

PHILOSOPHICAL INQUIRIES IN TEACHING AND LEARNING

By
Danielle L. Clevenger

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

Doctor of Philosophy
(Philosophy)

at the
UNIVERSITY OF WISCONSIN-MADISON
2026

Date of final oral examination: 04/20/2026

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

Harry Brighouse, Mildred Fish Harnack Professor of Philosophy, Philosophy

James Goodrich, Assistant Professor, Philosophy

David O'Brien, Assistant Professor, Educational Policy Studies

Elizabeth Brake, Professor, Philosophy

W. John Koolage, Professor, Philosophy

© Copyright by Danielle L. Clevenger 2026

All Rights Reserved

Acknowledgements

It is cliché to say that this dissertation would not have come to be without my advisors, but it is true. My chair, Harry Brighthouse, has advocated for me since the beginning of graduate school, and I remain deeply grateful for his encouragement and for helping me shape a dissertation genuinely aligned with my interests and values as a philosopher. Jimmy Goodrich's office often felt as much like home as my own. Though he would dispute it, his feedback, humor, and candid pep talks are chief among the reasons I made it to the finish line. David O'Brien and Elizabeth Brake offered insightful comments throughout, but also something rarer: the gift of making one feel intelligent, heard, and valued in philosophical conversation. I hope to carry myself with their grace and warmth in the future. John Koolage has my lasting gratitude for all he has done for me. Some people say they learned everything they needed to know about being a good person in kindergarten. I learned everything I needed to know about the kind of philosopher and teacher I hoped to become from John in his Introduction to Philosophy course.

To my partner and husband, Liam Thomas-Clevenger: I could write a million pages. Thank you is not enough, but it is still exactly right. Your boundless love and support, expressed in countless ways—including proofreading this entire dissertation—have made all the difference.

During my time in graduate school, I have been graced with the gift of being sustained by many kinds of family. First and foremost, my own family, especially my mom Julie, my sister Jordan, and my grandmother Kathy, carried me through the hardest parts of this process. Their faith in me held steady when my own did not. Also, shoutout to my niece and nephew for being the best possible comedic relief exactly when I needed it most. My gratitude also extends to my family-in-law, whose steady encouragement has meant so much.

I am especially grateful to David Vessel for his friendship and for being the world's best cornerman. His intelligence, work ethic, and loyalty have been a continual source of admiration and inspiration. I also owe him thanks for welcoming me into his family. To Madeleine Vessel, the whole Vessel Clan, plus honorary member Mark Walker: this dissertation is better, and its author steadier, because of your support.

Words cannot adequately express what Emily Fletcher's friendship has meant to me. Her beautiful spirit has shaped not only my life, but the philosophy department as a whole. My Madison family, including Umberto, Nebula, Amy, and Gary, have brought joy and comfort to these years in ways both large and small. Their questions, observations, and support left their mark on this dissertation.

No acknowledgment of friendship here could ever be complete. To name every friend who mattered to the writing of this dissertation would take far more pages than I have. I hope that, in thanking a few here, I may also express something of what I owe to many others. Michael Bruckner has been one of my closest friends since we began graduate school together, and his advice, empathy, and understanding have been indispensable. Crystal Seet and Matt Maxwell were the best office neighbors one could ask for, keeping me grounded and helping me remain true to myself throughout graduate school. Though farther away, the enduring friendship of Gabriel Stroe and Kelsey DeGuise has meant just as much. I am grateful that they have always been there when I needed them most and let me talk their ears off about matters both inside and outside philosophy.

I also want to acknowledge Lynne Prost for her generous mentorship. Working with her was a highlight of my time at Madison, and much of this dissertation bears the imprint of that

influence. I would be remiss not to mention the Philosophy Front Office, which has long been one of my favorite parts of the department. Wendy Crabb, Miriam Paris Fagan, and Charlene Sime, went above and beyond to make me, and all the students in the department, feel welcomed and supported. Finally, to all of my students: thank you for everything you have taught me. My very first students, Aksel Balkir, Avra Reddy, Wyatt Schoephoerster, and Nicolas Marchan, in particular, deepened my understanding of what it means to do philosophy. Your excitement and passion were, and remain, a beautiful gift.

Table of Contents

<i>Acknowledgements</i>	<i>i</i>
<i>Dissertation Introduction</i>	<i>1</i>
<i>Paper One: What is Active Learning? A Philosophical Audit</i>	<i>5</i>
Abstract for Paper One	5
1.I Introduction	5
1.II Philosophy of the Learning Sciences	7
1.III Concerning Trends in Active Learning Literature	8
1. IV What, exactly, is wrong with active learning?	12
1.IV.a Negation-Type Definitions	15
1.IV.b Impotent Definitions	22
1.IV.c No Definition	26
1.V Where do we go from here?	29
<i>Paper Two: Embodied Bootstrapping: An Embodied Account of Concept Learning</i>	<i>34</i>
Abstract for Paper Two	34
2.I Introduction	35
2.II Bootstrapping	36
2.III Hard-Route Bootstrapping	39
2.III.a Step One: Acquisition of a Placeholder Symbol Set (PSS).....	39
2.III.b Step Two: Computational Constraints Partially Interpret the PSS	40
2.IV Embodied Cognition and Embodied Bootstrapping	44
2.V An Embodied Bootstrapping Process	47
2.V.a Embodied Step One: Acquire a Placeholder Symbol Set (PSS).....	47
2.V. b Embodied Step Two: Embodied Constraints Partially Interpret the PSS.....	47
2.V.c Embodied Step Three: Full Interpretation of Placeholder Symbol Set via Modeling	50
2.VI Strengthening the Theory of Bootstrapping	53
2.VII The Deviant Interpretation Challenge	53
2.VII.b Learners Acquire Similar Concepts	59
2.VII.c Learners Hold Concepts Consistently	60
2.VIII Conclusion	61

Paper Three: Care, Justice, and the Duty of Pedagogical Competence	62
Abstract for Paper Three	62
3.I Introduction	63
3.II The Natural Alignment of Justice and Higher Education	64
3.III Grounding Teaching Obligations in Justice	67
3.IV Grounding Teaching Obligations in Care	73
3.V The Content of Teaching Duties of Care	77
3.V.a Caring for Students as Students	79
3.V.b Caring for Students as People	80
3.V.c Caring for Students as Individuals	81
3.VI Objection 1: It's All Style	83
3.VII Objection 2: It's Too Demanding	86
3.VIII Conclusion	91
Bibliography	93

Dissertation Introduction

Teaching and learning matter because they shape how people come to understand the world and how they are able to navigate within it. Learning is not merely the accumulation of information. It is developing the knowledge, skills, and dispositions that undergird autonomy and allow people to meaningfully set and pursue their own ends. This dissertation seeks to show that philosophy and teaching and learning are natural allies. Teaching and learning needs philosophy, and philosophy, in turn, needs teaching and learning.

The papers collected here are united by the thought that philosophy, the learning sciences, and the scholarship of teaching and learning should be in closer conversation than they often are. Each paper stages one such conversation. In each case, philosophical analysis is brought to bear on questions that arise within the study or practice of learning or teaching. At the same time, empirical pedagogical work is treated not as merely illustrative of philosophical concerns, but as something that can and should bear on philosophical theorizing itself. The result is not a single linear argument, so much as a set of philosophical interventions into a shared domain. Together, these papers explore some of the ways in which philosophy can clarify, challenge, and enrich our understanding of teaching and learning, and some of the ways in which attention to teaching and learning can sharpen and redirect philosophical inquiry.

The first paper, “What is Active Learning? A Philosophical Audit,” examines one of the most prominent categories in contemporary pedagogy—active learning. Active learning is widely promoted as a pedagogical “best practice,” credited with improving student outcomes and advancing classroom equity. Yet the empirical literature on AL is marked by inconsistency: studies frequently reach contradictory conclusions, and interventions that appear successful in

controlled contexts often fail to reproduce more broadly. This paper argues that these problems stem not merely from methodological shortcomings but from conceptual deficiencies in the category of active learning itself. Drawing on philosophy of science and feminist philosophy, this paper develops criteria for scientifically adequate concepts and applies them to prevailing definitions of AL. It shows that current conceptions fall into two problematic types: negation-style definitions, which render the concept untestable, and impotent definitions, which fail to support inductive generalizations. The widespread tendency to forgo definitions altogether compounds this instability. The paper concludes by considering implications for AL's status as a best practice and sketching directions for conceptual refinement.

The second paper, "Embodied Bootstrapping: An Embodied Account of Concept Learning," brings the dissertation's central theme into view from the other direction. Here the question is not how philosophy can help clarify educational inquiry, but how empirical work on learning can help assess and improve a philosophical theory. This paper argues that hard-route conceptual bootstrapping is significantly strengthened when it is formulated within an embodied view of mind, rather than a narrowly computational one. Empiricists appeal to bootstrapping to explain how learners acquire genuinely new conceptual primitives and, thereby, expand their expressive power beyond the mere recombination of concepts they already possess. However, traditional accounts of bootstrapping rely too heavily on computational assumptions and, as a result, fail to capture important features of how radical conceptual change actually proceeds. Drawing on recent work in embodied cognition, I develop an alternative account of bootstrapping that better explains how learners come to acquire new conceptual primitives through structured engagement with their bodies and environments. An embodied framework, I argue, both broadens the range of cases that bootstrapping can illuminate and provides a richer account of the mechanisms by

which new concepts are formed and stabilized. It also offers a principled response to the Deviant Interpretation Challenge, a prominent objection to bootstrapping theories. The paper thus develops an account of embodied bootstrapping, illustrates its operation through a concrete example, and argues that embodiment is not merely an auxiliary feature of learning but a theoretically significant component of bootstrapping itself. Understanding bootstrapping in this way has important implications for how we explain both the acquisition and the stability of genuinely new conceptual primitives.

The third paper, “Care, Justice, and the Duty of Pedagogical Competence,” turns from conceptual questions to normative ones. In some ways it builds from the earlier two papers to argue for ethical claims based, in part, on what is previously established regarding teaching and learning. Implementing good pedagogy in one’s teaching in higher education is morally significant. This paper argues that higher education instructors have a moral duty to bear the necessary and proportionate costs of becoming pedagogically competent. Such competence matters morally because higher education plays a central role in distributing socially and economically significant goods, while also helping to reproduce existing inequalities. Even so, concerns of justice cannot alone explain all instructors owe their students. Although it is crucial for understanding the broader institutional context, it does not fully account for the immediate, role-specific obligations that arise within the classroom itself. I argue that those obligations are grounded in care because the teacher-student relationship is structured by dependence, asymmetrical authority, and entrusted vulnerability. Drawing on Edward Hult Jr.’s account of pedagogical caring, I show that instructors must care for students qua students, as persons, and, where appropriate, as unique individuals. This account clarifies both the demands and the limits of pedagogical duty. Instructors are not required to become therapists or to take unlimited

responsibility for students' lives, but they are required to adopt reflective, revisable practices that support learning and minimize avoidable harms. The paper concludes by addressing two objections: (1) that good teaching is merely a matter of style rather than substantive design; (2) that the duty of pedagogical competence is overly demanding.

This dissertation demonstrates that some of the most important questions in philosophy and beyond begin precisely where disciplinary boundaries become porous. These papers aim less to close debates than to begin and deepen them and are offered as invitations to further conversation among philosophers, learning scientists, and Scholarship of Teaching and Learning scholars about questions that matter not only for theory, but for students' lives.

Paper One: *What is Active Learning? A Philosophical Audit*

Abstract for Paper One

Active learning is widely promoted as a pedagogical “best practice,” credited with improving student outcomes and advancing classroom equity. Yet the empirical literature on active learning is marked by inconsistency: studies frequently reach contradictory conclusions, and interventions that appear successful in controlled contexts often fail to reproduce more broadly. This paper argues that these problems stem not merely from methodological shortcomings but from conceptual deficiencies in the category of active learning itself. Drawing on philosophy of science and feminist philosophy, this paper develops criteria for scientifically adequate concepts and applies them to prevailing definitions of active learning. It shows that current conceptions fall into two problematic types: negation-style definitions, which render the concept untestable, and impotent definitions, which fail to support inductive generalizations. The widespread tendency to forgo definitions altogether compounds this instability. The paper concludes by considering implications for active learning’s status as a best practice and sketching directions for conceptual refinement.

1.1 Introduction

Empirical studies on active learning have mixed results. One of the most widely cited studies on active learning reports that students who were in traditional lecture-based courses were 1.5 times more likely to fail than students in active learning courses.¹ Studies have also found that when active learning is employed the gaps between privileged students’ outcomes and under-resourced

¹ Scott Freeman et al., (2014).

students' outcomes (often called "achievement gaps") shrink by 33% for exam pass rates and by 45% for course pass rates.² However, other studies have reported no association between active learning and student learning gains.³ Still other studies find much more modest effects on student learning.⁴ When empirical studies produce differing results regarding the same intervention, it raises concerns about the consistency, reliability, and generalizability of the findings. Such concerns make it challenging to form definitive conclusions about the intervention, and in this case specifically, provide clear recommendations for practicing teachers.

I argue that the inconsistencies found in the literature are not best explained by flaws in experimental design or by difficulties of classroom implementation. Rather, they point to a deeper problem: the concept of active learning itself is ill-suited for the kind of rigorous empirical inquiry it is expected to sustain. Although concerns about the concept of active learning are well documented in meta-reviews, what has not been clearly articulated is the precise nature of its conceptual deficiencies—a gap this paper aims to address.⁵ To demonstrate this, I categorize the most employed conceptions of active learning into one of two categories—negation-type definitions or impotent-type definitions. Each category reveals a distinct issue with using the definitions in its domain to understand active learning. Negation-type definitions result in an undefined category that makes any hypothesis which employs it untestable. Impotent-type definitions create a category of active learning that is too heterogenous to license a strong induction. In either case, these conceptions have stymied scientific inquiry into active learning's

² Theobald et al., (2020).

³ T. M. Andrews et al., (2011).

⁴ Wilke & Straits, (2001).

⁵ For example see: Drake, (2012); Doug Lombardi et al., (2021); Driessen et al., (2020); Bernstein, (2018).

efficacy and equity. As a result, the literature on active learning shows concerning trends of inconsistency and a lack of replicability.

This paper develops this claim further and charts a path towards improvement. Section I briefly primes the conceptual issues with active learning and speaks to philosophy's role in sorting out such issues. Section II surveys the empirical literature to highlight two concerning trends—contradictory conclusions and reproducibility failures. Section III uses criteria from philosophy of science and feminist philosophy to assess what makes a concept scientifically adequate and applies these standards to current definitions of active learning. This analysis reveals that existing conceptions fall into multiple flawed types. It also highlights the widespread tendency to forgo definitions altogether, which further compounds the problem. Finally, Section IV examines the implications of this instability for active learning's designation as a “best practice” and considers emerging alternatives and refinements. The aim is not to provide a definite new conception of active learning but to diagnose the deficiencies of current ones and to encourage more rigorous conceptual work to support empirical inquiry.

1.II Philosophy of the Learning Sciences

The learning sciences investigate which educational means conduce to learning and why. This paper is a contribution to the *philosophy* of the learning sciences. Like much philosophy of science, it begins from a problem that currently vexes the scientific field in question, and it aims to help resolve that problem by employing some tools of philosophical analysis to diagnose the source of the problem and help scientists in the field advance their work about it. In this case, the puzzle concerns the best practice of active learning.

As a way of intuitively motivating that there is a conceptual problem here, consider the following three activities: a think-pair-share; a reading exercise using manipulable figurines; and constructing a working trebuchet. The first activity, a think-pair-share, is a brief activity where students are given a moment to reflect on their answer to a question, which they then discuss with a nearby peer before answers are shared out to the whole class. The second activity, reading with manipulables, is a reading comprehension activity where students use figures to act out key actions or plot points of a story. For example, if a cow moves from the barn to the pasture, then the student might take a toy cow and move it across the table to a spot that has been designated as the pasture. The third activity, constructing a working trebuchet, is a sustained activity where engineering students work together to make a machine that meets certain qualifications. The differences in the time, effort, locomotion, and collaboration required by the three activities are considerable. Given these significant disparities, it is unclear why we ought to think they all belong in the same category. They are clearly very different learning activities employed for different reasons and in different settings. Yet, evidence drawn from investigation of these activities is used to draw conclusions about a whole category—active learning.

1.III Concerning Trends in Active Learning Literature

There are two concerning trends in the literature on active learning that, I argue, stem from inadequacies in the concept itself: inconsistent conclusions and replicability failures. This first section is dedicated to briefly establishing these issues. My claim of inconsistent conclusions is evident in the wide variability of reported outcomes regarding active learning's efficacy and

equity: some studies show significant positive effects⁶, others more modest ones⁷, and still others none at all.⁸ For instance, Michele Shuster and Ralph Preszler's study of two introductory biology courses (Bio 111 and Bio 112) exemplifies this inconsistency.⁹ Their investigation, like many others, sought to determine whether adopting active learning would enhance student outcomes and promote equity, yet the results echoed the literature at large by producing findings that were mixed and difficult to reconcile.

In the study, four different teaching models were employed over 16 semesters from Fall 2003 to Spring 2011. The teaching models included a control group (lecture) and three different models of active learning (LC, WRK, LCBC). What they found, is that while the active learning models generally performed better than the control group, *how much* better and *for whom*, varied significantly based on the specific model being employed. Relationships between course grades and course model, gender, and ethnicity were evaluated using two-way and three-way contingency table analyses. In all cases, if the probability associated with the Pearson χ^2 was

⁶ Adams & Dove, (2018); Deslauriers, Schelew, & Wieman, (2011); Di Vesta & Smith (1979); Dori & Belcher, (2005); Freeman et al., (2014); Hake, (1998); Hernández-de-Menéndez et al., (2019); Hewapathirana & Almasri, (2022); Johnson, Johnson, & Smith (1998); Knight & Wood, (2005); Laws, Sokoloff, & Thornton, (1999); McCarthy & Anderson, (1999); Redish, Saul, and & Steinberg, (1997); Ruhl, Hughes, & Schloss, (1987); Shuster & Preszler, (2014); Sokoloff, Laws, & Thornton, (2007); Wieman, a(2017) and b(2014).

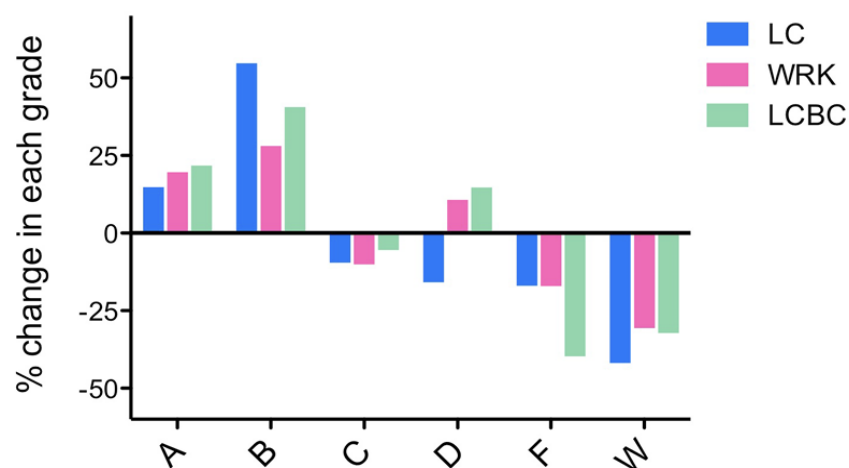
⁷ Butler, Phillmann, & Smart: (2001); Cant & Cooper (2010); Ebert-May, Brewer, & Allred, (1997); Felder et al., (2000); Gallagher, (1997); McKeachie, (1972); Michel, Cater III, & Varela, (2009); Nurbavliyev, Kaymak, & Sydykov, (2022); Richmond & Hagan, (2011); Ruiz-Primo et al., (2011); Shi et al., (2019); Sorcinelli, (1991); Stead, (2005); Turpen & Finkelstein, (2009); Venton & Pompano, (2021); Wilke & Straits, (2001).

⁸ Adams & Dove, (2016); Adams et al., (2011); Aguilon et al., (2020); Albanese & Mitchell, (1993); Andrews et al., (2011); Berkson, (1993); Colliver, (2000); Copeland, Scott, & Houska, (2010); Elliott et al., (2010); Goffe & Kauper, (2014); Hartling et al., (2010); Hilligoss, (1992); Hood et al., (2021); Kapur et al., (2022); Kirschner, Sweller, & Clark, (2006); Louis & Brown, (2017); Maries et al., (2020); Mayer, (2004); Rondon, Sassi, & De Andrade, (2013); Weltman & Whiteside, (2010).

⁹ Shuster & Preszler, (2014).

<0.01, the authors concluded that the variable(s) in question had a significant impact on student grades. The active learning models showed statistically significant (Pearson chi-squared $p < 0.001$) effects on student grades. Figure 1 shows the percent change of grade distribution based on the active model as compared to the lecture-based control group.¹⁰

Figure 1. Percent Change of Grade Distribution by Active Learning Model



The overall conclusion that the active models performed significantly better is apparent in the increased proportion of “As” and “Bs” and the decreased proportions of “Fs” and “Ws”.

However, note the differences in the strength, and in some cases, the direction of the changes between the three active models. Although all three conditions are considered active, their effects on grade distribution varied from condition to condition. Thus, if three different researchers studied one of these conditions, comparing it to a control group, they would come to different conclusions about the efficacy of active learning. This is especially true when such determinations become complicated by concerns of whether the improvements seen are worth the resources spent to achieve them.

¹⁰ Ibid, 13.

A similar pattern of inconsistency emerges when we turn from overall efficacy to questions of equity. Active learning is often championed as a best practice because it is thought to offer particular benefits for under-resourced student populations. Yet, this study found no significant differences in grade distributions between males and females when data were pooled across all course models (Pearson chi-squared $p = 0.75$).¹¹ On this basis, one might conclude that active learning does not provide disproportionate benefits for under-resourced groups. However, such averaging obscures important variation. When the data were disaggregated by active learning model, significant gender differences emerged (Pearson chi-squared $p < 0.001$).¹² Here, we see vividly how different approaches to analyzing the same data can yield contradictory conclusions: someone conducting a broad meta-analysis might conclude that active learning had no effect on equity, while a focused study of a single model would indicate that active learning does matter. This tension exemplifies the larger difficulty in interpreting the literature on active learning's equitability.

Beyond these contradictory conclusions, a second problem concerns a failure of reproducibility. While this challenge is familiar in other sciences, including psychology, it is especially pressing for active learning because it is so often promoted as a best practice—a designation reserved for teaching methods that have consistently demonstrated effectiveness across diverse contexts. The enthusiasm surrounding active learning is fueled by the expectation that it can be empirically validated as both effective and equitable. Yet, a study analyzing the largest random sample of college science instructors found no association between active learning and student learning

¹¹ In some disciplines, such as science, technology, engineering, and mathematics (STEM), women count as an under-resourced population. If active learning was more equitable, as is often claimed, then this result would tell against that.

¹² Ibid, 14.

gains.¹³ The authors suggest this discrepancy arises because prior studies were often conducted by education researchers studying their own courses, where fidelity to the intervention was high. Thus, when a random sample was obtained, the positive results achieved by education researchers studying their own courses, was not seen in the larger, more representative sample. This finding is consistent with other studies which demonstrate that instructors often inadvertently alter or remove crucial elements of new practices, rendering them ineffectual.¹⁴

The inconsistent findings and failures of reproducibility discussed above suggest a more fundamental problem. Rather than reflecting shortcomings in experimental design, these difficulties point to the inadequacy of “active learning” as a scientific concept—one too imprecise to sustain rigorous empirical inquiry. The next section explores this conceptual challenge.

1. IV What, exactly, is wrong with active learning?

Because my central claim is that active learning is unsuitable for rigorous empirical inquiry, it helps to begin by briefly considering what constitutes a good conception or definition. A good conception should possess clarity and precision to avoid vagueness and ambiguity, be non-circular, and maintain internal coherence. It should tie the concept to reality and differentiate it from other nearby concepts. Additionally, it should live in the goldilocks zone—not too restrictive as to discount legitimate instances of the concept, but not so permissive that anything and everything counts as a legitimate instance. The boundaries it draws should also help explain

¹³ Andrews et al., (2011).

¹⁴ Borrego et al., (2013); Daubenmire et al., (2015); Henderson & Dancy, (2009); Stains & Vickrey, (2017); Turpen, & Finkelstein, (2009).

why some examples are legitimate instances of the concept and why others are not. I take these criteria to be generally mundane and easily assented to by philosophers and scientists alike.

It can also be argued that a concept might be legitimately judged for its effectiveness in serving certain aims. For instance, active learning is (supposedly) a scientific concept. This means that conceptions of active learning can, and should be, assessed as to how well they facilitate fundamental scientific activities like hypothesis formation, experimentation, and modeling.¹⁵

There are some who take this idea even further and argue that not only should concepts be judged based on how well they facilitate certain ends but that the concepts themselves should be understood as being comprised, in part, by an epistemic goal.¹⁶ Taking a concept's aims into account is not a new proposition. Feminist philosophy exemplifies the progress achievable when an ameliorative approach to conceptual analysis is adopted.¹⁷ Sally Haslanger's landmark article, "Gender and Race: (What) Are They? (What) Do We Want Them To Be?" is notable for its explication of this methodological approach. She distinguishes between descriptive and analytical (or ameliorative) projects.¹⁸ The former reveals our actual use of concepts, while the latter suggests what concepts we should construct and use, given our inquiry's aims. For the purposes of this paper, the descriptive approach is helpful because it provides a starting point for diagnosing the failures of current active learning conceptions. However, once we have a clear grasp of why the current conceptions are inadequate, the ameliorative approach should be adopted. This framework clarifies the task ahead. After diagnosing the shortcomings of current

¹⁵ Carballo, (2020); 304.

¹⁶ Brigandt, (2010): 19–40.

¹⁷ This distinction was earlier referred to as "analytic" rather than "ameliorative".

¹⁸ Haslanger, (2000): 31–55.

conceptions of active learning, more pointed attention can be turned to creating a better version of the concept.

Current definitions of active learning suffer from two related, but distinct, issues. In order to flesh them out it is helpful to categorize the most common definitions used into one of two categories: Negation-Type Definitions and Impotent Definitions. Each category elucidates a distinct challenge with using definitions that fall under its domain to produce valid and generalizable findings regarding active learning. Before turning to these two categories in detail, it is worth first dispelling what I call the naïve reading of active learning. The naïve reading of active learning says: active learning was only ever meant to pick out those learning of activities that met the criteria of non-lecture and effective. This conception of active learning seems easily dismissed because, if this conception was adopted, it would deflate active learning's usefulness as a scientific concept. It is not capable of creating the kind of category that is useful to researchers. The enthusiasm around active learning is driven by the belief that it can be empirically proven to be highly effective and equitable. Thus, it is fundamental to the epistemic aims of the concept that it be testable, capable of supporting new observations, predictions, and experiments that advance our understanding of teaching and learning. In the learning sciences and Scholarship of Teaching and Learning (SoTL) discourse, active learning is treated as a category whose members share meaningful properties, not as a loose, heterogeneous collection of learning activities or methods. Thus, the naïve conception fails because it is incapable of producing a precise category. A clear, precise category is not only desirable for advancing the epistemic aims of active learning; it is also necessary if we are to make sense of the way it is used in scholarly discourse.

1.IV.a Negation-Type Definitions

Negation-type definitions are those that classify active learning as a negation—the opposite or absence of passive learning. Any understanding gained from the definition is made most salient or meaningful due to its contrast with the supposedly clearer concept of passive learning. Some examples of common definitions that fall into this category are:

1. “Active learning is, in short, any learning activity engaged in by students in a classroom other than listening passively to an instructor’s lecture.”¹⁹
2. “Active learning engages students in the process of learning through activities and/or discussion in class, as opposed to passively listening to an expert. It emphasizes higher-order thinking and often involves group work.”²⁰
3. “The term “active learning” is a concept that is known more for what it is not. The emphasis on active learning originated as a contrast to the perceived “passive” atmosphere of the lecture.”²¹
4. “Active learning simply means getting involved with the information presented, really thinking about it (analyzing, synthesizing, evaluating) rather than just passively receiving it and memorizing it.”²²

Definitions (2) and (3) might be thought of as enhanced versions of a negation-type conception. They still forefront a conception of active learning characterized by what it is *not*, but they do

¹⁹ Faust & Paulson (1998):4.

²⁰ Freeman et al., (2014): 8413-14.

²¹ Pelley, (2014): S13.

²² King, (1993): 31.

attempt to articulate some positive conceptual features as an addition. For example, Definition (2) adds at least one criterion (higher-order thinking) and gestures to an example (group work). Nonetheless, the emphasis is still placed on the contrastive element (as opposed to passively listening to an expert). An analysis of these definitions reveals a concerning trend. The SoTL and Learning Sciences literature are dominated by negation-type definitions, and the problem is that such definitions render the concept of active learning undefined and, therefore, empirically untestable.

To see why this matters, we need to consider what is required for a concept like active learning to function scientifically—namely, the ability to generate clear, testable hypotheses. To be empirically testable, a hypothesis must do two specific things: (a) it must make observable predictions and (b) it must be contrastable with another hypothesis.²³ If either of these conditions are not met, then the hypothesis fails to be the sort of thing that can be empirically investigated. Criterion (a) can be restated another way: testing a hypothesis requires that it make a prediction that can be checked by observation.²⁴ If a hypothesis fails to meet this requirement, it is not the sort of thing that can be empirically investigated. Consider Hypothesis one (H1) which says, “there are an even number of ornaments on my tree.” This hypothesis (H1) meets the first criteria of testability because it specifies a circumstance we are able to observe. Namely, it makes a prediction (even number of ornaments) that we can check by counting the number of ornaments on my tree. Compare this to Hypothesis two (H2) which says, “there are an even number of ornaments on my tree only when you aren’t looking at it.” (H2) fails to meet criterion (a) because

²³ Sober, (1999): 52.

²⁴ Ibid, 49.

the prediction it makes (even when not looking) cannot, in principle, be checked. There is no counting, no observation, that can be gathered to support or undermine this hypothesis.

Understanding criterion (b) is slightly more complicated than (a). There are two reasons that testability ought to be understood as an inherently contrastive activity. The first can be illustrated by a modified example. Hypothesis three (H3) says “there are an even number of ornaments on my tree.” Hypothesis four (H4) says “there are an odd number of ornaments on my tree.” (H3) and (H4) are testable because there are observations that will discriminate between them. In other words, these hypotheses are empirically testable because we can collect data, say by counting the ornaments on my tree, which will provide evidence that favors either (H3) or (H4). The evidence we collect will give us reason to adopt one of the two specific hypotheses.

However, consider the following modification: (H3’) “There are an odd number of ornaments on my tree only when someone is looking” and (H4) “There are an odd number of ornaments on my tree”. With this modification the hypotheses suddenly become untestable—there is no observation of any kind that could be made that would discriminate between the two. Any observations that could be made will support each hypothesis equally. If the hypotheses cannot, in principle, be distinguished on the basis of observation then it is not possible to test them empirically. Note that (H4) remains the same in both examples. This is why it is odd to say that (H4) itself is testable or untestable. It is better to conclude that certain questions are testable or untestable. When (H4) is employed in the question of odd versus even, then it is part of an empirically testable question. However, when employed in the question of odd versus odd only

when looking, it is not. Thus, we should think of questions as the fundamentally testable entity, hypotheses themselves are testable (or not) only derivatively.²⁵

The second reason lies in the fact that hypotheses make probabilistic predictions. They very rarely tell us what outcome *must* occur. Rather, they make a certain observation more or less *likely* to be observed. Now it can be argued that some hypotheses are capable of giving us precise probabilities, but even then, such results are not informative. For example, we know that observing 200 heads in 1,000 coin tosses is highly improbable *if* the coin is fair.²⁶ However, knowing that it is improbable to observe 200 heads in 1,000 coin tosses does not itself entail that the hypothesis that the coin is fair is false. Furthermore, there is no principled line at which an improbable observation should cause us to reject a hypothesis. This is why testing probabilistic hypotheses, which is the main activity of science, is an inherently contrastive activity. Such testing is only informative in that the evidence only points away from one hypothesis because it is pointing towards another. We can judge which hypotheses do better and which do worse only in competition.²⁷

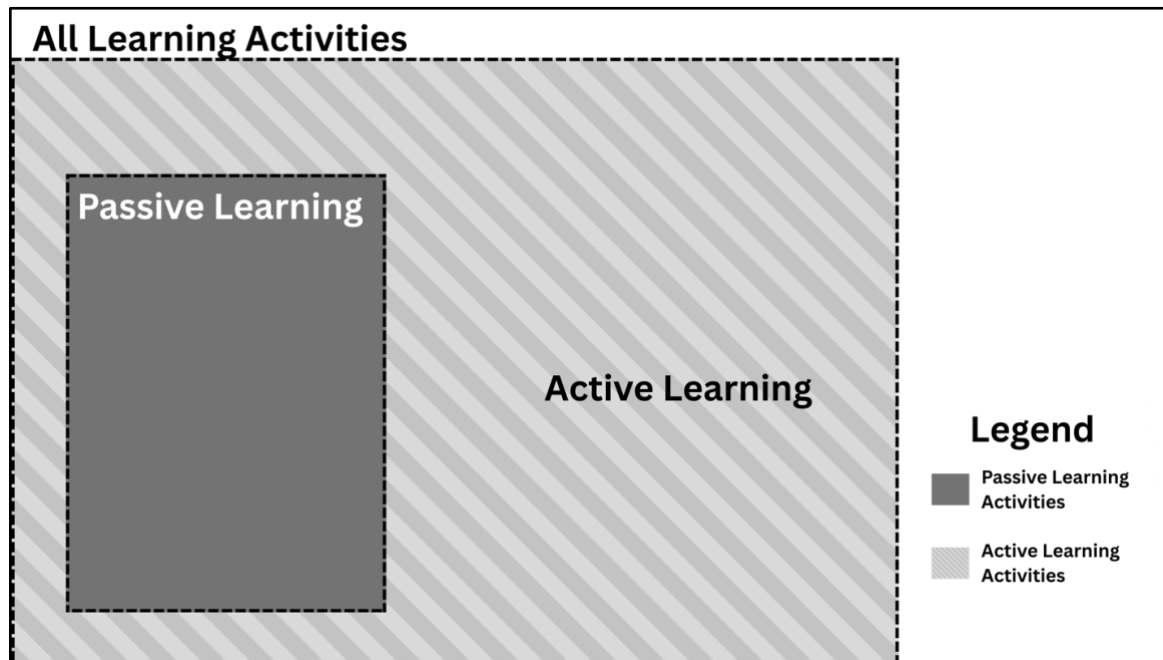
With this understanding, let us return to our negation-type definitions. Let us grant for a moment that passive learning itself is a well-defined concept. If we consider a state space that contains all possible learning activities, we can then put a box around a subset of those learning activities which are the passive ones (see Figure 2).

²⁵ Ibid, 53.

²⁶ Ibid, 58.

²⁷ Royall, (1997).

Figure 2. Active Versus Passive Learning State Space

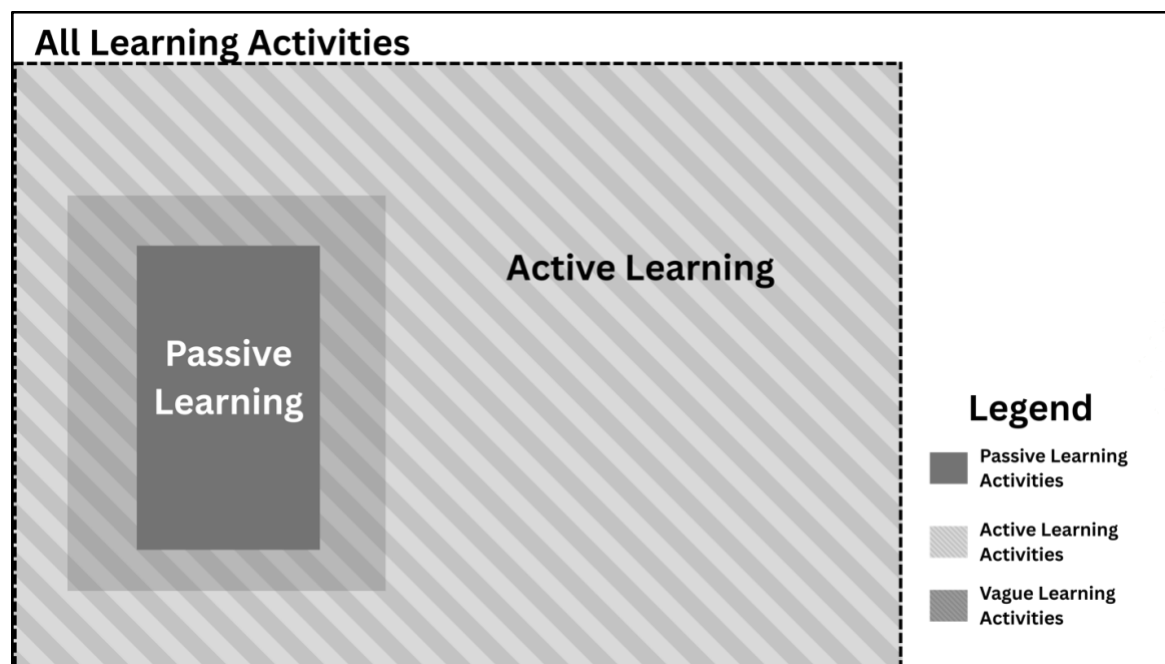


Per our negation-type definitions this would mean that all other learning activities would be active. This is problematic because it essentially creates a catch-all category. To draw conclusions about it you would need to test a string of conjunctions that includes every learning activity that is not passive. This would encompass all known and *all possible* learning activities except for those that are passive, including activities that might not have been developed yet and those that might never be developed but could still, theoretically, exist. It is important to highlight here, that this means we would need to be able to consider an unknown instance of active learning and predict what effects it would have on learning, which is impossible. This makes it, if not theoretically impossible, extremely difficult in practice to test because the category of active learning becomes an unwieldy conjunction of every learning activity that does not fall into the passive category. Thus, any hypothesis that invoked active learning would be on the hook for making predictions about all known and unknown learning activities, except for

those that fall into the passive space, to make any conclusions about the category. Drawing any logically licensed conclusions from testing this undefined concept would be sloppy science, as the homogeneity and precision necessary for a cogent induction have not been established. Furthermore, due to its catch-all nature, it is not a suitable contrastive hypothesis for passive learning. Thus, we see that negation-style definitions of active learning fail to be testable in the precise and nuanced ways we want.

The example shown above granted a major assumption that is most likely untenable. One thing a good conception should do is help identify instances of a concept. While it is not theoretically impossible for a negation definition to pick out instances of a concept that are testable, it requires a certain exhaustive relationship between the concepts. For example: dead and not dead or pregnant and not pregnant, are cases where the negation of the concept produces a suitably testable alternative. Above I argued that even if we do consider passive learning to be a well-defined concept, its negation does not create a testable alternative like not dead or not pregnant. However, the case becomes even more complicated in that passive learning itself is ill defined. It is generally agreed that straight and continuous exposition from an instructor counts as passive learning. However, there is no principled criterion for when exposition from an instructor transitions from being passive learning to part of an active learning activity. This would mean the state space would look something like the following.

Figure 3. Active Learning State Space with Vague Passive Boundary



To say that any exposition from the instructor creates a passive learning activity is surely too strong. Even staunch supporters of active learning will grant that some exposition from the instructor, such as answering questions or providing feedback, are at least permissible, if not desirable. The vagueness that plagues passive learning prevents it from creating a clear contrast class of concepts. Understanding active learning as its negation then does little to help us decide which activities are active and which are passive. Comparing the efficacy and equitability of passive versus active learning fundamentally necessitates being able to identify correctly and principledly a given learning activity as passive or active. It is difficult to observe coherently and draw conclusions about something if one is confused about the identity of that which one is supposedly observing. Negation type conceptions are very limited in their ability to identify instances of active learning in a way that is principled and not *ad hoc*. The result is a haphazard

guess as to which activities ought to count as active learning, with no clear characteristics having been articulated. This confusion is only amplified when the inadequate conception is employed in activities like experimental design, hypothesis formation, and observation. As a result, any empirical work predicated on a negation type concept is likely to be plagued by confusion and inconsistency, as we have seen with active learning.

1.IV.b Impotent Definitions

The second category of definitions, which I call *impotent*, differ from the negation types in that they do offer a positive conceptualization without any negation. Popular examples from the literature include:

5. “Instructional activities involving students in doing things and thinking about what they are doing”.²⁸
6. “The process of having students engage in some activity that forces them to reflect upon ideas and how they are using those ideas. Requiring students to regularly assess their own degree of understanding and skill at handling concepts or problems in a particular discipline. The attainment of knowledge by participating or contributing. The process of keeping students mentally, and often physically, active in their learning through activities that involve them in gathering information, thinking, and problem solving.”²⁹
7. “Active learning is an interactive and engaging process for students that may be implemented through the employment of strategies that involve metacognition,

²⁸ Bonwell & Eison (1991): 5.

²⁹ Collins & O’Brien, (2011): 5.

discussion, group work, formative assessment, practicing core competencies, live-action visuals, conceptual class design, worksheets, and/or games.”³⁰

I call such definitions impotent because despite offering positive conceptual features, they lack the power necessary to do any useful carving of the conceptual space. The positive conceptual features are so numerous and permissive that they fail to differentiate active learning from other nearby concepts. Remember the three learning activities introduced at the start of this paper (think-pair-share, reading with manipulables, building a trebuchet) and how different they were when closely considered. Yet, impotent definitions like these force us to label all three as active learning and treat them as if they compose one homogenous category, which creates a façade of unity over what is, at best, a loose collection of disjointed, heterogenous activities.

To demonstrate the over-permissiveness of this conceptualization, consider the typical lecture which is a paradigm example of passive learning. On any given day in colleges across the country, students pile into lecture halls where they sit and listen to what the educator says, occasionally writing it down. By showing up and listening to the lecturer, the student is participating in the lecture. In fact, by showing up and listening they have fulfilled their role in the lecturing activity. They have participated in exactly the manner and amount required by a lecture. If after the lecture they can parrot anything the lecturer said, then, according to the definition, they have “attain[ed] knowledge by participating”. If the goal of definitions is to help determine what ought to count as a legitimate instance of active learning and disqualify

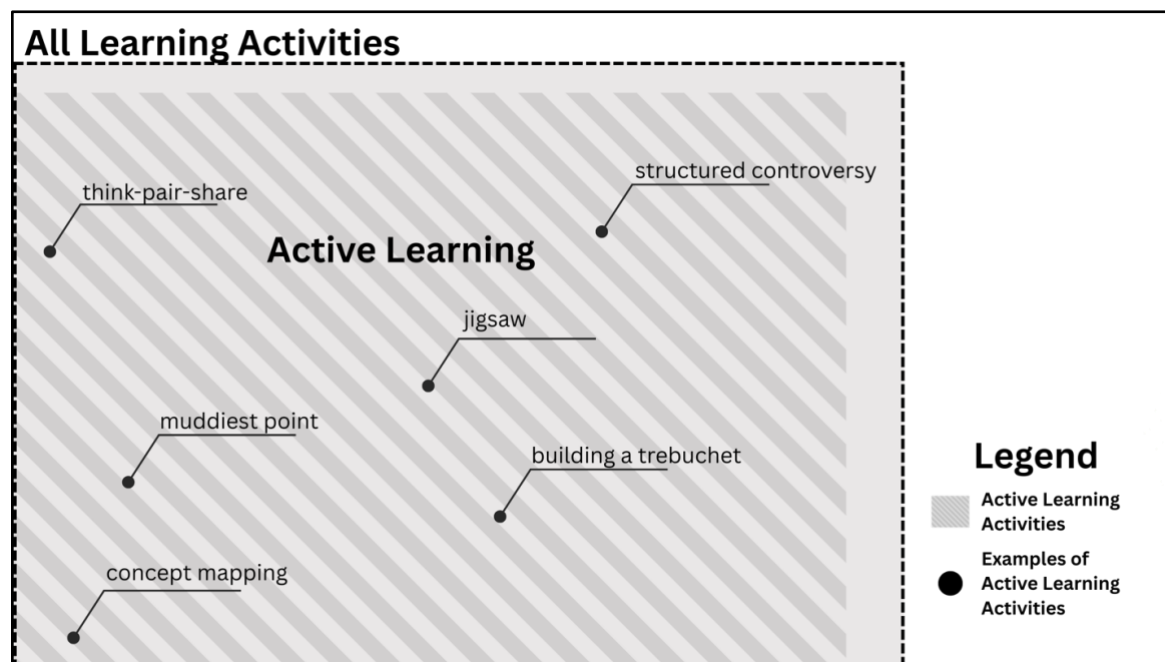
³⁰ Much of the contemporary empirical work on active learning is conducted by discipline-specific experts who specialize in teaching their discipline. STEM disciplines are particularly prolific in this sense, making it much easier to find studies, and even meta-analyses of studies, looking at questions of active learning in STEM fields. For example: Driessen et al, (2020): 6.

illegitimate instance, then impotent definitions permit an incredibly broad and diverse set of activities to count as active learning.

For example, definition (5) describes active learning as “instructional activities involving students in doing things and thinking about what they are doing.” The sole qualification on “instructional activities” in this definition is that the students are doing things. This leaves us to wonder: Does it matter what things the students are doing? Surely, whatever it is that they are doing should align with or otherwise facilitate the targeted learning outcome. Similarly, is it any thinking about what they are doing that matters, or should it be structured in some specific manner (e.g., reflective, analyzing, comparing, etc.)? Writing down what the lecturer has said verbatim requires both physical and mental activity that cannot be done without thinking. To copy down what was said the student must remember what was said, form the intention to write it down, and then proceed to physically write the words. Activities such as remembering or intending are modes of thinking. This makes rote transcription of what the lecturer is saying an active learning activity. To most, this will strike them as patently wrong—the thought process and physicality of transcription is not of the type we intuitively associate with active learning. Impotent-type definitions lend themselves to an understanding of active learning that is, by the design of the definition, incredibly broad to the point of encompassing nearly any classroom activity if some degree of reflection is involved. The lack of specificity regarding the nature and depth of the activity and reflection required means that many possible activities could be considered as instances of active learning, and yet they may not significantly differ from traditional learning practices. Put another way, short of the classic lecture paradigm, it is hard to find learning activities that would *not* count as active learning according to these definitions.

This can be visualized in Figure 4 where all the activities that fall under active learning are highlighted with paradigm examples noted.

Figure 4. Impotent Active Learning State Space



Impotent-type definitions rely strongly on examples to carry most of the explanatory burden. Relying on such a strategy is problematic. Given the many features of the example one might latch on to, the number of instances that become classified as active learning bloats. This is due to the vague and often ambiguous nature of descriptors used in Impotent-type definitions. The most flagrant problem caused by this bloating is that it impedes the production of a suitably homogenous concept. Recall at the beginning of this section the three learning activities (think-pair-share, reading with manipulables, building a trebuchet) that all qualified as active learning and how different they are along a number of dimensions. A theory or a concept is never itself tested. Rather, theories and theoretical terms are used to produce hypotheses that are then tested through experiments. The results of those experiments are then interpreted as evidence bearing

on the truth or accuracy of a theory. One common result of experiments is an inductive claim. After testing a sample from a larger group, observations made during the experiments are generalized to the broader context of the full group. Such claims are most truth tracking when there are some number of salient similarities between the instance(s) observed in the experiment and the broader group(s).

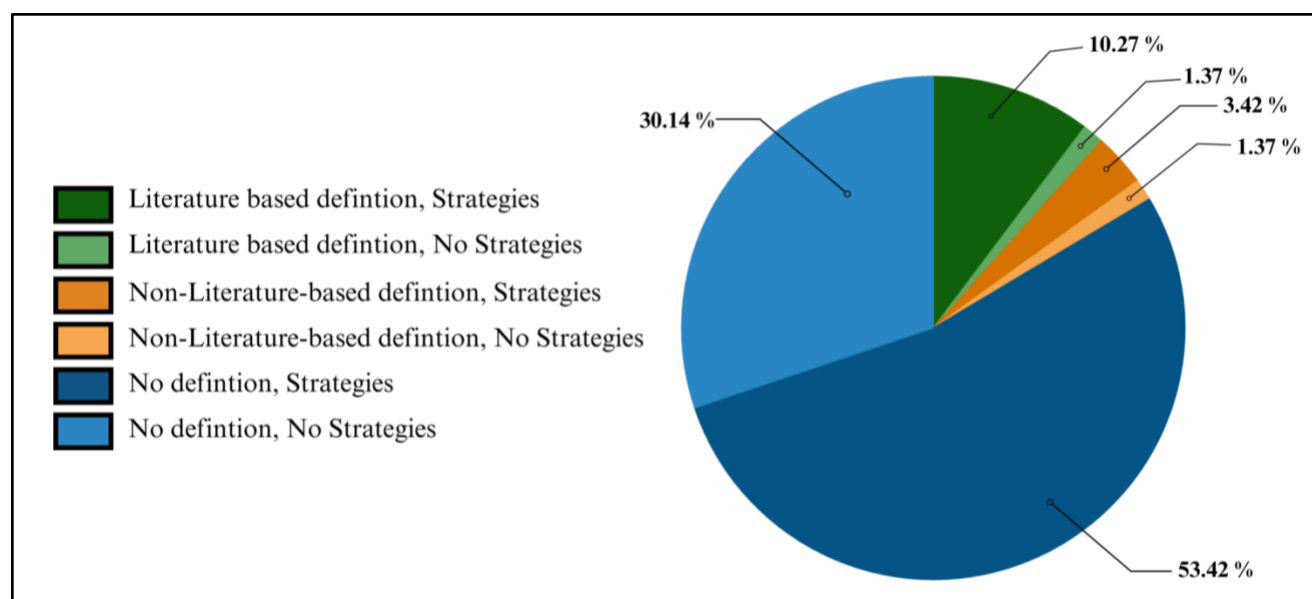
The problem with drawing sweeping generalizations about active learning on the basis of the success or failure of any particular active learning intervention is that the current conceptions of active learning cannot offer any reassurance that such generalization is warranted. Any testing done on one of the paradigm activities does not necessarily tell us anything about the group of active learning because there is no guarantee that the other instances share deep, meaningful similarities that would cause them to have a similar type or strength of effect on student learning. It is thus epistemically unwise to perform any inductive inferences or to draw any sweeping generalizations.

1.IV.c No Definition

While looking at what definitions are put forth in the literature is enlightening, so is noticing a dearth of definitions. It is telling to find a lack of definitions where they are reasonably expected to be found. One meta-analysis investigating the prevalence of active learning definitions in scholarly work published on active learning examined three top journals in Biology Education Research. Emily P. Driessen et al., “Demystifying the Meaning of Active Learning in Postsecondary Biology Education,” *CBE—Life Sciences Education* 19, no. 4 (2020): 1–9, <https://doi.org/10.1187/cbe.20-04-0068>. Any article using the term ‘active learning’ was placed into a category based on whether: (a) a definition of some sort was offered; (b) this definition was literature based; and (c) sample activities, termed strategies, were offered to gesture at some

conception of active learning. The results are best understood by looking at a pie chart showing the proportions of each category.³¹

Figure 5. Ways in which articles from LSE, the Journal of Microbiology & Biology Education, and CourseSource use the term “active learning.”



Note the two largest proportions are that of ‘no definition of any kind with example strategies’ at 53.42% and ‘no definition of any kind with no example strategies’ at 30.14%. *This means that 83.56% of peer-reviewed articles on active learning do not reference any kind of definition of it!* The scientists performing this study echo the concerns I have already outlined saying, “based on our investigation into the contemporary literature, it is apparent that people interpret active learning in a variety of different ways (i.e., interacting or engagement, not lecturing/listening, group work, scaffolding/constructivism, individual formative assessments, application/synthesis, problem solving, student centered, and evidence based)”.³² They are not alone in this analysis,

³¹ Driessen et al, 3.

³² Ibid, 5.

and such concerns have been documented by others (e.g., Prince 2004).³³ This lack of employing any definition of active learning is likely due to the fact that there is an immense proliferation of definitions and none of them seem to offer a distinct advantage over another. This only strengthens my argument that current conceptions of active learning are unsuitable for rigorous empirical inquiry and need serious attention.

In summary, current conceptions of active learning can largely be grouped into two categories—negation-style definitions and impotent-style definitions—each of which poses distinct problems for the science of teaching and learning. Negation-style definitions characterize active learning primarily by what it is not, which renders the concept too vague to generate clear, testable hypotheses. Impotent-style definitions, by contrast, offer descriptive features but fail to provide the epistemic resources necessary to license broader inductive generalizations about efficacy or equity. Many studies, moreover, bypass the issue entirely by forgoing any definition of active learning, further compounding the conceptual instability of the field. Both approaches are therefore ill-suited to support the kind of rigorous empirical inquiry expected of a scientific category. This is especially troubling because active learning is widely treated as such a category, one that is supposed to ground empirically supported claims about teaching effectiveness and educational equity. Yet the conceptual deficiencies of current definitions have gone largely unaddressed, leaving the literature marked by confusion and inconsistency. It is precisely this instability that explains why studies of active learning so often reach divergent conclusions about its efficacy and equity.

³³ This coheres with Michael Prince's metaanalysis, "Does Active Learning Work? A Review of the Research" in which he concludes it is impossible to find or provide a universally accepted definition of active learning because different researchers have interpreted the term differently. He argues that this makes global questions like "does active learning work?" unanswerable because the conclusions from various empirical sources cannot be aggregated.

1.V Where do we go from here?

One very concerning implication stemming from the unreliability of the empirical work on active learning is that active learning does not deserve its current designation as a best practice.

Remember, it is in virtue of demonstrated, repeated efficacy across broad contexts that a practice becomes a best practice. When a practice is designated as a best practice it receives a seal of approval from the community. It becomes something to promote to instructors, something we encourage them to adopt. Promoting active learning as a benevolent monolith is likely to mislead teachers about the type and strength of benefits they will see in their students' learning. Teachers are often forced to make difficult decisions deciding how to spend their finite resources (e.g., time, effort, funding). They want and need to make wise investments with their limited resources that maximize benefits for their students. Thus, if a teacher consults a list of best practices and sees the touted benefits of active learning they might think it a worthy investment of their resources.³⁴ Yet, because active learning is not a monolith, their calculations are likely to be misguided. While some activities might genuinely produce large learning gains and have positive consequences for classroom equity, this has not been demonstrated for every activity that qualifies under the current conceptualization of active learning. Without a more nuanced conversation the teacher might invest significant resources into an active learning activity that has only modest results. Furthermore, because the conception of active learning is so unclear it is hard to pinpoint exactly what is salient about certain active learning interventions—this is the problem of fidelity. The teacher must then avoid two pitfalls: investing more resources than appropriate based on the learning gains likely to be achieved and accidentally removing the salient features of the learning activity that yielded the learning gains in the first place. My

³⁴ Fosmire & Macklin, (2002).

ultimate worry is that continuing to promote active learning as a best practice could actually be a detriment to students. As it stands now, the empirical literature is not justified in giving a blanket endorsement of all the activities that currently count as active learning under the current deficient definitions. Allowing instructors to invest their resources as if every active learning activity is created equal is wasteful at best and harmful at worst. We do not want to rob students of opportunities for success because we failed to verify that the practices that are supposed to help them are reliably and intentionally doing so. This underscores the importance of continuing to address and improve the concept of active learning.

Some attempts have been made to create a better concept of active learning. One such reimagining is “deep active learning”.³⁵ Proponents of deep active learning argue that meaningful learning only takes place when the “active” part of active learning includes authentic engagement with the world writ large. Therefore, any definition of active learning must acknowledge this crucial feature. A more common attempt to remedy the situation is done by surveying stakeholders and constructing a conceptualization from their responses through aggregation and reduction.³⁶ This step, while admirable, is deeply misguided. Mere survey, without critical conceptual analysis, is vulnerable to mistaking popular opinion or agreement for truth. Even stakeholders can be mistaken about what features of their experiences are salient to a given endeavor, especially when they are novices. Furthermore, in the process of aggregating and reducing disparate responses it must also avoid creating a hybrid monster that no stakeholders are willing to endorse. Correspondingly, it must not go too far in the opposite direction by creating a conceptualization that, while endorsed by stakeholders, is so broad and

³⁵ Matsushita, (2018).

³⁶ Driessen et al., 5.

accommodating that it recreates the exact confusion it was meant to resolve. The motivation of these projects is sympathetic to my own, driven by the desire to produce a better concept of active learning. However, as of now, such endeavors have not successfully produced the more precise and homogenous concept of active learning needed to progress our theoretical or empirical understanding of active learning.

Another effort that has received significant attention is the ICAP framework. Michelene TH Chi and Ruth Wylie, “The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes,” *Educational Psychologist* 49 (2014): 219--243. This framework is directly responding to some of the concerns I have articulated about active learning. In addition to the categorizations of (p)assive and (a)ctive are (c)onstructive and (i)nteractive. Each category is associated with specific student engagement behaviors. This is advantageous in that it easily lends to the formation of empirically testable hypotheses and experiments. It also identifies a positive property, mode of student engagement, as the underlying site of salient similarities between activities. However, the ICAP framework is admittedly a behavioral metric. It relies on a yet unproven causal link between mode of student engagement and the underlying knowledge acquisition processes. Furthermore, it is admittedly heavily reliant on a computational understanding of the mind. While it might represent an improvement to current conceptions of active learning, it needs time for additional studies to substantiate that claim. It can also be argued that ICAP does not offer a different conception of active learning, rather it merely resituates the existing concept based on its connections to other additional concepts. It might, therefore, represent an alternative to active learning, but it does not offer a better conception of it.

One way to improve the concept of active learning might be to treat it the same way we have treated the concept of love, by making distinctions within the concept based on the different

elements of it. Species of love are differentiated, for example, based on the object being loved (e.g., self, others, inanimate objects). They are also differentiated based on the relationship between the lover and the beloved (e.g., parent-child, friendship, sexual). Making these distinctions within the larger category produces more usable sub-concepts that can yield more homogenous groupings. This ultimately makes them more meaningful and licenses generalized inferences from a sample population to the group, as well as from the group to novel instances. Another way to improve the concept would be to replace the notion of active learning as a binary property with a conception that understands active learning as a graded concept. This would relieve pressure to identify an activity simply as passive or active and allow for a more nuanced ranking of activities along a passive to active spectrum.

A final, potentially promising avenue can be found in the work being done on embodied cognition. Embodied cognition and education are natural allies whose connections are just beginning to be explored. Embodied research in the field of teaching and learning can serve as source and inspiration for possible, principled distinctions within the concept of active learning. One such example has already been mentioned earlier in this paper, reading with manipulables. The effect of manipulables on cognition, especially learning, is often a source of support for arguments in favor of embodied theories of cognition. Identifying learning activities that make use of manipulables gives a precise, clear criterion for identifying which activities belong in the subcategory and which do not. It creates a homogenous category based on a potentially salient property. It thus avoids two of the pitfalls of the overarching concept—heterogeneity and lack of explanatory power. Obviously, this is a fledgling example that would need significantly more development, but it is enough to serve as a proof of concept.

It would be disingenuous to pretend to have the answers for fixing active learning. I have argued that the root of active learning's troubles comes from the inadequacy of the concept itself, or at least our current conceptions of it. Developing a better conception of active learning is thus a necessary first step in addressing the inconsistencies found in the current empirical literature as well as those problems caused by a reliance on those findings. I have also suggested some ways we might approach the creation of a better concept of active learning. However, this only serves a small, tentative step forward. It is my hope that this paper will draw the attention of scholars from philosophy, the learning sciences, and SoTL alike and spark more purposeful and sustained efforts to produce a precise and nuanced concept of active learning that is capable of, among other things, guiding better scientific inquiry.

Paper Two: *Embodied Bootstrapping: An Embodied Account of Concept Learning*

Abstract for Paper Two

This paper argues that hard-route conceptual bootstrapping is significantly strengthened when it is formulated within an embodied view of mind, rather than a narrowly computational one. Empiricists appeal to bootstrapping to explain how learners acquire genuinely new conceptual primitives and, thereby, expand their expressive power beyond the mere recombination of concepts they already possess. However, traditional accounts of bootstrapping rely too heavily on computational assumptions and, as a result, fail to capture important features of how radical conceptual change actually proceeds. Drawing on recent work in embodied cognition, I develop an alternative account of bootstrapping that better explains how learners come to acquire new conceptual primitives through structured engagement with their bodies and environments. An embodied framework, I argue, both broadens the range of cases that bootstrapping can illuminate and provides a richer account of the mechanisms by which new concepts are formed and stabilized. It also offers a principled response to the Deviant Interpretation Challenge, a prominent objection to bootstrapping theories. The paper thus develops an account of embodied bootstrapping, illustrates its operation through a concrete example, and argues that embodiment is not merely an auxiliary feature of learning but a theoretically significant component of bootstrapping itself. Understanding bootstrapping in this way has important implications for how we explain both the acquisition and the stability of genuinely new conceptual primitives.

2.1 Introduction

Children live in a world that often outruns their understanding. Objects disappear when out of sight until they acquire object permanence. Liquid in a tall, skinny glass is more than the same liquid in a short, wide glass until they acquire conservation of matter. Yet, in a remarkably short time, children come to possess all types of concepts which organize this initially confusing field of experiences: concepts of object, number, causation, agency, biological kind, and fairness, among many others. The speed, scope, and depth of this achievement make concept acquisition a remarkable phenomenon in need of explanation.

Empiricists maintain that learners can expand their expressive power through the acquisition of new conceptual primitives. This claim is philosophically significant because it implies a form of conceptual change that cannot be reduced to mere recombination of concepts already possessed. If learners can acquire genuinely new primitives, then they can come to entertain thoughts that were not previously available to them. The central challenge for the empiricist, then, is to explain how such radical conceptual change is possible. One influential answer appeals to hard-route conceptual bootstrapping. Traditional accounts of bootstrapping, however, rely too heavily on computational notions of the mind.³⁷ I argue that they would do better to draw on recent work in embodied cognition in order to explain how bootstrapping works.

The central claim of this paper is that the theory of bootstrapping is significantly strengthened when it is formulated within an embodied view of mind rather than within a narrowly

³⁷ Jacob Beck, “Can Bootstrapping Explain Concept Learning?,” *Cognition* 158 (January 2017): 110–21; Susan Carey, “Concept Innateness, Concept Continuity, and Bootstrapping,” *The Behavioral and Brain Sciences* 34, no. 3 (2011): 152–62; Susan Carey, “On Learning New Primitives in the Language of Thought: Reply to Rey,” *Mind & Language* 29, no. 2 (2014): 133–66; Susan Carey, “Précis of *The Origin of Concepts*,” *Behavioral and Brain Sciences* 34, no. 3 (2011): 113–24.

computational framework. This allows it to capture a broader range of cases and provides richer resources for explaining how bootstrapping proceeds. Additionally, I argue that altering the theory of bootstrapping to adopt an embodied notion of the mind provides resources for offering a principled reply to the Deviant Interpretation Challenge, a prominent objection to theories of bootstrapping. The goal of this paper is to develop an account of how embodied bootstrapping works and provide an example of it. My conclusion is that embodiment is not merely an auxiliary consideration but is a theoretically significant component of bootstrapping, with implications for how we understand the acquisition and stability of new conceptual primitives.

2.II Bootstrapping

In this section, I clarify the idea of bootstrapping so that the concerns discussed later in the paper can be brought into sharper focus. An instructive example comes from the Indigo Bunting, a migratory bird that travels north to breed and raise its young. Steven Emlen discovered that Buntings are not born knowing where north is; rather, they acquire this concept through a neural mechanism that tracks the rotation of the night sky.³⁸ He demonstrated this by showing that when young Buntings placed in a planetarium were exposed to an artificial night sky that rotated around an arbitrarily chosen point, instead of the true night sky (which rotates around the North Star), they identified the arbitrarily chosen point as “north.” In essence, their domain specific learning mechanism, which takes perceptions of the night sky and yields a concept of north, was tricked. This is possible because the constraint implicitly coded in the learning mechanism was

³⁸ For a full description see: Stephen Emlen, "Migratory orientation in the Indigo Bunting, *Passerina Cyanea*: part I: Evidence for use of celestial cues," *The Auk* 84, no. 3 (1967): 309-342; Stephen Emlen, "Migratory orientation in the Indigo Bunting, *Passerina Cyanea*. Part II: Mechanism of celestial orientation." *The Auk* 84, no. 4 (1967): 463-489.

center-of-rotation. This rule is used on the Bunting's perceptions of the night sky indiscriminately. If the Bunting's perceptions of the night sky are skewed due to some manipulation (e.g., the planetarium experiments), then the conceptual content that fills the conceptual role prepared for "north" will be skewed correspondingly.

There are two crucial points to appreciate in this example. First, this is an example of bootstrapping because it demonstrates how the Indigo Bunting could acquire the new primitive "north" without it being composed of or defined using concepts the Bunting already possessed. All that is required are representations of the night sky and a dedicated learning mechanism implicitly coded with the computational constraint of center-of-rotation. This simple process produces the new primitive concept of north without any use of other concepts or definitions. Second, this is an example of what is known as the "easy route" of bootstrapping because the mental architecture, in this case the domain specific learning mechanism, was already present. One final thing to note about these types of domain specific learning mechanisms is that they often are time sensitive and atrophy very quickly. In the case of the Indigo Bunting, if the nestlings are deprived of ever seeing the night sky while they are young, no amount of exposure to the night sky later in life will ever allow them to learn the concept of north, and they will perish.

The Indigo Bunting case usefully illustrates the mechanics of bootstrapping, but it also represents a comparatively easy case.³⁹ Its relative simplicity lies in the fact that the bird's cognitive architecture already makes room for the concept of north, making it the so called "easy route" of

³⁹ Instances of the "easy-route" of bootstrapping can be found in domain-specific learning mechanisms, which are often seen in non-human animals; Carey, "On Learning New Primitives in the Language of Thought: Reply to Rey," 140.

bootstrapping. The deeper challenge concerns cases in which no such antecedent cognitive structure is available, making it the “hard-route” of bootstrapping.

A useful way to understand the difference between easy and hard-route bootstrapping is through the analogy of a dinner-party seating chart. Imagine that you have already set a table for six guests. Each place setting occupies a determinate position relative to the table and to the other settings—one seat is at the head of the table, two are next to each other, some are across from one another, and so on. As the host, you assign each guest to a seat on the basis of these relations. The guest of honor goes at the head of the table, partners are seated beside one another, and guests who should be kept apart are placed farther away. Because the seating structure is already in place, your task is simply to determine which guest best fits which seat.

Now imagine that, just before dinner begins, you unexpectedly need to accommodate an infant in a highchair. This case is different. There is no place already prepared for the infant, and it is not yet clear how the infant should fit into the existing arrangement. You must determine not only where the infant should go, but also how to modify the seating structure itself to make room for this new kind of guest. This makes deciding where to put the young guest much more difficult than it was to make your original seating chart.

In the analogy, the original place settings represent preexisting conceptual roles, while the guests represent contents that can fill those roles. In easy route bootstrapping, the relevant role is already present in the learner’s cognitive architecture. The problem is simply to determine what content belongs in that role. The infant case represents hard-route bootstrapping. Here there is no antecedently available role waiting to be filled. Instead, the learner must create a new conceptual role *and only then* determine what content belongs there. That is what makes the hard-route of

bootstrapping hard. It requires not just filling an available place but altering the structure itself to accommodate a genuinely new one.

2.III Hard-Route Bootstrapping

With this distinction in view, we can now turn to the hard-route of bootstrapping. These are the cases of greatest philosophical interest, since it is here that the deepest explanatory challenges arise. Hard-route bootstrapping is best thought of as proceeding in three stages, each of which is reviewed in a subsequent section.⁴⁰

2.III.a Step One: Acquisition of a Placeholder Symbol Set (PSS)

One of the first things we are taught as toddlers is how to count. Counting is among the earliest and most important structured practices children learn, in part because it gives them a way to keep track of the quantities that organize everyday life. Psychologist Susan Carey, a notable proponent of hard-route bootstrapping, treats this ordinary developmental achievement as psychologically and philosophically significant. She argues that children acquire the concept of an integer when they come to understand how counting implements the successor function, such that each next numeral denotes a cardinality exactly one greater than the last.

If you have ever taught a child to count, you know that the first thing they learn to do is memorize and parrot the count list (1...2...3...4...5...etc.,) In this case, language (the count list) serves as a set of symbols. However, these words (symbols) do not hold precise content for children yet. You ask them to bring you two cookies, and they bring you four instead; or you ask

⁴⁰ Jacob Beck, “Can Bootstrapping Explain Concept Learning?,” *Cognition* 158 (January 2017): 110–21.

for one cookie, and they bring you the whole package. At this stage, the placeholder symbols are empty to the learner. They can repeat the words ‘one’ and ‘two’ and so on in the correct order, but they do not yet know what these words, these symbols, refer to in the world. This demonstrates one way a set of empty symbols can be acquired through methods like mimicry and memorization. The child knows the symbols, but to them, the symbols are empty in a very meaningful sense.

In a more technical sense, the first step of hard-route bootstrapping is to acquire a placeholder symbol set (PSS). A PSS is simply a set of symbols that have been stripped of any conceptual content. Symbols are thought of as part of symbol-referent relationship. When you have the symbol, you are pointed towards its corresponding referent in the world. The symbol ‘Polaris’ points you toward the North star. Placeholder symbols lack a referent. They do not point towards something in the world, they just “hold space” so to speak. It is important to note that the process of bootstrapping crucially depends on the placeholder symbols being a set. By this I mean that there is more than one symbol being interpreted and the symbols bear specific, distinct relationships to one another. It is only a set of symbols that can produce the preserved relationships which are coopted as constraints for interpreting the new primitive in the second step.

2.III.b Step Two: Computational Constraints Partially Interpret the PSS

To understand the role constraints play in the bootstrapping story, Jacob Beck supplies a nice analogy by comparing it to a logical system.⁴¹

⁴¹ Jacob Beck, “Can Bootstrapping,” 116.

Just as any logic needs not only axioms but also rules of inference, any computational theory of mind that posits explicit representations also needs procedures that govern how those representations can be manipulated. Much of the power of a logical or mental system derives from its inference rules or procedures. But inference rules and procedures are also limited in a way that axioms and explicit representations are not. They cannot feature as premises or conclusions in reasoning, and so cannot be directly manipulated by other inference rules or procedures. They thus exist only as constraints on computations.

What makes bootstrapping a viable theory, rather than a piece of theoretical handwaving, is the role played by constraints. Constraints render the process one of partial interpretation by providing a clear route through which the symbols in the PSS can come to acquire new content. For that reason, the plausibility of bootstrapping depends on there being empirically credible ways of constraining, and as a result, partially interpreting placeholder symbols. It is worth noting here, that the only constraints that are considered are computational in nature. This is a point I will return to later.

To be clear, the interpretation of the PSS is only partial at this stage. The relevant constraints do not uniquely determine content. An analogy helps clarify this. Imagine Prince Charming trying to identify Cinderella after she has fled at midnight, leaving behind only her glass slipper. The slipper functions as a constraint on the range of possible candidates. Suppose Cinderella wears a size five slipper, and there are roughly 3,822,561,000 female persons in the world. If only 1.8% wear size five shoes, the slipper reduces the candidate pool to 68,806,098. That is a dramatic narrowing, but it still falls far short of uniquely identifying Cinderella. The same is true of the constraints generated by the preexisting relations among symbols in a PSS. Those relations do not fully determine the content of placeholder symbols, but they do substantially narrow the range of concepts that could appropriately occupy those positions.

This is precisely the role linguistic structure plays in Carey's account of integer acquisition. In her framework, language provides the relevant computational constraints. It does so because language consists of rule-governed operations over a finite set of symbols. Carey argues that children's initial interpretation of the count list depends on semantic resources already present in language, including the singular-plural distinction and the meanings of quantificational expressions such as "some" and "a." An important consequence is that children learning languages that encode numerical differences in other forms from English may generate different hypotheses about the partial meanings of number words. Carey's work thus offers a concrete example of how computational constraints can limit, organize, and partially interpret a PSS. Once those constraints have reduced the range of plausible candidate contents for the placeholders, the learner moves on to Step Three.

2.III.c Step Three: Full Interpretation of Placeholder Symbol Set via Modeling

In step three, the child uses what they already know about small exact quantities, the child recognizes a pattern: "two" is one more than "one," "three" is one more than "two," and "four" is one more than "three." By analogy, the child maps this pattern onto the count list and, through induction, concludes that each numeral names a quantity exactly one greater than its predecessor. This is how the count list comes to embody the successor function.⁴²

This final process is referred to as "modeling" and encompasses a number of non-demonstrative processes such as, "analogical mapping, thought experimentation, limiting case analysis, induction, and abduction."⁴³ It is in this step that the "correct" interpretation of the PSS

⁴² For more details see: Carey, "Bootstrapping & the Origin of Concepts," 63.

⁴³ Jacob Beck, "Can Bootstrapping," 112.

emerges from the narrowed field of candidate concepts. Interestingly, it is theorized that learners who are more able to employ these modeling strategies, particularly analogizing, acquire new concepts easier than learners who struggle to create such structures for themselves. Researchers have shown that the modeling process can be effectively facilitated by helping learners analogize from better known domains to the target domain. One particularly effective pedagogical activity for learning the concept density was having learners compare boxes with different numbers of dots in them. Learners were successful in understanding that a denser object is akin to a box that has more dots in it.⁴⁴ Beck suggest the efficacy of this type of intervention is due to the fact that analogizing helps learners construct new conceptual roles for their PSS, "...thereby endowing them with further content."⁴⁵

With step three, the hard-route bootstrapping process is complete. The learner has now come to possess a genuinely new concept; in Carey's case, this is the concept of integers. What is significant is that this concept is primitive in the relevant sense: it is not constructed through the mere recombination of antecedently available concepts, as "*unkind*" is constructed from "*un*" and "*kind*."

This is the essence of the bootstrapping account. Its empirical success depends fundamentally on the presence of real-world examples of bootstrapping. Empiricists bear the burden of showing that it actually occurs in our process of concept learning. Current examples of hard-route bootstrapping are in short supply. Carey's groundbreaking work only covers three examples, and

⁴⁴Carol Smith et al., "Using Conceptual Models to Facilitate Conceptual Change: The Case of Weight-Density Differentiation," *Cognition and Instruction* 9, no. 3 (1992): 221–8; Carol Smith and Chris Unger, "What's in Dots-Per-Box? Conceptual Bootstrapping With Stripped-Down Visual Analogs," *Journal of the Learning Sciences* 6, no. 2 (1997): 143–81.

⁴⁵ Jacob Beck, "Can Bootstrapping," 120.

only the acquisition of the concept of integers is worked out in detail. I argue that a theory of bootstrapping will be much more successful in supplying detailed examples of hard-route bootstrapping if the theory is altered to include a new factor: embodied constraints. The next sections are dedicated to elucidating how embodied constraints might play a role in the bootstrapping process and explaining why their addition deepens the explanatory resources of theory and helps address one of the most prominent objections to the theory.

2.IV Embodied Cognition and Embodied Bootstrapping

Until this point, the constraints in the bootstrapping process have been conceptualized as solely computational in nature. This is a mistake and unduly narrows the theory's explanatory resources and its empirical consequence class. The addition of embodied constraints to the bootstrapping process is inspired by insights from the embodied cognition literature.⁴⁶ This section explores how we can make sense of embodied constraints and sets the stage for explaining why broadening the theory of bootstrapping to include them strengthens the theory.

One area of embodied cognition worth focusing on for the sake of developing the notion of embodied constraints is that of conceptual metaphors. In general, metaphors enable us to understand one thing in terms of another. Consider the traditional metaphor, "the eyes are the window to the soul." Although eyes are not literally windows, our understanding of windows helps illuminate what this phrase is meant to convey. A window is an opening through which we

⁴⁶ Embodied cognition is not a single unified theory, so different versions assign different roles to embodiment and will therefore interact with the bootstrapping process in different ways. For discussion of the embodied research programme see: Larry Shapiro, "The embodied cognition research programme." *Philosophy compass* 2, no. 2 (2007): 338-346. Further elaborated in: Lawrence Shapiro, *Embodied cognition*, Routledge, 2019.

can see which would otherwise be hidden. Applying the properties of windows to someone's eyes helps us see how the metaphor presents the eyes as giving access to a person's inner life.

Metaphors work by exploiting a source-and-target structure. The source is used to interpret the target. In the case above, "windows" serves as the source domain and "eyes" as the target domain. Like traditional metaphors, conceptual metaphors also employ this same source-and-target structure. The way that researchers George Lakoff and Rafael Núñez describe conceptual metaphors is very similar to the way bootstrapping is described. They describe them as, "...grounded, inference-preserving cross-domain mapping[s]... that allow us to use the inferential structure of one domain (say, geometry) to reason about another (say, mathematics)." ⁴⁷ Similarly, the punchline of the bootstrapping process is that a placeholder symbol set creates a structure, based on the relationships between the symbols, that can then be used to partially interpret the empty symbols of the set. Conceptual metaphors create structures from relationships already familiar to our bodies and their movements. These structures are then exploited to help understand other domains.

For example, orientational metaphors, a subset of conceptual metaphors, reveal this process in a distinctly embodied form. Orientational metaphors are so pervasive in everyday language that we easily overlook their body centric nature. Lakoff and Johnson write: ⁴⁸

We call them "orientational" metaphors because most of them have to do with spatial orientation: UP-DOWN, FRONT-BACK, IN-OUT, ON-OFF, DEEP-SHALLOW, CENTRAL-PERIPHERAL. These spatial orientations arise from the facts that we have bodies of the sort we have and that they function as they do in our physical environment.

⁴⁷ George Lakoff and Rafael E. Núñez, *Where Mathematics Comes from: How the Embodied Mind Brings Mathematics into Being*, 1st ed (Basic Books, 2000), 6.

⁴⁸ George Lakoff and Mark Johnson, *Metaphors We Live By*, (University of Chicago Press, 1980), 14.

What is distinctive about these metaphors is that the source domain is not an external object like a window but instead our own bodily experience. Lakoff and Núñez argue metaphors with body-based source domains yield particularly strong connections and inferences because those connections and inferences are inherently meaningful. Some experiences need not be consciously interpreted because the meaning or inference structure is already apparent due to the nature of the learner's body and environment. This inherent or automatic meaning-making makes conceptual metaphors particularly valuable because they supply source domains that need not be explained in terms of something else.

For instance, orientational metaphors are instructive because particular orientations (e.g., up-down) are inherently meaningful to us due to the kinds of bodies we have and the ways we move through the world. In general, "up" is associated with positive states. This association is not arbitrary.⁴⁹ Rather, it is grounded in recurring features of our physical experience. When we are healthy, we are "up and about," able to stand, move, and act; when we are ill, we are often forced to lie down; when someone loses consciousness, they collapse; and in death, the body is no longer able to stand or move at all. In this way, health and vitality come to be associated with "up," while illness, unconsciousness, and death are associated with "down." This bodily pattern helps explain why a phrase like "things are looking up" conveys improvement or optimism. Because "up" is tied to beneficial bodily states, it acquires a positive valence that can then structure our understanding of other, more abstract domains. This is important for the story of

⁴⁹ Ibid.

embodied bootstrapping because it begins to show, in a concrete way, how constraints might be imposed from the experience of our bodies and their movement.⁵⁰

2.V An Embodied Bootstrapping Process

With an understanding of how conceptual metaphors aid in meaning-making, we can now describe an embodied bootstrapping process.

2.V.a Embodied Step One: Acquire a Placeholder Symbol Set (PSS)

The first step is to acquire a placeholder symbol set. On the embodied version of the view, this step broadens what can count as a PSS. The learner need not acquire a placeholder structure through language alone. A PSS might also be formed through stable, repeatable patterns in bodily orientation and movement. Recurrent experiences such as reaching, grasping, containing, balancing, moving toward, or moving away from things can generate structured relations that are available prior to full conceptual interpretation. In such cases, the learner acquires not merely an isolated symbol, but an organized pattern that can later be taken up in the bootstrapping process and used to constrain the interpretation of new content.

2.V. b Embodied Step Two: Embodied Constraints Partially Interpret the PSS

In the embodied version of bootstrapping, the second step is altered. Instead of using solely computational constraints to partially interpret the PSS, embodied constraints are used instead of, or in addition to, the computational constraints. Because the embodied constraints also serve to

⁵⁰ Linking metaphors provide an avenue for more complexity and greater abstraction and while body-based conceptual metaphors often require no explanation, linking metaphors are often explicitly taught such as the density and dots-per-box example.

dictate how the symbols in the PSS relate to each other, they also can be used to winnow down the candidate concepts appropriate for the PSS. To see how this might work in practice we will use this conceptual metaphor: Arithmetic is Collecting Objects.

Imagine a learner, a young child, who enjoys exploring her environment—her crib, her highchair, her toys (along with other things she shouldn't have). As she interacts with her environment, she will experience many things, such as receiving objects and having them taken away. We often hand children things like crackers, sippy cups, and toys. We also take things away from them, like objects that aren't actually toys, or their sippy cup when they need to be fastened in their car seat. The repetition of these actions, of acquiring and losing objects, becomes the source domain for Arithmetic is Collecting Objects. This is why body-based metaphors are described as highly intuitive and requiring little instruction. All that was needed to acquire the source domain of this grounded metaphor was to act like a normal child in a rich environment. Early Childhood Mathematics Education (ECME) research strongly supports the claim that children come into preschool and kindergarten already having developed an “everyday mathematics.”⁵¹ In a meta-article focused on identifying ECME best practices, the authors write: ⁵²

Everyday mathematics is an essential and even inevitable feature of the child's cognitive development, and like other aspects of the child's cognition... develops in the ordinary environment, usually without direct instruction. Indeed, everyday mathematics is so fundamental and pervasive a feature of the child's cognition that it is hard to see how children could function without it.

⁵¹ Arthur Baroody et al., “The Development of Young Children’s Early Number and Operation Sense and Its Implications for Early Childhood Education,” in *Handbook of Research on the Education of Young Children* (Routledge, 2014); Douglas Clements and Julie Sarama, “Effects of a Preschool Mathematics Curriculum: Summative Research on the Building Blocks Project,” pt. 1136-1163, *Journal for Research in Mathematics Education* 38, no. 2 (n.d.); Herbert P. Ginsburg et al., “Mathematical Thinking and Learning,” in *Blackwell Handbook of Early Childhood Development*, ed. Kathleen McCartney and Deborah Phillips (Blackwell Publishing Ltd, 2006).

⁵² Ginsburg, et. al., “Mathematical Thinking,” 3.

This research nicely dovetails with Lakoff and Núñez's work on grounded metaphors because it also recognizes that systematic inferences can occur from commonplace childhood activities. In fact, grounded metaphors are well suited to explain why, before any sort of instruction, children enter school having already developed an everyday mathematics.

Having identified the meaningful, body-based source domain of Arithmetic is Collecting Objects, we can begin to see how it partially constrains the PSS. The inferences the learner draws from collecting objects constrains the relationships of the symbols in the PSS. For example, because Arithmetic is Collecting Objects's source domain is objects (e.g., crackers, blocks, sippy cups) we see the physical qualities of the objects transferred to the target domain of mathematics. This explains why, even as adults, we think about and talk of numbers as if they are physical. Numbers are aligned with objects in the source domain. The physical properties of objects are part of the preserved inference structure used to constrain the PSS. For example, objects can be given and taken away; therefore, numbers can be given and taken away. Objects are inert, stable, and constant. Numbers are inert, stable, and constant. Like blocks, numbers don't smush together to form a new thing when collected; they are discrete. You can have more or less of them, and they can be bigger and smaller than one another. All these constraints on numbers originate from the body-based source domain and greatly affect how we conceptualize the way numbers relate to each other and how they can be manipulated.

Notice, however, that numbers need not function according to the same rules as physical objects, even though we often treat them as though they do. As abstract concepts, they do not have any physicality at all! Numbers cannot actually be collected or gathered because they are not tangible things. For the same reason, they cannot be given or taken away. They are not bigger and smaller in the same way different piles of blocks might be. Furthermore, they cannot be discrete or

constant for the same reasons blocks are discrete and constant. All of these constraints are the residue of using Arithmetic is Collecting Objects to partially interpret the PSS. Once the modeling phase provides further interpretation of these constraints, the learner will find that the fully interpreted PSS aligns with the conceptual primitive we term “addition.”

2.V.c Embodied Step Three: Full Interpretation of Placeholder Symbol Set via Modeling

The final step of the process is the modeling phase, where a number of methods such as analogical mapping, trial and error, induction, and abduction are used to fully interpret the PSS. Usually, the modeling phase for the primitive addition begins in earnest in the kindergarten classroom. ECME literature describes the general purpose of mathematics instruction as facilitating a transition from everyday mathematics into a rigorous scientific mathematics.⁵³ This is an excellent description of what happens during the modeling phase of embodied bootstrapping. Prior to formal instruction, our young learner has already developed a robust partial interpretation of the PSS for the concept of addition. A teacher must now guide the learner as they finish the interpretation process by creating learning activities with conditions that will help the learner acquire the “correct” interpretation of addition. In this case, our ECME teacher needs to facilitate proper analogical mapping between objects and numbers and between collecting and addition. When this is completed, the learner will have successfully bootstrapped the new primitive concept of addition.

Because teachers often help learners through this final stage, examples of effective instruction can illuminate how modeling in the embodied bootstrapping account works in practice. Earlier, I

⁵³ Lev Vygotsky, *Thought and Language*, ed. Alex Kozulin (MIT Press, 1986), 178.

noted that Carey had considerable success teaching density by first helping learners grasp the relevant structure of a concept in a more familiar source domain, such as the number of dots in a box. This was effective because it allowed learners to use the inferential structure of that source domain to reason about the less familiar target domain of density in physics. The following learning activity illustrates what step three might look like in the embodied case.

Below are the instructions for a common ECME learning activity focused on addition.

Step 1: Present the activity board with an empty green box at the top, an empty yellow box in the middle, and an empty blue box at the bottom. The boxes are aligned with the numerals 1, 2, and 3, and the learner is told to “Add It!”

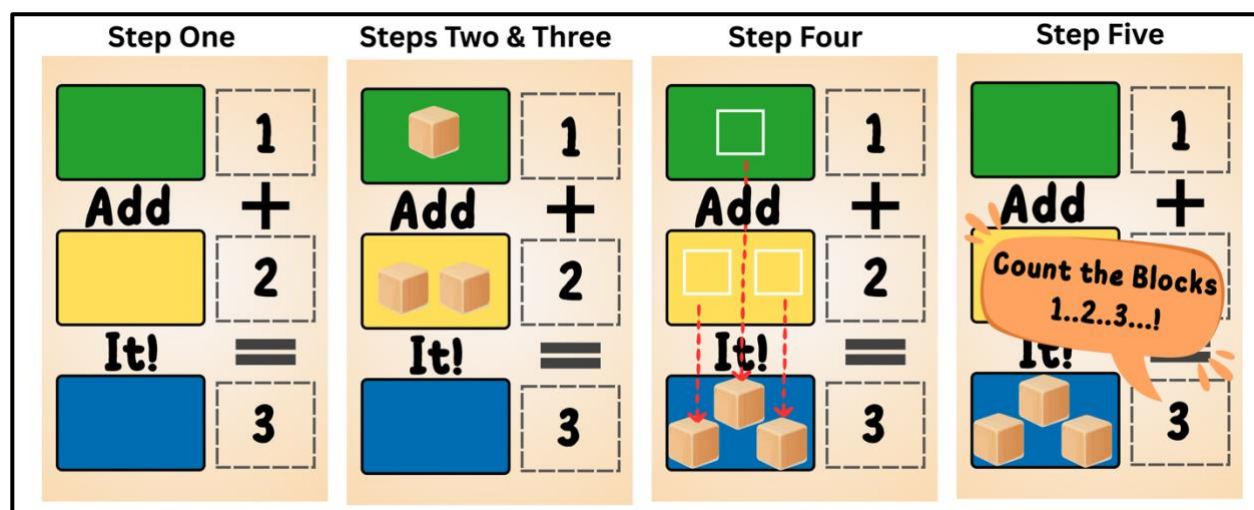
Step 2: Place one block in the green box.

Step 3: Place two blocks in the yellow box, while leaving the blue box empty. This sets up the two quantities that will be combined.

Step 4: Move the blocks from the green and yellow boxes into the blue box. The single block from the green box and the two blocks from the yellow box are brought together, so that the blue box now contains all three blocks.

Step 5: Ask the learner to count the blocks in the blue box: “1, 2, 3.” This helps the learner connect the combined set of blocks with the result of the addition problem.

Figure 6. “Add It!” An Early Childhood Mathematics Learning Activity.



The most important thing to notice here is how carefully every part of this learning activity aligned. This is an example of analogical mapping from a body-based domain to a mathematical domain. The blocks, which are from the grounded domain, are aligned with numbers in both verbal and written form. This alignment is literal. The blocks are placed in squares directly across from the corresponding numeral. This creates a connection between the physical block and numbers. Once this has occurred for both the green and yellow boxes, the learner then pushes or pulls all of the blocks into a pile in the blue box at the bottom of the card. This physical act of collecting all of the objects is then connected to the counting and tallying all of the objects to get a bigger number. This serves as a basic example showing how real teachers approach analogically mapping Arithmetic is Collecting Objects, piece by piece, into the domain of mathematics to help the learner build an understanding of addition.

2.VI Strengthening the Theory of Bootstrapping

If this example of embodied bootstrapping is viable, then the payoff is not limited to the case of arithmetic. Rather, it suggests that the empiricist has access to a much larger body of bootstrapping cases if they adopt an embodied bootstrapping theory rather than a purely computational account. The embodied view preserves the computational examples, since those remain genuine cases of constraint of the PSS, but it also captures further cases in which bodily experience helps structure and partially interpret a PSS. Just as embodied theories of cognition need not deny the existence of computational cognition, so can embodied bootstrapping accept the existence of embodied constraints without denying the existence of computational cognition. The work on conceptual metaphor is especially important here because it identifies a wide range of body-based source domains that can guide understanding in new target domains. Once those cases are admitted, bootstrapping begins to look less like a rare and fragile phenomenon and more like a pervasive feature of human learning. This strengthens the theory by making it more empirically plausible. It also broadens what the theory predicts, resulting in a more developed empirical consequence class, which is theoretically beneficial. And although conceptual metaphor provides the clearest illustration for present purposes, it is unlikely to be the only relevant insight from embodied cognition. Other lines of work, such as research on simulation, also appear promising and may reveal further ways in which embodiment can serve to constrain a PSS.

2.VII The Deviant Interpretation Challenge

One natural worry with the story of bootstrapping is that while the constraints placed on a PSS do genuinely narrow the number of candidate conceptions that can fill the symbols, they do not

explain how learners come to endorse the singular “correct” interpretation of the new primitive. Georges Rey challenges Carey with a version of this critique referred to as the Deviant Interpretation Challenge.⁵⁴ The Deviant Interpretation challenge can be thought of as concerned with the riddle of induction or underdetermination problems. The riddle of induction version of the challenge hinges on the fact that the learner is attempting to determine which predicates of the concept are lawlike and which are accidental. They need to do this in order to isolate the correct concept. However, as the riddle of induction emphasizes, attempting such a determination is impossible with limited observations. This leaves the learner stuck, with no principled reason for choosing the correct generalization over any of the other generalizations that also fit the narrowed PSS. Another way of interpreting the concern is to view it as an underdetermination problem. Limited data sets will always offer equal support for more than one hypothesis. The Deviant Interpretation Challenge is this kind of worry applied to bootstrapping. The Deviant Interpretation Challenge is problematic because we do think that learners come to endorse correct interpretations of concepts. If bootstrapping is to be the empiricist’s story about how learners acquire concepts, it needs to explain why learners come to endorse correct, as opposed to deviant, concepts.

In response to the Deviant Interpretation Challenge, Beck argues that such a challenge would be equally applicable to any ampliative account of learning, not just bootstrapping.⁵⁵ It is worth noting, however, that the Deviant Interpretation Challenge as formulated by Rey would only be

⁵⁴ For more on the Deviant Interpretation Challenge see: Georges Rey, “Innate *and* Learned: Carey, Mad Dog Nativism, and the Poverty of Stimuli and Analogies (Yet Again),” *Mind & Language* 29, no. 2 (2014); Georges Rey, “Concepts and Conceptions: A Reply to Smith, Medin and Rips,” *Cognition* 19, no. 3 (1985): 297–303.

⁵⁵ Jacob Beck, “Can bootstrapping?,” 121.

applicable if the theory of bootstrapping was concerned with how learners *ought* to proceed in acquiring new primitives. Goodman's worry was a normative one. For example, suppose you have Hypothesis A and Hypothesis B. Hypothesis A says that there are no fairies living in your garden; Hypothesis B says that there are fairies living in your garden, but they cease to exist when someone looks for them. Your data set consists of recorded observations where you check to see if there are fairies in your garden. Unfortunately, each time you look, there are no fairies to be found. Your observations fail to discriminate between Hypothesis A and Hypothesis B. This leaves you between a rock and a hard place. What principled reason could you have for selecting one hypothesis over the other? However, the theory of bootstrapping is not concerned with how learners should go about acquiring new concepts. Rather, it is trying to *describe* how learning a new primitive occurs. It is not beholden to the assumption that learners are rational agents who are always epistemically responsible. It can, therefore, sidestep this version of The Deviant Interpretation Challenge.

Even so, this does not mean that bootstrapping is completely free from these concerns. Human concept learning is likely systematic even if it is not ideal. There are three phenomena that, taken together, point towards this conclusion: (1) learners come to endorse the correct interpretations of new concepts; (2) learners endorse highly similar interpretations; (3) learners endorse the correct interpretation consistently over time. Thus, the Deviant Interpretation Challenge is more productively reformulated as a worry about what explanatory resources are available to explain these three phenomena, which computational bootstrapping struggles to do well.

The first phenomenon is that learners characteristically come to endorse the correct interpretation of a new concept. This suggests that the space of possible interpretations is not navigated at random. Although deviant interpretations may arise along the way, learners are not left

indefinitely among them. Rather, they tend to arrive at the interpretation that accords with successful use and instruction. Explaining this tendency is one of the central burdens any account of bootstrapping must bear.

The second phenomenon, that learners endorse highly similar conclusions, applies both to cases where the “correct” interpretation is endorsed and in cases where a “deviant” interpretation is temporarily endorsed. In many situations, teachers who have a great deal of experience teaching their subject, are able to predict what deviant interpretations their learners will endorse at certain points in the learning process. For example, in philosophy, beginner students often learn the concept of “validity” as it applies to arguments. A valid argument is one that has a successful deductive support relationship. However, because ‘valid’ is used colloquially as a synonym for correct or true, students will often label an argument as invalid if they believe the argument’s premises are false. This is mistaken because a valid argument can have untrue premises precisely because validity is only concerned with the support relation between the premises, not the content of the premises themselves.

The third phenomenon is that, once learners arrive at a particular interpretation, they do not typically abandon it or replace it at random. Human concept learning does not appear to involve constant, unprompted switching between competing interpretations. This is significant because if there were no systematic basis governing which interpretations learners endorse, we would expect much greater instability: concepts should shift frequently and without clear cause. But this is not usually what we observe. Instead, learners’ interpretations appear to be anchored in some underlying logic or system that gives them a degree of stability over time. An adequate account of bootstrapping must therefore explain not only how new interpretations are generated but also what stabilizes them once they are adopted.

As it currently stands, computational bootstrapping may be able to explain these three phenomena, but it does not explain them in a smooth or unified way. The view can describe how constraints narrow the range of possible interpretations and that modeling later yields the right concept. Still, that does not by itself explain why learners so often reach the “correct” concept. It also does not fully explain why their mistakes are often so similar or why their concepts remain fairly stable over time. To handle those points, the theory seems to need extra assumptions. Those assumptions can look ad hoc because they are not clearly supplied by the computational bootstrapping story itself. The problem, then, is not that computational bootstrapping fails. The problem, is that its current form does not have enough explanatory resources to account for the patterned and stable character of human concept learning.

I argue that altering the theory of bootstrapping to draw from an embodied theory of cognition makes all three phenomena unremarkable. This makes for a stronger theory of bootstrapping capable of fully responding to the Deviant Interpretation Challenge. Let us take each phenomenon in turn.

2.VII.a Learners Acquire Correct Concepts

Phenomenon one observes that learners come to endorse “correct” concepts. Throughout this paper, I have placed “correct” in scare quotes because it is not clear what, exactly, would make a concept correct, or what it would mean to specify truth conditions for one. Even so, we do have at least a partial answer. One reason we treat some concepts as correct is that they are useful. They help us navigate the world effectively.

This point becomes easier to see once embodiment is brought into view. Human beings have broadly similar bodies, and we inhabit broadly similar environments. As a result, many of the

things that matter for navigating the world matter to us in similar ways. Consider our ears. Human ears are structured similarly and function similarly. This makes our auditory experience and our sense of balance alike in important respects. Contrast this with certain birds. Some birds have small amounts of iron in their inner ears, which allow them to navigate by tracking the earth's magnetic field. For such creatures, the magnetic field is an important navigational resource. But human beings are not embodied in that way. We do not ordinarily perceive the earth's magnetic field, and so we do not form navigational concepts around it. The orientational metaphor UP IS GOOD makes a similar point. For a radially symmetrical creature floating in zero gravity, "up" might have little or no significance at all. Without gravity, and without a body organized around distinctions like head and feet or front and back, it is not clear what "up" would even amount to as a meaningful orientation. The general lesson is that the usefulness of a concept depends, in part, on the kind of body one has and the kind of environment one must navigate.

Although addition is more abstract than these literal cases of navigation, the same basic principle still applies. Addition is useful for creatures like us because it helps us keep track of objects, notice when quantities change, and avoid being cheated. It also supports more advanced forms of mathematics, which in turn make possible things like, safer buildings and complex economic systems. In that sense, addition helps sustain the kind of social world that creatures like us inhabit. Once this is recognized, it becomes less surprising that different people might independently endorse the same interpretations as "correct." What is useful for one creature like me will often also be useful for another creature like me.

Two further considerations make this convergence even less surprising. First, useful concepts are taught to us from a young age. We do not learn in isolation. Those who already possess these concepts can actively transmit them to others. Second, if we return to step three of the bootstrapping process, we can see that full interpretation does not emerge until the modeling stage is complete. In that stage, the learner tests a partially interpreted symbolic structure against possible candidate contents. If the fit is poor, the learner can revise the interpretation and try again. This process generates further feedback and makes it easier to determine which contents are good candidates for the structure. For that reason, it should not be especially surprising that we often converge on the same interpretations as useful—or, in the present sense, as “correct.” The final stage of bootstrapping allows for testing and revision until an adequate fit is achieved.

2.VII.b Learners Acquire Similar Concepts

In order to disentangle the idea that learners endorse similar concepts from the idea that learners endorse correct concepts, let us momentarily assume that learners never endorse correct concepts. Even on that assumption, it should not be surprising that learners arrive at similar interpretations. As we have seen, embodiment functions as a source of constraints on the interpretation of a PSS, and human bodies are broadly alike. We possess the same general sense organs, our bodily orientation is similar, and the basic dynamics of our movement through the world are similar as well. Because of these common features, we also share a wide range of experiences, such as collecting objects or feeling the pull of gravity. Any one of these embodied similarities could help explain why learners tend toward similar interpretations. Taken together, however, they make such convergence especially unsurprising. Given how much we share in bodily form and experience, it would in fact be far more puzzling if learners did not settle on similar interpretations.

2.VII.c Learners Hold Concepts Consistently

Phenomenon three observes that learners hold their concepts relatively consistent even across long spans of time. If there were no system anchoring interpretation, there would be little reason for a learner to hold on to one interpretation of a concept rather than another. We would expect endorsement to be unstable. Learners might move from one interpretation to the next without any clear basis for doing so. But this is not what we ordinarily observe. In general, learners do not switch interpretations at random.

The qualification “at random” is important. Embodied bootstrapping is not committed to the claim that the first interpretation a learner endorses must remain fixed for life. That would be incompatible with any plausible account of learning. Learning is a developmental process, and interpretations can change as new experiences arise, as old constraints weaken, or as new ones are added. A learner may therefore come to abandon one interpretation in favor of another. But that is very different from moving unpredictably and erratically between interpretations for no reason at all. The point is not that changes in how we interpret concepts never occurs. It is that such change is usually guided.

Once embodiment is brought into the picture, this stability becomes less mysterious. Our bodies are not randomly organized, nor do they change arbitrarily from moment to moment. They are structured in relatively stable ways and are adapted to carry out a common range of functions. Of course, our bodies develop over time and there is variation across individuals. But there is also enough continuity within a person and enough similarity across persons to give concept learning a stable basis. If embodiment helps constrain interpretation, then it is not surprising that learners tend to retain an interpretation over time unless they are given some reason to revise it. On this

view, stability is not accidental. It reflects the fact that concept learning is anchored by relatively enduring features of the body and its relation to the world.

2.VIII Conclusion

In this paper, I have argued that the theory of bootstrapping is significantly improved when it is situated within an embodied theory of the mind rather than a solely computational theory of the mind. Computational bootstrapping captures an important part of the story, but it does not exhaust the kinds of constraints that can guide the interpretation of a placeholder symbol set. Once embodied constraints are theorized and developed, bootstrapping becomes more empirically plausible because it can draw not only on the familiar computational cases but also on a wider range of cases grounded in bodily experience. Conceptual metaphors provide one especially clear illustration of how embodied constraints can structure and partially interpret placeholder symbol sets. And, the case of early arithmetic shows how this process can be completed through modeling in real pedagogical settings. This expanded account also places the theory in a stronger position to answer the Deviant Interpretation Challenge. If learners share similar bodies, similar environments, and similar patterns of interaction with the world, then it is far less mysterious that they tend to arrive at similar interpretations, to converge on interpretations that are useful enough to count as “correct,” and to retain those interpretations with considerable stability over time. What emerges from this discussion is a more expansive picture of bootstrapping itself. On that picture, embodiment is not peripheral, but a theoretically important part of explaining how the process is constrained and guided. Embodied bootstrapping is a promising way of strengthening the story of bootstrapping and allowing the empiricist to demonstrate how new conceptual primitives are acquired.

Paper Three: *Care, Justice, and the Duty of Pedagogical Competence*

Abstract for Paper Three

This paper argues that higher education instructors have a moral duty to bear the necessary and proportionate costs of becoming pedagogically competent. Such competence matters morally because higher education plays a central role in distributing socially and economically significant goods, while also helping to reproduce existing inequalities. Even so, concerns of justice cannot alone explain all instructors owe their students. Although it is crucial for understanding the broader institutional context, it does not fully account for the immediate, role-specific obligations that arise within the classroom itself. I argue that those obligations are grounded in care because the teacher-student relationship is structured by dependence, asymmetrical authority, and entrusted vulnerability. Drawing on Edward Hult Jr.'s account of pedagogical caring, I show that instructors must care for students qua students, as persons, and, where appropriate, as unique individuals. This account clarifies both the demands and the limits of pedagogical duty. Instructors are not required to become therapists or to take unlimited responsibility for students' lives, but they are required to adopt reflective, revisable practices that support learning and minimize avoidable harms. The paper concludes by addressing two objections: (1) that good teaching is merely a matter of style rather than substantive design; (2) that the duty of pedagogical competence is overly demanding.

3.I Introduction

In what is colloquially known as “the buried bodies cases,” a defendant being tried for murder revealed to his defense attorneys that he had committed two other, unrelated murders and disclosed the location of the bodies. For ordinary citizens, failing to report such a disclosure to relevant authorities would be morally repugnant. Yet, we do not necessarily draw the same conclusion of the attorneys who made the choice not to disclose this information. This is because being an attorney is a particular kind of social role and some social roles carry special permissions and special duties. My claim in this paper is that teaching in higher education is likewise a role with special moral duties. Specifically, I argue instructors have a duty of care to bear the necessary and proportionate costs of becoming pedagogically competent in the classroom.

This duty is not exhausted by the employment contract, though it is consistent with the fact that teaching obligations are also partly institutional and contractual. Instructors voluntarily occupy a role in which they exercise significant influence over whether students succeed academically, which in turn, affects students’ life prospects. Because of that role, students can rightfully expect more from their higher education instructors than subject-matter expertise alone. They can, and should, expect that their instructors will facilitate their learning of the material in ways that are competent, attentive, and developmentally responsive.

I also argue for a second, related claim. As higher education functions within societies marked by persistent inequality, instructors incur justice-based responsibilities as participants in institutions that shape the distribution of opportunity. These justice-based duties do not by themselves explain everything instructors owe to students. Even in a perfectly just society, teachers would still owe their students competent, attentive, and developmentally responsive teaching. But the

nonideal context of contemporary higher education intensifies the stakes of pedagogical failure and gives instructors additional reason to resist institutional arrangements that predictably disadvantage some students over others. However, justice-based obligations cannot and do not explain all that instructors owe to their students. Some other normative bridge is required to fully explicate all the duties that instructors incur in virtue of their role.

3.II The Natural Alignment of Justice and Higher Education

Before offering positive argument that higher education instructors' duties of pedagogical competence are grounded in care, it is instructive to consider why those duties cannot instead be grounded in the related framework of justice, which might initially seem better suited to concerns about equity in education.

Among the many roles higher education plays in a democratic society, two are especially salient for questions of justice. First, patterns of success in higher education are a use barometer of educational equity. The patterns of access, persistence, and graduation, reveal disparities across demographic groups. Second, higher education is a powerful mechanism for reproducing those same inequalities, insofar as it functions as a gatekeeper to valuable social and economic goods.⁵⁶

To briefly motivate these claims, consider the differences in postsecondary degree attainment by race and socioeconomic status, two of the most salient demographic categories across which inequality in American society are most commonly structured. In 2022, Asian and White adults were much more likely than Black/African American, American Indian/Alaskan Native, and

⁵⁶ Sandra R. Baum and Michael S. McPherson, *Can College Level the Playing Field? Higher Education in an Unequal Society* (Princeton University Press, 2022).

Hispanic/Latino adults to have attained a postsecondary degree.⁵⁷ Similarly, in 2022, estimated bachelor's degree attainment rates were almost 4 times higher for those coming from families in the highest income quartile than those in the lowest (58 percent vs. 16 percent).⁵⁸ Notably, disparities in socioeconomic status do not seem to account for racial gaps.⁵⁹ When this is considered alongside the fact that individuals with a bachelor's degree earn substantially more over their lifetimes than those with only a high school diploma and are significantly less likely to experience unemployment, we can begin to see how higher education achievement can be seen as a matter of justice.

Furthermore, the increased earnings associated with degree attainment are not merely abstract economic phenomena; they shape individuals' opportunities, security, health, and civic standing. Wealth is a safety net and facilitates access to more cost-effective goods that build security over generations. Those who are wealthy are afforded better chances to navigate personal, political, and natural disasters like major health crises, war, and hurricanes, while remaining relatively unscathed.⁶⁰ Additionally, the benefits of upward economic mobility are not limited to the tangible material benefits of wealth. As Sandy Baum and Michael McPherson put it:⁶¹

The equal opportunity ideal also extends beyond economic mobility, because it concerns not just making a lot of money but having the kind of work and life you

⁵⁷ American Council on Education, *U.S. Population Trends and Educational Attainment* (2022).

⁵⁸ Margaret Cahalan et al., *Indicators of Higher Education Equity in the United States 2024: 50-Year Historical Trend Report* (The Pell Institute for the Study of Opportunity in Higher Education, Council for Opportunity in Education (COE) and Alliance for Higher Education and Democracy of the University of Pennsylvania (Penn AHEAD)., 2024), 6.

⁵⁹ Anthony P. Carnevale and Jeff Strohl, *Separate & Unequal: How Higher Education Reinforces the Intergenerational Reproduction of White Racial Privilege* (Georgetown Public Policy Institute, 2013).

⁶⁰ This is even more pressing as global warming increases the number and strength of climate change events.

⁶¹ Baum and McPherson, "Can College," 34.

desire and are willing to work for. The American Dream is not only “anybody can grow up to be rich” but also “anybody can grow up to be president.”

Individuals with college degrees are more likely to have meaningful choices regarding the kind of work or career they ultimately pursue. While the social value of work is certainly not exhausted by credentialed professions, access to many careers regarded as especially meaningful or professionally rewarding, such as medicine or law, now depend upon obtaining a college degree. Those with a college degree are also more likely to vote, be otherwise civically engaged, and engage in volunteer work.⁶² Meaningful civic engagement is particularly crucial for justice in democratic societies because justice in such societies is highly dependent on citizens’ advocating for and affecting governmental decisions regarding their interests.

Additionally, persistence and retention rates are comparatively lower among under-resourced student populations.⁶³ This means that inequities in degree attainment are not simply a matter of disproportionate enrollment between privileged and under-resourced students; under-resourced students are also more likely to exit higher education altogether without returning. The consequences of leaving college without a credential are often substantial. Because many students must take out loans in order to attend college, those who withdraw before obtaining the credential linked to higher earnings remain responsible for repaying the debt they incurred. Accordingly, the burdens of noncompletion fall especially heavily on borrowers: within six years, 22% of those who left college without completing a degree had defaulted on at least one

⁶² Sandy Baum et al., *Education Pays 2013: The Benefits of Higher Education for Individuals and Society*, Trends in Higher Education (The College Board, 2013).

⁶³ See for example: <https://nscresearchcenter.org/persistence-retention/>

loan, whereas only 2% of borrowers who earned a bachelor's degree had done so.⁶⁴ Moreover, while borrowing has risen across all demographic groups, racial minority students have borne a disproportionate share of that burden.⁶⁵ Higher education, in this sense, functions as a high-stakes investment whose returns are contingent upon successful completion. The interpretation of the data presented in the section supports the claim that those most likely to reap the returns of higher education are individuals who are already socially privileged. Consequently, the benefits associated with earning a college degree accrue disproportionately to those who are comparatively better positioned to start.

3.III Grounding Teaching Obligations in Justice

I do not argue it is the case that failure in the classroom is always the cause of students choosing to withdraw from college. The environment (e.g., accessibility, feeling of belonging) and access to other supportive resources (e.g., tutoring, mental-health support, advising) most certainly can and do play a role in student persistence. However, what happens in the classroom is at least as crucial. We need only look at so called “gateway” courses to confirm this. Gateway courses are introductory, usually high-enrollment prerequisite courses that function as institutional threshold points, since success in them often determines whether students can advance within a major or persist in higher education at all. Because they occupy this gatekeeping role, they frequently reproduce broader educational inequities. Students from first-generation, low-income, and

⁶⁴ Lawrence Gladieux and Laura Perna, *Borrowers Who Drop Out: A Neglected Aspect of the College Student Loan Trend*, nos. 05–2 (The National Center for Public Policy and Higher Education, 2005), 8.

⁶⁵ Cahalan, et. al., *Indicators*, 6.

historically underrepresented backgrounds tend to experience disproportionately high rates of failure, withdrawal, and incomplete grades in such courses.⁶⁶

The Gardner Institute has had significant success in reducing the impact of gateway courses through an approach of faculty development and course redesign.⁶⁷ This makes it clear that instructors are able to make an impactful difference in higher education.

Instructors do not design the whole system, but they exercise significant discretionary power within it. They choose assessment strategies, attendance policies, modes of participation, grading practices, and pacing, of their courses. These choices affect who succeeds, who persists, and who comes to see higher education as a place where they can belong and flourish. Given this background, it is reasonable to think that what is happening in the higher education classroom ought to be directly within the reach of justice.

This line of thought is evident in the discourses of the learning sciences and the Scholarship of Teaching and Learning (SoTL), which are often structured around justice-based concepts. Concerns such as equality, fairness, equity, equal opportunity, which seem justicial in nature, permeate the SoTL literature. Global organizations such as the American Association of Colleges & Universities explicitly state that they are dedicated to the democratic purposes of and equity in higher education. However, it is not clear that justice can ground all the duties teachers owe their students. Whether or not one will be troubled by the concerns elaborated here is dependent on

⁶⁶ Andrew Koch and Brent Drake, “Digging into the Disciplines I: Accounting for Failure—the Impact of Principles of Accounting Courses on Student Success and Equitable Outcomes,” *John N. Gardner Institute for Excellence in Undergraduate Education*, 2018; Andrew Koch, *Transforming the Gateway Course Experience: A Call to Action for Higher Education*, 1st ed. (Routledge, 2024).

⁶⁷ For a brief summary see: “Rethinking Gateway Courses to Improve Outcomes, Equity,” *The Feed*, January 24, 2020.

which theory of justice they adopt. However, there are at least two concerns which might reasonably make one wary of grounding what teachers owe their students in justice.

The first concern is that justice is concerned primarily with institutions, not individuals' specific choices. Such a concern can be considered broadly Rawlsian. Rawls famously argues that the scope of justice is limited to what he calls the *basic structure* of society—the fundamental political, social, and economic institutions that shape the distribution of rights, liberties, opportunities, and resources. Rawls emphasizes that the principles of justice governing the basic structure, “must not be confused with the principles which apply to individuals and their actions in particular circumstances,” since, “these two kinds of principles apply to different subjects and must be discussed separately.”⁶⁸ This restriction of scope is central to Rawls's framework. Justice concerns the basic structure because a society is just only when its social arrangements are justifiable to each citizen. He takes the form of a society's basic institutions to be both necessary and sufficient for securing that kind of justification. This restriction seems to put the regulation of what individual teachers do in their own classrooms outside the purview of justice.⁶⁹

Another result of focusing on the arrangement of institutions is that when such institutions are suitably non-ideal, Rawls (among others) argues that one's duties morph. We are only required to comply with institutions that are just. When they are not, we are required to, “further just arrangements not yet established, at least when this can be done without too much cost to

⁶⁸ John Rawls, *A Theory of Justice*, (1971; Original ed, Belknap Press, 2005); 54-55.

⁶⁹ Gina Schouten has some work arguing that perhaps the basic structure is capable of more expansive intervention than initially thought. If such arguments work, then it is possible that individual, role-specific teaching duties could be grounded in justice. I take this to only strengthen the force of the duty. However, such arguments are not without controversy, so I have chosen to set them aside here. See: Gina Schouten, “Restricting Justice: Political Interventions in the Home and in the Market,” *Philosophy & Public Affairs* 41, no. 4 (2013): 357–88; Gina Schouten, “Gender and the Division of Labor,” in *Routledge Handbook of Philosophy, Politics, and Economic* (Routledge, 2022.).

ourselves.”⁷⁰ I argue the conditions of higher education are suitably unjust as to trigger this change in duty.

The research cited in the first section of this paper serves as strong evidence of higher education’s unjustness. Furthermore, the incentive structure of many institutions of higher education are suitably perverse to be considered non-ideal. Institutions of higher education often profess a tripartite commitment to research, teaching, and service. In practice, however, the prevailing incentive structure tends to privilege research productivity over the other two domains.⁷¹ Especially in R1s, hiring, tenure, promotion, and prestige are frequently tied most closely to publications and grant acquisition, creating pressures that can marginalize teaching and service, despite their central importance to the educational mission. While this is likely most prominent in research universities (e.g., R1s, R2s), I think there is good reason to suspect that similar features infect, small liberal arts colleges (SLACs), flagship state universities, and comprehensives, among others.

This imbalance is compounded by the asymmetry in professional preparation. Faculty typically receive extensive methodological and disciplinary training in research during graduate school, but comparatively little formal training as teachers. This preparation is in line with many universities’ incentive structures, which highly reward research. While teaching and service are valued in name, they often receive less recognition, fewer resources, and are not perceived as prestigious. When poor teaching is accepted, even implicitly encouraged by the incentive structure of the university, it is those who are already at a disadvantage that suffer. Those who are well resourced, who can afford tutors or have family members that have already completed a

⁷⁰ Rawls, *A Theory of Justice*, 115.

⁷¹ This is not true of every institution, but is true of many.

degree, are more likely to weather these non-ideal conditions. They will still succeed. However, for those who come in already at a disadvantage, (e.g., first-generation college students, students of color, students whose first language is not English) these conditions disproportionately affect their success for the worse. The system thus serves to further replicate and entrench the inequalities seen in society at large.

One might also believe that such a change is triggered by the fact that developing and employing good teaching practices might place an unfair burden on some instructors in light of what those around her are doing and the incentives of the institution in which she finds herself. Laura Valentini argues that, as existing institutions become increasingly unjust due to pervasive non-compliance, the natural duty of justice correspondingly becomes less demanding. This claim makes especially good sense if duties of justice are understood primarily as duties concerning the basic structure of society, since the justice of that structure depends on large-scale patterns of coordination that no single individual can secure alone.

On this view, many of the actions justice requires, such as creating, sustaining, or reforming fair institutions, can achieve their intended moral effects only when enough others also comply. Under conditions of widespread non-cooperation, however, any one individual's marginal contribution to the realization of just institutional arrangements becomes severely limited. In such circumstances, it would be unreasonable to hold that the natural duty of justice continues to impose equally demanding requirements on individuals. Justice cannot require agents to bear unlimited burdens in pursuit of structural outcomes that their actions alone cannot plausibly bring about. For this reason, Valentini concludes that, "Justice—a moral concern so pervasive in contemporary political philosophy—can do less moral work than assumed. Beneficence, I have

concluded, will often have to pick up the slack.” She emphasizes that this conclusion has significant implications for how political obligations should be understood under non-ideal conditions and cautions against the common tendency in contemporary political theory to treat all high-stakes moral questions as matters of justice alone. Moral principles such as beneficence, by contrast, may remain demanding even where the actions of others fall short, precisely because their force does not depend in the same way on collective compliance at the level of the basic structure.

If either the concern regarding the scope of justice or the transformation of one’s duties in non-ideal conditions are true, then we should not focus on what is happening in any given classroom. Rather, efforts should be focused on elucidating the duties of instructors to reform the institution of higher education.⁷² While it is likely true that instructors ought to work towards establishing more just institutions of higher education, to think that doing so exhausts all duties teachers owe their students is a mistake.

Consider two instructors, Teacher A and Teacher B. Teacher A satisfies these duties of justice by advocating for institutional reforms that would make higher education more equitable, yet she makes no effort to alter her own pedagogy in ways that would improve equity for the students presently in her classroom. Teacher B also advocates for institutional reform but, in addition, also makes reasonable pedagogical changes aimed at producing more equitable outcomes for her students. Teacher B is plainly morally preferable. Yet the difference between them should not be understood as the difference between fulfilling one’s obligations and performing a supererogatory act. The fact that an action would significantly benefit others is not, by itself,

⁷² What the content of these duties are and how instructors ought to go about fulfilling them is a matter for another paper.

sufficient to show that it is morally required, since some beneficial actions may impose extraordinary costs on the agent. Here, however, the relevant pedagogical efforts are achievable at comparatively modest personal cost. Given the substantial influence instructors exercise over student success, the importance of those outcomes for students' broader life prospects, and the relative modesty of the sacrifices involved, it is more plausible to conclude that Teacher B is fulfilling a moral requirement that Teacher A neglects. She is not going beyond duty; she is responding appropriately to obligations internal to the role of teacher. In the subsequent sections of this paper, I will argue that the duty instructors' duty to bear the necessary and proportionate costs of becoming pedagogically competent are grounded in an ethic of care and influenced by the special nature of the teacher-student role.

3.IV Grounding Teaching Obligations in Care

If we accept Valentini's suggestion that beneficence must "pick up the slack," the next question is how beneficence ought to be understood in the context of teaching and what, precisely, it requires of instructors. More specifically, we must ask both what grounds instructors' pedagogical duties and what the substantive content of those duties is.

One appealing way to ground an instructor's pedagogical duties would be to opt for some form of a utilitarian view. After all, teachers are dealing with large numbers of students who they ought to do right by. Utilitarianism would suggest that we maximize the good, whatever that may be, for the greatest number of people. But to do this would miss the key role that my students are, well, *my* students. It seems that I owe duties to them that I do not necessarily owe to my colleagues' students. I have certain commitments to my students that they should and can rightfully demand from me that Professor So and So's students cannot. And, if I was to neglect

my own students to benefit Professor So and So's, I have somehow failed, even misunderstood, my own duties. It is this relationship that grounds my duties to my students, not some abstract commitment to maximizing some universal good. What this suggests, and what several critics of utilitarianism have explicitly argued, is that moral reasoning cannot be reduced to the impartial maximization of good without obscuring the claims of those to whom we stand in particular relationships.⁷³ Furthermore, even where utilitarianism happens to endorse actions that align with a care-based view, this alignment is merely incidental and fails to recognize that the moral basis of such obligations lies first and foremost in the relationship itself.

A useful way to clarify the care-based grounding of pedagogical duties is to begin with a case outside education. Consider the parent–child relationship. On a care-ethical view, the moral demands that arise within this relationship are not best understood as derivative of abstract impartial principles alone, nor as exhausted by whatever arrangements the parent may have explicitly consented to. Rather, they arise from the concrete relation itself: a relation marked by dependence, vulnerability, asymmetries of power, and the child's need for sustained attention and support. This is why it would be morally unintelligible for a parent to claim that, although they have chosen to become a parent, they are under no obligation to feed, protect, or attend to their child because doing so would interfere with activities they personally prefer. The point is not merely that neglect would have bad consequences, though it would. It is that standing in this particular relationship (role) to someone else gives rise to responsibilities internal to that particular role.

⁷³ I think this is still accurate despite the he Peter Railton–Frank Jackson or Richard Yetter line. See: Philippa Foot, *Moral Dilemmas: And Other Topics in Moral Philosophy* (Oxford University Press, 2002); J. J. C. Smart and Bernard Williams, *Utilitarianism For and Against* (Cambridge University Press, 2008).

Care ethicists, such as Nel Noddings and Carol Gilligan, help make sense of this structure by emphasizing that moral life begins not from the standpoint of detached rule application but from encounters with concrete others whose needs call for response.⁷⁴ On such a view, caring relations are morally significant because they place agents in positions where attentiveness, receptivity, and responsiveness are required. These requirements are not wholly optional, nor are they entirely self-defined. Although parental duties are in an important sense voluntarist—one generally becomes a parent through choices for which one can be held responsible—the content of those duties is not simply set by preference. One may choose whether to enter the relationship, but one does not, having entered it, retain full discretion over what one owes within it. The child's dependency and the parent's power together constrain what responsible action can look like.

This feature of parental obligation makes it a useful analogical model for teaching. The instructor–student relationship is, of course, not identical to the parent–child relationship; students are not children, and teachers do not stand in an all-encompassing caregiving role. Still, the analogy is instructive because teaching likewise involves a structured relationship of dependence and asymmetrical authority. Students depend on instructors for access to knowledge, evaluation, guidance, and the conditions under which learning can occur. Instructors, in turn, exercise significant control over classroom norms, assessment structures, opportunities for participation, and the distribution of pedagogical attention. Once one voluntarily occupies the role of teacher, one does not simply acquire a set of contract-like duties whose content can be

⁷⁴ Nel Noddings, *Caring: A Feminine Approach to Ethics & Moral Education*, 2nd ed (University of California Press, 2003); Carol Gilligan, *In a Different Voice: Psychological Theory and Women's Development* (Harvard University Press, 2003).

minimized at will. Rather, one assumes responsibilities shaped by the nature of the relationship itself and by the needs of those entrusted to one's instruction.

Seen in this light, duties of pedagogical competence are not well understood as optional expressions of personal generosity or as merely instrumentally useful techniques for producing better aggregate outcomes. They arise because teaching is a morally thick practice in which one party is especially well positioned to affect another's intellectual development, chances of success, and, in many cases, broader life prospects. Just as a parent may not permissibly refuse basic care because other pursuits seem more appealing or rewarding, an instructor cannot simply dismiss the obligation to cultivate pedagogical competence on the grounds that doing so is inconvenient or competes with other uses of time. That claim need not imply that instructors must sacrifice without limit. It implies, rather, that the role itself carries substantive responsibilities, and, that at least some degree of attentiveness to how one's teaching affects students, is part of what it means to occupy that role responsibly.

This analogy also helps illuminate the non-contractual character of pedagogical duty. Teaching, like parenting, is voluntary in the sense that one typically chooses to enter the role. But once one has done so, the vulnerable position of the other party and the normative structure of the relationship place limits on what one may permissibly neglect. The moral force of these duties therefore comes not simply from general obligations to promote welfare, nor only from institutional job descriptions, but from the fact that teaching establishes an ongoing relation in which students' learning and flourishing are partly placed in the instructor's hands. A care-based account captures this especially well because it explains why these obligations are sensitive to

dependence, power, and need, while still allowing that their precise content must be interpreted in light of the specific practices and constraints of higher education as it is today.

Importantly, this grounding of pedagogical responsibility does not depend upon broader institutional failures or injustices within higher education. Even in a perfectly just educational system—one in which resources, opportunities, and institutional structures were distributed fairly—individual instructors would still owe their students competent, attentive, and developmentally responsive teaching. The obligations generated by care therefore operate independently of, and in addition to, duties of justice. Recognizing this point helps clarify why failures of basic pedagogical competence constitute not merely ordinary professional imperfections, but moral shortcomings.

3.V The Content of Teaching Duties of Care

Recognizing that care independently grounds pedagogical duties clarifies why failures of teaching can constitute moral failures, but it does not yet tell us what caring teaching requires in practice. One account that is particularly helpful in translating the theory of care ethics into the realm of education is Edward Hult Jr.'s account of pedagogical caring.⁷⁵ Hult's framework is especially useful because it treats care not as an undifferentiated emotional disposition, but as a role-specific moral orientation.⁷⁶ On his account, pedagogical caring must be understood in light of what teachers are for and of the distinctive responsibilities generated by the teacher-student

⁷⁵ This is not to say that others' work on care ethics is not useful or relevant to duties of teaching, this account is just particularly illustrative.

⁷⁶ Richard E. Hult, Jr., "On Pedagogical Caring," *Educational Theory* 29, no. 3 (1979): 237–43.

relationship. That starting point is important for the present argument. If instructors have duties to bear the necessary and proportionate costs of becoming pedagogically competent, those duties are not free-floating demands to maximize good wherever possible. They arise, rather, from the fact that one has assumed a role in which others' learning, development, and prospects are placed in one's hands.

Hult's role-based approach is particularly attractive because it avoids two opposite mistakes. On the one hand, it rejects the thin view according to which teaching is exhausted by content delivery and formal assessment. On the other hand, it also resists the thought that caring teaching requires the instructor to become a therapist, intimate confidant, or limitless emotional caretaker. Hult explicitly frames pedagogical caring as a role concept and warns against conflating pedagogical and psychotherapeutic responsibilities.⁷⁷ This is a crucial point for higher education. A care-based account of teaching duties should indeed be morally demanding, but it should not be unbounded. Instructors are responsible for students in a specific capacity—as learners situated within a pedagogical relationship structured by authority, dependence, and asymmetries of expertise. The content of their duties must therefore be specified in a way that is both morally serious and appropriately limited to their role. One useful feature of Hult's discussion is his appeal to R.S. Downie's levels of recognition (as individual, as person, as student) to further elucidate the content of the duties of care in teaching. For present purposes, however, those levels are best approached in reverse order, beginning with the most role-specific and moving outward.

⁷⁷ Ibid.

3.V.a Caring for Students as Students

At the most basic level, the instructor must recognize the student qua student. This is the minimal but indispensable form of pedagogical care. To recognize students in this way is to recognize their standing as learners whose ability to succeed is shaped by course design, clarity of expectations, forms of assessment, opportunities for practice, and the accessibility of instructional support. At this level, care requires more than good will. It requires pedagogical competence. An instructor who persistently teaches in ways that predictably confuse students, relies on assessments that mismeasure learning, or ignores obvious obstacles to participation, fails to care for students precisely as students. The wrong here is not simply professional underperformance. It is a moral failure to discharge duties incurred by the role one has chosen to occupy.

Understanding care at this level also clarifies the limits of the duty. Because the object of concern is the student as student, the teacher's obligation is not to remedy every dimension of a student's life. Hult's framework is helpful precisely because it protects against the familiar worry that a care-based conception of teaching would require impossible levels of personal knowledge or emotional resources. Such a requirement would be especially implausible in large lecture courses, where instructors may teach hundreds of students at once. Care cannot require what the role and circumstances make impossible. It can, and does, require that instructors make their courses as educationally responsible as they are able given their particular circumstances. It requires that they create clear pathways for participation, use transparent expectations, build opportunities for feedback, and refer students to appropriate professionals when their needs

exceed the instructor's competence. In this respect, the duty to become pedagogically competent is itself partly a duty of boundary-setting and role clarity. Caring well includes knowing what one is, and is not, equipped for as a teacher.⁷⁸

3.V.b Caring for Students as People

The second level of recognition requires recognizing students not merely as occupants of the learner role in the abstract, but as persons whose experience of the classroom may differ in morally significant ways. Students arrive with different histories of preparation, different relations to authority, and different vulnerabilities to common classroom structures. A caring instructor must therefore cultivate enough attentiveness to notice when routine pedagogical decisions are not affecting all students equally. Without reflection and self-correction, instructors are likely to mistake what is familiar to them for what is accessible to students.

This second level also helps explain why students must not be used as instruments of self-aggrandizement. An instructor may be tempted to treat the classroom primarily as a stage for displaying brilliance, securing admiration, or reaffirming personal authority. That temptation is real, especially in academic cultures that attach prestige to performance, charisma, and intellectual dominance. Yet from a care-based perspective, such an orientation is a distortion of the pedagogical relationship.⁷⁹ In a caring relationship, the instructor is called to attend and respond to the needs specific to that relationship. Self-aggrandizing pedagogy represents a failure of attentiveness. The teacher's focus has shifted from students' learning to the maintenance of the

⁷⁸ This plausibly includes appropriately communicating this to the student. How and when this communication should occur is ripe for further consideration but is outside the scope of this paper.

⁷⁹ Other ethical traditions (e.g., Kantian) could also support this point. I think care ethics does a particularly good job at explaining and supporting our intuitions about why the failure is a wrong in this case, so I place my focus there.

teacher's image. The wrong, then, lies in a failure of responsiveness within an entrusted relationship of asymmetrical dependence. Students are not merely disrespected in some general moral sense; they are neglected in the very capacity in which the teacher has a special duty to care for them.

3.V.c Caring for Students as Individuals

At the third level of recognition, pedagogical care involves recognition of the student as a unique individual self. Recognition, in this sense, is a recognition of something like students' distinctive talents, struggles, and aspirations. Hult argues that this sort of recognition is only appropriate in *some* teaching contexts, such as work with doctoral students who are being trained to be colleagues.⁸⁰ His argument does not imply that every instructor must cultivate deep personal familiarity with every student, nor that higher education classrooms should be transformed into sites of intensive emotional caretaking. Rather, it acknowledges that in some teaching contexts instructors should remain alert to the fact that students are whole persons whose intellectual development matters from within their own lives. The moral point is not that teachers must know everything about their students, but that they must resist treating them as faceless, fungible inputs to deposit knowledge into.

Once Hult's account is read in this way, it becomes easier to specify the duty to bear the necessary and proportionate costs of becoming pedagogically competent. The relevant duty is not merely to have caring motives; it is to undertake the reflective work required to make one's care intelligent, reliable, and educationally effective. A teacher who genuinely cares for students qua

⁸⁰ Hult, "On Pedagogical," 242.

students cannot remain indifferent to evidence that her teaching practices are not working well. For that reason, pedagogical reflection is not ancillary to care. It is part of its content. If one's role places one under an obligation to support students' learning responsibly, then one must be prepared to ask whether one's methods are in fact serving that end.

This aligns with many of the best practices that the Scholarship of Teaching and Learning has been advocating on behalf of for years. For example, student-centered learning can be understood as an expression of care at the second level of recognition because it treats students not merely as generic occupants of the learner role, but as distinct persons whose perspectives and needs matter pedagogically. To teach in a student-centered way is to acknowledge that students do not encounter course material from a uniform starting point, and that effective teaching, therefore, requires attentiveness to how different students experience the classroom. In this respect, student-centered pedagogy reflects care not simply by improving learning outcomes, but by listening to, respecting, and promoting the students' participation in the learning process. Another example is the process of continuous improvement. Hult argues that caring for students qua students, "implies provision of skill for the enhancement of student learning," and continuous improvement gives this requirement a concrete pedagogical form.⁸¹ As an iterative and context-sensitive process, it requires instructors to identify problems, test and refine responses through ongoing evidence collection, and adapt their practices over time in pursuit of sustainable change. In this way, continuous improvement exemplifies how care for students, at the level of their role as learners, demands not static goodwill but the ongoing development of pedagogical competence. Other concrete practices might include: keeping a teaching journal,

⁸¹ Hult, "On Pedagogical," 243.

collecting student feedback, seeking out classes or workshops on pedagogy, or staying informed of relevant teaching and learning literature. None of these practices are costless. They require time, humility, and sustained effort. But they are paradigmatic examples of the kind of necessary and proportionate costs that care-based pedagogical duty can require instructors to bear.

In short, Hult's account helps show that the content of teaching duties grounded in care is neither maximalist nor minimal. It does not require teachers to become therapists or to assume unlimited responsibility for students' lives. But neither does it permit instructors to retreat into a narrow conception of their role as mere transmitters of content. To care pedagogically is to recognize students first as students, then as individuals whose differences matter to learning, and finally, where appropriate, as unique selves whose development calls for humane attention. Once this is granted, the duty to develop pedagogical competence follows naturally. Instructors owe students not simply subject-matter knowledge, but reflective and revisable teaching that a caring discharge of the role requires.

3.VI Objection 1: It's All Style

One objection to the view I have developed here is that good teaching is merely a matter of style rather than a matter of substantive moral consideration or objective effectiveness.⁸² On this objection, pedagogical differences are analogous to differences in temperament, presentation, or personal flair: one instructor is more animated, another more reserved; one favors discussion,

⁸² This line of argument has been often levied by faculty around me who are skeptical of, unaware of, or uninterested in the work the of SoTL and learning sciences scholars.

another lecture; one is warm and informal, another distant and highly structured. If teaching is understood in this way, then some faculty are inclined to conclude that there is no meaningful standard by which one could say that an instructor ought to teach differently, only that different instructors have different styles. The objection is plausible in part because teaching does in fact admit of considerable variation, and because many instructors can be effective while embodying very different classroom personas.

In response, it is illustrative to return to the earlier analogy of parenting. As there are many ways to be a good parent, so are there many ways to be a good teacher. However, this does not mean that good parenting, or good teaching, is a free-for-all. As with parenting, certain practices are much more appropriate, effective, and moral than others (e.g., time outs instead of beating one's child). What practices one should adopt depends in part on one's aims, one's "children", and the material facts of one's circumstances.

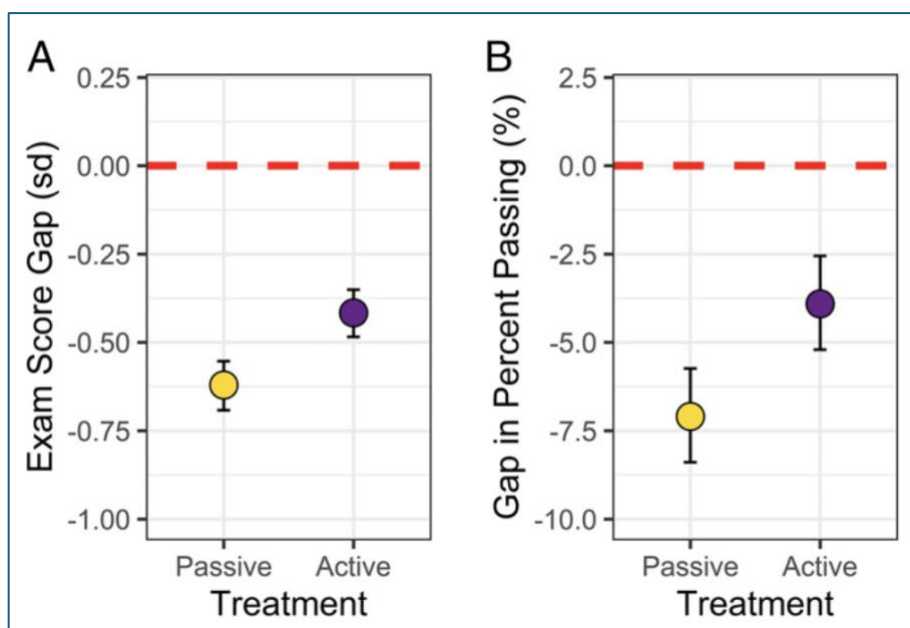
For example, a substantial body of empirical research indicates that certain pedagogical designs are more effective, on average, than traditional lecture-based approaches.⁸³ A widely cited meta-analysis by Freeman et al. found that active learning strategies in undergraduate Science, Technology, Education, and Mathematics (STEM) courses were associated with significantly improved examination performance and lower failure rates compared to traditional lecturing.⁸⁴ In another particularly illustrative study, the exam pass rates and course pass rates for underserved students were compared across traditional lecture courses and active learning courses. As shown

⁸³ See earlier in this dissertation: "*What is Active Learning? A Philosophical Audit.*"

⁸⁴ Scott Freeman et al., "Active Learning Increases Student Performance in Science, Engineering, and Mathematics," *Proceedings of the National Academy of Sciences* 111, no. 23 (2014): 8410–15.

in Figure 4, the differences between the groups were measured in standard deviations from parity.⁸⁵ Students in traditional lecture courses exhibited a larger negative achievement gap (approximately -0.6 SD), whereas courses using active learning show a smaller gap (approximately -0.4 SD), indicating improved relative performance for underserved students. As the values closer to zero represent greater equity in outcomes, the results suggest that active learning environments are associated with a meaningful reduction in the exam score and course-passing gaps.

Figure 7. Course and Exam Pass-Rates Gap for Under-resourced Students by Class Form.⁸⁶



Beyond decreasing the achievement gap for a single course (a worthy goal itself), success in the first year of college is often predictive of degree attainment. This might be due to the fact that

⁸⁵ Elli J. Theobald et al., “Active Learning Narrows Achievement Gaps for Underrepresented Students in Undergraduate Science, Technology, Engineering, and Math,” *Proceedings of the National Academy of Sciences* 117, no. 12 (2020): 6479.

⁸⁶ Ibid.

those students who come to college more prepared are more likely to succeed in their first year and eventually graduate. However, the success seen from interventions like first-year seminars suggests that something more nuanced might be happening.⁸⁷ If students experience good teaching and classroom success early on, it can influence how they feel about themselves and higher education in general, which in turn affects their rates of persistence and degree attainment. This is particularly important for under-resourced student populations, for whom doubts of belonging are more likely to undermine their confidence and achievement.⁸⁸

Crucially, findings such as these do not establish that every instructor must adopt identical practices. They do, however, undermine the view that pedagogy is purely a matter of stylistic preference. When credible evidence indicates that certain designs improve both effectiveness and equity, ignoring that evidence becomes morally salient. Teachers ought not adopt all of the same practices, but to pretend that all pedagogical choices are equally good is to ignore the real differences they make in shaping students' opportunities to learn, persist, and succeed.

3.VII Objection 2: It's Too Demanding

Even if no one wishes to defend it in writing, there will be a concern from many reading this paper that this view is too extreme, too costly. Higher education instructors might object that

⁸⁷ Jayne Brownell and Lynn Swaner, "Five High-Impact Practices: Research on Learning Outcomes, Completion, and Quality," *Peer Review* 14, no. 3 (2012); George Kuh, "High-Impact Educational Practices: What They Are, Who Has Access to Them, and Why They Matter," *Peer Review* 14, no. 3 (2012): 3–18.

⁸⁸ For evidence of this claim see: Gregory M. Walton and Geoffrey L. Cohen, "A Question of Belonging: Race, Social Fit, and Achievement.," *Journal of Personality and Social Psychology* 92, no. 1 (2007): 82–96.

they already carry a heavy load (especially in some institutional contexts) and that expecting them to become specialists in the scholarship of teaching is demanding too much of them. Meaningful course redesign and pedagogical development are often costly undertakings. They demand substantial investments of time, energy, and institutional support, all of which are finite and therefore must be drawn away from other responsibilities. For instructors, these resources are effectively zero-sum. Increased attention to teaching may come at the expense of research, service, administrative labor, mentorship, and responsibilities beyond academic life. The burden is compounded by the fact that, within many (but certainly not all) institutions of higher education, professional advancement, prestige, and even job security, are tied more closely to research output than to teaching quality.

To highlight the pervasiveness of such beliefs, let me turn for a moment to my home discipline of philosophy. Recently, the philosopher Stephen Bloch-Schulman describe his experience as a philosopher who focuses on pedagogy:

I am the opposite of surprised when, in conversation with other philosophers, my mentioning that [philosophical pedagogy] is a focus of my work as a philosopher is a great way to encourage the person I am talking to to shift uncomfortably, scrunch up their nose as if I have an unpleasant smell about me, and find an excuse to go talk to others. To the “real” philosophers.”⁸⁹

⁸⁹ Stephen Bloch-Schulman, “Making the Case for the Value of Teaching through Self-Authorship: From ‘What Does the Discipline Want of Me?’ To ‘Who Do I Want to Be?’” in *Innovations in Teaching Philosophy: A Toolkit for the 21st-Century Classroom* (2026), 325.

There is a pervasive belief, even if no one wants to explicitly own up to holding it, that being a good teacher and being a good researcher are fundamentally at odds.⁹⁰ And, to prioritize one's teaching over one's research is often career suicide, if one wants to be a prominent scholar in their field. There is way of reading this view which requires instructors not only to develop an additional skill set, but to possibly do so at the expense of other personal goals they might hold. This objection deserves to be taken seriously, since faculty workloads, institutional resources, and access to pedagogical training vary widely across institutions and appointments. Many instructors are asked to teach while managing large course loads, significant research expectations, and little meaningful support for developing as teachers.

Again, returning to the parenting analogy is instructive. We do not completely excuse parents in non-ideal circumstances from providing a minimal level of care to their child.⁹¹ We should not completely excuse instructors in non-ideal conditions. This does not mean that instructors are required to achieve pedagogical perfection. Nor are they required to sacrifice all research or other commitments in favor of teaching. However, I argue they are required to meet a minimum threshold of teaching competence, below which they are morally culpable for failing to meet their duties to their students.

Developing this competence is not as costly as faculty often make it out to be. The point is not that instructors must become specialists in pedagogy, but that they must take seriously the needs

⁹⁰ This example is drawn from my own discipline of philosophy. However, I do not think the critique I am making is limited solely to philosophy, but is present throughout the academy, though likely to varying degrees in different disciplines and venues.

⁹¹ We might think they are not necessarily blameworthy for failing to care for their child, but we do not accept their incompetence. We have major social agencies dedicated to determining when an individual is fit to have custody of their child, and when they are unfit the child is taken away.

of students as dependent learners whose prospects are shaped by the quality of instruction they receive—the quality of choices the instructor makes. Instructors should be required to:

1. Engage seriously with relevant scholarship on teaching and learning.
2. Cultivate a reasonable level of pedagogical competence.
3. Employ empirically supported practices where feasible and appropriate.

It is true that instructors in higher education receive very little training in how to teach.⁹² And, what they do often receive, is less thorough or helpful than we might hope. However, those teaching in higher education often have extensive training as researchers, as scholars. As a group, we have spent countless hours reading and evaluating scholarly work in our fields. As a graduate student, it is expected that we become intimately familiar with the bodies of work that inform our research and scholarly outputs, and we are judged on our facility with the literature in our field. Clearly, then, instructors in higher education, have the skill necessary to fundamentally inquire into how to better their teaching. However, they stop short of applying it to their own practice. This should strike us as negligent, for most faculty, teaching will constitute a major portion of their role. Yet, for some, reason, there is a pervasive tendency to view ourselves as researchers first and foremost, and as teachers only secondarily. This tendency is misguided because it lends to obscuring the moral significance of the teaching we do.

Fields such as the Scholarship of Teaching and Learning and the Learning Sciences already provide extensive insight into how students learn, the barriers that commonly impede them, and the teaching practices most likely to promote their success. As a result, instructors are not

⁹² David W. Concepción et al., “The State of Teacher Training in Philosophy,” *Teaching Philosophy* 39, no. 1 (2016): 1–24.

required to reinvent the wheel or arrive at sound pedagogy through trial and error alone. What is required, rather, is that they use their previously developed research skills to engage a body of literature, that has often been deliberately synthesized for practical use, and then apply its findings thoughtfully within their own classrooms while reflecting on how those choices affect student learning.

To understand the force of the claim that neglecting pedagogical obligations can constitute a moral failure, we must distinguish among different forms of pedagogical shortfall. A care framework is especially useful here because it resists treating all failures alike and instead asks whether an instructor has responded appropriately to the needs and vulnerabilities of those in their care. On this view, at least three categories matter: ignorant failure, negligent failure, and conscientiously constrained practice. Ignorant failure occurs when an instructor never meaningfully engages with the scholarship of teaching and learning and is thus unaware of widely available evidence relevant to student learning. An example might be an instructor who thinks teaching is merely a matter of personal style. There is also negligent failure, in which an instructor is aware of the scholarship on teaching, or at least aware that pedagogical inquiry bears on student success, but makes no serious effort to develop minimal teaching competence. An example of instructors in this category might be those that primarily conceptualize themselves as disciplinary researchers who happen to teach. They have rigorous standards for their work as researchers, but they do not apply the same effort or rigor to their teaching. Finally, there is conscientiously constrained practice, in which an instructor attempts to teach responsibly and improve in good faith, but their progress is less than ideal due to their material circumstances. Unlike the duties that arise from justice, those arising from care are not lessened or morphed by the noncompliance of others. Even so, this does not mean that every pedagogical shortfall is

equally blameworthy. A care-based account still requires us to distinguish between failures that reflect a lack of attentiveness or responsiveness and those that arise despite sincere efforts under nonideal conditions.

These distinctions matter because moral criticism is appropriate only where there is a failure of attentiveness or responsiveness, not wherever ideal care is not attained. Faculty are not blameworthy for every pedagogical imperfection, since care is always exercised under conditions of limitation, dependency, and unequal support. An instructor who is trying to respond well but is constrained by institutional circumstances may still fall short of minimal teaching competence without thereby manifesting a moral defect. Ignorant and negligent failures are different. Where instructors refuse to engage even minimally with credible evidence about how students learn or knowingly disregard evidence about practices that would better support student learning, they fall short not merely of professional aspiration but of the moral demands attached to their role.⁹³ This is particularly clear when the evidence concerns inequitable patterns of exclusion or predictable barriers faced by underserved students, since indifference in such cases is especially morally salient.

3.VIII Conclusion

The argument of this paper has been that higher education instructors occupy a role with distinctive moral demands, and that among those demands is a duty to bear the necessary and

⁹³ While some might recognize professional obligations as a subset of moral obligations, others (especially those who are not professional ethicists) often tend to view professional obligations as distinct from, or less serious than, other moral obligations. This is the view I intend to contrast when I say that such shortcomings are moral, not professional.

proportionate costs of becoming pedagogically competent. Because higher education functions as a gatekeeper to unequally distributed social and economic goods, what happens in the classroom is morally consequential in ways that extend far beyond course grades alone. Justice helps explain why these stakes are so high and why instructors also incur responsibilities to resist institutional arrangements that predictably reproduce inequality. But, justice cannot fully explain what instructors owe the students who are already in their classrooms. For that, we need a care-based account. Care better captures the direct, role-specific obligations that arise from the teacher-student relationship itself: obligations of attentiveness, responsiveness, and competence owed to students as dependent learners whose prospects are shaped, in part, by the quality of the instruction they receive.

Once pedagogical duty is understood in these terms, the obligation to develop a threshold level of teaching competence follows naturally. This obligation does not require instructors to become therapists, to achieve pedagogical perfection, or to sacrifice every competing professional aim in the name of teaching. Nor does it imply that every pedagogical shortfall is equally blameworthy. But it does require that instructors engage seriously with available scholarship on teaching and learning, reflect on the effects of their pedagogical choices, and make reasonable efforts to adopt practices that better support student learning and reduce predictable forms of exclusion. Where instructors remain merely ignorant or negligently indifferent to these responsibilities, they fail not only professionally, but morally. If teaching in higher education is as impactful as I have argued, then pedagogical competence cannot remain an optional adornment to scholarly life. It must instead be recognized as part of what responsible participation in the role of higher education instructor requires.

Bibliography

- Adams, Caleb, and Anthony Dove. 2016. *Flipping Calculus: The Potential Influence, and the Lessons Learned*. 10 (3): 154–64.
- Adams, Caleb, and Anthony Dove. 2018. “Calculus Students Flipped Out: The Impact of Flipped Learning on Calculus Students’ Achievement and Perceptions of Learning.” *PRIMUS* 28 (6): 600–615. <https://doi.org/10.1080/10511970.2017.1332701>.
- Adams, Deanne M, Richard E Mayer, Andrew MacNamara, Alan Koenig, and Richard Wainess. 2011. “Narrative Games for Learning: Testing the Discovery and Narrative Hypotheses.” *Journal of Educational Psychology*, 104 (1): 235–49.
- Aguillon, Stephanie M., Gregor-Fausto Siegmund, Renee H. Petipas, Abby Grace Drake, Sehoia Cotner, and Cissy J. Ballen. 2020. “Gender Differences in Student Participation in an Active-Learning Classroom.” *CBE—Life Sciences Education* 19 (2): 1–10. <https://doi.org/10.1187/cbe.19-03-0048>.
- Albanese, Mark, and Susan Mitchell. 1993. “Problem-Based Learning: A Review of Literature on Its Outcomes and Implementation Issues.” *Academic Medicine* 68 (1): 52–81.
- Alqasa, Khaled Mohammed Ahmed, and Jehad Abdallah Atieh Afaneh. 2022. “Active Learning Techniques and Student Satisfaction: Role of Classroom Environment.” *Eurasian Journal of Educational Research* 98: 85-100.
- American Council on Education. *U.S. Population Trends and Educational Attainment*. 2022. <https://www.equityinhighered.org/indicators/u-s-population-trends-and-educational-attainment/educational-attainment-by-race-and-ethnicity/>.
- Andrews, T. M., M. J. Leonard, C. A. Colgrove, and S. T. Kalinowski. “Active Learning Not Associated with Student Learning in a Random Sample of College Biology Courses.” *CBE—Life Sciences Education* 10, no. 4 (2011): 394–405. <https://doi.org/10.1187/cbe.11-07-0061>.
- Ballen, Cissy J., Carl Wieman, Shima Salehi, Jeremy B. Searle, and Kelly R. Zamudio. 2017. “Enhancing Diversity in Undergraduate Science: Self-Efficacy Drives Performance Gains with Active Learning.” *CBE—Life Sciences Education* 16 (4): ar56. <https://doi.org/10.1187/cbe.16-12-0344>.
- Baroody, Arthur, Lai Meng-lung, and Kelly Mix. “The Development of Young Children’s Early Number and Operation Sense and Its Implications for Early Childhood Education.” In *Handbook of Research on the Education of Young Children*. Routledge, 2014.
- Baum, Sandra R., and Michael S. McPherson. *Can College Level the Playing Field? Higher Education in an Unequal Society*. Princeton University Press, 2022.
- Baum, Sandy, Jennifer Ma, and Kathleen Payea. *Education Pays 2013: The Benefits of Higher Education for Individuals and Society*. Trends in Higher Education. The College Board, 2013.
- Beck, Jacob. “Can Bootstrapping Explain Concept Learning?” *Cognition* 158 (January 2017): 110–21. <https://doi.org/10.1016/j.cognition.2016.10.017>.
- Berkson, Laeora. 1993. “Problem-Based Learning: Have the Expectations Been Met?” *Academic Medicine* 68 (10): S79-88.
- Bernstein, Douglas A. “Does Active Learning Work? A Good Question, but Not the Right One.” *Scholarship of Teaching and Learning in Psychology* 4, no. 4 (2018): 290–307. <https://doi.org/10.1037/stl0000124>.
- Berry Jr., Lemuel. 1991. *Collaborative Learning: A Program for Improving the Retention of Minority Students*.

- Bloch-Schulman, Stephen. "Making the Case for the Value of Teaching through Self-Authorship: From 'What Does the Discipline Want of Me?' To 'Who Do I Want to Be?'" In *Innovations in Teaching Philosophy: A Toolkit for the 21st-Century Classroom*. 2026.
- Bonwell, Charles C, and James A Eison. 1991. *Active Learning: Creating Excitement in the Classroom*. ERIC Clearinghouse on Higher Education. <http://files.eric.ed.gov/fulltext/ED336049.pdf>.
- Borrego, Maura, Stephanie Cutler, Michael Prince, Charles Henderson, and Jeffrey E. Froyd. 2013. "Fidelity of Implementation of Research-Based Instructional Strategies (RBIS) in Engineering Science Courses." *Journal of Engineering Education* 102 (3): 394–425. <https://doi.org/10.1002/jee.20020>.
- Brigandt, Ingo. "The Epistemic Goal of a Concept: Accounting for the Rationality of Semantic Change and Variation." *Synthese* 177, no. 1 (2010): 19–40. <https://doi.org/10.1007/s11229-009-9623-8>.
- Brownell, Jayne, and Lynn Swaner. "Five High-Impact Practices: Research on Learning Outcomes, Completion, and Quality." *Peer Review* 14, no. 3 (2012).
- Butler, Adam, Kayah-Bah Phillmann, and Lona Smart. 2001. "Active Learning within a Lecture: Assessing the Impact of Short, In-Class Writing Exercises." *Teaching of Psychology* 28 (4): 257–59. https://doi.org/10.1207/S15328023TOP2804_04.
- Cahalan, Margaret, Nicole Brunt, Terry Vaughan III, et al. *Indicators of Higher Education Equity in the United States 2024: 50-Year Historical Trend Report*. The Pell Institute for the Study of Opportunity in Higher Education, Council for Opportunity in Education (COE) and Alliance for Higher Education and Democracy of the University of Pennsylvania (Penn AHEAD)., 2024. <https://www.pellinstitute.org/wp-content/uploads/2024/05/05.01.24ExecutiveSummary-MCNBHHMC-Indicators-2024.pdf>.
- Cant, Robyn P., and Simon J. Cooper. 2010. "Simulation-based Learning in Nurse Education: Systematic Review." *Journal of Advanced Nursing* 66 (1): 3–15. <https://doi.org/10.1111/j.1365-2648.2009.05240.x>.
- Carballo, Alejandro Pérez. "Conceptual Evaluation: Epistemic." In *Conceptual Engineering and Conceptual Ethics*, 1st ed., edited by Alexis Burgess, Herman Cappelen, and David Plunkett. Oxford University Press/Oxford, 2020. <https://doi.org/10.1093/oso/9780198801856.003.0015>.
- Carey, Susan. "Bootstrapping & the Origin of Concepts." *Daedalus* 133, no. 1 (2004): 59–68. <https://doi.org/10.1162/001152604772746701>.
- Carey, Susan. "Concept Innateness, Concept Continuity, and Bootstrapping." *The Behavioral and Brain Sciences* 34, no. 3 (2011): 152–62. <https://doi.org/10.1017/s0140525x10003092>.
- Carey, Susan. "On Learning New Primitives in the Language of Thought: Reply to Rey." *Mind & Language* 29, no. 2 (2014): 133–66. <https://doi.org/10.1111/mila.12045>.
- Carey, Susan. "Précis of *The Origin of Concepts*." *Behavioral and Brain Sciences* 34, no. 3 (2011): 113–24. <https://doi.org/10.1017/S0140525X10000919>.
- Carnevale, Anthony P., and Jeff Strohl. *Separate & Unequal: How Higher Education Reinforces the Intergenerational Reproduction of White Racial Privilege*. Georgetown Public Policy Institute, 2013.
- Chi, Michelene, and Ruth Wylie. 2014. "The ICAP Framework: Linking Cognitive Engagement to Active Learning Outcomes." *Educational Psychologist* 49 (4): 219–43. <https://doi.org/10.1080/00461520.2014.965823>.
- Clements, Douglas, and Julie Sarama. "Effects of a Preschool Mathematics Curriculum: Summative Research on the Building Blocks Project." Pt. 1136-1163. *Journal for Research in Mathematics Education* 38, no. 2 (n.d.).

- Collins, John William, and Nancy P. O'Brien, eds. 2011. *The Greenwood Dictionary of Education*. 2nd ed. Greenwood.
- Colliver, Jerry A. 2000. "Effectiveness of Problem-Based Learning Curricula: Research and Theory." *Academic Medicine* 75 (3): 259–66.
- Concepción, David W., Melinda Messineo, Sarah Wieten, Catherine Homan, and Philosophy Documentation Center. "The State of Teacher Training in Philosophy." *Teaching Philosophy* 39, no. 1 (2016): 1–24. <https://doi.org/10.5840/teachphil201621942>.
- Copeland, David E., Jenna R. Scott, and Jeremy Ashton Houska. 2010. "Computer-Based Demonstrations in Cognitive Psychology: Benefits and Costs." *Teaching of Psychology* 37 (2): 141–45. <https://doi.org/10.1080/00986281003626680>.
- Dall'Alba, Gloria, and Søren Bengtsen. 2019. "Re-Imagining Active Learning: Delving into Darkness." *Educational Philosophy and Theory* 51: 1477–1489.
- Daubemire, Patrick L., Diane M. Bunce, Carolyn Draus, Meredith Frazier, Austin Gessell, and Mary T. van Opstal. 2015. "During POGIL Implementation the Professor Still Makes a Difference." *Journal of College Science Teaching* 44 (5): 72–81.
- Deslauriers, Louis, Ellen Schelew, and Carl Wieman. 2011. "Improved Learning in a Large-Enrollment Physics Class." *Science* 332 (6031): 862–64. <https://doi.org/10.1126/science.1201783>.
- Di Vesta, Francis J., and Deborah A. Smith. 1979. "The Pausing Principle: Increasing the Efficiency of Memory for Ongoing Events." *Contemporary Educational Psychology* 4 (3): 288–96. [https://doi.org/10.1016/0361-476X\(79\)90048-1](https://doi.org/10.1016/0361-476X(79)90048-1).
- Dori, Yehudit Judy, and John Belcher. 2005. "How Does Technology-Enabled Active Learning Affect Undergraduate Students' Understanding of Electromagnetism Concepts?" *Journal of the Learning Sciences* 14 (2): 243–79. https://doi.org/10.1207/s15327809jls1402_3.
- Driessen, Emily P., Jennifer K. Knight, Michelle K. Smith, and Cissy J. Ballen. "Demystifying the Meaning of Active Learning in Postsecondary Biology Education." *CBE—Life Sciences Education* 19, no. 4 (2020): 1–9. <https://doi.org/10.1187/cbe.20-04-0068>.
- Ebert-May, Diane, Carol Brewer, and Sylvester Allred. 1997. "Innovation in Large Lectures: Teaching for Active Learning." *BioScience* 47 (9): 601–7. <https://doi.org/10.2307/1313166>.
- Eddy, Sarah L., and Kelly A. Hogan. 2014. "Getting Under the Hood: How and for Whom Does Increasing Course Structure Work?" *CBE—Life Sciences Education* 13 (3): 453–68. <https://doi.org/10.1187/cbe.14-03-0050>.
- Elliott, Lisa Jo, Stephen Rice, David Trafimow, Laura Madson, and Malisa F. Hipshur. 2010. "Research Participation versus Classroom Lecture: A Comparison of Student Learning." *Teaching of Psychology* 37 (2): 129–31. <https://doi.org/10.1080/00986281003626862>.
- Emlen, Stephen. "Migratory orientation in the Indigo Bunting, *Passerina Cyanea*. Part II: Mechanism of celestial orientation." *The Auk* 84, no. 4 (1967): 463–489.
- Emlen, Stephen. "Migratory orientation in the Indigo Bunting, *Passerina Cyanea*: Part I: Evidence for use of celestial cues," *The Auk* 84, no. 3 (1967). 309–342.
- Faust, Jennifer L, and Donald R Paulson. 1998. "Active Learning in the College Classroom." *Journal on Excellence in College Teaching* 9: 3–24.

- Felder, RM, DR Woods, JE Stice, and A Rugarcia. 2000. "Teaching Methods That Work." *Chemical Engineering Education* 34 (1): 26–39.
- Fitzgerald, Debbie. 2013. "Employing Think–Pair–Share in Associate Degree Nursing Curriculum." *Teaching and Learning in Nursing* 8 (3): 88–90. <https://doi.org/10.1016/j.teln.2013.01.006>.
- Foot, Philippa. *Moral Dilemmas: And Other Topics in Moral Philosophy*. Oxford University Press, 2002. <https://doi.org/10.1093/019925284X.001.0001>.
- Fosmire, Michael, and Alexius Macklin. "Riding the Active Learning Wave: Using Problem-Based Learning as a Catalyst for Creating Faculty-Librarian Partnerships." *Libraries Research Publications*, 2002.
- Fredericksen, Elaine. 1998. "Minority Students and the Learning Community Experience: A Cluster Experiment." *Annual Meeting of the Conference on College Composition and Communication*.
- Freeman, Scott, Sarah L. Eddy, Miles McDonough, et al. "Active Learning Increases Student Performance in Science, Engineering, and Mathematics." *Proceedings of the National Academy of Sciences* 111, no. 23 (2014): 8410–15. <https://doi.org/10.1073/pnas.1319030111>.
- Gallagher, Shelagh A. 1997. "Problem-Based Learning: Where Did It Come from, What Does It Do, and Where Is It Going?" *Journal for the Education of the Gifted* 20 (4): 332–62. <https://doi.org/10.1177/016235329702000402>.
- Gilligan, Carol. *In a Different Voice: Psychological Theory and Women's Development*. Harvard University Press, 2003.
- Ginsburg, Herbert P., Joanna Cannon, Janet Eisenband, and Sandra Pappas. "Mathematical Thinking and Learning." In *Blackwell Handbook of Early Childhood Development*, edited by Kathleen McCartney and Deborah Phillips. Blackwell Publishing Ltd, 2006. <https://doi.org/10.1002/9780470757703.ch11>.
- Gladieux, Lawrence, and Laura Perna. *Borrowers Who Drop Out: A Neglected Aspect of the College Student Loan Trend*. Nos. 05–2. The National Center for Public Policy and Higher Education, 2005.
- Goffe, William L., and David Kauper. 2014. "A Survey of Principles Instructors: Why Lecture Prevails." *The Journal of Economic Education* 45 (4): 360–75. <https://doi.org/10.1080/00220485.2014.946547>.
- Haak, David C., Janneke HilleRisLambers, Emile Pitre, and Scott Freeman. 2011. "Increased Structure and Active Learning Reduce the Achievement Gap in Introductory Biology." *Science* 332 (6034): 1213–16. <https://doi.org/10.1126/science.1204820>.
- Hake, Richard R. 1998. "Interactive-Engagement versus Traditional Methods: A Six-Thousand-Student Survey of Mechanics Test Data for Introductory Physics Courses." *American Journal of Physics* 66 (1): 64–74. <https://doi.org/10.1119/1.18809>.
- Hartling, Lisa, Carol Spooner, Lisa Tjosvold, and Anna Oswald. 2010. "Problem-Based Learning in Pre-Clinical Medical Education: 22 Years of Outcome Research." *Medical Teacher* 32 (1): 28–35. <https://doi.org/10.3109/01421590903200789>.
- Haslanger, Sally. "Gender and Race: (What) Are They? (What) Do We Want Them To Be?" *Noûs* 34, no. 1 (2000): 31–55. <https://doi.org/10.1111/0029-4624.00201>.
- Henderson, Charles, and Melissa H. Dancy. 2009. "Impact of Physics Education Research on the Teaching of Introductory Quantitative Physics in the United States." *Physical Review Special Topics - Physics Education Research* 5 (2): 020107. <https://doi.org/10.1103/PhysRevSTPER.5.020107>.
- Hernández-de-Menéndez, Marcela, Antonio Vallejo Guevara, Juan Carlos Tudón Martínez, Diana Hernández Alcántara, and Ruben Morales-Menendez. 2019. "Active Learning in Engineering Education. A Review

- of Fundamentals, Best Practices and Experiences.” *International Journal on Interactive Design and Manufacturing (IJIDeM)* 13 (3): 909–22. <https://doi.org/10.1007/s12008-019-00557-8>.
- Hewapathirana, Gertrude Iranganie, and Firas Almasri. 2022. “Active Learning Compared With Lecture-Based Pedagogies in Gender and Socio-Cultural Context-Specific Major and Non-Major Biology Classes:” In *Advances in Higher Education and Professional Development*, edited by Jared Keengwe. IGI Global. <https://doi.org/10.4018/978-1-7998-9564-0.ch014>.
- Hillgoss, Tonya. 1992. “Demystifying ‘Classroom Chemistry’: The Role of the Interactive Learning Model.” *Teaching Sociology* 20 (1): 12. <https://doi.org/10.2307/1318542>.
- Hood, S., N. Barrickman, N. Djerdjian, et al. 2021. “‘I Like and Prefer to Work Alone’: Social Anxiety, Academic Self-Efficacy, and Students’ Perceptions of Active Learning.” *CBE—Life Sciences Education* 20 (1): 1–12. <https://doi.org/10.1187/cbe.19-12-0271>.
- Huggett, Kathryn N., and William B. Jeffries. 2021. “Overview of Active Learning Research and Rationale for Active Learning.” *IAMSE Manuals*, 1–7. https://doi.org/10.1007/978-3-030-62916-8_1.
- Hult, Jr., Richard E. “On Pedagogical Caring.” *Educational Theory* 29, no. 3 (1979): 237–43. <https://doi.org/10.1111/j.1741-5446.1979.tb00856.x>.
- Johnson, David W., Roger T. Johnson, and Karl A. Smith. 1998. “Cooperative Learning Returns To College What Evidence Is There That It Works?” *Change: The Magazine of Higher Learning* 30 (4): 26–35. <https://doi.org/10.1080/00091389809602629>.
- Kapur, Manu, John Hattie, Irina Grossman, and Tanmay Sinha. 2022. “Fail, Flip, Fix, and Feed – Rethinking Flipped Learning: A Review of Meta-Analyses and a Subsequent Meta-Analysis.” *Frontiers in Education* 7: 1–19. <https://doi.org/10.3389/educ.2022.956416>.
- King, Alison. “From Sage on the Stage to Guide on the Side.” *College Teaching* 41, no. 1 (1993): 30–35.
- Kirschner, Paul A., John Sweller, and Richard E. Clark. 2006. “Why Minimal Guidance During Instruction Does Not Work: An Analysis of the Failure of Constructivist, Discovery, Problem-Based, Experiential, and Inquiry-Based Teaching.” *Educational Psychologist* 41 (2): 75–86. https://doi.org/10.1207/s15326985ep4102_1.
- Knight, Jennifer K, and William B Wood. 2005. “Teaching More by Lecturing Less.” *Cell Biology Education* 4 (4): 298–310. <https://doi.org/10.1187/05-06-0082>.
- Koch, Andrew, and Brent Drake. “Digging into the Disciplines I: Accounting for Failure—the Impact of Principles of Accounting Courses on Student Success and Equitable Outcomes.” *John N. Gardner Institute for Excellence in Undergraduate Education*, 2018.
- Koch, Andrew. *Transforming the Gateway Course Experience: A Call to Action for Higher Education*. 1st ed. Routledge, 2024. <https://doi.org/10.4324/9781003444961>.
- Kovarik, Michelle L., Jill K. Robinson, and Thomas J. Wenzel. 2022. “Why Use Active Learning?” In *ACS Symposium Series*, edited by Thomas J. Wenzel, Michelle L. Kovarik, and Jill K. Robinson, vol. 1409. American Chemical Society. <https://doi.org/10.1021/bk-2022-1409.ch001>.
- Kuh, George. “High-Impact Educational Practices: What They Are, Who Has Access to Them, and Why They Matter.” *Peer Review* 14, no. 3 (2012): 3–18.
- Lakoff, George, and Mark Johnson. *Metaphors We Live By*. University of Chicago Press, 1980.
- Lakoff, George, and Rafael E. Núñez. *Where Mathematics Comes from: How the Embodied Mind Brings Mathematics into Being*. 1st ed. Basic Books, 2000.

- Laws, Priscilla W, Pamela Rosborough, and Frances Poodry. 1999. "Women's Responses to an Activity-Based Introductory Physics Program..." *American Journal of Physics* 67 (7): S32–37.
- Laws, Priscilla, David Sokoloff, and Ronald Thornton. 1999. "Promoting Active Learning Using the Results of Physics Education Research." *UniServe Science News* 13: 14–19.
- Lombardi, Doug, Thomas F. Shipley, Astronomy Team, Biology Team, Chemistry Team, Engineering Team, Geography Team, Geoscience Team, and Physics Team, et al. "The Curious Construct of Active Learning." *Psychological Science in the Public Interest* 22, no. 1 (2021): 8–43.
<https://doi.org/10.1177/1529100620973974>.
- Lorenzo, Mercedes, Catherine H. Crouch, and Eric Mazur. 2006. "Reducing the Gender Gap in the Physics Classroom." *American Journal of Physics* 74 (2): 118–22. <https://doi.org/10.1119/1.2162549>.
- Louis, W, and J Brown. 2017. "Reports of a Failed 'flipping' Experiment: Lower Course Satisfaction, and under-Performance among Low Achievers." *Vancouver International Conference on the Teaching of Psychology*.
- Maries, Alexandru, Nafis I. Karim, and Chandralekha Singh. 2020. "Active Learning in an Inequitable Learning Environment Can Increase the Gender Performance Gap: The Negative Impact of Stereotype Threat." *The Physics Teacher* 58 (6): 430–33. <https://doi.org/10.1119/10.0001844>.
- Matsushita, Kayo, ed. 2018. *Deep Active Learning*. Springer Singapore. <https://doi.org/10.1007/978-981-10-5660-4>.
- Mayer, Richard E. 2004. "Should There Be a Three-Strikes Rule Against Pure Discovery Learning?" *American Psychologist* 59 (1): 14–19. <https://doi.org/10.1037/0003-066X.59.1.14>.
- McCarthy, J. Patrick, and Liam Anderson. 1999. "Active Learning Techniques Versus Traditional Teaching Styles: Two Experiments from History and Political Science." *Innovative Higher Education* 24 (4): 279–94. <https://doi.org/10.1023/B:IHIE.0000047415.48495.05>.
- McConnell, David A., LeeAnna Chapman, C. Douglas Czajka, Jason P. Jones, Katherine D. Ryker, and Jennifer Wiggen. 2017. "Instructional Utility and Learning Efficacy of Common Active Learning Strategies." *Journal of Geoscience Education* 65 (4): 604–25. <https://doi.org/10.5408/17-249.1>.
- McKeachie, Wilbert. 1972. "Research on College Teaching." *Educational Perspectives* 11 (2): 3–10.
- Michael, Joel. 2006. "Where's the Evidence That Active Learning Works?" *Advances in Physiological Learning* 30 (4): 159–67. <https://doi.org/10.1152/advan.00053.2006>.
- Michel, Norbert, John James Cater III, and Otmar Varela. 2009. "Active versus Passive Teaching Styles: An Empirical Study of Student Learning Outcomes." *Human Resource Development Quarterly* 20 (4): 397–418. <https://doi.org/10.1002/hrdq.20025>.
- Noddings, Nel. *Caring: A Feminine Approach to Ethics & Moral Education*. 2nd ed. University of California Press, 2003.
- Nurbavliyev, Omarbek, Serkan Kaymak, and Bakhyt Sydykov. 2022. "The Effect of Active Learning Method on Students' Academic Success, Motivation and Attitude towards Mathematics." *Journal of Language and Linguistic Studies* 18. <http://www.jlls.org/index.php/jlls/article/view/4409>.
- Pelley, John. 2014. "Making Active Learning Effective." *Medical Science Educator* 24 (S1): 13–18. <https://doi.org/10.1007/s40670-014-0087-1>.
- Prince, Michael. 2004. "Does Active Learning Work? A Review of the Research." *Journal of Engineering Education* 93 (3): 223–31. <https://doi.org/10.1002/j.2168-9830.2004.tb00809.x>.

- R. Drake, John. "A Critical Analysis of Active Learning and an Alternative Pedagogical Framework for Introductory Information Systems Courses." *Journal of Information Technology Education: Innovations in Practice* 11 (2012): 039–052. <https://doi.org/10.28945/1546>.
- Rawls, John. *A Theory of Justice*. 1971. Original ed. Belknap Press, 2005.
- Redish, Edward F., Jeffery M. Saul, and Richard N. Steinberg. 1997. "On the Effectiveness of Active-Engagement Microcomputer-Based Laboratories." *American Journal of Physics* 65 (1): 45–54. <https://doi.org/10.1119/1.18498>.
- Rey, Georges. "Concepts and Conceptions: A Reply to Smith, Medin and Rips." *Cognition* 19, no. 3 (1985): 297–303. [https://doi.org/10.1016/0010-0277\(85\)90037-X](https://doi.org/10.1016/0010-0277(85)90037-X).
- Rey, Georges. "Innate and Learned: Carey, Mad Dog Nativism, and the Poverty of Stimuli and Analogies (Yet Again)." *Mind & Language* 29, no. 2 (2014): 109–32. <https://doi.org/10.1111/mila.12044>.
- Richmond, Aaron S., and Lisa Kindelberger Hagan. 2011. "Promoting Higher Level Thinking in Psychology: Is Active Learning the Answer?" *Teaching of Psychology* 38 (2): 102–5. <https://doi.org/10.1177/0098628311401581>.
- Rondon, Silmara, Fernanda Chiarion Sassi, and Claudia Regina Furquim De Andrade. 2013. "Computer Game-Based and Traditional Learning Method: A Comparison Regarding Students' Knowledge Retention." *BMC Medical Education* 13 (1): 1–8. <https://doi.org/10.1186/1472-6920-13-30>.
- Royall, Richard. 1997. *Statistical Evidence: A Likelihood Paradigm*. 1st ed. Routledge. <https://doi.org/10.1201/9780203738665>.
- Ruhl, Kathy L., Charles A. Hughes, and Patrick J. Schloss. 1987. "Using the Pause Procedure to Enhance Lecture Recall." *Teacher Education and Special Education: The Journal of the Teacher Education Division of the Council for Exceptional Children* 10 (1): 14–18. <https://doi.org/10.1177/088840648701000103>.
- Ruiz-Primo, Maria Araceli, Derek Briggs, Heidi Iverson, Robert Talbot, and Lorrie A. Shepard. 2011. "Impact of Undergraduate Science Course Innovations on Learning." *Science* 331 (6022): 1269–70. <https://doi.org/10.1126/science.1198976>.
- Schouten, Gina. "Gender and the Division of Labor." In *Routledge Handbook of Philosophy, Politics, and Economic*. Routledge, 2022.
- Schouten, Gina. "Restricting Justice: Political Interventions in the Home and in the Market." *Philosophy & Public Affairs* 41, no. 4 (2013): 357–88.
- Shapiro, Lawrence. *Embodied Cognition*. Routledge, 2019.
- Shapiro, Lawrence. "The embodied cognition research programme." *Philosophy Compass* 2, no. 2 (2007): 338–346.
- Shi, Yinghui, Huiyun Yang, Jason MacLeod, Jingman Zhang, and Harrison Hao Yang. 2019. "College Students' Cognitive Learning Outcomes in Technology-Enabled Active Learning Environments: A Meta-Analysis of the Empirical Literature." *Journal of Educational Computing Research* 58 (4): 791–817. <https://doi.org/10.1177/0735633119881477>.
- Shuster, Michele I., and Ralph Preszler. 2014. "Introductory Biology Course Reform: A Tale of Two Courses." *International Journal for the Scholarship of Teaching and Learning* 8 (2): 1–30. <https://doi.org/10.20429/ijstl.2014.080205>.
- Smart, J. J. C., and Bernard Williams. *Utilitarianism For and Against*. Cambridge University Press, 2008.

- Smith, Carol, and Chris Unger. "What's in Dots-Per-Box? Conceptual Bootstrapping With Stripped-Down Visual Analogs." *Journal of the Learning Sciences* 6, no. 2 (1997): 143–81. https://doi.org/10.1207/s15327809jls0602_1.
- Smith, Carol, Joseph Snir, and Lorraine Grosslight. "Using Conceptual Models to Facilitate Conceptual Change: The Case of Weight-Density Differentiation." *Cognition and Instruction* 9, no. 3 (1992): 221–83. https://doi.org/10.1207/s1532690xci0903_3.
- Sober, Elliott. "Testability." *Proceedings and Addresses of the American Philosophical Association* 73, no. 2 (1999): 46–47.
- Sokoloff, David R, Priscilla W Laws, and Ronald K Thornton. 2007. "RealTime Physics : Active Learning Labs Transforming the Introductory Laboratory." *European Journal of Physics* 28 (3): S83–94. <https://doi.org/10.1088/0143-0807/28/3/S08>.
- Sorcinelli, Mary Deane. 1991. "Research Findings on the Seven Principles." *New Directions for Teaching and Learning* 1991 (47): 13–25. <https://doi.org/10.1002/tl.37219914704>.
- Springer, Leonard. n.d. *Effects of Small-Group Learning on Undergraduates in Science, Mathematics, Engineering, and Technology: A Meta-Analysis*.
- Stains, Marilyne, and Trisha Vickrey. 2017. "Fidelity of Implementation: An Overlooked Yet Critical Construct to Establish Effectiveness of Evidence-Based Instructional Practices." *CBE—Life Sciences Education* 16 (1): rm1. <https://doi.org/10.1187/cbe.16-03-0113>.
- Stead, David. 2005. "A Review of the One-Minute Paper." *Active Learning in Higher Education the Journal of the Institute for Learning and Teaching* 6: 118-131.
- The Feed*. "Rethinking Gateway Courses to Improve Outcomes, Equity." January 24, 2020.
- Theobald, Elli J., Mariah J. Hill, Elisa Tran, et al. "Active Learning Narrows Achievement Gaps for Underrepresented Students in Undergraduate Science, Technology, Engineering, and Math." *Proceedings of the National Academy of Sciences* 117, no. 12 (2020): 6476–83. <https://doi.org/10.1073/pnas.1916903117>.
- Turpen, Chandra, and Noah D. Finkelstein. 2009. "Not All Interactive Engagement Is the Same: Variations in Physics Professors' Implementation of Peer Instruction." *Physical Review Special Topics - Physics Education Research* 5 (2): 020101-1–18. <https://doi.org/10.1103/PhysRevSTPER.5.020101>.
- Venton, B. Jill, and Rebecca R. Pompano. 2021. "Strategies for Enhancing Remote Student Engagement through Active Learning." *Analytical and Bioanalytical Chemistry* 413 (6): 1507–12. <https://doi.org/10.1007/s00216-021-03159-0>.
- Vygotsky, Lev. *Thought and Language*. Edited by Alex Kozulin. MIT Press, 1986.
- Walton, Gregory M., and Geoffrey L. Cohen. "A Question of Belonging: Race, Social Fit, and Achievement." *Journal of Personality and Social Psychology* 92, no. 1 (2007): 82–96. <https://doi.org/10.1037/0022-3514.92.1.82>.
- Wang, Xueli, Ning Sun, Seo Young Lee, and Brit Wagner. 2017. "Does Active Learning Contribute to Transfer Intent Among 2-Year College Students Beginning in STEM?" *The Journal of Higher Education* 88 (4): 593–618. <https://doi.org/10.1080/00221546.2016.1272090>.
- Weltman, David, and Mary Whiteside. 2010. "Comparing the Effectiveness of Traditional and Active Learning Methods in Business Statistics: Convergence to the Mean." *Journal of Statistics Education* 18 (1): 1–10. <https://doi.org/10.1080/10691898.2010.11889480>.

- Wieman, Carl E. 2014. "Large-Scale Comparison of Science Teaching Methods Sends Clear Message." *Proceedings of the National Academy of Sciences* 111 (23): 8319–20.
<https://doi.org/10.1073/pnas.1407304111>.
- Wieman, Carl. 2017. *Improving How Universities Teach Science: Lessons from the Science Education Initiative*. Harvard University Press.
- Wilke, R. Russell, and William J. Straits. "The Effects of Discovery Learning in a Lower-Division Biology Course." *Advances in Physiology Education* 25, no. 2 (2001): 62–69.
<https://doi.org/10.1152/advances.2001.25.2.62>.