For example, say that I am not motivated to send a check to Oxfam to help alleviate extreme poverty. There may be any number of reasons why I am not so motivated. Perhaps I’m not motivated because I do not believe that I am obligated to do so. Or perhaps it is because I am consumed with worry about my own financial security that despite believing that I’m obligated I cannot bring myself to write checks month after month. Nevertheless, I am not off the hook if I don’t intend to send the check, nor am I off the hook if I can’t will myself to send the check because either way, were I to try my best, I’d surely succeed at writing those checks. So by including “trying” language, Estlund believes that he analyzes “can  $\phi$ ” such that lacking the motivation to  $\phi$ , regardless of the reason, never gets in the way of being able to  $\phi$ .

Wiens, however, argues that there *are* cases where an agent can’t  $\phi$  because she lacks the motivation. Specifically, he thinks that there are some cases where an agent gives up trying to do something—and in this sense, ultimately lacks the motivation to do it—and

---

<sup>33</sup> Estlund 2011, 213.

because of this, is unable to do that thing. Wiens characterizes these cases as ones where an agent makes what he calls *a good faith effort* to do something, but is overwhelmed by countervailing motivations against doing it—motivations that ultimately lead the agent to give up. You may be moved to vocalize your dissent against the cruel dictatorship (to use his example), but each time the opportunity arises you become aware of the threat to your life that your protest will incur. Time and again, you fail to express your dissent, despite efforts to suppress your fear of the likely consequences. By Wiens’ lights, we should count these kinds of cases as cases of inability. He thus offers a revision of Estlund’s analysis:

[OIC<sub>WE</sub>] A person is able to (can)  $\phi$  if and only if, were she to repeatedly make good faith attempts to complete a sequence of acts that conduces to  $\phi$ -ing, she would tend to successfully  $\phi$ .<sup>34</sup>

The idea, then, for Wiens is that the “not giving up” component of Estlund’s analysis permits too many abilities. Wiens acknowledges that Estlund is primarily moved by cases where someone initiates an attempt to do something, but does so half-heartedly or otherwise gives up too quickly. Maybe you flirt with the idea of protesting the regime, but quickly change your mind when you see the perks of joining ranks. And he agrees with Estlund that we should say that giving up so easily like this is not sufficient for being unable to do something. But, Wiens says, “the never give up clause” would also capture cases where an agent cannot help but give up and so it is too permissive a conception of ability. We should instead adopt the “good faith effort” qualifier.

For the purposes of this paper, we will accept Wiens’ revision of Estlund’s analysis (hereafter, “WE”).<sup>35</sup> Again, we do so because it provides a good test case for our thesis about the difficulty of forming beliefs about our abilities. So we will accept (roughly) that

---

<sup>34</sup> Wiens 2015, 8.

<sup>35</sup> We identify it this way rather than as Wiens’ analysis because he doesn’t endorse it as the correct analysis of *ability*; rather, he thinks that *if* you are going to take on the handful assumptions about an *ability* analysis that Estlund takes on, then this is the most plausible version of such an analysis.

you can act just in case you tend to succeed at acting when you make a good faith effort to do so.

#### 4. Believing whether I can

Now that WE gives us a better idea of the conditions under which an agent counts as able or unable to do something, we are in a position to assess how difficult it can be for an agent to form justified beliefs about her own abilities in a particular instance of moral-practical deliberation.<sup>36</sup> To aid in this task, let's consider one of Wiens' examples: Claudia, whose goal is to write a book.

When Claudia first attempts to write a book, she realizes that doing so is much harder than she anticipates. But she doesn't give up. She continues to write, developing a character here, a plot point there. After a year of writing, there is still no novel. Claudia becomes increasingly dejected, her self-confidence plummets. Eventually she abandons her book. But then a few years later, Claudia comes up with a new book idea and she renews her intention to write. Just as before, she starts out motivated and makes small progress here and there. She asks her friend in the publishing business to give her some feedback, which is largely critical. After this, she once again begins to feel hopeless and defeated, and these feelings of despair lead her to abandon the project once and for all.

According to WE, Claudia is unable to write a book. She has repeatedly made good faith attempts to complete a sequence of acts that conduces to writing a book, but nevertheless continues to fail. Ultimately her countervailing motivations—hopelessness and a crisis of self-confidence—lead her to give up. So though it is the case that were she to try and not give up, she would write a book, it is also true that she cannot help but give up. But now let us imagine Claudia when she is initially deliberating about the possibility of writing a book. She wants to decide whether she should adopt such an intention. She judges book-writing to be a sufficiently valuable endeavor and she can tell of herself that she *wants* to write a book.

---

<sup>36</sup> We will always assume that an agent must have justified beliefs about her abilities in order to appropriately deploy OIC, but from this point forward, for ease of exposition, I will drop the “justified” modifier.

But she also wants to know: *Can* she write a book? Claudia recognizes that she shouldn't intend to do *all* of the things that she'd like to do—e.g., that she shouldn't intend to dunk a basketball from the three-point line—since some of these things seem obviously impossible to her. But for many of the things that she wants to do, it is just not clear to her whether she could succeed. How can she figure out which are the things that she should attempt?

By WE, there are at least three things that Claudia needs to take stock of in order to figure out whether she can write a book: (a) the sequence of actions conducive to  $\phi$ -ing; (b) whether she can perform each of those acts in the sequence that conduces to  $\phi$ -ing; and, (c) how trying will bear on her chances of successfully completing the sequence of actions that conduces to  $\phi$ -ing. Let's consider these in order.

*a. The sequence of actions conducive to  $\phi$ -ing*

According to WE, an agent's ability to  $\phi$  depends in part on her repeated good faith attempts to complete a sequence of acts that conduces to  $\phi$ -ing. In order for the agent to justifiably infer that there is some sequence that she can complete, she must first be able to identify—at least roughly—what some enabling sequence looks like. For any given  $\phi$ -ing, there are likely to be a number of different sequences that are conducive to  $\phi$ -ing. But while it is true that believing that you can  $\phi$  requires you to identify only one such sequence, for some  $\phi$ s, it can be very difficult to grasp even just one such sequence (even if there is one). If you cannot in the first place grasp *how* some action can be performed, then you cannot address the ultimately important question of whether you can perform it.

There are, of course, many cases where we do not have any trouble identifying how to do something. For example, consider the set of sequences of actions that conduce to satisfying my current hunger. There are many: There is a sequence in which I walk over to

the drawer where I keep the take-out menus, realize that I crave Indian more than Chinese, and then telephone the restaurant. Alternatively, there is another (slightly more complicated) sequence where I boil a pot of water and then add spaghetti, slice an onion, simmer some tomato sauce, grate some Parmesan, and mix it all together in an ordered fashion. And there are others still. Alleviating my hunger is the kind of  $\phi$ -ing for which there is no problem coming up with several sequences of events that I take to be conducive to its successful performance.

But now consider Claudia's case. Claudia needs to have some idea of what sequence of actions conduces to writing a book, but book-writing isn't like satisfying hunger in this regard. For the complete novice, there is not an obvious step-by-step guide that, if followed, will guarantee or come close to guaranteeing success in writing a book. There might be some actions that seem to be an obvious part of any sequence conducive to this particular action: "identify a plot," "develop characters," "write the first chapter," etc. But there are likely to be others that are less obvious and perhaps idiosyncratic to the agent attempting to  $\phi$ : "learn to be patient with writer's block," "determine your best sources of inspiration," "build relationships with other writers in order to learn more about the writing process." Since it is plausible that Claudia has no idea which of the many possible steps are essential to the successful writing of a book for her, it's plausible that she won't be able to find one even if, as a matter of fact, she would be successful were she to try. That is, not only are there many ways to go about writing a book, but of those many ways, it might be very difficult for Claudia at the outset of her deliberation to identify which is likely conducive to *her* success.

So here is our first lesson: In order to form a judgment about whether she can  $\phi$ , an agent needs to identify a sequence of actions that likely conduces to her successfully  $\phi$ -ing, and this can be tricky (even when there is such a sequence).

*b. Whether the agent can perform each action in the sequence that conduces to  $\phi$ -ing.*

Now let us imagine that Claudia manages to find a sequence that she believes will result in her writing a book. Perhaps she spends time reading autobiographies of writers and attending meet-ups for local writers, so that she has more than just an idea for a book. She has a possible plan. She believes that if she writes four hours a day, shares her work and gets feedback, commiserates with fellow writers when she is stuck or feeling defeated, and so on, then she will succeed. But now Claudia is faced with yet another hurdle. For each of these actions in the sequence, she must establish that these, too, are things that she can do. Can she really sit for four hours a day, uninterrupted? Can she muster the courage to ask for feedback on the parts that she has written, even though she might find this process excruciating? These are hard questions, too! And in order for Claudia to find a sequence that serves as grounds for believing that she can write a book, she not only must identify a sequence that would very probably result in the successful completion of the book, but also a sequence whose many steps are executable.

In effect, what (b) requires of an agent is that she applies the analysis of ability to each action in the sequence that conduces to  $\phi$ -ing, so according to WE, whether she would tend to succeed at each action in the sequence, were she to make a good faith effort to do so. The process of figuring out whether you are able to perform some action is an iterative one. More abstractly, imagine that  $\psi$ -ing is part of a sequence of actions conducive to  $\phi$ -ing. In order to determine whether one can  $\psi$ , one must think about whether there is a sequence of actions that conduce to  $\psi$ -ing. And perhaps in this process, she might identify  $\zeta$ -ing as conducive to  $\psi$ -ing, which then will lead her to think about whether she is able to  $\zeta$ . We could run out of Greek letters!



In theory, this process may proceed infinitely; but in practice, there's probably no infinite regress to fear. I suspect that we usually eventually arrive at actions for which there are no mysterious or complicated sequences of further actions required in order to accomplish the action in question. For example, eventually, in the course of this iterative procedure, we will reach actions such as “boil a pot of water,” “choose a name for the protagonist,” or “show up to the protest.” But it is hard work deconstructing one's sequence so that the tasks one faces are these *actionable* actions. And for many possible plans, it is highly likely that we'll reach an obstacle somewhere in the inferential process before we arrive at any set of actionable actions that we believe that we can complete.

Our second lesson, then, is that once an agent has identified an appropriate sequence of actions that conduces to  $\phi$ -ing (or something roughly close to an appropriate sequence), she must reflect on her abilities with respect to the further actions in this sequence. Like the task of finding a sequence conducive to  $\phi$ -ing, this task is not impossible. But it is yet another tall order.

*c. How trying bears on success*

A third thing that an agent has to think about as she assesses her ability to  $\phi$  has to do with the “tends to succeed” clause of WE—viz., how the very act of trying repeatedly to  $\phi$  bears on the tendency to successfully  $\phi$ . Wiens and Estlund cash out the “tends to succeed” clause in terms of meeting a threshold of possible worlds where one successfully  $\phi$ s. So for WE, you tend to succeed *iff* of the worlds where you make a good faith effort, you successfully  $\phi$  in sufficiently many of them.<sup>37</sup> It might be perfectly reasonable for an agent in

---

<sup>37</sup> You might worry that ability—or at least, some abilities—does not require “sufficiently many” of the good-faith efforting worlds to be successfully  $\phi$ -ing worlds. Instead, one might think that all that is required for (some) ability is that, of the worlds in which an agent makes a good faith effort, there

her initial deliberation about whether to adopt some intention to judge that, *as she is now*, she is not able to  $\phi$ , or able to do all of the  $\psi$ s that comprise the sequences that are conducive to  $\phi$ -ing. Or to translate this point into possible-worlds speak, it may very well look to her as if, even if she were to try, there simply are not enough of those worlds in which she succeeds. But this judgment is not sufficient to justify the agent's thinking that she cannot  $\phi$  *tout court*. Why? Because repeatedly trying can itself improve and increase one's abilities. Or perhaps more specifically, an agent can change what she is able to do through *practice*. So though an agent may not be able to do all of the  $\psi$ s that comprise the sequences that are conducive to  $\phi$ -ing right now, she can come to be able to do those things by practicing them. Given the effect that practice can have on an agent's abilities, it would be unjustified for an agent to judge that she cannot  $\phi$  without accounting for how practicing might help her chances of  $\phi$ -ing. Claudia, for example, may rightly judge that she cannot right now sit for four hours at a time and write. But this judgment does not license a judgment that she will *never* be able to perform such an action. The nature of sitting for long periods of time to write requires concentration, patience, and endurance (among other things). These capacities plausibly can be *developed* over time—usually by practicing them over and over and over again.

It is not always the case that trying increases the likelihood of our success. No matter how many times I try to draw a jack of spades from a well-shuffled standard complete deck of cards, the chances of my success are only 1 in 52. And I would be deeply confused about

---

is *at least one* in which she succeeds in  $\phi$ -ing. For example, we might think that there are relatively few close worlds in which injury-ridden Kirk Gibson gets called to pinch hit in the bottom of the ninth and hits a playoff-game-winning home run; but, nevertheless we may want to say that he was able to do just that. Similarly, we might think that there are only a few worlds in which Keri Strug, sprained ankle and all, sticks the landing of her vault to clinch the gold medal for Team USA, but we still want to say that she was able to do so. After all, they did such things, so they were able. Here, we simply adopt the “sufficiently many” threshold for the sake of argument. For an argument against the idea that a single success is sufficient for an ability, see Southwood and Wiens 2016.

the way things work if I believed that my trying had some effect on this chance of success. But with respect to how trying relates to success, we would be wrong to assume that all actions are of the choosing-a-jack-of-spades variety. In fact, a lot of the actions that humans are interested in attempting do require capabilities that evolve over time. I won't be able to complete the 32 fouettés in the Black Swan variation when I first begin ballet lessons, but with time, effort, and practice, my chances of success improve. Likewise, it might be years and years before I can sit peacefully in an undisturbed hour-long meditation, but eventually, after doing the right work, I can. Importantly for our purposes, in deliberation, it can be really hard to tell how our trying will bear on our success. In part, this is because it can be so hard to imagine being able eventually to do something that we right now cannot. Sometimes it isn't until we actually have engaged in some amount of practice that we gain a clear idea of whether it can help bring us to success.<sup>38</sup>

The agent also faces another, related difficulty. She faces the difficulty of judging whether she can mount a good faith effort. Though Wiens does not offer a more precise conception of what he takes a good faith attempt to be, it seems to have at least two key components. The first is a condition of sincerity: an agent must (i) be attempting to  $\phi$  because both she genuinely believes  $\phi$ -ing to be worthy of pursuit and (ii) genuinely intends to successfully  $\phi$ . I do not see any difficulty in an agent being able to tell of herself whether she meets both components of the sincerity condition—that is, whether she can tell of herself whether she genuinely believes  $\phi$ -ing to be worthy of pursuit and genuinely intends to successfully  $\phi$ .

---

<sup>38</sup> The relationship between trying and doing gets even more complicated, since sometimes trying can get in the way of succeeding. Think of cases where we think that someone has failed because they tried too hard, exerted too much effort—for example, the elite athlete who botches the play or the match by “being too much in their head.”

But in addition to this condition of sincerity, a good faith effort requires persistence: for a number of attempts up to some unspecified threshold, in the event that an agent is unsuccessful, she will initiate a further attempt. (Wiens makes this condition explicit in WE by including the “repeatedly” qualification.) And it is much more difficult for an agent to know of herself whether she can satisfy this condition. In a lot of the difficult and complex actions that humans undertake, they face setbacks and obstacles. Many of the actions in a sequence (or a sequence within a sequence) that conduces to  $\phi$ -ing will be actions that pertain to the agent’s responding to such setbacks and obstacles—e.g., *braving* negative feedback, *withstanding* the frustration of writer’s block, etc. Determining whether an agent can perform these setback-responding actions requires her to estimate her own intellectual, emotional, and motivational resources. She, in particular, must assess the extent to which she has *grit*—i.e., what Sarah Paul and Jennifer Morton have characterized as the willingness to persevere in the face of evidence that one lacks the ability to succeed.<sup>39</sup> This kind of self-knowledge can be hard to come by—especially prior to actually facing the setback or obstacle. Sure, Claudia can anticipate that crushing feedback will be terrible to get, but before getting it, she hasn’t actually felt that despair. And without having felt that despair, it will be hard for her judge whether she has what it takes to overcome it. What’s more, even if she can identify various kinds of coping mechanisms ahead of time, she can still be in the dark about whether they will work for her.

So our third lesson is two-fold. First, when faced with deliberation about one’s abilities, an agent must think about whether trying to  $\phi$  will affect the likelihood of successfully  $\phi$ -ing. More specifically, she must now approximate whether trying will sufficiently increase her chances of success. Recall the “tend to succeed” clause of WE.

---

<sup>39</sup> Paul and Morton 2019.

Whatever the threshold for “sufficiently many worlds” is, an agent counts as able to  $\phi$  only when sufficiently many of the worlds in which she tries are also worlds where she succeeds. The effect of one’s trying, then, only matters to the extent that it pushes her past this threshold. Determining whether practicing can push her past the threshold prior to actually practicing can be difficult. Second, in addition to thinking about how trying matters for her success, an agent also has to determine whether she has what it takes to overcome adversity in the course of her of trying. Once again, knowing whether you have what it takes to overcome adversity prior to actually facing that adversity is yet another tall order for an agent to fill as she sets out to determine whether she is able to  $\phi$ .

\*\*\*\*\*

We now have a rough idea of what we need to consider when forming beliefs about our abilities. But remember that when OIC is used to putatively block some normative requirement, the line is that we are not obligated to  $\phi$  when we *cannot*  $\phi$ . Until now, what I’ve argued is that it is often hard to establish that you *are* capable of some action—i.e., that you *can* take some action. But I have not argued that it’s equally hard to establish of some action that you *cannot* take it. It may seem to some that this knowledge is more easily gained. When it comes to determining that you can take some action, you must have reason to believe that *all* of the aforementioned requirements are satisfied. But when it comes to determining that you cannot take some action, you must discover only that *one* of the aforementioned requirements is not satisfied. Thus the difficulties that come with inferring that you *cannot*  $\phi$  seem to pale in comparison to the difficulties that come with inferring that you can  $\phi$ . Or so the story might go.

But I think that this line of reasoning is mistaken. Though it’s true that an action’s failing one of these requirements is sufficient for its being something that we cannot do, it is

typically much harder to confirm that an action has failed a requirement than that it has passed a requirement. More concretely, in order to justifiably infer that you can  $\phi$ , you must find *some* plan that passes the relevant tests, but in order to justifiably infer that you *cannot*  $\phi$ , you must justifiably infer that there is *no* plan that passes one of the relevant tests. Since there are typically so many possible plans to check (perhaps even infinitely many), this latter task is harder. For even if you've established that quite a few plans don't pass muster, it could be that you've just missed that one white raven of a plan that enables you to  $\phi$ . Though we have not said anything thus far about the case of collective action, it is worth pointing out even just briefly that this point will hold, too, in the collective case. We might be tempted to think that it is only easier to establish that we cannot do something than to establish that I cannot do something because the nature of collective actions is so much more complex than individual actions. But the involvement of many agents not only increases the complexity of a given collective action, it also increases the number of possible available plans. And this increase in the number of available plans raises the bar for judging that we cannot since there are that many more sequences conducive to  $\phi$ -ing that we must rule out, one by one.

##### 5. The demands of justice, the methodology of political philosophy

The above considerations provide the bulk of the case against the deliberative usefulness of OIC. For clarity's sake, we can state the general argument *a bit* more succinctly:

- (1) It is permissible for an agent to use OIC to rule out some moral requirement,  $R$ , on the basis of her inability to  $\phi$  only if she has a justified belief that she cannot  $\phi$ .
- (2) An agent has a justified belief that she cannot  $\phi$  only if she has surveyed the all of the plans that conduce to  $\phi$ -ing and justifiably inferred that there is none that she would tend to successfully complete were she to make a good faith effort to do so.

- (3) For many possible moral requirements, an agent cannot survey all of the plans that conduce to  $\phi$ -ing and justifiably infer that there is none that she would tend to successfully complete (in the sense specified in 3).
- (4) Therefore, many possible moral requirements are such that they cannot permissibly be ruled out on OIC grounds (in the sense specified in 1).

The argument, of course, hinges centrally on the claim that it can be hard to discern your abilities and your inabilities, and that in cases where one confronts this difficulty with respect to some purported moral requirement, there are no grounds for using OIC to eliminate that requirement.

One might object that this case against the deliberative usefulness of OIC is overstated because it isn't as hard to discern your abilities and inabilities as often as I've made it seem. You know that you cannot run faster than the speed of light or go back in time to save Franz Ferdinand. And you know that you can finish reading this sentence. When you easily know that you can do something because the action's execution is uncomplicated, you don't need to consider the elements discussed in section 4. And of course, when you easily know that you can do something, there should be no temptation to use OIC as requirement-blocking. So, too, one might insist, for the cases when it seems straightforward that you cannot  $\phi$ —e.g., the cases that look as if they require you to defy the laws of physics. Here, it seems as if the reason that it is hard to find a sequence of action that conduces to successful  $\phi$ -ing is that there isn't one, and no amount of deliberation is required to make us more confident that this is so. Such requirements—those seemingly satisfiable only by impossible actions—must be obvious candidates for the OIC chopping block, right?

Though there are surely *some* cases where OIC can be applied unproblematically in this way—e.g., when dismissing some putative requirement that could *only* be fulfilled by travelling at a super-luminous speed—I suspect that unproblematic cases are much more

rare than many are wont to think. This is because many of the cases that might *seem* like good candidates for elimination by OIC are ones where we actually have the ability to fulfill the requirement, but just don't know it yet. After all, human history is chock-full of humans having done things that were once (possibly even justifiably) thought to be impossible—voyaging oceans and skies and space, curing diseases, etc. Humans continue to try to do things that many great minds think are impossible, such as drag a 125-million-ton iceberg from Antarctica to South Africa in order to resolve a water-shortage crisis.<sup>40</sup> What's more, the candidate moral requirements that many believe there are good reasons to think are genuine do not require seemingly impossible actions. Not even Singer (1972) thinks that we should give money to East Bengal at superluminal speeds.

This point is worth keeping in mind. The claim is not that generally it is difficult to have justified beliefs about our own (in)abilities. It is rather that when it comes to the kinds of actions that are called for by many plausible moral requirements—e.g., to give up eating factory-farmed meat, to forgo air travel, to devote fewer resources to your own children, to register voters, to build community, etc.—it is difficult to have justified beliefs that we are unable to do these things.

How does the argument bear on the debate between Estlund and Wiens and on the greater debate between the ideal and nonideal theorists? Quite straightforwardly, the point forces us to acknowledge that even if OIC is true, it might not matter (or matter as much as perhaps some have thought) for figuring out what justice (and morality, generally) requires of us. This is because OIC is not very helpful as we attempt to figure out what these demands are (and especially, what they are not).

---

<sup>40</sup> See <https://www.bloomberg.com/news/features/2019-06-06/towing-an-iceberg-one-captain-s-plan-to-bring-drinking-water-to-4-million-people>, last accessed 10 July 2019.



There is a methodological lesson for political philosophers here. Presumably, when we propose and defend various principles of justice that we take to be action-guiding, it may be as difficult for us philosophers to have justified beliefs about agents' abilities as it is difficult for agents themselves. This means that any attempts that philosophers themselves make to use OIC to argue against various proposed requirements will have to satisfy the epistemic challenges laid out here. Nonideal theorists have been especially forceful in their insistence that political philosophers take more seriously considerations of feasibility, especially the ways in which human motivational limitations determine what is feasible. But if I am right, then it's at least very difficult to dismiss any putative normative requirement on these grounds (because of the tremendous difficulty that goes into establishing that some action is infeasible). Thus even if we follow the nonideal theorist in granting serious force to considerations of "cannot", there is still a significant burden for any philosophers wishing to allege that some "ought" fails on OIC grounds.

#### 6. What we ought do when we are uncertain about our abilities

The above discussion draws our attention to the difficult issue of how agents should think about their obligations when uncertainty about their ability to  $\phi$  persists—that is, when they cannot settle one way or the other on the matter of their ability to  $\phi$ . While this is not the occasion to address this question in its entirety, I have a few preliminary thoughts. Consider some purported requirement  $R$ . Assume that all of the other normative considerations line up such that a verdict of "being able" is sufficient for justified acceptance of  $R$  and that a verdict of "being unable" is sufficient (on grounds of violating OIC) for justified rejection of  $R$ . Also assume that the agent knows that ultimately she must either

accept or reject  $R$ .<sup>41</sup> It seems that in these circumstances, there are at least two obvious deliberative principles that an agent might opt for:

- (7) Given uncertainty about my abilities, I should act as if I am able to fulfill  $R$  and so accept  $R$ . That is, I should be *aspirational* about my abilities.

Or,

- (8) Given uncertainty about my abilities, I should act as if I am unable to fulfill  $R$  and so reject  $R$ . That is, I should be *circumspect* about my abilities.

It seems plausible that political philosophers may divide on this question according to standard party lines. The ideal theorist might say something like the following: “We have weighty moral reasons to adopt (7) as our deliberative principle. Why? Because when it comes to betting on our abilities, false negatives are morally worse than false positives. That is, it would be really bad to assume that we are unable to  $\phi$  when we are, in fact, able to do so—worse than to assume that we are able to do something when we are, in fact, unable to do that thing. It would be a grave moral mistake to fail to realize some moral ideal because we assumed that we were unable to do so. By contrast, there is not a similarly grave moral mistake to be made in spending our time pursuing some moral ideal even if, as things turn out, the successful realization of this requirement is just not in the cards for us. Only in the former case did we squander an opportunity to achieve something morally worthy.”

There is something important to be said for this line of reasoning. But it also clear what the nonideal theorist will say in reply: “It is a mistake to believe that the morally worse option is to bet on being unable to do something rather than to bet on being able to do it. It is true that when we bet on being unable, we may miss out on the realization of some very attractive moral ideal. However, that bet also comes with great risk—the risk of wasting

---

<sup>41</sup> There are other alternatives here. For example, we could adopt some graded attitude towards  $R$  that does not readily translate into acceptance and rejection. This is not the place to develop a framework in these terms, so we set this possibility aside for now. See Hicks 2018 for extensive discussion of graded attitudes in contexts like these.

one's time pursuing something that one is, realistically, probably going to fail to achieve, and (more importantly) the risk of failing to achieve some perhaps less attractive but nonetheless important moral progress. Only when we are conservative about our abilities are we more likely to make *some* moral progress. And morally speaking, we ought try to do what is good and more likely rather than what is better but less likely. Thus, we are morally required to deliberate with (8) instead of (7)."

There is also a certain force to the nonideal theorist's *abilities conservatism*. If we're circumspect about our abilities, we are hedging our bets in favor of getting things done. We can be more efficient and productive on this approach. But there is a different kind of value to be had in going the aspirational route. Adopting an aspirational deliberative principle like (7) is valuable insofar as it at least makes it possible for things that we have already determined to be normatively worthy to be brought about. By contrast, if we are circumspect, there is no chance of those things happening (barring some cosmic miracle) because we will never try in the first place. Perhaps there is weighty reason to keep hope for the realization of such possibilities alive. This would be a kind of moral case for being aspirational about one's abilities in deliberation under persistent uncertainty.

Relatedly, perhaps we should be aspirational because it is better to have tried and failed than to have never tried at all. In what sense better? Better in that being aspirational about one's abilities would promote one's flourishing. This would be a prudential case for adopting a principle like (7). Though I cannot make either the moral or the prudential case for (7) here, we can see what are the kinds of considerations one could marshal in its favor.

Can we opt for a mixed strategy to deal with persistent uncertainty? Surely. We could adopt (7) in some circumstances where we cannot settle whether we are able to do something and (8) in others. The trick then would be to identify principles that can direct

these adoptions. For example, maybe when it comes our collective actions in demanding high-stakes scenarios, being circumspect is best. But when it comes to my individual actions, being aspirational is best.<sup>42</sup> I can't defend any particular mixed response here, but mixed responses deserve our attention.

Apart from the important question of whether and when (7) or (8) is the best deliberative principle to adopt, the fact that political philosophers may divide on this question according to party lines is a telling observation. Why? Because if I am right, then the ideal and nonideal theorists may misunderstand the nature of their disagreement. For the most part, their dispute have been sold as methodological or meta-philosophical dispute—a dispute about how the political philosopher ought theorize about the nature of principles of justice. But as things appear now, the dispute between the ideal and the nonideal theorist is, at heart, a substantive dispute about how we are morally required to proceed when we are not sure what we can do. Perhaps if we recognize this point, we can better understand our engagement with one another and the aims of our dialectic.

## 6. Conclusion

In this paper, I have drawn our attention to the fact that for many epistemically demanding actions, it is quite a challenge for any agent to form justified beliefs about one's ability to successfully complete such actions. This point is especially important for agents as they think about whether to accept or reject various moral requirements since such requirements will likely consist in or require an agent's attempting epistemically demanding

---

<sup>42</sup> There is also an interesting third option that is worth briefly mentioning: Our moral reasons ultimately weigh in favor of us initially acting as if we are able, but only up until some point. This option would require that we have some theory about when it is rationally and morally permissible to give up, or stop trying, which itself strikes me as another interesting and hard question.

actions. Thus, the principle “ought implies can” seldom (if ever) allows the agent to eliminate some putative normative requirement on the grounds that it is too demanding.

If I am right that rational agents may be uncertain about their own abilities much of the time, then philosophers should be interested in the question of what agents ought to do under conditions of uncertainty of this kind. There are of course considerations of rationality that will bear on the question of what we ought do in conditions of uncertainty about our abilities, but there are likely also moral considerations. Political philosophers should recognize that their disagreements over the respective merits of ideal and nonideal theory may very well stem from a substantive disagreement over what our moral considerations tell us to do when we are not sure what we are able to do.

### Chapter 3: Authenticity and Imagination in Transformative Choice<sup>†</sup>

**Abstract:** Sometimes we can choose whether to do something that we haven't done before. Philosophers such as L. A. Paul have worried that we cannot make these choices well given our lack of experience. In this paper, I argue that there is no problem with making such choices rationally, but that there might be a problem when it comes to making these choices authentically. I then consider whether and how we ever choose authentically in the dearth of experience and argue that authentic choice is sometimes possible in such circumstances if we make adequate use of our imagination.

#### 1. Introduction

Though you face two choices today—one whether to have lentils for dinner tonight, the other whether to become a philosopher—they are very different kinds of choices. First, the stakes are different. Your life won't be made considerably worse should you choose your dinner poorly, but it may be should you choose your career poorly. Thus you could really regret the latter in a way that you won't the former. Relatedly, your career choice is more personal than your dinner choice. We typically think that your line of work says more about who you are than what you eat for dinner on some particular night. Another difference (perhaps a consequence of the first two) is that these choices *feel* different as you make them. Choosing your career feels harder and you may agonize over it more than choosing your dinner. It is also true that while you will choose what to eat for dinner thousands of times in the course of your life, you will only choose a career once, maybe twice. This means that you may choose a food that you have eaten countless times, but when it comes time to choosing a career, you will in some sense be choosing whether to experience something that you have never done before.

Some have thought that this last feature of the choice of career threatens our ability to make such choices well. Without any direct acquaintance with the possible outcomes, how

can you *really* understand what you are choosing? Experience, they say, is a vital resource in choosing well. There are no suitable substitutes. In the first part of this paper, I address the question of how best to understand the worry about choosing what we haven't experienced. I suggest that we should think about the worry as one about the authenticity of such choices. En route to the formulation of the problem that I think better than any formulation on offer in the literature, I consider why other previously considered versions fail to be gripping. Then, in the second part of the paper, I argue that though a lack of experience can sometimes threaten one's ability to choose authentically, it needn't always. Why? Because sometimes when you cannot choose on the basis of your past experience, you can authentically choose on the basis of what you've imagined about what you haven't experienced.

## 2. Rational transformative choice

Sometimes we have new experiences. Having grown up in a small Midwestern town in the U.S., my first taste of the tropical fruit durian at 30 was a new experience. So, too, was my first New York City subway ride at 20. As ordinary wisdom would have it, experience changes us—*new* experiences change us. When we try durian or ride the subway for the first time, we undergo an epistemic change. Before tasting or riding, we didn't know what those experiences were like. But afterwards, we do. Let us follow L. A. Paul in saying that all new experiences are in this sense *epistemically transformative*. Occasionally new experiences can change us in other ways, too. They can, for example, change us personally by setting the course of our life in this direction or that, or by altering how we think of ourselves or how others see us. These experiences are what Paul calls *personally transformative*. Her central example is that of becoming a parent. Others that usually come to mind include marrying

your partner or divorcing them, taking a career or losing it, moving to a new country, or having a religious conversion. For now, we'll operate with our ordinary (perhaps vague) understanding of when some experience counts as new and what these transformations consist in.

Though it is not always up to us whether we have a new experience, sometimes it is. Sometimes we face a choice where at least one of our options is experientially unknown to us—that is, we face *a transformative choice*. And so we must choose whether to possibly undergo some epistemic and perhaps personally transformative experience. When there is the possibility of a personal transformation, it will be what Edna Ullmann-Margalit succinctly calls a *big* decision.

Now, we all think that deciding is the kind of thing that one can be better or worse at doing, and, barring some peculiar circumstances, the kind of thing that we'd like to do well. This goes for all decisions, but probably especially for those big ones that matter so much for what shape our life takes. Philosophers (among others) have spilt a lot of ink over what the standards are for making decisions well, and, arguably, rationality is second to none: To choose well is to choose rationally. Moreover, there is a generally agreed upon approach for choosing rationally. Treat your choices (the advice goes) like the choice whether to pack an umbrella in the morning. You can take the umbrella or leave it. And either it will rain or it won't. How annoying will it be to carry an umbrella around all day only not to have to use it? More or less irksome than getting caught in a downpour with no protection? Consider how much you value each of the possible outcomes and how likely you think each is in the event you take either choice. And then with a few quick decision-theoretic calculations, voilà, you



have the rational choice. This standard recipe requires you to come up with just two simple ingredients—your *credences* and your *valuings*.<sup>43</sup>

By and large, theorists of rational choice have thought of all decisions as on a par—whether transformative or not, big or small—in the sense that every decision can be made well using this same recipe.<sup>44</sup> But a handful of philosophers—I will here count Paul, Ullmann-Margalit, as well as Agnes Callard—have recently alleged that this thought is too quick. In particular, they caution that the standard recipe may be inadequate for our big transformative decisions. There is ultimately a bouquet of worries that these three have, but here we shall focus on one due especially to Paul.<sup>45</sup> To get the flavor, think again of the choice whether to take an umbrella to work. But this time, imagine someone who has never

---

†I'm grateful to Monika Betzler, Christoph Fehige, Anca Gheaus, André Grahle, and Chris Register for their thoughtful comments on earlier drafts of this paper. Thanks to audiences at Ludwig Maximilian University, University of Saarbruecken, and the pre-APA conference on Transformative Experience for their feedback. And thanks to Russ Shafer-Landau, Mike Titelbaum, and especially Reuben Stern for their feedback on the very many versions of this paper.

<sup>43</sup> I'll talk of valuings even though it does some violence to our ordinary language instincts to use this noun form. I'll talk as well of "assigning a value," "valuing," "knowing the value," to refer to a general phenomenon that includes both the act of settling one's views about how valuable something is as well as the mental state that results from this settling. A valuing has an object—the thing valued—but is itself an attitude. I leave it open exactly what character this attitude has—for example, whether it is more belief-like or more desire-like—but I do take it to be an attitude that is *not* immediately apparent to us in the way that, say, our hunger sometimes is. In other words, I don't take it to be especially feeling-like (even though it may be determined by our feelings or produce subsequent feelings in us). Why is this assumption important? Because when it comes to assigning values for the sake of choosing rationally, we must imagine that when an agent is faced with the question, "How do I value X?" it is not the same for her as being faced with the question, "Am I hungry?"

<sup>44</sup> To be clear, not all decision theorists think that decision theory should be used by agents as a deliberative tool. For instance, some think of decision theory as merely descriptive, as a way of inferring an agent's beliefs and desires from her revealed preferences. Others use decision theory merely evaluatively—i.e., as a way to *assess* the rationality of an agent's choices. These camps do not provide a recipe, in a strict sense of "recipe." The group I'm speaking of, then, is those decision theorists who want their theory to provide normative guidance to an agent as she deliberates about what she ought to do.

<sup>45</sup> I focus on Paul's account, but I think that she and Callard have similar concerns, albeit couched in different vocabularies. Paul approaches things from within a decision-theoretic framework and focuses on choice. Callard's approach is more action-theoretic and she focuses on one's reasons for action. But I think these differences are ultimately superficial and that it's fair to think of them as gesturing towards the same worry. In what follows, I end up discussing things in Paul's language rather than Callard's, but there's no principled reason for this.

experienced rain or seen an umbrella. A Martian, perhaps. The Martian amasses all of the usual sorts of evidence in order to come to a view about how likely it will be that it rains or doesn't. He next begins the process of evaluating his outcomes—e.g., judging (on the basis of what he's heard about life on Earth) that it is better to be encumbered by a pesky umbrella on a dry day than to be sopping wet and umbrella-free. Once done with this step, he then proceeds to the usual calculations. But we might balk at the initial evaluation step: Aren't those evaluations a bit peculiar coming from someone who has never experienced rain or umbrellas for himself? We can draw the thought out further by comparing the Martian's valuing to the Seattleite's. Don't the former's evaluations of the outcomes seem odd in a way that the latter's do not? Fans of rain may think that the Seattleite is wrong about what's most valuable, but their criticism likely won't go any further than that. When it comes to the Martian's choice, on the other hand, these rain-lovers may additionally contest that the Martian has no ground on which to stand despite the vast pool of evidence at his disposal.

What explains this asymmetry? Paul seems to be right that there is a difference in their respective evaluations, and, moreover, that this difference is plausibly chalked up to the difference in their experience. What's tricky about rationally choosing whether to taste durian for the first time? To appreciate new kinds of music? To become a parent?<sup>46</sup> Paul argues that in order to choose rationally, you need to be able to assess the value of these things, but that you can't give such an appraisal of their value without experiencing them for yourself. To adopt Paul's language, since you don't *know what they're like*, you cannot know what their value is for you. Paul regards this as a central problem with the standard recipe for rationally choosing in the case of transformative choice (whether big or small). We face an

---

<sup>46</sup> Tasting durian and becoming a parent are Paul's examples. Choosing to listen to a new kind of music is adapted from Callard's example of deciding to become a music aficionado.

experience deficit in such choices and this deficit gets in the way of us acquiring one of the two ingredients necessary for the recipe—the valuing. To make this iteration of the problem a bit more precise:

1. Rational choice requires that we are able to assign values to all of the possible outcomes of our choice.<sup>47</sup>
2. In order to assign values to those outcomes, we have to have the *requisite experience* of them.<sup>48</sup>
3. In transformative choice, we don't have the requisite experience.
4. So we can't make transformative choices rationally.

Let's think of this as the *problem of rational transformative choice* (or, PORT, for short) and of premise 2 as the crux of the problem. It's not intended to be sound as stated, and Paul would not say otherwise. But it will be instructive to use this argument as a starting point for where we want to go, especially because something like this argument is the one that many others have taken to be representative of the worry. In the nascent literature that has resulted in response to Paul's work, the problem has been seen by and large as a problem about whether we can make these choices rationally, given the dearth of experience.<sup>49</sup>

---

<sup>47</sup> For Paul, Callard, and their interlocutors, the claim here is not that an agent in some more literal sense cannot assign a value to the outcomes that she hasn't experienced. Of course, she can always assign values by flipping a coin or deploying some other random procedure. But all parties of the debate do not count this as sufficient for being able to assign a value *that aims to be correct*. So we join them in making this assumption. It is worth noting that our ability to make *rational* choices may not depend at all on our ability to assess the values of outcomes in this way, because rationality may provide instructions only *given* our credences and valuing, regardless of whether our credences and valuing are justified or correct. This point previews what we will discuss below in section 2.

<sup>48</sup> In this paper, I'll use "requisite experience" to pick out the experience of the very thing that we're choosing and not some related or similar experience. So, if I am deciding whether to taste durian, whether I have had the requisite experience is whether I have tasted durian before. This designation is important since it might be tempting to read (2) as saying that one can assign values to an outcome only if one has any experience that is associated with the relevant outcomes. For example, there is a sense in which a male OBGYN has experience of childbirth, but (2) does not say that this is the kind of experience necessary for assigning value to some outcome involving childbirth.

<sup>49</sup> The bulk of papers in the *Res Philosophica* 2015 special issue on transformative choice focus on a version of PORT.

We should keep in mind that PORT isn't *the*, but merely *a* possible problem for the rationality of such choices. When transformative choices are of the personal sort, there are additional questions about their rationality. If a personal transformation means becoming a different person (where that is understood as shifting from one to another set of beliefs and desires), then we may wonder whether an inconsistency between these two sets makes such a shift irrational. This would be a question about the diachronic rationality of such choices.<sup>50</sup> Additionally, a personal transformation could raise questions about the metaphysics of selfhood, including what it means exactly to become a different person or which person—earlier you or later you—is the real or authoritative you. These questions might also affect how such choices can be made rationally. But here I set these concerns aside and focus on the problem outlined above. In part, this is because the problem outlined above rests centrally on the feature of transformative choices that has to do with lacking experience. These other questions may arise even when the agent in question has all of the requisite experience.

There is, of course, always uncertainty when we deliberate about what to do. We don't know what the future will bring—whether it will in fact rain or whether the durian I taste will in fact be ripe. And sometimes when we feel dim about the prospects of assigning values to the possible outcomes that we face, it may be due to this, let's say, *ordinary* uncertainty. It can feel difficult to estimate the value of becoming a parent because there are just so very many ways in which one's parenting—an extended and complex event—could manifest itself, and the value seems highly variable depending on which of the myriad possibilities becomes actual. Due in part to the extended and complex nature of the outcomes that we face in our *big* decisions, these choices especially have a great deal of

---

<sup>50</sup> This is the primary question that vexes Ullmann-Margalit. See her 2006, 2007. See also, Pettigrew Manuscript, which addresses this question at length from within the decision-theoretic framework.

ordinary uncertainty, which may contribute to their often *feeling* hard. But we must take care to keep the nags of ordinary uncertainty at bay so that we don't let their difficulty eclipse Paul's worry. The worry posed (via premise 2) in PORT is not that we cannot assign values to our unexperienced outcomes because we do not know what will happen. It is (again) that we cannot assign the requisite values because, whatever might happen, we do not know what it will be like for us.<sup>51</sup>

There is one additional stage-setting remark to make. For Paul, experience is not only necessary in order to acquire the knowledge needed for assigning values for rational choice, it is (probably) the only way that we can come by such knowledge. Or put another way, she does not think that several of the other sorts of epistemic sources that we ordinarily turn to when our experience runs out, e.g., testimony or imagination, can get the job done.<sup>52</sup> By Paul's lights, you can get all of the testimony that you like about what it's like to have a child, gather all of the scientific data on the matter you like, try as hard as you like to imagine it; but, none of these options will suffice to produce the requisite knowledge. This may be puzzling to you, but for the moment let's put a pin in that thought.

## 2. What rationality requires

As we entertain the question of whether some choices are especially difficult to make well because we lack the requisite experience, the cases that we've considered are meant to

---

<sup>51</sup> One way of isolating the worry is to imagine that you are given by a partial oracle all information of a certain sort about the possibilities that you face—e.g., to have a child or not. This partial oracle tells you exactly what will happen from the point that you don't have a child to the end of your life, as well as what will happen from the point that you do have a child to the end of your life. The only thing that the partial oracle doesn't tell you is how it will be for you to live through these outcomes—that is, how what will happen will affect your feelings, emotions, attitudes, beliefs, understanding of yourself and the meaning your life has. The point for Paul is that even with all of this information from the partial oracle, you would not be able to assign a value that will allow you to make a rational choice.

<sup>52</sup> On this point, she and Callard are in agreement.

have some intuitive force that pushes us in the direction of an affirmative answer. There *is* something odd about your assessing the value of durian, rain, or parenting when you haven't experienced those things for yourself. But good philosophers that we are, the first thing that we should do after being moved by these cases is search for counterexamples—i.e., for cases where agents can choose well despite their lack of experience with the possible outcomes. Along these lines, consider the case of a slave choosing whether to escape her capture. Having been born into slavery, she has never experienced what it is like to be a freed person. According to PORT, this person cannot make a rational choice about whether to escape, because she cannot appraise the value of her freedom. Hopefully this will strike you (as it has me) as wrong. Surely it is possible for this choice to be a rational one. And more importantly, surely such a person can appraise the value of her freedom. In fact, she may very well stand in a superior position to judge its value precisely because she has been deprived of the requisite experience.

Our judgment of the slave case directs us to further scrutinize (2)—that is, whether experience is needed in order to get the second ingredient in the rational choice recipe. As it turns out, there has been a chorus of responses to Paul expressing skepticism about (2). Chang 2015, Kauppinen 2015, Harman 2015, and Dougherty et al. 2015 all in some fashion or another contend that our ability to assess the overall value of the outcomes that we face doesn't require us to have experience of those outcomes. They argue that once we clarify the metaphysics of value, the epistemology can also be made clearer. They maintain further: the value of our outcomes does not depend—at least not in any troubling way—on the nature of our experience of those outcomes. Instead, it derives primarily from their non-experiential features. If this is right, then it seems clear that we can come to know the value of an outcome sans any experience of that outcome; we can simply acquire this knowledge of its

non-experiential features via familiar epistemic routes like testimony and counterfactual reasoning. These objections are each forceful in their own right, and taken together constitute a serious challenge to the worry about transformative choices thus far construed. Does this—the *value-theoretic response to PORT*—preempt any further discussion? Unsurprisingly, I think that there is more to say. In order to say it, it'll help us to consider other ways in which Paul has talked about the difficulty posed by transformative choice. Here are two additional ways that she articulates her worry:

“In particular, what is necessary for *the right sort of epistemic acquaintance* is that you represent the nature of the lived experience of the outcome to yourself under the subjective mode of presentation.”<sup>53</sup>

“But the agent’s inability to grasp and represent the propositions concerning the nature and character of the outcome, or, we might say, *her inability to grasp these propositions under the right mode of presentation*, implies that she cannot be aware of the subjective value of the outcome.”<sup>54</sup>

The italicized phrases point us to a different kind of problem from the one that we’ve been considering. To get a grip on this other problem, think once again about the Martian’s choice of whether to bring an umbrella. The initial gloss on the case was that the Martian could not assign values to facing rain without an umbrella, etc., because he cannot know, e.g., what it will subjectively feel like to be caught in the rain, sopping wet. In our recipe metaphor, the Martian could not get one of the two ingredients—the valuing—necessary to make a rational choice. But the alternative gloss is that the Martian can get both ingredients. He can get all sorts of information about rain and umbrellas, and can subsequently make a value judgment based on that information. But those valuing will be *defective* or *lacking* somehow. Why? Because the outcomes are not *understood* in the right way. On this suggestion, an agent’s *grasp* of the possible outcomes can differ in some relevant way,

---

<sup>53</sup> Paul 2015b, 480.

<sup>54</sup> Paul 2015b, 501.

depending on *the way* that she becomes acquainted with those outcomes—e.g., whether it is through testimony or through experience. Thus two agents armed with the same valuing and credences might make the same choice, but differ with respect to whether they choose well in accordance with their distinct grasps of the values.

We can give a little more precision to this worry—let’s call it the *problem of rational transformative choice, modified* (or, PORT-M)—like so:

1. Choosing rationally requires the right sort of valuing of the possible outcomes of our choice.
2. We have the right sort of valuing of our outcomes only if we know what these things are like in the right way.
3. We know what our outcomes are like in the right way only if we have the requisite experience of them.
4. In transformative choice, we do not have the requisite experience.
5. So we can’t make transformative choices rationally.

Understanding the problem of transformative choice in the manner above helps to explain not just certain features of Paul’s view that have struck many (including me) as odd, but also to explain more generally some of our folk ways of talking about and treating experience.

Imagine the following scene:

To learn about rain and umbrellas, the Martian reads everything he can that may be relevant to his value assignment. For example, he reads about wetness, about human reactions to rain, and about human neuropsychology (in order to form an opinion of whether human reactions are likely to be like his own). After reading, he declares his disdain for being caught in the rain without an umbrella. But we scoff at this declaration. “Martian,” we say, “you don’t really know what you value. You don’t what it’s like to be caught in the rain without an umbrella!” The Martian responds by rattling off a list of reasons to hate rain from his studies. Indeed, he gives a surprisingly rich and nuanced picture of what it is in fact like. But we dig in our heels and insist—“But, Martian! You don’t *really* know what it is like!”

To my ears, the scene plays out quite ordinarily (save the Martian). It illustrates our tendency to privilege experience-based declarations of what things are like and what their value is, as



well as our tendency to distinguish these privileged declarations by using really-talk.<sup>55</sup> Paul herself also uses this kind of talk:

“But in a case where you really don’t know what it’s like to taste the menu item, you can’t use maximin, or maximax, or any other decision-under-ignorance rule to rationally make a decision based on what you think it will taste like.”<sup>56</sup>

And,

“[T]ransformative experiences] are experiences in which we learn something new. We learn what something is (really) like. Upon learning what that thing is like, we gain abilities to remember, imagine, and recognize.”<sup>57</sup>

And,

“[P]eople tell stories about how they knew what people had told them (about parenting, or going to war, or moving to a country with a very different culture, etc.) beforehand, but there was still a distinctive and extremely important sense in which that did not prepare them for what the experience was *really* like.”<sup>58</sup>

The Martian’s scene above also illustrates the judgment that we sometimes have that in the absence of experience, testimony and other epistemic sources are not suitable replacements. I cannot learn all that I know about parenting from *What To Expect When You’re Expecting*, and then assert that I *really* know what parenting is like. Even if I know beforehand that durian is stinky because reality food TV star Andrew Zimmern told me so, there is, we often think, something relevantly different about *the way* that I know durian’s stench after I’ve actually smelled it for myself. Earlier on we noted that Paul, too, insists that extra-experiential epistemic resources such as testimony cannot make up for what is lost given our dearth of experience when we’re making transformative choices. In the context of PORT,

---

<sup>55</sup> Cath 2019 offers an admirably clear articulation of the ways in which “knows what it’s like” ascriptions are ambiguous, as well as the sense in which our privileged ascriptions (or “gold-standard” ascriptions, as he calls them) pick out knowledge that is based on having had the requisite experience.

<sup>56</sup> Paul 2015a, 16.

<sup>57</sup> Paul 2014, 10.

<sup>58</sup> Paul 2015b, 487, emphasis hers.

that claim seemed false. But in the revised version of the problem, it begins to seem more plausible.

So, as we aim to formulate a more credible version of a problem for transformative choice, PORT-M is a step in the right direction. Nevertheless, it is not yet *quite* the formulation that we desire because it suffers from another problem. Like its predecessor, PORT-M makes the problem of transformative choice one about rationality. It says that transformative choices cannot be made rationally because we are unable to value our outcomes in the right sort of way (premise 1), where thus far what we have understood about “the right sort of way” is that it has to do with how we *grasp* the outcome. Even without a further specification, this claim should strike us as suspect. As we’ve observed, in order to choose rationally according to standard accounts of rational choice (e.g., classical formulations of decision theory), we can choose rationally simply on the basis of our confidence and value assignments to the outcomes. Paul raises a worry that we cannot know the values of some outcomes without experiencing them, but—as convincingly argued in the value-theoretic response to PORT—there are other ways to come to know the value of an outcome (e.g., testimony, science, etc.).

By contrast, PORT-M insists that there is *more* to rational choice than choosing on the basis of these two ingredients. Again, it says that rational choice additionally and more specifically requires that our valuing is known to us in a particular way. Though decision theorists do believe that there are some constraints on the value assignments that rational agents make, these constraints typically ensure the coherence of those value assignments—for example, that they do not violate transitivity. Coherence constraints place limits on *what* value you can assign to a given outcome, given what your other values are. The constraint issued by PORT-M, however, would not prohibit an agent from assigning any particular

value to a given outcome. Rather, it says that when an agent fails to grasp, or stand in the right cognitive relation to the outcome, then the value that she assigns to it—whatever it happens to be—cannot issue a rational choice. This is certainly not a requirement of rationality that decision theorists have discussed before. But how can we tell whether it is, indeed, a requirement that they have simply overlooked?

Though we can't here mount a case that would rule decisively against the general idea that an agent's valuing is rational only if she stands in a particular cognitive relation to the object of that valuing, we can offer at least some reason to think that we should be wary of such a requirement. Think back to the Martian and the Seattleite. We imagined that they have exactly the same valuing—that is, that they both value being dry and encumbered over being without an umbrella but soaked. It is hard to see why the Martian's but not the Seattleite's valuing should count as *irrational*. This valuing allows the Martian to engage in coherent and effective means-end reasoning, just as the Seattleite's valuing does. Given that the Martian values as he does, plugging his value assignments into the decision-theoretic calculus will enable him to maximize expected utility. Additionally, the Martian's valuing is informed by reasons, just as the Seattleite's is: when asked, he can provide rain- and umbrella-related justifications of why he values the way that he does. So without some richer conception of rationality that would illuminate for us why the Seattleite's but not the Martian's value can issue a rational choice (*even* when they assign the same values), there is not much reason to believe that rationality puts this kind of pressure on agents. Again, this is not to say that there certainly isn't such a kind of rationality that makes this demand of

agents. But Paul doesn't explicate such a conception and there doesn't appear to be an obvious one already on offer.<sup>59</sup>

Thus, while PORT-M is a more plausible formulation of the problem of transformative choice than PORT because the former does not make the doubtful claim that without the requisite experience, agents have no resources to help them assign values to the possible outcomes in their choice, PORT-M still isn't perfect. This is because it is unlikely that *rationality* requires us to stand in any particular cognitive relation to the outcomes that we value. In the next section, we consider whether an alternative normative standard of choice more plausibly does issue such a requirement.

### 3. What authenticity requires

There are lots of ways that we might judge a choice to be a good one. We might, looking to Robert Frost on his less travelled path, admire unorthodox choices. Or, looking to the hemlock-drinking Socrates, admire principled choices. Or we might prize obedient choices, putting Abraham on a pedestal. Plenty of us think moral choices are the good ones.<sup>60</sup> Sometimes choosing well is choosing authentically. What standard of good choice other than rationality could help us make the best sense of the problem of transformative choice?

---

<sup>59</sup> Of course, if we really wanted to rescue the idea that there is trouble for making rational transformative choices, we could articulate such a conception. But for me, the more interesting route to pursue, or anyway, the route that we'll pursue here is to locate the problem elsewhere.

<sup>60</sup> Thinking of rationality in the narrow sense of being average utility-maximizing, some of these standards will be extra-rational, while others will be complementary to the standards of rationality. There are broader conceptions of rationality, including the one that says acting rationally is just to act for some reason. If we adopt this broader conception, maybe any choice standard that we adopt will fall under the heading of being a kind of rational standard.

As it turns out, Paul herself provides one answer to this question: authenticity. Though it is the part of her view that is the least developed and the most ignored, she thinks that transformative choices pose a special problem for choosing authentically. She writes:

“A way of emphasizing the importance of introspective reflection in these sorts of cases is to say that high-stakes choices like this can be subject to an “authenticity norm.” The authenticity norm concerns the way you make life choices in concert with your first-personal understanding of who you are and what you want from life. An agent who authentically understands herself first-personally grasps her defining nature and values from the inside, that is, she knows who she is under the subjective mode of presentation.”<sup>61</sup>

So for Paul, the pressure to have the right sort of valuing—that is, a valuing based on the right grasp of one’s outcomes—is pressure that comes from the standard of authenticity.<sup>62</sup> Let’s see whether adopting *this* as the relevant standard for choice can give us a more convincing formulation of our problem. Here is what the *problem of authentic transformative choice* (or PAT) looks like:

1. Choosing authentically requires the right sort of valuing of the possible outcomes of our choice.
2. We have the right sort of valuing of our outcomes only if we know what these things are like in the right way.
3. We know what our outcomes are like in the right way only if we have the requisite experience of them.
4. In transformative choice, we do not have the requisite experience.
5. So we can’t make transformative choices authentically.

I do think that the first premise re-cast as a requirement of authenticity rather than of rationality gives PAT more credibility. In order to see this, it is helpful to consider an example in which two choices intuitively differ with respect to their authenticity.

On the one hand, there’s Mitch, who has never tried mapo tofu, but chooses not to eat it when offered the chance because he has heard tales of its extreme spiciness. On the

---

<sup>61</sup> Paul 2015b, 483.

<sup>62</sup> In her most explicit moments, Paul argues that one cannot rationally *and* authentically make a rational choice. She does not say more about how she understands the standards of rationality and authenticity to be related to one another.

other hand, there's Marlon, who opts not to eat mapo tofu because he has tasted it before and finds it too spicy. We can imagine that Mitch and Marlon dislike worlds in which they ingest mapo tofu equally, and even that Mitch and Marlon are in shared agreement about precisely which worlds are preferable (and not preferable) to those worlds in which they eat the mapo tofu. But still, it seems that Mitch's choice is defective in a sense that Marlon's is not because of the *way* that Mitch disvalues the tofu. It isn't that Mitch's valuing lacks sincerity or is not fully endorsed by him. Nor is it that it is not reflective of his true or real self (whatever that could possibly mean). Rather, there is a legitimate claim against Mitch that his disvalue isn't *felt* in the sense that he does not know what it's like to taste mapo tofu for himself, and in this way, it fails to be connected to *him*. There seems no such legitimate claim against Marlon. It seems reasonable to construe this difference in terms of authenticity in this very general sense of not being appropriately reflective of or tied to oneself. Mitch's choice is inauthentic because it's not on the basis of what he has felt for himself, while Marlon's is because he chooses on the basis of what it has felt like for him to be contact with the thing being valued.

It would unfortunately take us too far afield to fully flesh out this conception of authenticity. But this certainly a very rough sketch seems to both have some credibility as a conception of authenticity as well as be a credible way of cashing out our intuitive judgments in cases like Mitch's and Marlon's, like the Martian's and Seattleite's, as well as those that Paul focuses on, like the choice to become a parent. With the hope that this sketch is sufficient to make us receptive to the idea that the problem of transformative choice is an authenticity problem, we adopt it going forward.

With these considerations of authenticity in hand, then, let us briefly pause to take stock. We started by granting that something seems odd about the Martian's choice to bring

the umbrella. His relative valuing of carrying an umbrella with no rain and getting caught in the rain without an umbrella seems *off*, somehow lacking in a way that a Seattleite's similar valuing does not. We've chalked this off-ness up to a likely difference in the nature of the Martian's and the Seattleite's valuing, which is in turn due to a difference in the *way* that the Martian and the Seattleite understand the outcomes of their choice. Without having experienced rain and umbrellas, the Martian's choice is inauthentic in the sense that he cannot appreciate what it feels like to be soaked, no matter what he has learned before coming to Earth. Initially, we wondered if the kind of difficulty that results from this failure to have the right sort of appreciation should be understood as a difficulty in choosing rationally. But we found that there is some reason to think that rationality does not require a special appreciation of our outcomes; a value assignment based on any old grasp will do. Thus we shifted from PORT (and PORT-M) to PAT, and came to understand the challenge posed by transformative choice in terms of authenticity, where a choice qualifies as *authentic* only when the agent's value appraisals based on an appropriate understanding of their outcomes. Where to from here?

In the rest of this paper, I argue that PAT ultimately fails on the grounds that there are some transformative choices where we do know what our outcomes are like in the right way absent experience—namely, those choices where our valuing are based on our having appropriately imagined the possible outcomes that we face. So I reject premise 3 in PAT. Let's turn now to that view.

#### 4. Choosing what we have already experienced

In order to make the case against premise 3 of PAT, we have to make sense of what it means to “know what an outcome is like in the right way,” or to have what we can call *an*

*appropriate appreciation* of that outcome. One possibility, of course, is that the appropriate appreciation of one's outcomes just *is* an appreciation based on or generated by one's requisite experience. Then the knowledge that you need for authentic choosing just *is* experience-based knowledge. Without any argument to this end, however, this option not only begs the question against those who argue that transformative choices aren't problematic, but also fails to illuminate. And Paul does not supply such an argument. Exactly what we want to figure out is *why* the requisite experience purportedly facilitates a valuing that can issue an authentic choice. What is it that experience allows us to do as we think about our outcomes in deliberation that other epistemic resources such as testimony do not?

Paul suggests some possible answers (which we saw hints of in her remarks considered earlier). Perhaps appropriately appreciating one's outcomes is being able to represent them under the subjective mode of presentation, to first-personally grasp them, or (as she mentions twice in other places) to be able to *de re* represent them.<sup>63</sup> These suggestions aren't question-begging, but because Paul does not elaborate on them, they are not very illuminating either. That doesn't mean that we should rule them out as possible contenders, but if we consider them directly, I think we'll quickly find ourselves in the weeds. (In part, we're likely to find ourselves lost because we would have to do too much speculating about what Paul thinks but has not said.) Instead, I suggest that we proceed by reflecting more carefully on those ordinary choices that we make when we have experience of our options. Since we know that, by Paul's lights, experience facilitates an appropriate appreciation, thinking about what is characteristic of how we grasp our possible outcomes when we have

---

<sup>63</sup> For the claim that what we need is to be able to represent our outcomes under the subjective mode of presentation, see pp. 480, 481, 483, 484, 486, and 493 of 2015b. For the claim that we need is a *de re* representation of our outcomes, see 478, 486fn11, 515, and 516 of *ibid.* Her claim that what is necessary is a first-personal representation is the most widespread, but see, e.g., 481, 483, 535 of *ibid.*



experienced them can help us build a cluster concept of the notion this notion—“an appropriate appreciation”—that we seek clarity about.

Let’s consider the following choices:

6. Choosing to taste a green apple after having tasted it before
7. Choosing to have a child after having had one already

Here’s *one* plausible way in which my deliberation about these choices could go: I think about what tasting a green apple (or refraining), having a child (or not) will be like. To do this, I build in my mind’s eye (as the metaphor goes) a representation of each possible future outcome that I might live. This representation is a scene of sorts, one where I am the central character and which depicts some way in which an event happens to me. I use my judgments about the various likely ways that things might go to tinker with various features of the scene—I see one where my second child is a boy, and the other where it is a girl—and look for what else changes in the scene once I make these particular adjustments. I then let these representations (in all of their various iterations) guide my judgment about how valuable each of my possible outcomes is, and eventually make my choice on the basis of these evaluations.

Let’s call such a representation a *representation for deliberation*, or *RFD*.<sup>64</sup> Now, what kind of character my RFD has will depend on what epistemic resources I can draw from as I construct it. One central resource that I can use is my past *direct* experience. Relying on memory, I recall what the thing in question was like in the previous time(s) that I’ve experienced it. I then use those past representations (*experience-sourced representations*, let’s say) to develop a similitude. To illustrate: When I’m thinking about what green apples taste like as I choose whether to have this apple or an orange, I recall my previous experiences with

---

<sup>64</sup> Of course, we probably do not always deliberate using RFDs. For instance, we may not when we have to deliberate with some greater degree of expediency.

eating green apples in order to put a picture in mind of what this future green apple eating will be like.<sup>65</sup>

Given that I model RFDs after my experience-sourced representations, when those are available, let us think about what the latter are like. (We ordinarily call these things perceptions.) Among other things, these representations are often *pictorial* in the sense that they are of a thing (e.g., an event or an object) whose various features come *composed*—that is, arranged and ordered (e.g., spatially and causally) in such-and-such ways with its various components presented *tout ensemble*. For example, when I have the experience of eating a green apple, the various parts of that experience—the greenness, the roundness, the tough-skinned-ness, the tartness—come together to give me a representation *green apple*. In fact, so long as I possess the usual sorts of sensory equipment, I cannot usually have a perception of this experience that does not have these features (unless I have some special intervention by which I can block one of my modes of sensing).

My experience-sourced representations have additional features: They are *perspectival* in the sense that there is always a subject to which the information that appears is appeared to. This feature determines *how* the picture is composed (for example, whether the glint on the apple’s skin from the daylight appears on the right or left). These representations are *passive* in that we don’t do much “online” cognitive work to obtain them; our various modes of sensory perception do much of the work for us. (We exercise some agential control occasionally and to the extent that we are able to direct what in particular we attend to, but much of the time even our attention is not something we control.) In addition, our experience-given representations are *abundant* in that they are especially rich in content. That

---

<sup>65</sup> See Nanay 2016 for a similar model of decision-making that emphasizes the role of imagination. Nanay argues that this model is more empirically plausible than the “pure” belief-desire models that seem to be assumed by decision theories and some theories of action.

is, experience supplies us with a vast wealth of details about the thing experienced. (If a picture is worth a thousand words, an experience is worth a million.) And, partly constitutive of this richness, these representations are *felt*—i.e., they engage us so that the sensations, feelings, emotions, attitudes and other reactions that we have to the thing being experienced become a part of our representation of it.

When it comes to building an RFD, we can usually preserve at least some of these features, to some degree. We produce pictorial representations and the pictures or scenes that we produce are modeled after the details of the original. They are perspectival—there is a point of view to which they appear. They can also be abundant and felt. Think about your own past experiences of eating a green apple. Can you envision an apple before your mind's eye? Or a scene of yourself, perhaps yesterday at lunch with your spouse, eating an apple after your peanut-butter sandwich? Can you conjure the sour taste? (Is your mouth watering?) Do greenness and roundness automatically present themselves before your mind? Can you feel smooth and firm skinned-ness, perhaps even the sticky-wetness of apple juice trickling down your chin?

There are, of course, differences between the original perception and your RFD. Though in perception we can't help but be the focal point, the thing towards which everything that is represented is represented to, when we build an RFD, we have more flexibility in which perspectives we can take. We can recall the original scene as it appeared to us, but also, in an act of abstraction, can consider it from other focal points. Or as the familiar metaphor goes, we become able to see things (including ourselves) *from the outside*. Our RFDs may also be abundant and felt, but probably less so than our original perceptions. They are, it seems, inevitably subject to a certain amount of *fading*. I can recall *green apple sourness* to some degree, but it doesn't have nearly the intensity as when I'm actually eating a

green apple. To what extent the richness and feltness fade will be a matter of a handful of considerations, including, *inter alia*, the fastidiousness of our memory and the nature of the things remembered. Finally, while we are passive in perceiving, we are not when building RFDs. We are in the latter exercise, after all, *constructing* representations, and so we exert a great deal more control over *what* we represent.

The takeaway point here is that when we are deploying these mental simulations in the service of deliberation and we have requisite past experience to rely on, our simulations, or RFDs, are perception-*like* in character. Among other things, this means that the thoughts that we have in this kind of deliberation—those thoughts, remember, whose character facilitates authentic valuing—probably have some or all of the various properties that Paul focuses on. They are first-personal (at least some of the time). They are *de re* in that they are thoughts of things.<sup>66</sup> And they are obviously subjective in the sense that they are presented to me as a part of my conscious experience (which is the sense of “subjective” that Paul, I think, has in mind).

But what, exactly, gives these thoughts their authenticity-conferring capabilities? Arguably, it is their feltness. This is the feature that contains information about what the thing represented is like *for you*. It is less clear how the other features—being imagistic, being first-personal, etc.—directly promote authenticity. Rather, they seem important for authenticity only to the extent that they facilitate feltness.<sup>67</sup> So when we ask why we can

---

<sup>66</sup> Though it may seem obvious that we have *de re* thoughts of non-occurrent or non-actual things, these kinds of cases are not paradigmatic and thus often not discussed by theorists of the *de re*; it is not always straightforward how various theories of the *de re* could or should be extended to such cases. Jeshion 2002 is one example that I know of that gives a theory of the *de re* that is explicitly intended to extend to such cases.

<sup>67</sup> Whether and how they do is, for us here, a matter of speculation. It seems plausible that features like being image-centric (which may be constitutive of being *de re*) facilitate feltness because given the kinds of cognitive creatures that we are, images engage our affective faculties in a way that image-less thoughts (usually) do not. Likewise, abundance (which may itself be facilitated by image-centrality)

choose authentically in cases like the choice to eat a green apple again or to have a child again, the explanation seems to be this: Our requisite past experience allows us to construct an RFD that is felt. When we assign a value to our possible outcome on the basis of felt RFDs, we assign a value based on an appropriate appreciation of those outcomes. This is plausibly the sense in which we value our outcomes in the way that allows this valuing to confer authenticity on our subsequent choice.<sup>68</sup>

### 5. What imagination can (and cannot) do

If it turned out that the only way to get a felt RFD is by having the requisite prior experience, then PAT would go through. But we can get felt RFDs in other ways, and so the argument fails. When we don't have the requisite experience, sometimes we can use imagination instead. As it turns out, we have already granted the role of imagination in deliberation, since the representation-constructing activity that we have described above is itself best understood as an act of imagination. Imagining, generally, is a cognitive capacity for constructing representations of non-occurrent things, where often those representations are centrally imagistic (as opposed to centrally propositional.)<sup>69</sup> When we construct outcomes on the model of previous experience, we are *imagining* what the future will be like based on what the past was like by constructing a centrally imagistic representation. And we have

---

seems to facilitate feltness because the latter seems to be triggered only by a critical mass of detail. It seems less plausible that first-personal perspectival-ness is critical to feltness since observing scenes can produce feeling just as being the subject of them can. But perspectival-ness generally may be important—again, perhaps the extent to which it is bound up with a thought's being image-centric. Since all we really care about is feltness, the devil actually isn't in these details. Though, as we continue, we may use these additional features like being *de re* and being abundant as *good evidence* for a thought's being felt, and thus, authenticity-conferring.

<sup>68</sup> If someone like Paul, who is skeptical of our ability to make authentic choices rationally, disagrees with this account, the onus is on her to defend an alternative one.

<sup>69</sup> For helpful discussion of the difference between image-centric imaginings and proposition-centric imaginings, see McGinn 2009 and van Leeuwen 2013.

already granted that these imaginings have a particular character—pictorial, perspectival, abundant, and (importantly for us) felt. So we have seen already that imagination is poised to give us the kind of thought that we need in order to authentically value. The question we really need to address is whether imagination can give us this right kind of thought even when we are trying to imagine things that we haven't experienced.

In this kind of case, by hypothesis, we have never had the requisite perception on which to model our RFD. So we must rely on other epistemic resources in order construct our scene. Perhaps the first possibility that comes to mind is some other *related* experience of my own. Maybe I haven't eaten a Granny Smith, but I have eaten a Braeburn, so I can use my memories of the latter to build an idea of the former. They are, after all, both apples. Relying on prior related experience is a relatively easy way for us to get an RFD of the right sort, since we deploy a significantly similar procedure to the one that we do when we have had the requisite experience: we construct a scene on the model of a memory. The central difference is that we usually need additional information to help us understand the nature and degree of resemblance between the two kinds of experience in question. This information can help us augment our imagining in ways that signify our recognition of the fact that these experiences are more like siblings or cousins than twins. Since we ourselves won't have this information (because we're missing one of the comparison classes), we'll have to seek out descriptions—of the requisite experience, but also of how it compares to the ones that we have had. Below I'll say something about how descriptive information bears on the character of our RFDs. But for now the important point is this: there is no in-principle reason why related experience should give us an RFD different *in kind* from that facilitated by a requisite experience.

There are pitfalls to having only related experience to rely on. Perhaps it's not Braeburns that I've eaten, but oranges, and I try to build an idea of what it will be like to eat a Granny Smith by thinking about what it was like for me to eat an orange. They are, after all, both fruits. Perhaps, further, I haven't had the good fortune of knowing anyone who has had Granny Smiths or, ideally, Grannies and oranges. In this case, I will run into a problem about how *accurate* my RFD is—that is, how accurate the details of the scene are, including, importantly, the details about what kind of reaction I would in fact have to such an experience. There are two kinds of inaccuracy I can incur: my RFD may, first, fail to resemble in various ways actual instances of the things in question (in this case, green apples) and so, second, may fail to resemble the future item that I might choose and so experience (in this case, the next green apple).

The possibility of an inaccurate representation does not threaten authentic choice for several reasons. First, even given inaccuracies, we still have a thought with the same features as our direct experienced-based thoughts—pictorial, felt, etc. These are the kinds of thoughts that deliberative imagination can't help but produce given its perception-like functioning. Second, since inaccuracy does not threaten an RFD's authenticity-conferring status in the case of direct experience-based representations, it should not do so in the case of indirect experienced-based ones, either. As we noted earlier, the abundance and feltness of my original perceptions are subject to fading when I construct RFDs modeled on these perceptions due to various lapses in memory. Memory not only fades, but distorts. The point might be tricky to see in the case of green apples, but think about a choice to take a vacation to Europe. If I haven't been to Europe in over three decades, my imagining what Europe *will* be like for me based on what it *was* like for me might be drastically different from what it in fact will be like for me. Surely Europe has changed a lot in the last 30 years. I certainly have.

In cases where we know the future will not resemble the past perfectly, we tend to rely on the same augmenting strategy that we use to adjust for differences between our related experience and the target experience. But this strategy is not foolproof and we are often left with an inaccurate idea. If this isn't a problem for the authenticity-conferring status of our direct experience-based RFDs, it can't be for other RFDs either.

In addition to past, related experience, we also can draw on testimony to help us imagine what previously unexperienced outcomes will be like. When you are given a piece of testimony, you are given a description. A description may not necessarily facilitate the production of an RFD that is pictorial, perspectival, abundant, and felt, since there are two different ways in which we can cognitively register descriptions. Consider the following two choices:

8. Choosing to taste a green apple having never tasted it before, but on the basis of your friend's testimony
9. Choosing to have a child after reading Ina May Gaskin's *Guide to Childbirth*

You *could* go about making these choices as follows. Your friend's testimony might just be that she likes green apples; you could, in your deliberation, reason that you resemble your friend in some respects and so probably like what she likes, and so decide to go for the green apple. Alternatively, your friend's testimony might be a description of what green apples taste like. You hear that they are tart, juicy, and crisp, reason that tart, juicy, crisp foods are tasty, and so, as before, you decide to take the plunge. At no point do you run through a scene in your mind of you eating a green apple. You *merely* have *dicta* in mind, no *res*. Similarly, according to Ina May's description of childbirth, it is magical, ecstatic, one of the most profound experiences that you can have as a woman. Magical, ecstatic, profound experiences are a critical part of flourishing, you reason, so you go in for having a kid. Again,



you skip any attempt to play through any scenes in your mind of you birthing or raising a kid.

If you made your choices in this way, I submit that you are not making authentic choices. This is so both in an intuitive sense, but also in the slightly more elaborate sense that we've sketched above—you aren't appropriately grasping your possible outcomes. Rather, your representations are flat, scant, impotent. You seem to register these testimonies *merely propositionally*, as items to be added to the stock of propositions that you already possess and that you can deploy for the purpose of serving as the bases of inferences. Your engagement with them seems to be nothing more than an impersonal reasoning in order to infer what is valuable. Your thought reflects neither the quality of living these outcomes, nor how it would be for *you* to be the one to live them.

But using the descriptions provided by testimony in this way is not the only way that we can use testimony. And it is not using testimony in the service of imagination. Though we *can* simply use these descriptions as propositions to add to our mental stock, we are *also* capable of using them to imagine. When we imagine on the basis of descriptions, we *translate* our descriptions so that we can build an image of the thing that they describe. This translating is the characteristic activity that we engage in any time that we read a story. Take, for instance, John Kennedy Toole's description of a slovenly, overweight man whose "green hunting cap squeezed the top of the fleshy balloon of a head" and whose "full, pursed lips protruded beneath the bushy black moustache and, at their corners, sank into little folds filled with disapproval and potato chip crumbs." As I read this description, I picture Ignatius J. Reilly.<sup>70</sup> Can we translate descriptions to imaginings even when they are descriptions of a thing that I have not experienced? Of course we can. Ignatius is a case in point. Can I

---

<sup>70</sup> Toole 1980, 1.

translate descriptions to an imagining that is abundant and felt of a thing that I have not experienced? Surely. Every time I am *moved* this way and that throughout the course of reading a story, I am engaging in this activity. When Ignatius's slobbish, crumb-specked face enters my mind and I feel a hint of disgust, I have a felt imagining of something that I haven't experienced.

There are hurdles to relying on descriptions. How easy it will be for us to get felt RFDs from descriptions will depend a great deal on the quality of those descriptions. Getting the right sort of picture or scene in mind is challenging when I only have impoverished descriptions to work with. This is why it may be important to rely on descriptions not just from others who have had the requisite experience, but others who have had the requisite experience *and* are skilled enough in introspecting and articulating to describe that experience richly enough. This is no small feat. Getting a felt description-based RFD also depends on the quality of *my* imagination, or more specifically, on the extent of my own creativity. This act of translating, of assembling scenes from descriptions of those scenes, is also challenging. The more so when these scenes have to be vivid and complex enough to engage our affective faculties. It requires practice. But these particular hurdles don't show that description-based, experience-less RFDs cannot be felt, only that they may be more difficult to come by than their experience-based relatives.

There is an additional worry, again about accuracy. How accurate your description-based RFD is will depend on the accuracy of the descriptions you are given, including how accurate they are *of you*. But once again, ending up with an inaccurate RFD isn't a problem, for the same reasons that it was not a problem earlier. A felt but inaccurate RFD is still a felt RFD. When it comes to the issue of having a representation that gives you an authentic valuing, accuracy doesn't matter, since experienced-based RFDs are also vulnerable to being

inaccurate. This next green apple might be different from those previous green apples (maybe it's a new variety!); this second kid could be extremely different from the first (maybe she rather than he arrives via C-section). Relying on my past experience alone to imagine what those future experiences will be like will result in my having imagined something different from what in all likelihood will come to be. If these RFDs can be authenticity-conferring nevertheless, so too can an RFDs built inaccurately from testimony.

Another, more pressing worry for description-based RFDs is whether there are some features of what an experience is like that simply resist being described in a manner that can be translated into a felt imagining. For example, if I really concentrate, I think that I can produce very faintly the tartness of a green apple without actually eating an apple right now. Similarly, replaying the moment that I first held my son in the delivery room, I experience again that same warm and tender joy. However, I cannot, no matter how hard I try, reproduce at will the particular visceral pains of childbirth. It is, let's say, *translation-resistant*.<sup>71</sup> Does this threaten our ability to rely on descriptions to get authenticity-conferring RFDs? Can I authentically choose to have a child even if I cannot feel those particular pains outside of the moments of actually having a child?

I think the response to this worry is two-fold. First, there is again a parity argument to be had. Even as I undertake to imagine what it will be like to have a second child, I cannot construct a representation of that possibility that contains the corresponding pains of childbirth. This feeling, like a lot of pains, cannot be entertained at will. When there are features of my past experience that resist recall, I usually include them in my representation of those past experiences by deploying descriptions that I myself supply. I supplement the

---

<sup>71</sup> There are two, not mutually exclusive ways that something can be translation resistant. First, something can be particularly difficult to describe, so they are resistant in the experience-to-description direction. Second, something can be particularly difficult to imagine given even a rich description, and thus resistant in the description-to-experience direction.

imagistic content of my representation with some propositional content to mark that the thing being represented has such-and-such a feature.<sup>72</sup> I know from the experience of my first birth that it will be painful, I may even have a way of describing that pain that seems accurate. But the way that I use that experience in forming an idea of what my future possible child-bearing will be like is by constructing a scene with whatever qualitative content I can assemble and then assigning to that scene the relevant descriptive content. So, even in cases where I have the requisite experience, I have to rely on descriptions as I construct RFDs of my future possible outcomes. If relying on descriptions to represent the translation-resistant content does not block an RFD's authenticity-conferring status in these cases, then it should not in other cases.

The second response to this worry is to point out that, while there may be certain features of experience that cannot be entertained at will and that resist translation, many of the aspects of experience that are important for authentic valuing are not problematic in these ways. While pains and other particular bodily sensations seem to be particularly resistant to our attempts to exert agency over them, our feelings and moods are much less so. Joy, fear, dejection, agitation, jealousy, etc., *are* the sorts of things that we can be coaxed to feel through thought alone. Our engagement with reading fiction is, again, evidence of this. There are even less concrete feelings—like the to-be-avoided-ness that accompanies pains—that seem to be readily brought to mind. So though there may be some things that we can never feel by imagining, there still is a great deal that we can. To the extent that having a robustly felt understanding of our future possible outcomes is important for our

---

<sup>72</sup> Peter Kung thinks that this phenomenon of stipulating or “assigning” content (as he refers to it) is a regular feature of imagining, not just a feature of imagining in this special case where one has translation-resistant content. See Kung 2016, especially section 4.

authentically valuing those outcomes, it does not appear that a small class of translation-resistant features will get in the way of obtaining this robustness.

## 6. Conclusion

This, then, is a first sketch of how we can imagine what we haven't experienced in such a way that allows us to authentically value. There is a lot more that can be said—more about how we rely on previous experience and testimony, and more, too, about additional resources that we can use in concert with our imagination.<sup>73</sup> I won't, however, say more now. I hope that the picture outlined so far suffices to make a strong initial case for the powers of imagination, and thus a convincing case against the claim that we cannot appreciate our outcomes in a way that allows us to authentically value them unless we have experienced those outcomes.

To close, let's return once more and briefly to our Martian. Given our discussion above, what we should say about the Martian is this: it is at least *possible* for him to authentically value encumbered, rain-less days more than umbrella-less, soaked days (or vice versa!) despite his rain-less, umbrella-less existence. But to do so, he would have to do slightly different cognitive work than the Seattleite. He would have to gather enough of the relevant information (from others, from his past experiences) to be able to vividly imagine these possibilities so that he comes to have an idea of how it might be for him. When he assigns values to these possibilities on the basis of these ideas, these valuing can be authentic.

---

<sup>73</sup> One other that comes to mind now is when we observe others' experience, as when I develop a picture of what it is like to have a kid based on what I've witnessed of my brother's parenting experience. Dougherty, Horowitz, and Sliwa 2015 discusses this method.

It is certainly possible that things turn out differently from how the Martian imagined. Things may turn out so differently, in fact, that the Martian changes his mind about what he values. But this has no more bearing on the authenticity of his choice than it would on the rationality of his choice. Choosing well has never required us to be omniscient. It is also possible that the Martian will make inauthentic value judgments. He might not know the right people—i.e., people who have had the requisite experience and are good at describing it—or there might not be the right books on Mars. Or, as Nomi Arpaly has put it, he might just have a crappy imagination.<sup>74</sup> So much the worse for the Martian. But we now understand correctly why he has fallen short. It can be difficult to authentically choose what you haven't experienced, but this difficulty does not arise because we are inexperienced. It arises because we run into trouble imagining.

What is there to say about our initial, pre-theoretical comparative judgment of the Martian and Seattleite's respective valuing of getting stuck in the rain without an umbrella? Remember, we thought there was something odd about the Martian's valuing, an oddness not shared by the Seattleite's. The lesson from this paper is that we should be careful with this kind of judgment. Perhaps we tend to think of the person making value-judgments about what she hasn't experienced as doing so in an inauthentic manner à la the way that we, above, initially described choosing to eat a green apple on the basis of a friend's testimony or choosing to have a child after reading Gaskin's *Guide to Childbirth*. This tendency would explain our sense that experience-less value judgments are odd, but experienced-based ones

---

<sup>74</sup> Arpaly 2016. Arpaly's discussion of her lackluster imagination presses on a strange implication of my view, which is that you might not often choose authentically if you have a lackluster imagination. This implication is uncomfortable, not least because for some people it may not be within their power to improve their imaginative abilities to the point where they can produce the relevant kinds of representations. Fully addressing this worry would take us too far afield, but I do see it as a serious worry. But part of the response may be to point out that there are many kinds of authenticity and so even if there are some people that cannot authentically choose in the sense deployed in this discussion, there may be other kinds of authenticity that they can achieve.

are not. But now we know not to neglect to possibility that someone without experience has adequately used their imagination to make a value judgment. Maybe the Martian didn't do this. But maybe he did.

## Bibliography

- Arpaly, Nomi. 17 April 2016. What it's like to have a crappy imagination. *Special APA Pre-conference: Themes in Transformative Experience*. Westin Hotel, Seattle, WA. Conference presentation.
- Beard, J. 1990. *Personality correlates of the imposter phenomenon: An exploration in gender differences in critical needs*. Unpublished masters's thesis, Georgia State University, Atlanta.
- Brennan, Geoffrey & Southwood, Nicholas. 2007. Feasibility in action and attitude. In J. Josefsson D. Egonsson ed., *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*.
- Brighthouse, Harry and Swift, Adam. 2014. The place of educational equality in educational justice. In Kirsten Meyer (ed.), *Education, Justice and the Human Good: Routledge international studies in the philosophy of education*. London; New York: Routledge, 14-33.
- Brouwer, Frederick E. 1969. A difficulty with 'ought implies can'. *Southern Journal of Philosophy* 7/1: 45–50.
- Callard, A. 2016. Proleptic reasons. *Oxford Studies in Metaethics* 11. Oxford University Press. doi: <https://dx.doi.org/10.1093/acprof:oso/9780198784647.003.0006>
- Cath, Yuri. 2018. Knowing What it is like and testimony. *Australasian Journal of Philosophy* 97/1: 105–120.
- Chang, Ruth. 2015. Transformative Choices. *Res Philosophica* 92/2: 237-282.
- Clance, P. R., & Imes, S. A. 1978. The imposter phenomenon in high achieving women: Dynamics and therapeutic intervention. *Psychotherapy Theory, Research, and Practice* 15, 241–247.
- Clance, P. R., & O'Toole, M. A. 1987. The imposter phenomenon: An internal barrier to empowerment and achievement. *Women and Therapy* 6, 51–64.
- Cozzarelli, C, & Major, B. 1990. Exploring the validity of the impostor phenomenon. *Journal of Social and Clinical Psychology* 9, 401–417.
- Dougherty, Tom; Horowitz, Sophie & Sliwa, Paulina. 2015. Expecting the unexpected. *Res Philosophica* 92/2: 301–321.
- Estlund, David. 2011. Human nature and the limits (if any) of Political Philosophy. *Philosophy and Public Affairs* 39/3: 207–237.
- Estlund, David. 2016. Reply to Wiens. *European Journal of Political Theory* 15/3: 353–362.
- Flewelling, A. L. 1985. *The impostor phenomenon in individuals succeeding in self-perceived atypical professions: the effects of mentoring and longevity*. Unpublished master's thesis. Georgia State University, Atlanta



- Gheaus, Anca. 2015. Three Cheers for the Token Women! *Journal of Applied Philosophy* 32/2:163–176.
- Gilabert, Pablo & Lawford-Smith, Holly. 2012. Political feasibility: A conceptual exploration. *Political Studies* 60/4: 809–825.
- Graham, P. A. 2011. 'Ought' and ability. *Philosophical Review* 120/3 :337—382.
- Harman, Elizabeth. 2015. Transformative experience and reliance on moral testimony. *Res Philosophica* 92/2: 323–339. <http://dx.doi.org/10.11612/resphil.2015.92.2.8>.
- Harvey, J. C. 1981. *The impostor phenomenon and achievements: A failure to internalize success*. Doctoral dissertation. Retrieved from Dissertation Abstracts International, 42, 4969–4970B.
- Harvey, J. C., & Katz, C. 1985. *If I'm so successful, why do I feel like a fake?: The impostor phenomenon*. 1st ed. New York, NY: St. Martin's Press.
- Hawley, Katherine. 27 July 2016. Philosophical reflections on Imposter Syndrome. *Munich Center for Mathematical Philosophy Summer School for Female Students*. Munich, Germany. Keynote address. Available online at: <https://videoonline.edu.lmu.de/en/sommersemester-2016/8140>
- Hicks, Amelia. 2018. Moral uncertainty and value comparison. In Russ Shafer Landau (ed.), *Oxford Studies in Metaethics, Volume 13*. Oxford, UK: 161-183.
- Jackson, Frank. 1986. What Mary didn't know. *Journal of Philosophy* 83/5: 291–295.
- Jeshion, Robin. 2002. Acquaintanceless De Re belief. In Joseph Keim Campbell, Michael O'Rourke & David Shier (eds.), *Meaning and Truth: Investigations in Philosophical Semantics*. Seven Bridges Press, 53-74.
- Kauppinen, Antti. 2015. What's so great about experience? *Res Philosophica* 92/2: 371-388.
- Kolligian, Jr., J., Sternberg, R. J. 1991. Perceived fraudulence in young adults: Is there an “imposter syndrome”? *Journal of Personality Assessment* 56/2, 308-326.
- Kung, Peter. 2016. You really do imagine it: Against error theories of imagination. *Noûs* 50/1: 90-120.
- Lawford-Smith, Holly. 2013. Understanding political feasibility. *Journal of Political Philosophy* 21/3: 243-259.
- Lawford-Smith, Holly. 2013. Non-Ideal accessibility. *Ethical Theory and Moral Practice* 16/3: 653-669.
- Lawler, K. N. 1984. *The impostor phenomenon in high achieving persons and Jungian personality variables*. Doctoral dissertation. Retrieved from Dissertation Abstracts International, 45B, 2693.

- Leslie, Sarah-Jane et al. 2015. Expectations of brilliance underlie gender distributions across academic disciplines. *Science* 347, 262-265. doi: 10.1126/science.1261375
- McGinn, Colin. 2009. Imagination. In Brian McLaughlin, Ansgar Beckermann & Sven Walter (eds.), *The Oxford Handbook of Philosophy of Mind*. Oxford University Press.
- Meyer M, Cimpian A and Leslie S-J. 2015. Women are underrepresented in fields where success is believed to require brilliance. *Front. Psychol.* 6: 235. doi: 10.3389/fpsyg.2015.00235
- Murphy, Mary C. and Carol S. Dweck. 2010. A culture of genius: How an organization's lay theory shapes people's cognition, affect, and behavior. *Personality and Social Psychology Bulletin* 36/3, 283–296.
- Nanay, Bence. 2016. The role of imagination in decision-making. *Mind and Language* 31/1: 126-142.
- Toole, John Kennedy. 1980. *A confederacy of dunces*. New York: Grove Press.
- Olberding, Amy. 6 March 2018. The outsider. *Aeon*. Last retrieved from <https://aeon.co/essays/how-useful-is-impostor-syndrome-in-academia> on 22 June 2018.
- Paul, L. A. 2014. *Transformative experience*. Oxford University Press.
- Paul, L. A. 2015a. What you can't expect when you're expecting'. *Res Philosophica* 92/2: 1-23.
- Paul, L. A. 2015b. Transformative choice: Discussion and replies. *Res Philosophica* 92/2: 473-545.
- Paul, Sarah & Morton, Jennifer. 2018. Grit. *Ethics* 129/2: 175-203.
- Saka, Paul. 2000. Ought does not imply can. *American Philosophical Quarterly* 37/2: 93-105.
- Pettigrew, Richard. Unpublished manuscript. *Choosing for Changing Selves*.
- Rawls, John. 1971. *A Theory of Justice*. Harvard University Press.
- Schwan, Ben. 2018. What ability can do. *Philosophical Studies* 175/3: 703-723.
- Simmons, Dana. 2016. Impostor Syndrome, a reparative history. *Engaging Science, Technology, and Society* 2, 106-127.
- Sosis, Cliff. 5 October 2017. Rebecca Tuvel. *What it's like to be a philosopher?* Last retrieved from <http://www.whatisitliketobeaphilosopher.com/rebecca-tuvel/> on 22 June 2018.
- Southwood, Nicholas & Wiens, David. 2016. "Actual" does not imply "feasible". *Philosophical Studies* 173/11: 3037-3060.
- Tanesini, Alessandra. 2016. Intellectual humility as attitude." *Philosophy and Phenomenological Research*. Version of Record online: 4 AUG 2016 | DOI: 10.1111/phpr.12326

- Ullmann-Margalit, Edna. 2006. Big decisions: opting, converting, drifting.” *Royal Institute of Philosophy Supplement* 58: 157–172. <http://dx.doi.org/10.1017/S1358246106058085>.
- Ullman-Margalit, Edna. 2007. Difficult choices: To agonize or not to agonize?” *Social Research* 71/1: 51–78.
- Vranas, Peter B. M. 2007. I ought, therefore I can. *Philosophical Studies* 136/2: 167-216.
- Wiens, David. 2012. Prescribing institutions without ideal theory. *Journal of Political Philosophy* 20/1: 45-70.
- Wiens, David. 2016. Motivational limitations on the demands of justice. *European Journal of Political Theory* 15/3: 333-352.
- Winter, Caroline. Towing an iceberg: One captain’s plan to bring drinking water to 4 million people. *Bloomberg News*. <https://www.bloomberg.com/news/features/2019-06-06/towing-an-iceberg-one-captain-s-plan-to-bring-drinking-water-to-4-million-people>. Last accessed 15 June 2016.