

ESSAYS IN RISK TAKING, BELIEF FORMATION, AND SELF-DECEPTION

by

NATHAN R. ADAMS

A DISSERTATION

Presented to the Department of Economics
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

June 2018

DISSERTATION APPROVAL PAGE

Student: Nathan R. Adams

Title: Essays in Risk Taking, Belief Formation, and Self-Deception

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Economics by:

| | |
|------------------|------------------------------|
| Glen Waddell | Chair |
| Michael Kuhn | Core Member |
| William Harbaugh | Core Member |
| William Terry | Institutional Representative |

and

| | |
|----------------|---|
| Sara D. Hodges | Interim Vice Provost and Dean of the Graduate School |
|----------------|---|

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded June 2018

© 2018 Nathan R. Adams
All rights reserved.

DISSERTATION ABSTRACT

Nathan R. Adams

Doctor of Philosophy

Department of Economics

June 2018

Title: Essays in Risk Taking, Belief Formation, and Self-Deception

In this dissertation, I examine changes in risk-taking behavior, beliefs, and self-deception induced by changes in policy and behavior. Specifically, Chapter II examines player performance and risk-taking behavior in tournament environments which include eliminations in the middle of the tournament. I find that when players face elimination, they perform better and take risks more often. In addition, when facing elimination, players are more likely to have those risks pay off. Turning to the interaction between public policy and personal beliefs, Chapter III explores how public policy affects beliefs in the context of same-sex marriage. Exploiting the timing of the legalization of same-sex marriage, I find that legalization induces an increase in the proportion of people who have strong beliefs on same-sex marriage. I also find a substantial increase in measured state-level polarization due to legalization. Finally, Chapter IV presents the results of an experiment designed to uncover how self-confidence and self-deception change after performing dishonest behavior. In an online experimental laboratory, participants who cheated have higher confidence in their ability even when the opportunity to cheat is not present. In addition, participants who cheated, and were rewarded for cheating with a high reward, had higher beliefs in their ability.

This dissertation includes unpublished co-authored material.

CURRICULUM VITAE

NAME OF AUTHOR: Nathan R. Adams

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene, Oregon
Brigham Young University, Provo, Utah

DEGREES AWARDED:

Doctor of Philosophy, Economics, 2018, University of Oregon
Master of Science, Economics, 2014, University of Oregon
Bachelor of Arts, Economics, 2013, Brigham Young University

AREAS OF SPECIAL INTEREST:

Applied Microeconometrics
Labor Economics
Behavioral and Experimental Economics

PROFESSIONAL EXPERIENCE:

Graduate Employee, University of Oregon, 2013-2018
Adjunct Professor, Brigham Young University-Idaho, 2016-2018

GRANTS, AWARDS AND HONORS:

Graduate Teaching Fellowship, University of Oregon, 2013-2018
Teaching Award, University of Oregon, Department of Economics, 2017
Best Econometrics Performance Award, University of Oregon, 2014
Kleinsorge Fellowship, University of Oregon, 2013
Magna cum Laude, Brigham Young University, 2013
Heritage Scholarship, Brigham Young University, 2010-2013

ACKNOWLEDGEMENTS

I have benefited greatly from the helpful advice and guidance that I have received from the members of my committee. I am especially grateful to Glen Waddell and Michael Kuhn for mentoring me and for helping to improve this dissertation beyond what I could have done without them. I also thank seminar participants at the University of Oregon and at the Western Economic Association Conference.

I also wish to thank my family, peers, colleagues, and friends for their support. Most importantly, I am grateful to my wife, Kacy, for her constant support and encouragement.

TABLE OF CONTENTS

| Chapter | Page |
|--|------|
| I. INTRODUCTION | 1 |
| II. PERFORMANCE AND RISK TAKING UNDER THREAT OF ELIMINATION | 4 |
| Introduction | 4 |
| Other related literature | 6 |
| Background and data | 9 |
| Identification | 12 |
| Results | 18 |
| Baseline Performance Results | 18 |
| Baseline Risk-Taking Results | 21 |
| Sensitivity, Heterogeneity, and Bandwidth Considerations | 24 |
| Additional Heterogeneity | 32 |
| Conclusion | 41 |
| III. THE MALLEABILITY OF BELIEF: VARYING RESPONSES TO “SHOULD SAME-SEX COUPLES HAVE THE RIGHT TO MARRY?” | 42 |
| Introduction | 42 |
| Background and Data | 44 |
| General Social Survey | 44 |
| Policy Variation and Identification | 46 |
| Related Literature | 49 |

| Chapter | Page |
|--|------|
| Results | 51 |
| Support for Same-Sex Marriage | 51 |
| Legalization and Strength of Belief | 52 |
| Heterogeneous Responses to Legalization | 56 |
| Failed Attempts to Legalize Same-Sex Marriage | 62 |
| Direct Measures of Polarization | 67 |
| Longevity of Changes in Beliefs and Polarization | 69 |
| Conclusion | 71 |
| | |
| IV. CONFIDENCE AND CONTRITION: IS CHEATING INTERNALIZED IN PERFORMANCE ASSESSMENTS? | 73 |
| Introduction | 73 |
| Experiment 1 | 78 |
| Design | 78 |
| Results | 80 |
| Experiment 2 | 94 |
| Design | 95 |
| Results | 96 |
| Discussion | 99 |
| | |
| V. CONCLUSION | 102 |
| | |
| APPENDIX: SUPPLEMENTAL FIGURES AND TABLES | 104 |

| Chapter | Page |
|----------------------------|------|
| REFERENCES CITED | 108 |

LIST OF FIGURES

| Figure | Page |
|--|------|
| 1. How many players make the “70 plus ties” cut | 10 |
| 2. Deviations from tournament cut (<i>Rank</i>) | 16 |
| 3. How predictable is the elimination threshold? | 17 |
| 4. Does responsiveness change as elimination approaches? | 26 |
| 5. Mean rank of players | 28 |
| 6. The estimated discontinuity by number of threshold crossings | 31 |
| 7. Heterogeneity by hour of play | 34 |
| 8. Responses to “Homosexual couples should have the right to marry one another,” 2004–2016 | 46 |
| 9. Timing and method of same-sex-marriage legalization, by state | 48 |
| 10. Changes in belief regarding same-sex marriage induced by same- sex-marriage legalization | 55 |
| 11. Changes in belief regarding same marriage induced by same-sex- marriage legalization, by generation | 59 |
| 12. Changes in belief regarding same marriage induced by same-sex- marriage legalization, by education | 60 |
| 13. Changes in belief regarding same-sex marriage induced by same- sex-marriage legalization, by race | 61 |
| 14. Changes in belief regarding same-sex marriage induced by same- sex-marriage legalization, by average state belief in 2004 | 63 |
| 15. Changes in belief regarding same-sex marriage induced by same- sex-marriage legalization, by method of legalization | 64 |
| 16. Impact of failed attempts to legalize same-sex marriage | 66 |
| 17. Permanence of changes in belief induced by the legalization of same-sex marriage | 70 |

| Figure | Page |
|---|------|
| 18. Permanence of changes in state-level polarization induced by same-sex marriage legalization | 71 |
| 19. Objects Reported by Gender, Experiment 1 | 84 |
| 20. Distribution of Investment by Cheating, Experiment 1 | 91 |
| 21. Cheating propensity, by incentive to cheat, Experiment 1 | 93 |
| A1. Support for same-sex marriage, 1988 & 2004–2016 | 104 |
| A2. Task 1 | 106 |
| A3. Task 2 | 107 |

LIST OF TABLES

| Table | Page |
|--|------|
| 1. The performance gains under pending elimination | 19 |
| 2. Risk taking under pending elimination | 22 |
| 3. Bandwidth sensitivity by mean rank of player | 29 |
| 4. Heterogeneity by average ability of the field of competitors | 35 |
| 5. Heterogeneity by player ability | 36 |
| 6. Heterogeneity by par of hole | 38 |
| 7. Heterogeneity by cut outcome of most-recent tournament | 40 |
| 8. Support for same-sex marriage around same-sex-marriage legalization, 2004–2016 | 53 |
| 9. Intensity of chosen positions and same-sex-marriage legalization | 57 |
| 10. Failed attempts to legalize same-sex marriage | 65 |
| 11. Polarization around same-sex-marriage legalization | 68 |
| 12. Determinants of Cheating | 82 |
| 13. Confidence and Propensity to Cheat | 86 |
| 14. Cheating and Stated Confidence | 87 |
| 15. Cheating and Revealed Confidence | 88 |
| 16. Testing Stated vs. Revealed Preference (Table 14 vs. Table 15) | 89 |
| 17. Impact of Cheating on Stated and Revealed Confidence | 94 |
| 18. Impact of Rewards from Cheating on Stated Confidence | 97 |
| 19. Impact of Rewards from Cheating on Revealed Confidence | 98 |
| 20. Testing Stated vs. Revealed Preference (Table 18 vs. Table 19) | 99 |

| Table | Page |
|---|------|
| A1. Strength of opinion on same-sex marriage around same-sex-marriage legalization, 2004-2016 | 105 |

CHAPTER I

INTRODUCTION

The study of labor economics is the study of decisions and policies that affect the workplace. In this dissertation, I add to current understanding by exploring three of these labor market issues: the changes in performance and risk-taking behavior induced by a specific workplace incentive structure, the impact of public policy innovation on normative beliefs, and the consequences of cheating in the workplace.

In Chapter II, we revisit the incentive effects of elimination tournaments with a fresh approach to identification. These tournaments are common in the workplace, but many of the effects of this incentive structure on decision making are yet unknown. We consider how player performance and risk-taking behavior in the PGA Tour changes under this structure—when elimination is pending. Previous studies of this topic have often suffered from an inability to understand what role risk-taking plays in any changes in performance. However, our identification allows us to separately measure changes in performance due to risk-taking and changes in performance due to changes in focus and effort. Our results strongly support that performance improves under the threat of elimination, and 23 percent of the improvement in performance is due to productive increases in risk taking. These effects are concentrated among those closest to the margin of elimination and among lower-ability competitors. These results help to inform managers about the effects of using elimination structures as an incentive technique. This chapter is co-authored with Glen Waddell.

Next, I explore how beliefs may be affected by policy itself in Chapter III. Often, unintended consequences accompany the implementation of policy and this chapter deals with one of these consequences. It is important to fully understand the full effects of policy as a proper weighing of the costs and benefits requires fully understanding what those costs and benefits may be. I exploit the timing of the legalization of same-sex marriage to identify the impact of legalization on support for same-sex marriage among a representative sample of Americans. I demonstrate that the strength with which beliefs are held increases with legalization, though legalization does not induce changes in belief between support and opposition. As support shifts to stronger support, and opposition shifts to stronger opposition, I also find that legalization induces larger differences between races and educational groups. I then estimate that legalization accounts for an increase in measured state-level polarization of roughly 65 percent of a standard deviation. These results also apply to the workplace, where policies are implemented frequently, often without any regard to the unintended consequences on future human behavior.

Finally, in Chapter IV, we investigate the consequences of cheating on confidence in the workplace. Self-assessed ability matters in the workplace, presumably motivating self-advancement—whether an employee asks for a raise or puts their name forward for a promotion, solicits external offers, and how they self-assess their performance. Rewards for past successes can serve as signals of ability, so long as the relationship between ability and those rewards is understood and accounted for. In this chapter, we examine whether rewards obtained by cheating at a task influence self-assessed ability at that task. We design an experiment that allows us to estimate the relationship between past cheating and both stated beliefs of ability and, in a version of the task without the potential for cheating, revealed

confidence. Our results are suggestive that cheaters have both higher stated beliefs in their ability and will reveal themselves to be more confident in their ability even when they cannot cheat again. We also experimentally vary the *ex-post* reward to cheating, and find that a larger reward from cheating may also cause an increase in stated beliefs, but it has no impact on revealed confidence. This chapter is co-authored with Glen Waddell and Michael Kuhn.

Chapter V offers concluding comments and further policy implications of the three substantive chapters.

CHAPTER II

PERFORMANCE AND RISK TAKING UNDER THREAT OF ELIMINATION

This chapter was co-authored with Glen Waddell.

Introduction

Although elimination tournaments and similarly discrete outcomes of competitive environments are quite common, opportunities to consider individual behavior under the threat of elimination are rare. Having coincident measures of risk taking makes this opportunity all the more rare. In this paper, we separately identify changes in risk-taking behavior and performance due to the threat of elimination.

Of course, elimination tournaments are a particular form of contract convexity, which we might generally expect to increase risk taking. Quantifying changes in behavior due to pending elimination therefore informs an understanding of contracts somewhat more-broadly than would be implied by a strict interpretation of elimination tournaments. The shape of stock-option contracts, for example, also exhibits strong convexities in their “up or out” implications. The higher is the exercise price on the option, the more likely it is that payment will only be realized when the upside is realized and, thus, the more appealing risk-taking becomes. In fact, the vast majority of stock options are granted with exercise prices equal to the grant-date stock price (Barron and Waddell, 2008), thus implying that the only realization of monetary return is conditional on the stock price increasing—very much mimicking the “up or out” nature of elimination tournaments. Moreover, if the stock price does not exceed the exercise price, it

does not matter at the margin by how much it falls short.¹ We will see empirical regularities consistent with this in the tournaments we consider, as our analysis also suggests that both performance and risk taking increase under threat of elimination.

We use hole-level Professional Golf Association (PGA) records of player performance, inclusive of objective measures of *ex ante* risk taking with *ex post* realizations, which enables hole-by-player-by-tournament-by-year analysis of performance on both sides of an objectively determined discontinuity in the expectation of elimination. With players repeatedly observed on either side of the threshold, we measure the systematic variation (within-player) in both performance and risk taking that is explained by that discontinuity.

With this fresh approach to identification, our results strongly support that performance improves under the threat of elimination and suggest a sizable role for risk taking as part of the mechanism. Where we can separately identify changes in performance and risk taking, our estimates suggest that at least 23 percent of the improvement in performance induced by potential elimination is due to productive increases in risk taking—productive in the sense that risks taken under threat of elimination are paying off with higher probability. In a world where all of the additional risk-taking opportunities (from among those the PGA flags as risk-taking opportunities) induced by the threat of elimination pay off, risk taking would account for up to 49 percent of the increase in performance. These effects are concentrated among those closest to the margin of elimination, among lower-ability competitors, and diminish as elimination approaches—we actually

¹Barron and Waddell (2008) interpret the implications of such convexity as inducing a sort of “work hard not smart” strategy. In their context, unabated, this may even leave agents prone to excessive risk taking, as though there is nothing to lose.

find performance declines in the last few opportunities for a player to escape elimination. Interestingly, risk taking seemingly plays no role in this decline.

We consider related literatures in Section II, followed by background information and data description in Section II. In Section II, we present our empirical strategy, which we follow with the main results and supplemental analysis in Section II. We offer concluding remarks in Section II.

Other related literature

A large empirical literature has developed since Lazear and Rosen (1981) first demonstrated the efficacy of tournaments in promoting effort in a second-best world.² In a collection of papers looking at the incentive effects in a tournament environment, Ehrenberg and Bognanno (1990a,9) and Orszag (1994) together find mixed evidence of player performance responding to monetary payoffs.³ Exploiting variation in the design of the National Basketball Association (NBA) player draft, Taylor and Trogdon (2002) offer strong evidence of declining *ex post* performance on the elimination side of tournaments, identifying that teams having just lost the chance of a playoff birth lose significantly more often than teams that are still at the margin of making it into the playoff tournament.⁴ In these ways, we anticipate that margins of elimination matter to performance.

As we use data on professional golf tournaments, we implicate several other pieces of literature. For example, Guryan, Kroft, and Notowidigdo (2009)

²See Prendergast (1999) for a summary of the early literature.

³In related work, Knoeber and Thurman (1994) compared a tournament scheme to a pay scheme that combines relative rankings with information about absolute productivity differences. They find that changes in prize levels that leave the prize spreads unchanged have no impact on performance in tournaments.

⁴In this context, it is argued that eliminated teams turn their attention to the pending player draft, which rewards lower performance.

exploits random pairings of golfers to identify potential peer effects, finding no such relationship. However, Brown (2011) does find that the performance of non-superstars declines in the presence of superstars—a “Tiger Woods effect.” Pope and Schweitzer (2011) also find that professional golfers exhibit loss aversion, putting less accurately when at the margin of achieving a below-par score on a hole.

A somewhat large literature analyzes risk taking, generally, and often implicates areas of finance and the behavior of “C-level” executives. There are large incentives for executives, for example, to take on risk in order to make up for poor past performance. Imas (2016) summarizes the literature on risk taking after a loss and in a lab experiment finds that the effect of loss on risk taking depends on the timing of the realization of the loss. Participants who face the loss immediately after it occurs take on less risk than those who do not face the loss until the end of the experiment. Chevalier and Ellison (1997) considers mutual funds investment strategies, and find that mutual funds adopt riskier strategies when nearing the end of the calendar year in order to obtain a stronger end-of-year performance.

Other work examines the implications of up-or-out environments on risk taking in particular. In an experimental setting, Oprea (2014) explores the interaction between profit maximization, risk-taking, and survival. He finds that when profit maximization and survival can occur simultaneously, participants will choose strategies close to the optimum. However, when survival and profit maximization imply different strategies, subjects will choose survival over profit maximization, exposing themselves to greater risk in profits.⁵ As another example, Cabral (2003) presents a model set in a story of investments in research and development where, in equilibrium, the industry leader chooses a safe strategy

⁵Note that in our context, survival and profit maximization imply the same strategy—to make it past the cut.

while followers choose risky strategies. This model is tested in Mueller-Langer and Versbach (2013), where the second game in sequences of two-game soccer tournaments is used as a measure of whether teams play differently when they've lost the first game. They find no such evidence that pending elimination changes behavior. Of course, as a team sport, soccer may introduce an aggregation problem that challenges identification of the causal parameter of interest.⁶

Grund, Hocker, and Zimmermann (2013) considers risk taking in the NBA, demonstrating that teams who are losing near the end of the game take more three-point shots, but that these riskier shots do not translate into a higher score. In weightlifting competitions, where participants choose what they intend to lift, Genakos and Pagliero (2012) finds an inverted-U relationship between participant rank and the weight they intend to lift. (Higher- and lower-ranking lifters choose to attempt heavier lifts, which is the riskier strategy.) Performance is also lower for higher-ranked lifters, which suggests that risk taking may map into realized performance differently across player ability.⁷

Ozbeklik and Smith (2017) consider risk taking in golf tournaments and find that players with lower world ranking (OWGR) are more likely to take risks in match-play golf tournaments, as are those who are playing poorly compared to their contemporaneous opponent. However, Ozbeklik and Smith (2017) defines risk as *ex post* variability of score. We, instead, separately identify *ex ante* risk taking and the *ex post* outcome of having taken that risk. In particular, we use PGA-defined measures of risk-taking potential on each hole played on the PGA Tour,

⁶Taylor (2003) proposes a model that accounts for general-equilibrium effects, arguing that mutual-fund managers may best respond to the risk-taking incentives faced by other managers—those with nothing to lose—by taking more risks themselves in order to stay ahead.

⁷Increased tournament incentives in NASCAR leads to more accidents (Becker and Huselid, 1992), especially when closely ranked drivers are nearby (Bothner, han Kang, and Stuart, 2007).

and conditional on this sub-sample of holes, consider whether a player takes that risk or does not.⁸

Background and data

Tournaments on the PGA Tour can vary in format and scoring system. The most-common scoring system is called stroke-play—it is by far the scoring system most are thinking of when they think of golf. We use only stroke-play tournaments in our analysis.⁹ The winner of a stroke-play tournament is the player who completed all days of the tournament with the fewest cumulative number of strokes. However, in most stroke-play tournaments on the PGA Tour, it is also customary to cut players at the end of the second day of competition—with 18 holes in each round, that implies that elimination occurs after two times around the same 18-hole course. It is this pending threat of elimination that provides our identifying variation.

The elimination criterion used most often by PGA tournaments is to cut to 70 players, plus all ties. In a typical tournament, the “70 plus ties” rule falls in a fairly fat part of the distribution of player, so ties are not uncommon, and the number of players who actually make the cut thus varies.¹⁰ In Figure 1 we


⁸McFall and Rotthoff (2016) uses stroke-level data from golf tournaments, where they define risk as a player being near the green on a shot earlier than would be expected given the par of the hole. This is potentially confounded with, for example, unobserved player ability, but is interpreted as evidence of increased risk taking in response to the presence of superstars, and evidence of reference bias—players tending to take more risk when their current rank is further away from their OWGR. Their measure of risk also differs from our measure.

⁹In a match-play tournament, players compete one-on-one for a win on each hole. The player who has fewer strokes on the most holes is the winner of this type of play. Garcia and Stephenson (2015) examines player performance under the Stableford scoring mechanism—a very small number of tournaments use this system, where points are awarded for a player’s score relative to par—and finds no evidence that risk taking increases in response to this convexity.

¹⁰The purpose of instituting a cut to the field of competitors is primarily to speed up play, allowing players to play the third and fourth days of competition in pairs instead of in threes.

capture the realized number of players making the cut on average, across all 526 tournaments in our sample.

Figure 1. How many players make the “70 plus ties” cut



Notes: Given ties, the number of players who make the cut after 36 holes of play can exceed 70. In this figure, we plot the histogram of the number of players who make the cut in a given tournament, across all stroke-play tournaments on the PGA Tour, 2002-2016.

It is also the case that every player who makes the cut shares some portion of the total purse, while no player missing the cut receives any portion of the purse. In a tournament with the median purse of \$6,000,000, a last-place finish will yield \$12,000. More generally, however, prizes asymptote to 0.2 percent of the purse on average. As we are identifying off of the discontinuity created at the elimination

In 2008, an additional cut rule was added to the PGA Tour. If more than 78 players make the day-two cut, a second “70 plus ties” cut is held at the end of day three, to reduce the field of competitors going into the last day of competition. These players are said to have made the cut, but not finished. Players who are cut at the end of day three receive a share of the tournament purse that would be consistent with having played all four days and finishing in last place.

cut, we thus have in mind that the \$12,000 prize is part of the explanatory to any systematic differences in behavioral we observe around the elimination margin.

In our analysis we use the hole-level panel data provided by the PGA Tour’s ShotlinkTM—every hole played by every player in all PGA Tour events. We restrict our sample to stroke-play tournaments with four rounds of scheduled and completed play, and a “70 plus ties” elimination after the second day of competition.¹¹ We also restrict our analysis to tournaments that were not significantly influenced by weather (e.g., we drop tournaments where rounds were completely eliminated or where multiple rounds were played on one day instead of two). As identification is achieved around the elimination rule, we restrict our sample to where we have identification—all player-tournament-holes strictly within the first 36 holes of each tournament. We also restrict our analysis to players who completed 36 holes, reflecting that players have no obligation to complete each round, and anything falling short of two full days of competition may introduce problematic sample selection.

Our sample includes data on 2,630 players across 526 tournaments, all of them held between 2002 and 2016 inclusive. Our data include information about each player’s performance on each hole. Across courses, the PGA Tour defines certain holes as “going-for-it” holes, which we use to determine whether golfers systematically take more or less risk when elimination is pending. Risk taking—“going for the green,” as it would be called—is therefore defined by an attempt to reach the green in fewer strokes than would be suggested by the par on the hole. For example, on a par-five hole, instead of taking three shots to get to the green, a player might attempt to hit the green in only two strokes. We observe whether

¹¹For example, we discard the Master’s Tournament which has a cut at “50 plus ties,” but also has the provision that players within 10 strokes of the leader make the cut.

players indeed took the riskier strategy on these holes, and whether they were successful in their attempt to reach the green. Players who fail to reach the green often land in some sort of hazard, leading to higher scores than would be expected if the risk had not been attempted.

To better control for player ability we include players' world rankings from the previous year, which also facilitates our consideration later of heterogeneity across player ability. In short, this OWGR ranking is a weighted measure of tournament success in each player's two most-recent years of competition, with points awarded according to finishing placement in any tournament and more weight given to more-difficult tournaments.¹²

Identification

The fundamental source of variation we exploit is the discontinuity in player expectations of making the cut, introduced by the elimination of competitors that will occur at the end of 36 holes. Specifically, the PGA Tour's "70 plus ties" rule initiates a notion of pending elimination on all holes after the first. We will therefore ask whether there are identifiable differences in performance or risk taking when playing from the elimination side of this rule on each of the holes $h \in \{2, 3, \dots, 36\}$. After identifying average effects, we will explore heterogeneity in this relationship. For example, among other things, we will consider how it might change as the threat of elimination approaches, and how the threat of elimination might induce changes in performance and risk taking differentially for those who were eliminated in their most-recent tournament.

¹²Recency is also given more weight, considering that golf is a game where ability is highly varying across time, so recency may better reflect current rank. These rankings are not limited to PGA Tour players, but include every professional tour, and includes the top-200 players through 2006, and the top-300 players thereafter.

The main threat to identifying the effect of a potential elimination on player performance and risk taking is that unobserved player ability will simultaneously affect both player rank (i.e., their rank relative to the elimination discontinuity) and outcomes (i.e., strokes taken). In particular, as players who perform worse are more likely to be cut, we would expect to find lower average performance (i.e., higher scores) on the elimination side of the discontinuity. Identifying off of within-player variation will protect identification from this potential confounder—in our preferred specifications we include player-by-year-by-tournament fixed effects, addressing the concern that retrieving the causal parameter is hampered by unobservable ability. Formally, we therefore allow for player ability to vary other than within given tournaments. In the ideal experiment, we would compare a player to himself on the same hole in the same year in the same tournament with the same cumulative number of strokes, but at a differently ranked position due to the (exogenous) performance of the competitors he faced that weekend. This reveals the fundamental econometric problem, of course, as each golfer plays each tournament-hole only once. We do get close to the ideal experiment, however, by comparing a player to himself across holes in the same tournament, where those holes are played while at differently ranked positions relative to the cut. Restricting the sample to players who are closest to the elimination margin likewise mitigates this concern, which we will do as part of our bandwidth-sensitivity analysis.

The econometric specification we are describing can be written,

$$\begin{aligned}
 Y_{ihty} = & \beta_1 \mathbb{1}(\text{Rank}_{ihty} \geq E_{hty}) + \beta_2 \text{Rank}_{ihty} \\
 & + \beta_3 \text{Rank}_{ihty} \times \mathbb{1}(\text{Rank}_{ihty} \geq E_{hty}) + \delta \text{Par}_{ihty} + \gamma_{ity} + \epsilon_{ihty},
 \end{aligned} \tag{2.1}$$

where Y_{ihty} is a placeholder for the outcome of interest (e.g., total strokes, putts, risk taken) of player i on hole h of tournament t in year y , Par_{ihty} is the par of the

hole, $Rank_{ihty}$ is player i 's rank in the field of competitors (after having played hole $h - 1$ but before playing hole h), and E_{hty} is the elimination threshold.¹³ Player-by-year-by-tournament fixed effects are captured in γ_{ity} . We are primarily interested in the role of $\mathbb{1}(Rank_{ihty} \geq E_{hty})$ —with the threshold player i faces on hole h determined only by lagged performance, $\hat{\beta}_1$ identifies the difference in player performance that is systematic with being on the elimination side of an exogenous threat of elimination.

We consider four main outcomes in our analysis: the player's stroke total on the hole, a measure of how many putts were taken on a given hole, the binary choice of whether a player “Went for it” (on the subsample of holes the PGA officially designates as “going-for-it” holes), and a variable indicating whether the player hit the green after taking that available opportunity for a riskier strategy. Any difference in total strokes attributable to the pending threat of elimination we interpret as some change in effort or focus, or to playing the hole differently by taking more or less risk. Any variation in putts alone, however, cannot be attributable to risk taking—there are no options for taking risky putts, *per se*—which will help in identifying whether movement in total strokes is entirely attributable to risk. (It will not be.) With respect to the return to risk taking itself, we will interpret the player hitting the green after taking the risk as a measure of success.


Although the actual cut occurs at the end of the second day of competition, we observe each player's performance on each hole of the tournament, and can

¹³For 94 percent of the tournaments in our sample, half of the players will randomly be assigned to start on hole 10 in order to speed up play in the first few days of competition. We define h as the hole sequence for player i , such that that the player who starts on the course's first hole and the player who starts on the course's tenth hole start the day at $h = 1$. The results are robust to including an out-of-order indicator, including tournament-by-year-by-order fixed effects, and to the inclusion of player-by-tournament-by-year-by-order fixed effects.

therefore recreate the status of any pending elimination that would have been faced on approach to each hole. Given the “ties” included in the PGA Tour’s “70 plus ties” elimination rule, E_{hty} is a tournament-specific elimination threshold. The indicator variable $\mathbb{1}(Rank_{ihty} > E_{hty})$ therefore captures any player i with a rank worse than the “70 plus ties” cut as he approaches hole h and therefore faces elimination without some improvement. We normalize E_{hty} to zero in all figures and tables below, after accounting for ties. While elimination is according to ordinal ranking, we will also respect cardinal relationships when predicting player performance on either side of the elimination rule and, thus, define $Rank_{ihty}$ in deviations from the *stroke* total that would imply elimination (within tournament, of course, and recalculated for the entire field of players after each hole). In Figure 2, we report the distribution of rank over the pooled sample of all player-holes. In the end, if $Rank_{ihty} > 0$, then $Rank_{ihty}$ is equal to the number of strokes i must pick up in order to make the cut. If $Rank_{ihty} < 0$, then $|Rank_{ihty}|$ is equal to the number of strokes i could drop before he failed to make the cut. The interaction term in the model identifies the slope parameter on $Rank_{ihty}$ for those who face elimination as of hole h .

To the question of whether the elimination threshold is relevant to competitors, in Figure 3, we present a histogram of the relative-to-par elimination threshold. Across all tournaments in our sample, there is seemingly a high degree of predictability in the threshold, even before the tournament has begun. Players also have tournament-specific knowledge in the formation of their beliefs, suggesting an even tighter distribution of deviations around the threshold than is implied in the figure. Thus, it seems reasonable to assume that players are aware of the threshold. However, as is typical in tournaments, players move across the

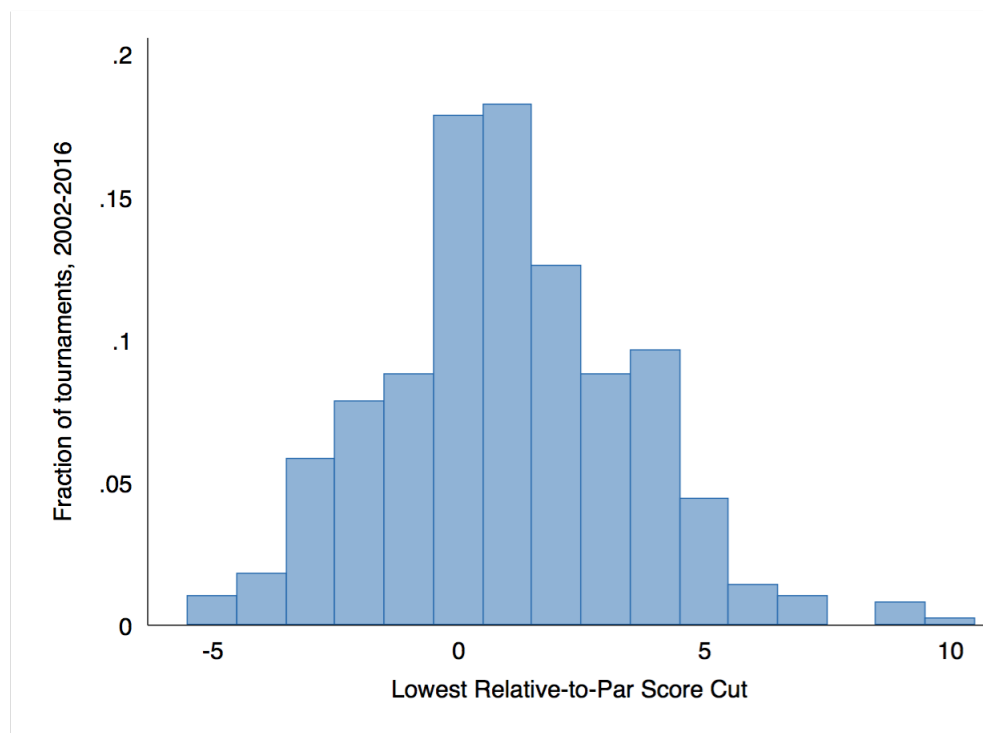
Figure 2. Deviations from tournament cut ($Rank$)



Notes: Given deviations from the implicit elimination threshold on each hole of play on the PGA Tour (E_{hty}), we plot the histogram of player rank at each hole (measured in strokes) relative to the elimination threshold, across all stroke-play tournaments on the PGA Tour, 2002-2016.

elimination threshold because they are playing better or worse *relative to their peers*—there is not a fixed threshold for success. As such, identifying player-specific differentials in performance and risk taking while on one side of the threshold or the other is contributed to by both own and competitor performance, and we assume that players anticipate their competitors’ best response.¹⁴

Figure 3. How predictable is the elimination threshold?



Notes: In this figure we plot the histogram of strokes *relative to par* that constituted the “70 plus ties” cut across all stroke-play tournaments on the PGA Tour, 2002-2016.

Our specifications are estimated using ordinary least squares, where γ_{iyt} indicates tournament specific controls for player heterogeneity, and the estimation

¹⁴Though not reported, we consider whether players differentially responded to being on the elimination side of the threshold as a function of their most-recent play—whether they had recently picked up or lost strokes relative to par—and find no systematic variation. In the heterogeneity analyses below, we do report the results of considering the outcome of players’ most recent tournament, the number of threshold-crossings players make in the current tournament, and their average deviation from the threshold.

of ϵ_{ihty} allows for clustering at the level of player-by-year-by-tournament. As the estimated coefficients are implying changes in *hole-level* performance, note that a small change in hole-level performance can amount to sizable changes in 36-hole performance over two days of competition.

Results

Baseline Performance Results

In Column (1) of Table 1, we report estimates of a baseline specification of hole-level performance on either side of the discontinuity in players' expectations of survival that is introduced by the "70 plus ties" elimination threshold. Without including controls for player ability, the positive slope parameter on *Rank* in Column (1) is consistent with better players tending to take fewer strokes to complete a hole, on average. In Column (2), we absorb this player-specific time-invariant heterogeneity into the error structure, which has the effect of reversing the sign of the estimated slope parameter. However, golf being the game it is, with players arguably experiencing hot and cold spells, there is reason to anticipate that player ability can vary across time in ways that would then escape player fixed effects. Thus, in Column (3) we allow for tournament-by-year player heterogeneity. To the extent players have good and bad weekends idiosyncratically, identifying the difference in player i 's performance using variation within a given weekend of competition, when he is in and out of facing elimination over the course of that tournament, will be our preferred specification.¹⁵ To the extent we have not controlled for varying player-specific heterogeneity, we anticipate the main threat to

¹⁵It is not uncommon for the best professional golfers to have bad weekends. For example, Jordan Spieth, the world-number-one golfer at the end of 2015, failed to make it through to the third day of competition in four of the 25 PGA Tour events he entered in 2015.

identification continuing to work against finding performance improvements on the elimination side.

Table 1. The performance gains under pending elimination

| | Strokes per hole | | | Putts only |
|---|-----------------------|-----------------------|-----------------------|-----------------------|
| | (1) | (2) | (3) | (4) |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.012*** (0.001) | 0.002 (0.001) | -0.050*** (0.002) | -0.029*** (0.001) |
| <i>Rank</i> ^a | 0.0002 (0.0003) | -0.002*** (0.0003) | -0.063*** (0.0004) | -.031*** (0.0003) |
| <i>Rank</i> × $\mathbb{1}(\text{Rank} \geq 70)$ | 0.0154*** (0.0005) | 0.006*** (0.0005) | 0.013*** (0.0007) | 0.0097*** (0.0005) |
| <i>Par</i> | 0.827*** (0.0007) | 0.827*** (0.0007) | 0.816*** (0.0007) | -0.035*** (0.0006) |
| Player FE | | Yes | | |
| Player-by-year-by-tournament FE | | | Yes | Yes |
| Observations | 2,519,650 | 2,519,650 | 2,519,650 | 2,519,650 |
| Number of groups | | 2,555 | 71,990 | 71,990 |
| R^2 | 0.377 | 0.382 | 0.405 | 0.010 |
| Mean | 3.96 | 3.96 | 3.96 | 1.60 |
| Implied change in strokes (per <i>t</i>) | | | -1.785 | -1.039 |

Notes: Standard errors in parentheses, allowing for clustering at the player-by-year-by-tournament level. *** p<0.01, ** p<0.05, * p<0.1. ^a In stroke play, lower integer ranks are better.

In our preferred specification, the estimated discontinuity in performance at the elimination margin is therefore -0.05, suggesting that players perform relatively better (than themselves, on the same weekend) when they face elimination, and thus have nothing to lose. As the estimated parameters are per-hole performance

measures, it is noteworthy to consider that the 36-hole equivalent yields a pre-cut marginal effect of -1.785 strokes—a meaningful improvement in performance given the margins that often make the difference between a player failing to make the cut and playing through to the final day of competition. The percent of players missing the cut by one stroke varies (across tournaments) from 2 to 16.5 percent, with 8.7 percent of the field of players in the average tournament missing the cut by one stroke.

We will shortly turn to consider the elasticity of risk taking with respect to potential elimination. Before doing so, we estimate in Column (4) our preferred specification but with the number of putts as the dependent variable. As putting is not subject to any choice of risk-related strategy, this result serves to establish our prior that risk is surely not able to explain the entire increase in performance. (Some debate exists about whether putting affords any risk taking or not. However, the PGA—as does the consensus opinion, it seems—does not acknowledge risk-taking opportunities with respect to putting.)¹⁶ The estimated discontinuity is -0.03, with the associated 36-hole equivalent of this marginal effect of -1.039. Increased performance on putts thus explains 58 percent of the improvement in overall score, which supports that players do indeed perform better when on the elimination side of the upcoming cut in ways that are independent of their risk taking. Without yet an available appeal to risk taking to explain this increase in performance, this improvement in score is most likely explained by increased focus and determination when the threat of elimination is more salient.

¹⁶Contrary to this, Pope and Schweitzer (2011) defines “risk-averse” putts as shots where the player does not aim for the hole, but with the objective of setting up a less-risky subsequent shot. However, this strategy is not optimal nor likely to occur as there is no downside risk to putting past the hole. Unlike “going for it” holes, where hazards make missing the shot risky, in putting it only matters how far away from the hole you are—it is optimal to try to get the ball in the hole. This is further reflected in any number of amateur and professional “how to” videos.

Baseline Risk-Taking Results

In Table 2, we adopt our preferred specification from above but restrict the sample to those holes designated by the PGA as “going-for-it” holes. In this table, we model the variation in three risk-related outcomes: whether players went for the green on risk-taking holes, and as indications of success, whether players hit the green and the distance to the hole they faced subsequent to having taken that risk.¹⁷ The sample size varies across columns of Table 2 as the risk-success outcomes are conditional on having taken the risk.

In general, our results in Column (1) suggest that there is a positive relationship between player rank and risk taking—a one-stroke decline in performance (an *increase* in *Rank*) is associated with a roughly 1-percent increase in the propensity to take a risk, all else equal. However, those on the elimination side of the cut choose the risky strategy 2.7 percentage points more often—a 5.4-percent increase in the probability of taking a risk when possible. Given an average of 7.2 holes (out of 36) on which it is possible to take this type of risk, this represents a potential improvement in performance of 0.19 strokes over the first two days of competition (i.e., $7.2 \times .027$). Thus, at the upper bound where all risk taking pays off, taking the risk on “going for it” holes can explain up to 49 percent of the gains in measured performance.¹⁸

¹⁷If we repeat the preferred specification of strokes, but restrict the sample to those holes designated by the PGA as “going-for-it” holes, we find similar point estimates even though there are no par-3 holes in this sub-sample, since it is expected that all golfers will always attempt to hit the green from the tee. This explains part of the decline in sample size going from Table 1 to Table 2.

¹⁸Restricting Table 1 Column (3) to a sample of “going-for-it” holes yields a point estimate discontinuity of -.055. Thus, where all additional risks taken reward the player with a one-stroke improvement, the fraction of gains attributable to additional risk taking is $.027/.055=.49$.

Table 2. Risk taking under pending elimination

| | $\mathbf{1}(\text{Went for it})$ | $\mathbf{1}(\text{Went for it}) = 1$ | |
|--|----------------------------------|--------------------------------------|------------------------|
| | | $\mathbf{1}(\text{Hit green})$ | Distance remaining |
| | (1) ^a | (2) ^a | (3) |
| $\mathbf{1}(\text{Rank} \geq 70)$ | 0.027*** (0.003) | 0.012*** (0.003) | -13.372*** (3.340) |
| <i>Rank</i> | 0.011*** (0.0006) | 0.009*** (0.0008) | -5.557*** (0.727) |
| <i>Rank</i> \times $\mathbf{1}(\text{Rank} \geq 70)$ | 0.003*** (0.0009) | -0.003** (0.001) | 1.974* (1.106) |
| Par | 0.037*** (0.002) | 0.130*** (0.002) | -146.008*** (4.648) |
| Player-by-year-by-tournament FE | Yes | Yes | Yes |
| Observations | 445,158 | 221,865 | 221,865 |
| Number of groups | 63,423 | 58,613 | 58,613 |
| R^2 | 0.005 | 0.016 | 0.122 |
| Mean (depvar) | 0.498 | 0.242 | 673.323 |

Notes: Standard errors in parentheses, allowing for clustering at the player-by-year-by-tournament level. *** p<0.01, ** p<0.05, * p<0.1. ^a Linear-probability models, though binary response models yield qualitatively similar results.

Having established playing at the margin of elimination increases risk taking (on the elimination side), in Column (2) we then consider the *ex post* realization of risk taking. Where players take risky strategies, the PGA records “success” quite simply—hitting the green on such a shot is considered success, which we capture with a binary outcome. Conditional on the risky strategy having been chosen, we find a one percentage-point increase in the probability of a successful outcome on the elimination side. When players are on the elimination side of the cut on a hole that affords the option to take a risk, that option is being chosen with higher likelihood. If those induced risks are similar in their riskiness (and are just opted for more often) then we would expect coincident declines in rates of success on the same side of the cut. Yet, we find *increases* in the probability of hitting the green on those risks taken from the elimination side of the elimination threshold. Assuming that selection into risk taking is monotonically increasing in the riskiness of the opportunity, the low-risk options should be played more often than the high-risk options. Risk taking therefore explains 23 percent of the gains in measured performance on “going-for-it” holes.¹⁹

As one last attempt to capture variation in player performance around risk taking, in Column (3) of Table 2 we consider the remaining distance to the hole after having taken a risky shot. While the PGA records failing to hit the green after choosing the risky strategy as failure, any systematic variation in the remaining distance to the hole may similarly point to improved performance. Conditional on the distance their risky shot is taken from, we find that when players face elimination they land their risky shots 13-inches closer to the hole,

¹⁹Restricting Table 1 Column (3) to a sample of “going-for-it” holes *on which risk was actually taken* yields an estimated discontinuity of -.053. Thus, accounting for rates of success, the fraction of gains attributable to additional risk taking is $.012/.053=.23$.

on average, than when they do not face elimination. Again, we interpret the data as suggestive that the threat of elimination is increasing “productive” risk taking.

Sensitivity, Heterogeneity, and Bandwidth Considerations

In the analysis above we identify the average effect of a pending threat of elimination. Below, we wish to consider whether the elasticity of performance and risk taking is evidently different as the threat of elimination becomes more salient, and the potential sensitivity of our results to the selection of players who contribute to identification. As it turns out, these are all “bandwidth” considerations, in a way, so we group them together below.

Responsiveness as Elimination Approaches

Here, we will explore the effect of pending elimination as the opportunities to influence outcomes slip away—as there are few holes remaining and the threat of elimination becomes more salient. We accomplish this by re-estimating our models while sequentially dropping the earliest holes played in each tournament.

In Figure 4, for each of our four outcome variables, we report the estimated effects of being on the elimination side of the elimination threshold across this metric. Moving to the right on the figure, we restrict the sample to fewer holes—those holes closer to the actual tournament cut at the end of the 36th hole.²⁰ In Panel A, the estimated discontinuity in strokes is negative throughout most of the specifications, though it attenuates as we discard early holes from the model, and actually flips sign as we identify only off of the threat of elimination over

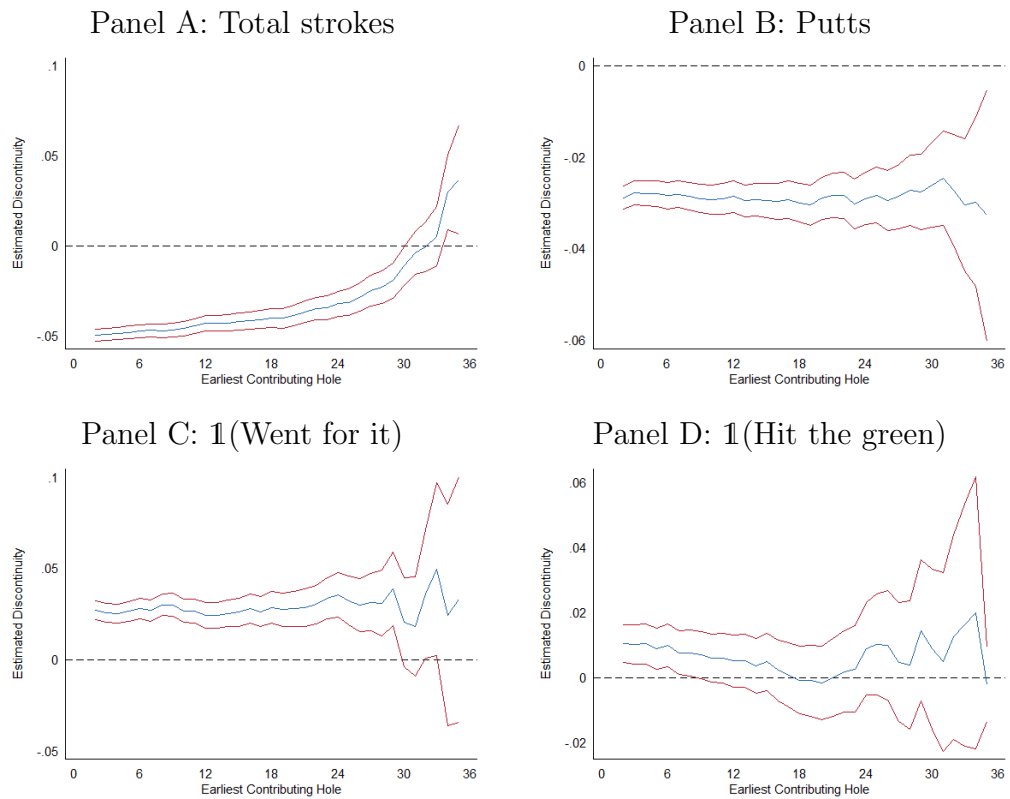
²⁰Since we are using player-by-tournament-by-year fixed effects, the last hole we can actually discard from our sample is hole 34. Recall also that hole 1 is not identified since players are not ranked prior to posting a stroke total.

the last few holes, where it is most salient. That is, on holes 34 and 35, players are performing *worse* when on the elimination side of the cut. (Recall that it is *reductions* in score that equate with *improvements* in performance.) This would be consistent with players losing focus as the opportunities for success slip away, or evidence that players have limited capacities to maintain performance levels with the stress of elimination being felt so strongly.

Though precision is lost, the estimated discontinuity in putts, in Panel B, is seemingly invariant across the same sequence of specifications. In Panel C we plot the estimated discontinuities for risk taking, which also prove very stable throughout the 36 holes, as does the differential probability that the player hits the green on risky shots as elimination approaches (Panel D). To the extent risk taking behavior is especially salient in the second day of the tournament, and players have better information about the field, the second day is when one might want to take risks. There is a tentative evidence in favor of this argument in Panel C of Figure 4—after being relatively stable in the first 18 hole (the first day), the estimated discontinuity in “going for it” starts increasing (as does, though imprecise, the corresponding discontinuity for “hitting the green” in Panel D).

Overall, then, we find consistent performance improvements associated with potential elimination that are in part driven by induced increases in successful risk taking. However, the improvements induced by the contract convexity represented in the threat of elimination do not withstand the immediacy of that threat, where significant *declines* in performance become evident as elimination approaches and there are few remaining opportunities to avoid elimination. Interestingly, however, risk taking and putting performance play little role in this eventual decline. This

Figure 4. Does responsiveness change as elimination approaches?



Notes: Point estimates and 95-percent confidence intervals from repeated estimations of Equation (2.1), restricting the sample to holes successively closer to elimination.

reversal is then most-easily attributed to differences in performance in the earlier shots of each hole.

The Responsiveness of Players Who Are Closer to the Cut on Average


Here we consider parameter sensitivity to the systematic removal of players (not holes) from the sample. Players are amassed around the cutoff for the first few holes, of course, and the contributions of those player-holes introduced to the identifying variation by those who will quickly separate from the field (in either direction) may not be the ideal experimental variation off of which to identify.²¹ For example, even after controlling for player ability, even the tournament's eventual winner could have easily played the first hole at one or two strokes over par, and would thus contribute one or two observations to the identifying variation, in ways that have little to do with any responsiveness to pending elimination. (Such would yield negative bias, as he subsequently performed better on the elimination side, which is what led him to eventual victory.) Clearly, the strongest and weakest players in the field, who will necessarily contribute at least a few observations around the cutoff, are hardly the marginal observations off of which we should identify. We would prefer, in the sense of the ideal experiment, to find these players randomly on each side of the elimination threshold, which we might best approximate by considering the sensitivity of the point estimate to collapsing on those players who are closer to the cutoff *on average* and are therefore most likely to experience the quasi-random play from both side of the cut.

In Figure 5 we plot the histogram of players' average within-tournament ranks, revealing how different players can be in their average deviation from the

²¹On average, 61 percent of the field is on the margin of being cut at hole 2.

threshold. The truly marginal players are arguably those clustered around the cut (at zero). In Table 3 we therefore report estimates of the discontinuity as we remove players with average (tournament-specific) ranks farthest away from the threshold—roughly, then, it is both the best and worst players being eliminated as we tighten up the estimation around the discontinuity. We consider the entire sample in the first column, those within ten strokes of the cut in the second column, through to using only those with a mean rank within one stroke of the cut. Total strokes and estimates of putting responsiveness are very stable as the bandwidth collapses in this way, restricting observations to players who are truly middling in each tournament.

Figure 5. Mean rank of players



Notes: Given deviations from the implicit elimination threshold on each hole, we plot the histogram of *average* player rank (measured in strokes) relative to the elimination threshold, across all stroke-play tournaments on the PGA Tour, 2002-2016.

Table 3. Bandwidth sensitivity by mean rank of player

| | Absolute deviation (in strokes) from elimination threshold | | | |
|--|---|--------------------------|-------------------------|-------------------------|
| | $ \mu_r \leq 20$ (1) | $ \mu_r \leq 10$ (2) | $ \mu_r \leq 5$ (3) | $ \mu_r \leq 1$ (4) |
| <i>Panel A: Strokes</i> | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.050*** (0.002) | -0.048*** (0.002) | -0.039*** (0.002) | -0.050*** (0.003) |
| Observations | 2,519,650 | 2,516,885 | 2,352,525 | 715,680 |
| <i>Panel B: Putts</i> | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.029*** (0.001) | -0.028*** (0.001) | -0.023*** (0.001) | -0.020*** (0.002) |
| Observations | 2,519,650 | 2,516,885 | 2,352,525 | 715,680 |
| <i>Panel C: $\mathbb{1}(\text{Went for it})$</i> | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | 0.027*** (0.003) | 0.026*** (0.003) | 0.021*** (0.003) | 0.021*** (0.005) |
| Observations | 445,158 | 444,684 | 415,562 | 126,400 |
| <i>Panel D: $\mathbb{1}(\text{Hit the green})$ conditional on $\mathbb{1}(\text{Went for it})=1$</i> | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | 0.012*** (0.003) | 0.011*** (0.003) | 0.010*** (0.003) | 0.015*** (0.005) |
| Observations | 221,865 | 221,671 | 206,636 | 62,382 |

Notes: All specifications include par, and absorb player-by-year-by-tournament unobserved heterogeneity into the error structure. Standard errors in parentheses, allowing for clustering at the player-by-year-by-tournament level. *** p<0.01, ** p<0.05, * p<0.1

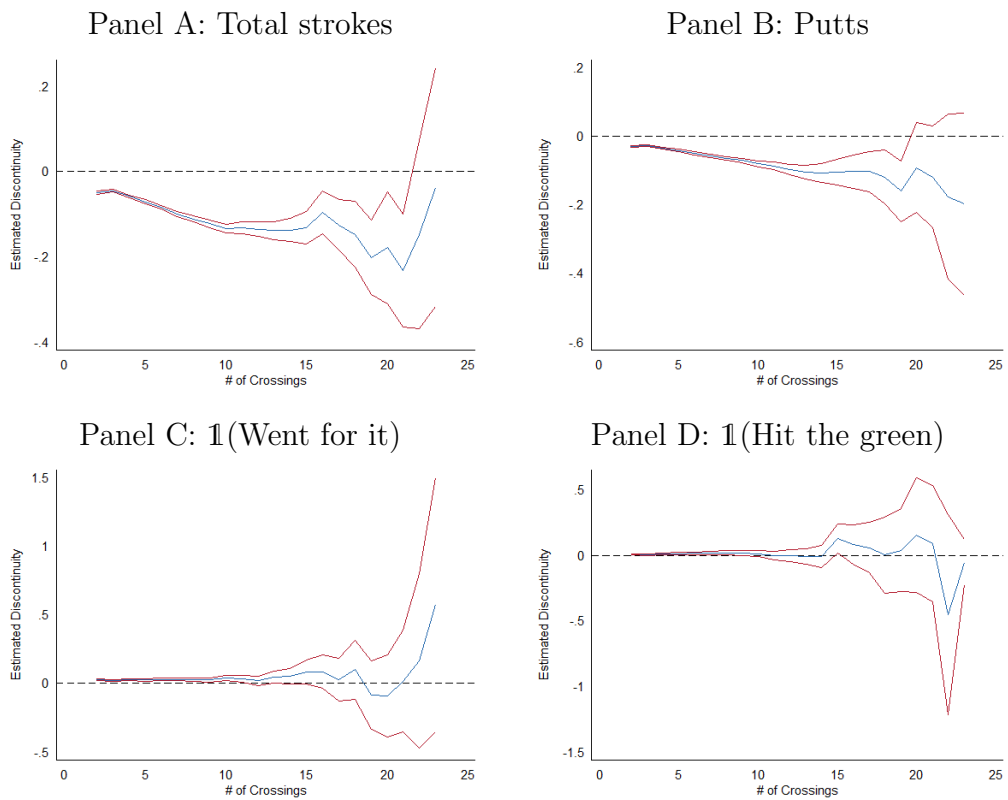
The Responsiveness of Players Who Are Around the Cut More Often

As one last alternative, rather than removing players based on their average rank over the tournament, we can collapse on players who are close to elimination most often. In Figure 6 we report the estimated discontinuities across increasingly restricted samples. The x-axis of the figure represents the minimum number of times contributing players crossed the threshold in either direction during the first 36 holes of competition. As we move to the right—this has the sample increasingly consist of players who crossed the threshold more frequently—the estimated discontinuity in performance (in Panels A and B) more than doubles, though loses precision in the expected way. This suggests that, if anything, players most often at the margin of elimination are more sensitive to which side of the threshold they are on. The propensity to take risks and to succeed at those risks remains relatively constant (in Panels C and D).

Why Not Restrict the Sample by Disaggregated Player-Rank?

Another possible bandwidth-sensitivity exercise would increasingly restrict the sample to tournament-year-player-holes where player ranks were nearest the cut. Given the importance of capturing unobserved player heterogeneity, in our preferred specifications we identify only off of within-player variation. As such, restricting the sample to only those player-holes where the player was close to elimination (not just close on average) quickly decreases our ability to detect any behavior of interest. More fundamentally, though, the objective in collapsing around the identifying threshold is to increase the comparability of the “treatment” and “control” groups, which this does not accomplish. For example, as previously discussed, a *Rank*-based bandwidth restriction gives increasing weight

Figure 6. The estimated discontinuity by number of threshold crossings



Notes: Point estimates and 95-percent confidence intervals from repeated estimations of Equation (2.1), restricting the sample to those “closer” to elimination, defined as having more crossings of the elimination threshold.

to players who will only pass through the treatment margin in the first few holes of the tournament, before they separate to either the front or back of the field of competitors. These players are in no way representative of the sort of marginal player off of which we wish to identify.

Additional Heterogeneity

Time of Day

We continue with our consideration of heterogeneity in other dimensions by first including an analysis of what could constitute a threat to identification. Though likely to impart only attenuation bias, the time-staggered play across the field of competitors may matter insofar as early groups are less informed about the stroke totals that will contribute to the “70 plus ties” elimination threshold (i.e., the E_{hty} above).

To partially address this, we stratify the model by the hour each hole was finished. To avoid conflating the time of day and hole sequence, we control for hole sequence directly. In Figure 7 we report the estimated discontinuities, noting again that the confidence interval naturally widens around observations at the end of the day, where there are fewer observations. In the end, however, we find the estimates quite robust to time of play. Even though morning players are arguably less informed than afternoon players, those in the morning are not seeming to react to potential elimination any less than those in the afternoon. This is also consistent with professional golfers knowing about where the cut line is going to be and reacting to that expectation, as we suggested in our discussion of Figure 3.

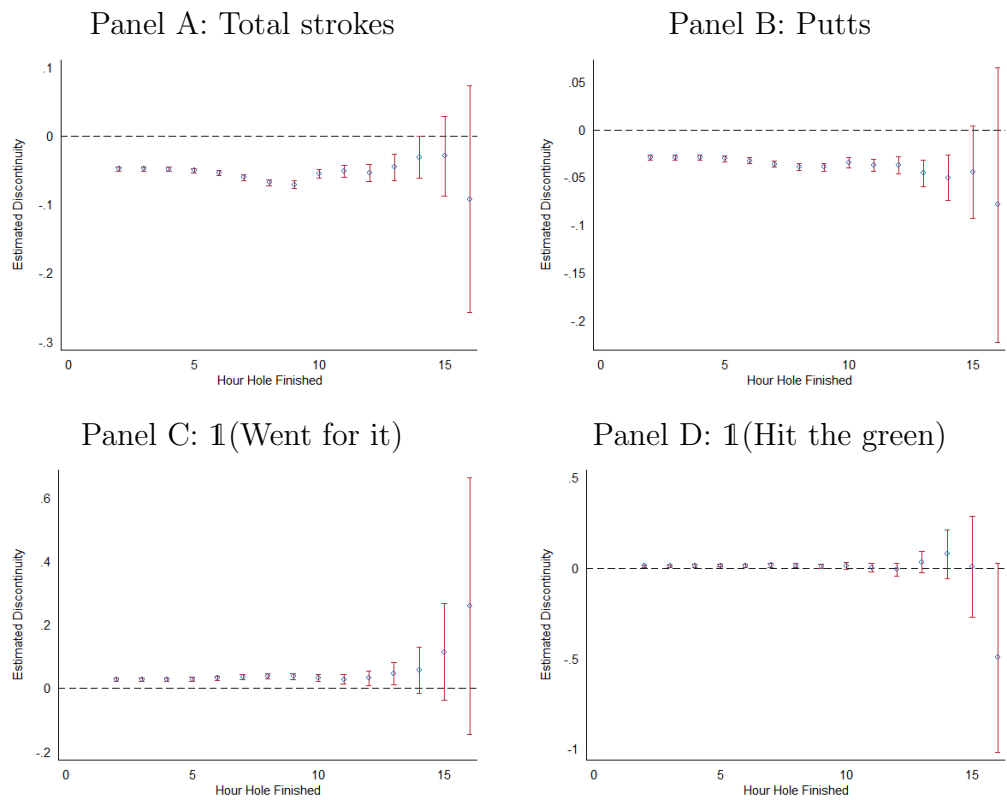
Field- and Player-Specific Ability

We next explore heterogeneous responses to the threshold. As players select into tournaments, it is interesting to consider heterogeneity by tournament, which we do by separately considering tournaments for which at least 30 percent, 50 percent, and 70 percent of the field is world ranked (according to the OWGR), respectively. Estimated discontinuities across these very different average-ability levels reveal no differential responsiveness to finding oneself on the elimination side of the cut, enough so that we are inclined to believe that differential selection into tournaments is not contributing to our results. These estimates are reported in Table 4.

We also consider heterogeneity at the player level. In Table 5 we stratify the sample into ranked and unranked players, and then separately estimate the discontinuity for unranked players, and then for stronger and stronger pools of ranked players. Doing so reveals that higher-ranking players respond less at the margin to potential elimination. The estimated discontinuity on stroke totals for the highly ranked players is about half the size as it is for the unranked players and the estimated discontinuity on putts is also smaller.

Moreover, the risk-taking behavior and subsequent risk success of highly ranked players does not differ with pending elimination. This pattern is consistent with differences in experience—more-experienced players know better or are more comfortable with their style of play, thus being less sensitive to conditions. Essentially, a sign of maturity and expertise may well yield lower elasticities with respect to pending elimination.

Figure 7. Heterogeneity by hour of play



Notes: Point estimates and 95-percent confidence intervals from repeated estimations of Equation (2.1), stratified by the (player-specific) hour play was completed.

Table 4. Heterogeneity by average ability of the field of competitors

| | Full sample (1) | % of field with OWGR ranking | | |
|--|----------------------|------------------------------|----------------------|----------------------|
| | | ≥ 30 (2) | ≥ 50 (3) | ≥ 70 (4) |
| <i>Panel A: Strokes</i> | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.050*** (0.002) | -0.050*** (0.002) | -0.054*** (0.002) | -0.057*** (0.003) |
| Observations | 2,572,654 | 2,354,695 | 1,649,900 | 645,015 |
| <i>Panel B: Putts</i> | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.029*** (0.001) | -0.029*** (0.001) | -0.030*** (0.002) | -0.028*** (0.003) |
| Observations | 2,572,654 | 2,354,695 | 1,649,900 | 645,015 |
| <i>Panel C: $\mathbb{1}(\text{Went for it})$</i> | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | 0.027*** (0.003) | 0.029*** (0.003) | 0.031*** (0.003) | 0.031*** (0.005) |
| Observations | 445,158 | 419,542 | 301,078 | 127,078 |
| <i>Panel D: $\mathbb{1}(\text{Hit green})$ conditional on $\mathbb{1}(\text{Went for it})=1$</i> | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | 0.012*** (0.003) | 0.013*** (0.003) | 0.015*** (0.004) | 0.019*** (0.005) |
| Observations | 221,865 | 208,883 | 153,307 | 66,793 |

Notes: All specifications include par, and absorb player-by-year-by-tournament unobserved heterogeneity into the error structure. Standard errors in parentheses, allowing for clustering at the player-by-year-by-tournament level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5. Heterogeneity by player ability

| | Unranked | OWGR-ranked players | | | |
|--|----------------------|----------------------|----------------------|----------------------|----------------------|
| | players | 1-300 | 1-200 | 1-100 | 1-50 |
| | (1) | (2) | (3) | (4) | (5) |
| <i>Panel A: Strokes</i> | | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.051*** (0.002) | -0.036*** (0.002) | -0.038*** (0.003) | -0.030*** (0.004) | -0.021*** (0.005) |
| Observations | 1,096,270 | 1,423,380 | 1,183,455 | 635,495 | 325,220 |
| <i>Panel B: Putts</i> | | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.032*** (0.002) | -0.020*** (0.002) | -0.021*** (0.002) | -0.020*** (0.003) | -0.021*** (0.004) |
| Observations | 1,096,270 | 1,423,380 | 1,183,455 | 635,495 | 325,220 |
| <i>Panel C: $\mathbb{1}(\text{Went for it})$</i> | | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | 0.029*** (0.004) | 0.022*** (0.004) | 0.020*** (0.004) | 0.006 (0.005) | 0.006 (0.008) |
| Observations | 188,045 | 257,113 | 214,492 | 116,356 | 59,776 |
| <i>Panel D: $\mathbb{1}(\text{Hit green})$ conditional on $\mathbb{1}(\text{Went for it})=1$</i> | | | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | 0.009* (0.005) | 0.009** (0.004) | 0.007 (0.005) | 0.009 (0.006) | 0.012 (0.008) |
| Observations | 89,001 | 132,864 | 112,793 | 64,634 | 34,574 |

Notes: All specifications include par, and absorb player-by-year-by-tournament unobserved heterogeneity into the error structure. Standard errors in parentheses, allowing for clustering at the player-by-year-by-tournament level. *** p<0.01, ** p<0.05, * p<0.1

It is also the case that as we collapse on higher-ranking players, we are collapsing on players who are increasingly likely to make a given cut, and thereby secure prize winnings with greater likelihood and in larger amount.

Among the highly ranked, the performance increase in strokes due to the pending elimination is similar in magnitude to the performance increase in putts, implying that increased putting concentration and focus can completely explain the overall performance increase.

Hole Difficulty

Opportunities to gain strokes vary according to the hole's par, necessitating the use of par as a control variable in our preferred specification. It is somewhat natural to then consider the potential heterogeneity across par. In Table 6 we report changes in our outcome variables across par-3, par-4, and par-5 holes separately.²²

The estimated discontinuity in all four outcomes is somewhat larger in magnitude as the par of the hole increases. That is, players are more likely to gain strokes in response to elimination on the longer par-5 holes, where there are more opportunities to gain strokes.²³ While some of this is surely mechanical—the standard deviations in strokes taken on par 3s, 4s, and 5s, are 0.623, 0.670, and 0.727—we see evidence of performance responding in similar ways even on the shorter par-3 holes, which again reflects that risk taking is not fully accounting for changes in performance. Explicitly accounting for these varying standard

²²As no par-3 holes are ever classified as “going-for-it” holes, the two risk-dependent variables cannot be estimated on par-3 holes.

²³This is evident in data, with 8,559 instances of players scoring an eagle on par-5 holes and only 212 eagles (i.e., holes in one) on par-3 holes.

Table 6. Heterogeneity by par of hole

| | Par 3 (1) | Par 4 (2) | Par 5 (3) |
|--|----------------------|----------------------|----------------------|
| <i>Panel A: Strokes</i> | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.036*** (0.003) | -0.047*** (0.002) | -0.056*** (0.004) |
| Observations | 578,327 | 1,496,835 | 444,470 |
| <i>Panel B: Putts</i> | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.020*** (0.003) | -0.029*** (0.002) | -0.031*** (0.003) |
| Observations | 578,327 | 1,496,835 | 444,470 |
| <i>Panel C: $\mathbb{1}(\text{Went for it})$</i> | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | | 0.014*** (0.004) | 0.026*** (0.003) |
| Observations | | 86,117 | 359,041 |
| <i>Panel D: $\mathbb{1}(\text{Hit green})$ conditional on $\mathbb{1}(\text{Went for it})=1$</i> | | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | | 0.0001 (0.003) | 0.014*** (0.004) |
| Observations | | 51,592 | 170,273 |

Notes: All specifications absorb player-by-year-by-tournament unobserved heterogeneity into the error structure. Standard errors in parentheses, allowing for clustering at the player-by-year-by-tournament level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

deviations, the effect sizes for par 3s, 4s, and 5s are 5.8, 7.0, and 7.7 percent of a standard deviation, again suggesting this isn't purely mechanical. Calculating effect sizes for these estimated coefficients, Players are also more risk-responsive to pending elimination on par-5 holes than on par-4 holes. Concerning returns to risk, when on the elimination side of the cut players are also more likely to hit the green after taking a risk on par-5 holes.²⁴

Does a Player's History of Elimination Matter?

In Table 7, we stratify by whether the player's most-recent tournament ended with him being eliminated after 36 holes, or not. Doing so suggests that those who were eliminated in their most-recent tournament were, if anything, less responsive to playing under pending elimination. As we have absorbed player-specific fixed effects into the error structure of our preferred specification, this implies that those coming off of an elimination play more similarly on either side of the pending cut than do those who successfully made it to the third day of competition in their most-recent tournament. This slight increase in responsiveness to the threat among those with recent success holds across total strokes, putts, risk taking, and in the probability that risk end in success. While we have no strong priors, this is consistent with heightened expectations from recent success interacting with the current threat of elimination to produce more of a motivating device.

²⁴This can't be explained by differences in hole difficulty as more players succeed in hitting the green on par-4 holes than par-5 holes. Across all players, there is a 23-percent risk-success rate for par-5 holes and 25-percent risk-success rate for par-4 holes.

Table 7. Heterogeneity by cut outcome of most-recent tournament

| | Outcome in most-recent tournament | |
|--|-----------------------------------|-----------------------|
| | Eliminated (1) | Not eliminated (2) |
| <i>Panel A: Strokes</i> | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.036*** (0.002) | -0.048*** (0.002) |
| Observations | 1,406,195 | 1,031,415 |
| <i>Panel B: Putts</i> | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | -0.024*** (0.002) | -0.026*** (0.002) |
| Observations | 1,406,195 | 1,031,415 |
| <i>Panel C: $\mathbb{1}(\text{Went for it})$</i> | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | 0.023*** (0.004) | 0.027*** (0.004) |
| Observations | 251,455 | 179,957 |
| <i>Panel D: $\mathbb{1}(\text{Hit green})$ conditional on $\mathbb{1}(\text{Went for it})=1$</i> | | |
| $\mathbb{1}(\text{Rank} \geq 70)$ | 0.009** (0.004) | 0.011** (0.005) |
| Observations | 127,377 | 86,759 |

Notes: All specifications include par, and absorb player-by-year-by-tournament unobserved heterogeneity into the error structure. First tournament for all players discarded. Standard errors in parentheses, allowing for clustering at the player-by-year-by-tournament level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Conclusion

The PGA records enable hole-by-player-by-tournament analysis around objectively determined cutoffs, where players are routinely observed on both sides of the threshold of elimination from tournament competition. We approximate the ideal experiment by comparing a player to himself across holes in the first-two days of a single tournament—on some holes, he will be on the elimination side of the threshold, while on other holes he will be on the safe side.

We exploit an opportunity to jointly observe performance, *ex ante* risk taking, and the *ex post* realization of risk. Collectively, we paint a picture of performance, risk taking, and rates of success on risks taken, each being higher when players play from the elimination side of the elimination threshold. Our results are robust to a battery of sensitivity exercises, and strongly suggest that sensitivity to elimination is stronger in lower-ability players. In particular, while measured performance does improve, there is no evident increase in risk taking among those of highest ability.

Overall, we suggest that 23 percent of the improvement in performance induced by potential elimination is due to productive increases in risk taking. These effects are most-evident among those closest to the margin of elimination, among lower-ability competitors, and diminish as elimination approaches.

CHAPTER III
THE MALLEABILITY OF BELIEF: VARYING RESPONSES TO “SHOULD
SAME-SEX COUPLES HAVE THE RIGHT TO MARRY?”

Introduction

In this paper, we consider the potential for policy innovation itself to lead to changes in belief. Specifically, we exploit variation in the timing of state-level legalizations of same-sex marriage to identify systematic changes in related belief—the “policy feedback” effect of same-sex marriage legalization. Given this policy feedback, we will then consider the extent to which such feedback explains polarization.

Understanding the dynamics of normative belief is important for economists to consider insofar as such belief may temper the acceptance of policy innovations or otherwise motivate behavior on margins that influence future policy. Martin and Yurukoglu (2017) and DellaVigna and Kaplan (2007) suggest that media-consumption choices influence political beliefs, for example—specifically that exposure to Fox News induces more-conservative belief. By extension, then, it is not difficult to imagine that political offices and public policy are shaped by changes in underlying belief.

More generally, there is a healthy segment of the population that identifies social issues as themselves important margins. Over the last 75 years, for example, roughly one in six people have identified a social issue as “the most-important problem” facing the United States (Heffington, Park, and Williams, 2017), with as

much as 74 percent of people (in 1989) flagging a social issue as the most important problem.¹

Others have considered the potential for feedback from policy to beliefs around the legalization of same-sex marriage. In strictly pre-post designs (Bishin, Hayes, Incantalupo, and Smith, 2015; Kreitzer, Hamilton, and Tolbert, 2014) report mixed evidence.² Flores and Barclay (2016) matches “treated” to “untreated” individuals in the American National Election Study (ANES) panel around four state-level changes to same-sex marriage (three state-ballot measures in 2012, and the 2013 California legalization). They find transition probabilities of support for gay rights suggestive of increasing support. However, unlike the GSS, the ANES lacks the dexterity to separately identify beliefs beyond broad opposition, support, or “ambivalence.”³ Our analysis of the data will suggest that there is much more movement evident when the survey instrument allows for movements in strength of belief.

One fear, of course, is that innovations in social policy may well play a role in increasing polarization in society. Since 1960, for example, the difference in ideological position of the median democrat and the median republican in the U.S. Congress has increased by 53 percent (Poole, 2005). Moreover, U.S. citizens

¹Examples include gender equality, civil liberties, discrimination and racism, same-sex rights, generational issues, crime, abortion, religion, family issues, or other values and morals. In general, the propensity for social issues to displace economic issues is strongly related to booms in the business cycle.

²Bishin, Hayes, Incantalupo, and Smith (2015) collects opinion data on same-sex marriage before and after relevant court rulings, and are unable to detect changes in opinion. Similarly, Kreitzer, Hamilton, and Tolbert (2014) cites that public approval of same-sex marriage increased in Iowa after its 2009 legalization. Although conditioned on some cofounders, these studies lack control-group responses, and may well be confounded by increasing public support for same-sex marriage, as we demonstrate in Figure 8.

³Regarding same-sex marriage, the authors define ambivalence toward same-sex marriage as support for civil unions for same-sex couples, but opposition toward same-sex marriage

are themselves increasingly more politically polarized (Pew Research Center, 2014). Some 92 percent of those who identify as “republican” are now measurably more conservative than the median democrat, while 94 percent of those who self identify as “democrats” are now more liberal than the median republican. Only twenty years ago, these same metrics were 64 percent and 70 percent, respectively.

In Section III, we provide a brief history of same-sex marriage in the US, and describe the data. It is in this section that we provide the specific context for identifying the effect of policy variation on public opinion. In Section III, we describe our empirical strategy, and report results of the basic model, where we will argue for a more-nuanced interpretation of the data than has the existing literature, as well as directly examining changes in polarization. We offer some discussion of the broader implications for research and policy in Section III.

Background and Data

General Social Survey

Since 1972, the General Social Survey (GSS) has collected the sentiment of Americans on such issues as national-spending priorities, crime and punishment, intergroup relations, and confidence in institutions. The GSS, facilitated by the National Opinion Research Center (NORC) at the University of Chicago, is the main source of data for our analysis, collected almost every year between 1972 and 1994, and every other year since 1994. Although the GSS first asked about beliefs on same-sex marriage in 1988 as a part of an international survey of beliefs,

we discard data from that year and only utilize data starting in 2004, when the question became part of the permanent series.⁴

The relevant responses we will track in the GSS are respondent's reflections on the question, "Do you agree or disagree? Homosexual couples should have the right to marry one another." A standard five-point scale was offered to respondents (i.e., Strongly disagree, Disagree, Neither agree nor disagree, Agree, and Strongly agree), which will become important to identifying some nuanced movement in belief over time. In Figure 8, we reproduce the proportion of respondents who selected each of the five options across the 2004-2016 time series as fractions of one.⁵ Although we will reveal a somewhat more nuanced story with the data than that of a simple movement of opinion toward acceptance, "average" support for same-sex marriage has indeed increased in the US over the last 12 years. Capturing central tendencies from distributions of opinion is non-trivial, however, evident in the figure is the transition, for example, from the modal response being "strong disagreement" in 2010 and earlier, to "agreement" in 2012, and then to a modal response of "strong agreement" by 2014.⁶ Similarly, the median response shifts from "disagree" to "agree" over the time series as the mass generally shifts from opposition toward support. In fact, the proportion of respondents who either "agree" or "strongly agree" with legal same-sex marriage doubles between 1988 and 2004 (not evident in the figure, but support increases from 12.4 to 29.7 percent) and doubles again by 2016 (from 29.7 to 59.3 percent). Clearly, being careful about


⁴The question was asked in 1988 as a part of the International Social Survey Programme (ISSP).

⁵While we do not consider the 1988 responses as part of our econometric exercise, in Appendix Figure A1 we do display the long trend, which is similarly toward greater support.

⁶Beliefs on other policies are seemingly trending more slowly, with the exception of inclinations toward marijuana use.

the role of trends in our empirical specifications will be important. Yet, as we do, we should also note the *distribution* of belief, which has the potential to reveal any potential polarization of opinion in recent years.

Figure 8. Responses to “Homosexual couples should have the right to marry one another,” 2004–2016



Notes: We plot the proportion of GSS respondents in each categorical response to the GSS question “Do you agree or disagree? Homosexual couples should have the right to marry one another,” as fractions of one.

Policy Variation and Identification

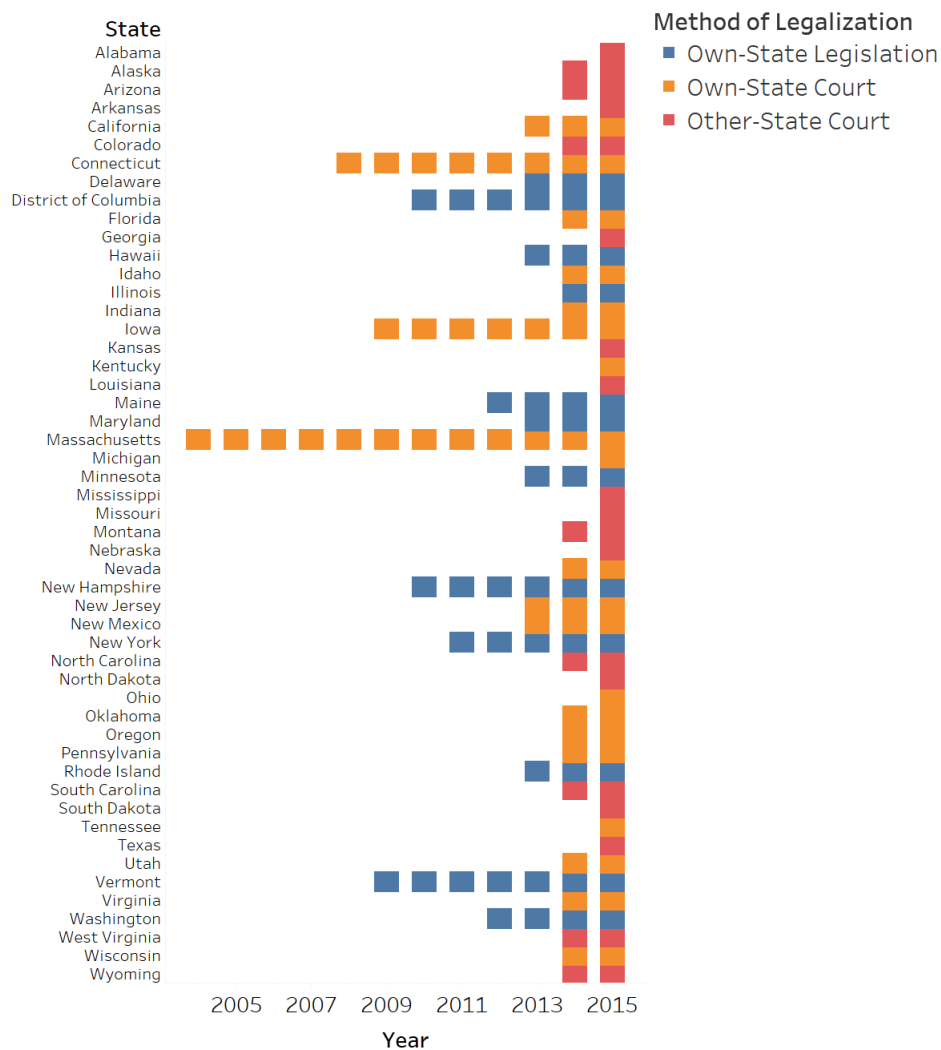
Of fundamental interest are the dynamics of belief around legalization. As such, we augment the GSS data with publicly available data on legislative and judicial decisions relating to same-sex marriage, at state and federal levels. In May 2004, Massachusetts became the first state to legalize same-sex marriage.

In response, over the subsequent two years, 23 other states passed constitutional amendments barring same-sex marriages. By the end of 2012, 31 states had constitutional bans. Meanwhile, still other states had joined Massachusetts in moving toward legalization, with eight states and the District of Columbia legalizing same-sex marriage by the end of 2012. In 2013 and 2014, a rash of court cases overturned many of these constitutional amendments and, in 2015, the U.S. Supreme Court legalized same-sex marriage nationwide.

In Figure 9, we present the timing and method of legalization across states. Much of the early movement was through legislative action to legalize same-sex marriage, either through referenda or state legislation, with seven of the first ten states legalizing same-sex marriage through one of these two methods. However, most individuals in the US experienced the legalization of same-sex marriage judicially, some through actions taken by their own state judiciary, but many others through federal circuit courts or the U.S. Supreme Court itself. In such cases, legalization is through what is referred to as binding precedent.⁷ There is an argument to be made, of course, that such imposed legality, given that it is a step removed from local public opinion, is in some sense cleaner identification. However, the timing of these—all of them happening between the 2014 and 2016 GSS surveys—does not easily allow for their full consideration. As such, we will present these results knowingly sacrificing power for added exogeneity, and be careful in the resulting inference.

⁷A lower court is under a binding precedent when a higher court has established case law that supersedes any potential lower court ruling. In the case of same-sex marriage, this most often occurred when a federal district court overturned a same-sex marriage ban, meaning that any courts under their jurisdiction would need to follow that ruling on any future same-sex marriage bans. This usually resulted in same-sex marriage bans being quickly overturned in states under a binding precedent.

Figure 9. Timing and method of same-sex-marriage legalization, by state



Notes: We present the timing of the legalization of same-sex marriage across states, color coded by method of legalization. We stratify legalization method into three categories: own-state legislation, which includes both legalization by popular referendum and action by the legislative branch; own-state court action; and other-state court action which also includes the Supreme Court national action.

Related Literature

Policy feedback specifically related to the legalization of same-sex marriage has been explored in a few domains. Harris (2015) uses a difference-in-differences model and a synthetic control to examine judicial retention and political participation after same-sex marriage legalization in Iowa. She finds that the legalization of same-sex marriage in Iowa led to the defeat of the justices who ruled in the case and to increased political participation in subsequent elections. Dee (2008) exploits variation in same-sex marriage legislation in Europe to determine how the transmission of sexually transmitted infections changed following legalization. He finds that infection rates decreased after legalization, attributing the decrease to lower rates of risky sex. In the public health literature, legalization of same-sex marriage has been associated with improved mental health among gay people (Hatzenbuehler, Keyes, and Hasin, 2009; Tatum, 2017).

Policies have been shown to affect beliefs in other arenas. In education, Cantoni, Chen, Yang, Yuchtman, and Zhang (2017) analyzes the impact of a gradual rollout of a politically motivated curriculum change, demonstrating that students' political beliefs under the new material more closely matched beliefs taught under the newer curriculum. Additionally, Clots-Figueras and Masella (2013) and Friedman, Kremer, Miguel, and Thornton (2016) explore the impact of education on beliefs, finding changes in political and social beliefs due to changes in educational curriculum. Di Tella, Galiani, and Schargrodsky (2007) uses exogenous variation in land ownership to identify increasing support for the free market, among other beliefs.⁸ In the political science literature, Erikson and Stoker (2011)

⁸For a summary of the economics literature on the interaction between private and public institutions and beliefs, see Alesina and Giuliano (2015).

exploits random variation in Vietnam draft lottery numbers and demonstrates that, after controlling for actual military service, men with lower draft numbers (and therefore more likely to be drafted) had opinions that were more anti-war and more-in-line with the democratic party. Pacheco (2013) considers state-year variation in workplace smoking bans on beliefs regarding the acceptability of smoking. Other examples here include welfare reform (Hetling and McDermott, 2008; Hetling, McDermott, and Mapps, 2008; Soss, 1999; Soss and Schram, 2007), and health-care policy (Barabas, 2009; Campbell, 2011; Gusmano, Schlesinger, and Thomas, 2002).

In early examinations of the consequences of group differences, ethnic diversity has been shown to frustrate economic growth. For example, Easterly and Levine (1997) uses cross-country comparisons to link the effects of ethnic fractionalization to economic growth, finding that higher ethnic diversity contributes to lower economic well-being. Alesina and Ferrara (2005) summarizes the large literature that has evolved around the connection between growth and measures of fractionalization or polarization. Exploring other consequences of polarization, other authors have found links between ethnic diversity and civil war (Montalvo and Reynal-Querol, 2005), genocide (Montalvo and Reynal-Querol, 2008), corruption (Papyrakis and Mo, 2014), and government quality (La Porta, Lopez-de Silanes, Shleifer, and Vishny, 1999). In these papers, it is common to attribute the slower growth to an inability of different groups to agree on policy or on public-good provision (Alesina and Ferrara, 2005). In addition, there is a wide literature on the effects of classroom diversity, suggesting mixed effects of diversity on achievement (Carrell and Hoekstra, 2010; Hanushek, Kain, and Rivkin, 2009; Hoxby, 2000). In more recent literature in political science, polarization has

been linked with both decreased legislative productivity and increased political participation and campaign investment (Nivola and Brady, 2008; Van Weelden, 2015).

Results

Support for Same-Sex Marriage

In considering the relationship between private beliefs and legalization, we will approach modeling as a difference-in-differences exercise, identifying off of the policy variation induced by the timing of legalization within and across states. In general, then, we will be considering specifications such as,

$$\mathbb{1}(Supportive_{isy}) = \beta \mathbb{1}(Legal_{sy}) + \delta_s + \gamma_y + \epsilon_{isy}, \quad (3.1)$$

where $Supportive_{isy}$ represents the individual belief of person i in state s in year y . In the GSS, belief is captured on a five-point scale. However, as a first pass, we define $\mathbb{1}(Supportive_{isy})$ as an indicator variable capturing whether the respondent takes the supportive position—that is, either “agrees” or “strongly agrees” with legalizing same-sex marriage. (This follows existing literature, though we will soon relax this restriction to analyze the five-fold responses in a multinomial logit specification.) We define δ_s and γ_y as state and year fixed effects, respectively, and $\mathbb{1}(Legal_{sy})$ as an indicator variable capturing whether same-sex marriage is legal in state s in year y . We estimate Equation 3.1 by ordinary least squares, allowing for ϵ_{isy} to capture any clustering at the state level. Point estimates can be interpreted as percentage-point changes in the probability of supporting the legalization of same-sex marriage.

In Table 8, we first report estimates directly associated with this baseline specification of Equation (3.1). In Column (1), we see what broader literatures typically interpret as evidence of causality running from the legalization of same-sex marriage to positive support for same-sex marriage—legalization inducing a 7.3-percent increase in the probability a respondent associates positively with same-sex marriage.⁹ However, given the prospect that states that legalized same-sex marriage may be the same states in which belief is trending upward faster, one might be concerned that such an interpretation of the data is not robust to allowing for state-specific trends, for example.

When we allow for state-specific time trends, in Column (2), we see how sensitive any such inference is, with the point estimate on legality attenuating and losing statistical significance. In Column (3), we add individual level controls, which further attenuates the point estimate on $\mathbf{1}(Legal)$.¹⁰ This exercise suggests that there is little support for a causal claim that legalization brings with it any significant movement between opposition and support.

Legalization and Strength of Belief

As the restriction to a simple notion of agree/not agree hides potential movements between the strength of beliefs and movement of those who do not have a strong opinion either way, we turn to examining changes in the strength of belief. We model the five-category dependent belief variable with a multinomial logit model.¹¹ In Figure 10, we report the implied changes in predicted probabilities

⁹See Flores and Barclay (2016)

¹⁰Individual level controls include age fixed effects, gender, race, employment status, income, religion, and educational attainment

¹¹The ordinal logit model is another natural model, but it is not appropriate for this context. The ordinal logit model depends on the “parallel regressions” assumption, which fails to be

Table 8. Support for same-sex marriage around same-sex-marriage legalization, 2004–2016

| | $\mathbb{1}(\textit{Supportive})$ | | |
|------------------------------|-----------------------------------|------------------|------------------|
| | (1) | (2) | (3) |
| $\mathbb{1}(\textit{Legal})$ | 0.033* (0.019) | 0.018 (0.020) | 0.009 (0.021) |
| State FE | Yes | Yes | Yes |
| Year FE | Yes | Yes | Yes |
| State-specific linear trend | No | Yes | Yes |
| Individual controls | No | No | Yes |
| Observations | 10,539 | 10,539 | 10,539 |
| Mean (dep var) | 0.457 | 0.457 | 0.457 |
| Impact size (at mean) | 0.073 | 0.039 | 0.019 |

Notes: In all specifications, the dependent variable is equal to one where respondents either agreed or strongly agreed that homosexual couples should be allowed to marry, and zero otherwise. Individual-level controls include sex, race, age, work status, income, religion, and education. Standard errors in parentheses, allowing for clustering at the state level. *** significant at 1%; ** significant at 5%; * significant at 10%.


across the five categorical responses permitted in the GSS and the associated impact sizes (given the different belief levels across categories). We derive these estimates and confidence intervals from a single multinomial logit model, which we interpret as the changes in mass at each of the five outcomes due to legalization. As before, we include state and year fixed effects, state-specific linear trends, and individual level controls. As with all such analyses, we are unable to identify where the mass in any one category is likely to have come from in the counterfactual world without treatment, but only how it has predictively changed. For example, and noteworthy here, we see that the estimated change in the probability that a respondent chooses “strongly agree” is roughly four percentage points higher after legalization. Relative to the underlying propensity to strongly agree among the control group (.213), this represents an impact of roughly 20 percent. Similarly, we see a *decrease* in the predicted probability that a respondent would choose “agree” of about the same magnitude and impact size. (We provide the multinomial logit estimates in Table A1 in the appendix.)

Although not statistically significant, there is also what appears to be an increase in the propensity to respond with “strongly disagree” and a decrease in the propensity to respond with “disagree” following legalization.¹² One interpretation of the data generating process is that mass is moving toward the “tails” of the distribution, in a roughly “U-shaped” manner. However, note also the slight increase in mass in the “Neither agree nor disagree” associated with legalization, as if the data are suggesting more of a “W-shaped” response to legalization.

satisfied here. However, the multinomial logit model can be thought of as a less-restrictive (and less-parsimonious) version of the ordinal logit model.

¹²The “strongly disagree” effect becomes statistically significant at the ten-percent level when using a less-conservative approach to our errors, when clustering at the state-year level instead of at the state level.

Figure 10. Changes in belief regarding same-sex marriage induced by same-sex-marriage legalization



Notes: On the left of each category (in blue), we plot the change in predicted probability with a 95% confidence interval from the multinomial logit model of the five-fold categorical response to “Do you agree or disagree? Homosexual couples should have the right to marry one another.” Point estimates are interpreted as the average change in the predicted probability of a respondent selecting each of the five different opinion choices due to legalization. On the right of each category (in red), we plot the implied impact—the percent change in likelihood of occupying that category, evaluated at its (control) mean.

This movement is suggestive of respondents holding their beliefs more strongly in response to legalization. This is also consistent with Table 8, in which respondents do not switch between opposition and support, but move within these categories.

In Table 9, we consider two different binary dependent variables with the intention of informing this notion of “U” or “W” shaped responses. In Column (1), we ask whether it is more likely that respondents choose either “Strongly disagree” or “Strongly agree,” rather than any of the interior categories available. We find a 0.092 increase in the probability that one responds in one of these two ways—an 18.7-percent increase. In Column (3), we include “Neither agree nor disagree” in our indicator variable and ask a similar question—are respondents more likely to choose either tail or the “neither” category after legalization, than they would have been prior to legalization? We find a 0.092 increase in the probability that one responds in one of these three ways—a 14.8-percent increase. With legalization, respondents are 24.2-percent less likely to select the “weakly” held positions of “disagree” or “agree.”

Heterogeneous Responses to Legalization

We next consider heterogeneity in these changes in belief. As age is an important predictor of support for same-sex marriage, it is important to consider the differential impact of legalization across age groups.¹³ We explore heterogeneous responses to legalization across different age cohorts. We stratify the sample into generations: greatest/silent generations (1915-1945), baby boomers (1946-1964), generation X (1965-1979), and millennials (1980-1988).¹⁴ We report the estimated

¹³The concern that generational concerns may bias results motivates our inclusion of age fixed effects in our main specification.

¹⁴In our sample, the youngest person surveyed was born in 1988.

Table 9. Intensity of chosen positions and same-sex-marriage legalization

| | $\mathbb{1}(SD, SA)$ (1) | $\mathbb{1}(SD, N, SA)$ (2) |
|-----------------------|-----------------------------|--------------------------------|
| $\mathbb{1}(Legal)$ | 0.092*** (0.024) | 0.092** (0.034) |
| Observations | 10,539 | 10,539 |
| R^2 | 0.036 | 0.022 |
| Mean (dep var) | 0.49 | 0.62 |
| Impact size (at mean) | 0.187 | 0.148 |

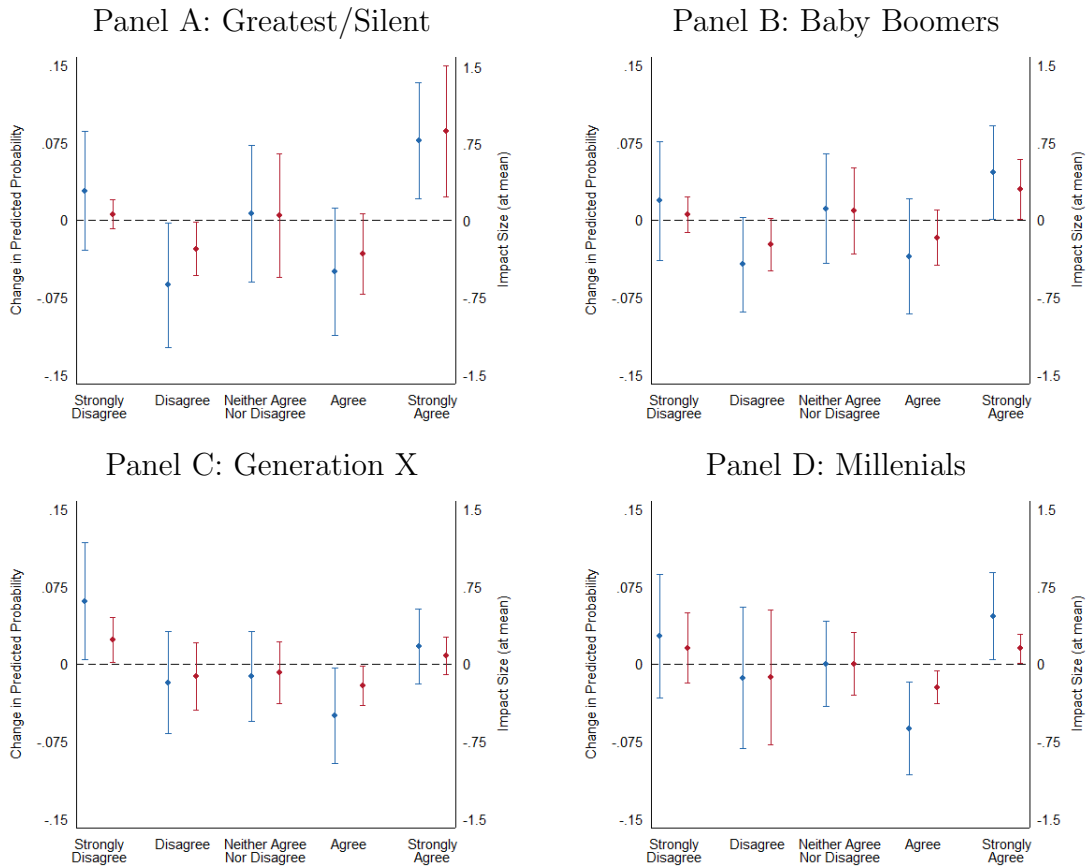
Notes: Specifications include state and year fixed effects, state-specific linear trends and individual-level controls. $\mathbb{1}(SD, SA)$ is an indicator variable for whether the respondent selected “strongly agree” or “strongly disagree.” $\mathbb{1}(SD, N, SA)$ is an indicator variable for whether the respondent selected “strongly agree,” “neither agree nor disagree,” or “strongly disagree.” Standard errors in parentheses, allowing for clustering at the state level. *** significant at 1%; ** significant at 5%; * significant at 10%.

coefficients and impact sizes from interacting legalization with these generational groups in Figure 11. Note that the same general pattern appears among all cohort groups—there is an increase in the estimated probability that people in each cohort will select the “strong” outcomes and a decreased probability that people will select the “weaker” outcomes. The consistency among these estimates helps alleviate concerns that these effects are being driven by changes in cohort composition over time.

In figures 12 and 13 we report similar systems of estimated impacts, having stratified the sample by education (Figure 12) and race (Figure 13). Notably, we see that it is the more-educated who are most inclined to *increase* their strong support of gay couples marrying with legalization, and the less-educated who are most inclined to *decrease* their support, though the general “W-shaped” responses are still generally evident across all groups. With respect to race and ethnicity, we again see this “W-shaped” pattern, though note the significant differences between whites, who are 26-percent more likely strongly agree after legalization, and Hispanic respondents, who are 41.6-percent more likely to strongly disagree after legalization. As white and more educated respondents have higher levels of support initially, this represents a divergence in opinion. As a result of the legalization of same-sex marriage, these educational and racial groups have increasingly-different beliefs from each other.

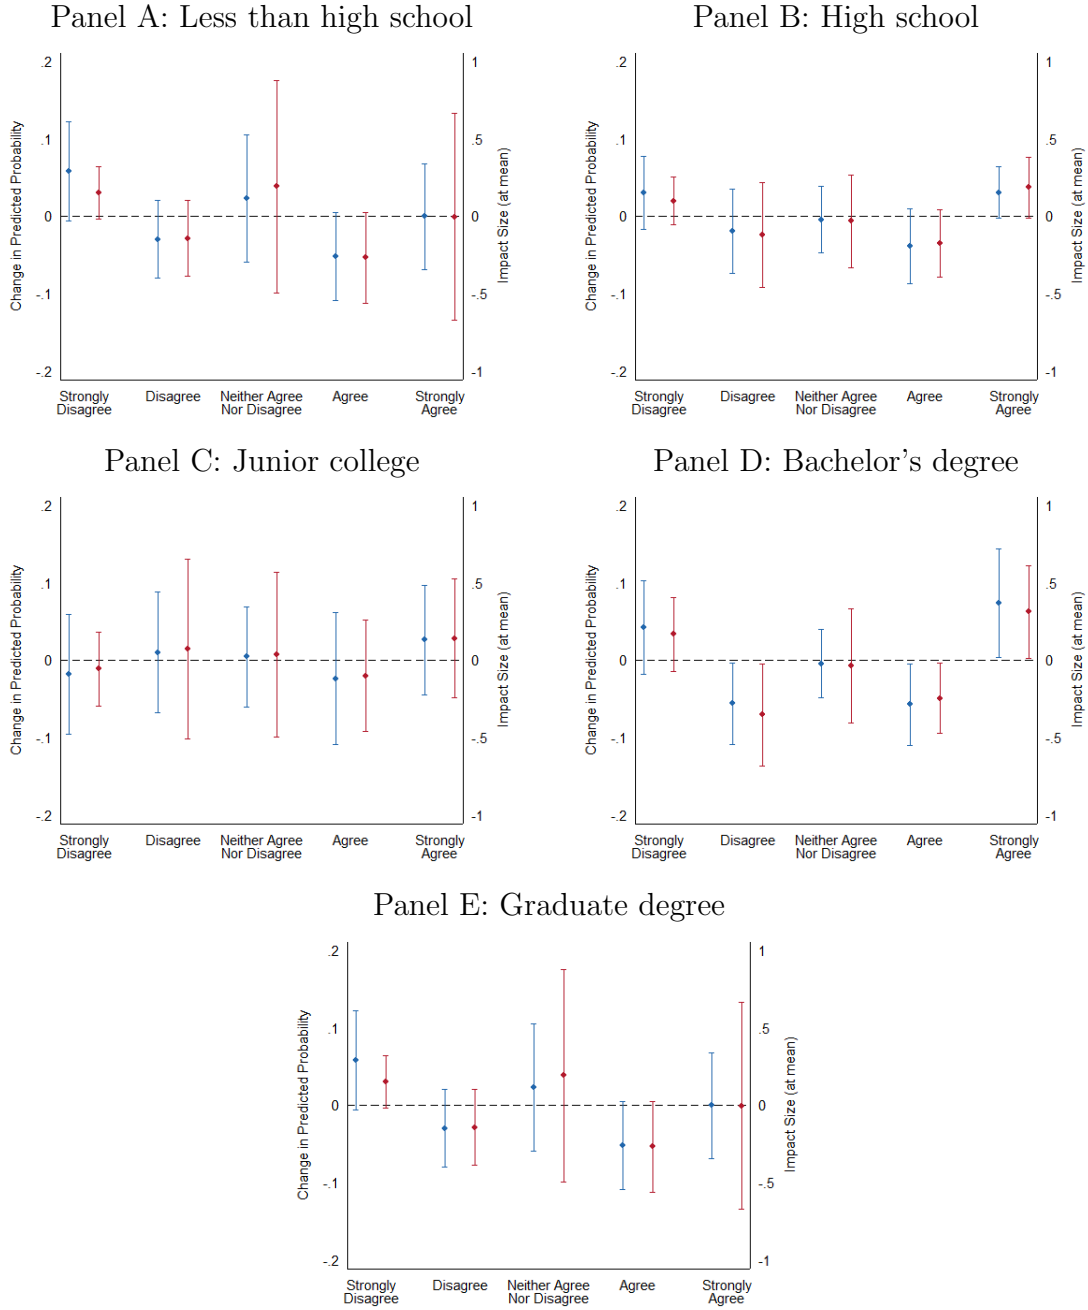
Finally, we explore heterogeneity across states with differential beliefs. We stratify states into terciles based on the average belief level of respondents in those states in 2004. We report the estimated effect of legalization on belief for these terciles in Figure 14. The “W” shape is again present across these panels, but

Figure 11. Changes in belief regarding same marriage induced by same-sex-marriage legalization, by generation



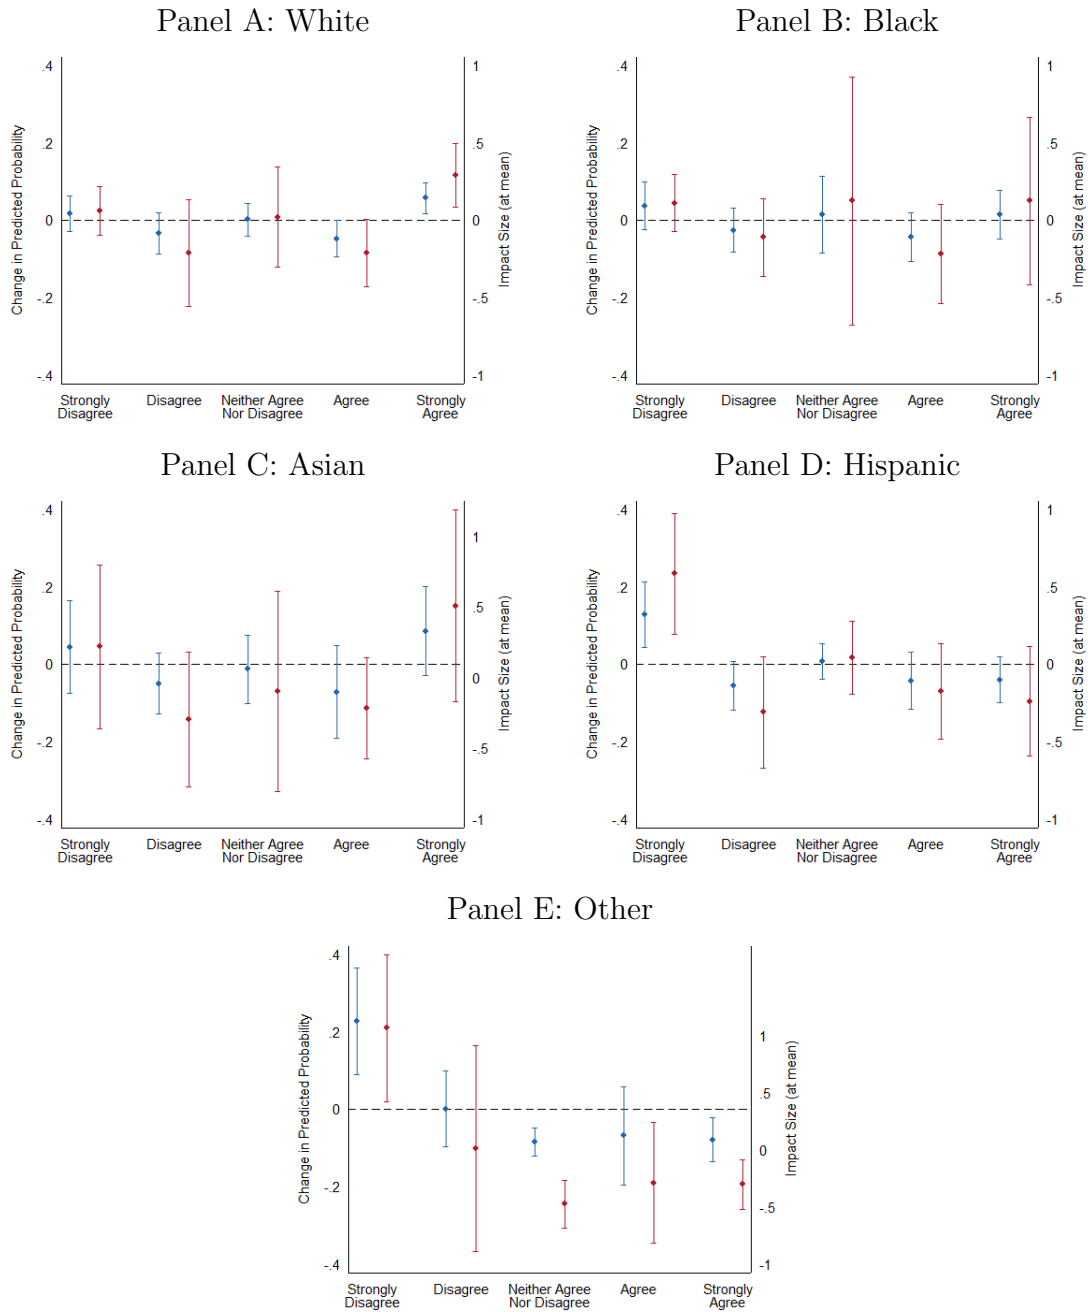
Notes: On the left of each category (in blue), we plot the change in predicted probability with a 95% confidence interval from the multinomial logit model of the five-fold categorical response to “Do you agree or disagree? Homosexual couples should have the right to marry one another.” Point estimates are interpreted as the average change in the predicted probability of a respondent selecting each of the five different opinion choices due to legalization. On the right of each category (in red), we plot the implied impact—the percent change in likelihood of occupying that category, evaluated at its (control) mean.

Figure 12. Changes in belief regarding same marriage induced by same-sex-marriage legalization, by education



Notes: On the left of each category (in blue), we plot the change in predicted probability with a 95% confidence interval from the multinomial logit model of the five-fold categorical response. On the right of each category (in red), we plot the implied impact, evaluated at its (control) mean.

Figure 13. Changes in belief regarding same-sex marriage induced by same-sex-marriage legalization, by race



Notes: On the left of each category (in blue), we plot the change in predicted probability with a 95% confidence interval from the multinomial logit model of the five-fold categorical response. On the right of each category (in red), we plot the implied impact, evaluated at its (control) mean.

notably, the “neither agree nor disagree” estimate is negative in Panel B. This is suggestive of the middle beliefs being emptied in those states.

We perform a similar exercise in Figure 15 by examining heterogeneity by the method of legalization. Again, the general “W-shaped” pattern appears across all three methods, although the movement is most pronounced in Panel A, among those states that legalized same-sex marriage through a popular referendum or through state legislative action. However, even among those states under a binding precedent in Panel C—those states in which legality is imposed most exogenously—we see a similarly-sized movement toward the extreme positions of “strongly agree” and “strongly disagree”, although we lack power to make strong inference.

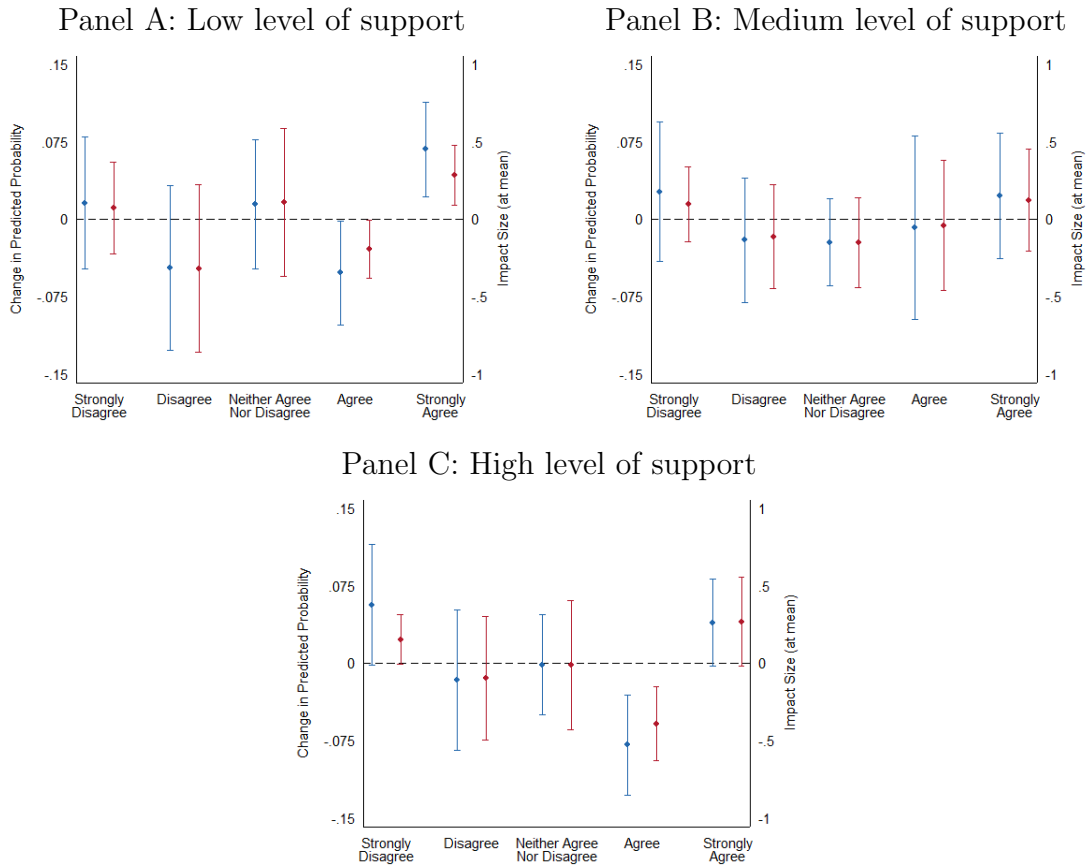
Failed Attempts to Legalize Same-Sex Marriage

Given the chaotic political environment surrounding the legalization of same-sex marriage, could it be that changes in beliefs are not due to the passage of legislation, but due to the political environment? To disentangle the political atmosphere from the passage of the law, we exploit failed attempts to legalize same-sex marriage. In these attempts, the news media attention, debate on social media, the salience of the law passing in one’s own state, etc. are presumably similar, with the outcome of the attempt being the only difference. In this way, we can differentiate between the activity surrounding legalization and the effect of the law.

We focus on attempts to legalize same-sex marriage through courts, as the bulk of the policy variation in the successful attempts occurred through the courts.¹⁵ We also focus on failed attempts to legalize same-sex marriage as opposed

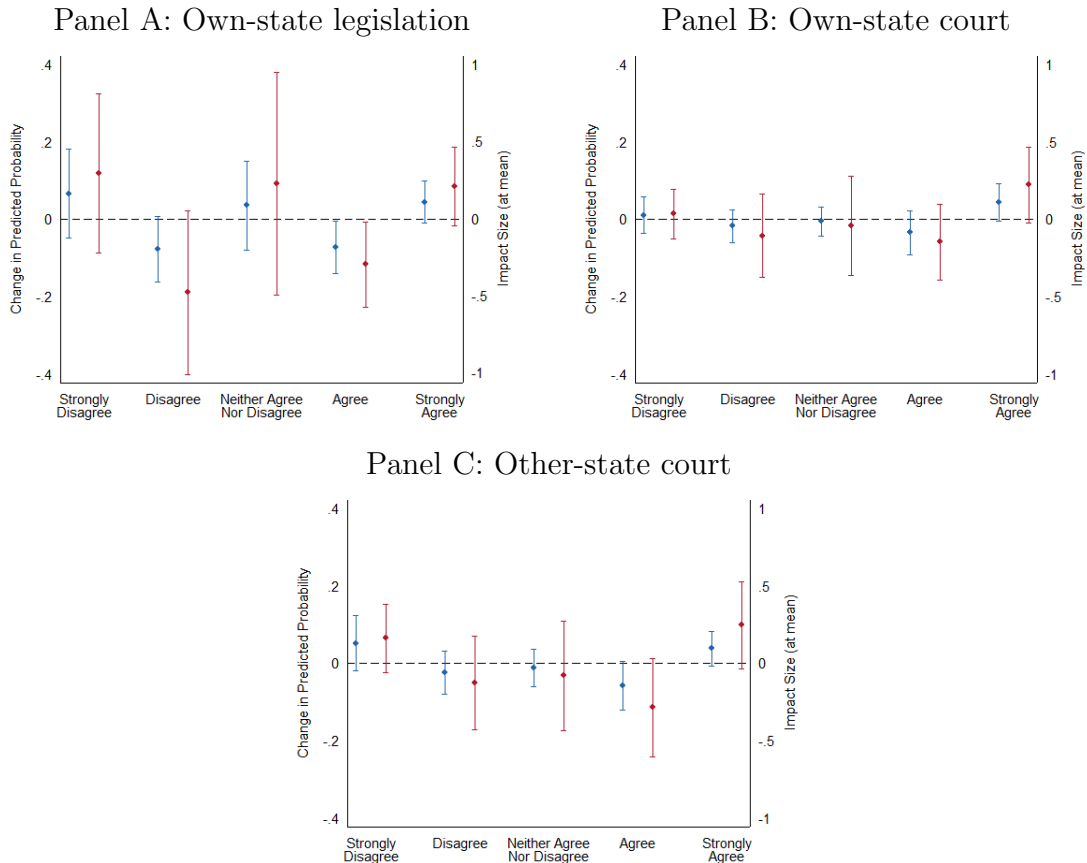
¹⁵Including other types of attempts does not meaningfully affect the results.

Figure 14. Changes in belief regarding same-sex marriage induced by same-sex-marriage legalization, by average state belief in 2004



Notes: On the left of each category (in blue), we plot the change in predicted probability with a 95% confidence interval from the multinomial logit model of the five-fold categorical response to “Do you agree or disagree? Homosexual couples should have the right to marry one another.” Point estimates are interpreted as the average change in the predicted probability of a respondent selecting each of the five different opinion choices due to legalization. On the right of each category (in red), we plot the implied impact—the percent change in likelihood of occupying that category, evaluated at its (control) mean.

Figure 15. Changes in belief regarding same-sex marriage induced by same-sex-marriage legalization, by method of legalization



Notes: On the left of each category (in blue), we plot the change in predicted probability with a 95% confidence interval from the multinomial logit model of the five-fold categorical response to “Do you agree or disagree? Homosexual couples should have the right to marry one another.” Point estimates are interpreted as the average change in the predicted probability of a respondent selecting each of the five different opinion choices due to legalization. On the right of each category (in red), we plot the implied impact—the percent change in likelihood of occupying that category, evaluated at its (control) mean.

to successful attempts to ban same-sex marriage in order to keep the environment as identical as possible to the successful attempts to legalize same-sex marriage. We report the list of failed attempts, along with the corresponding success in Table 10.

Table 10. Failed attempts to legalize same-sex marriage

| State | Failed attempt | Successful attempt |
|-----------|----------------|--------------------|
| Indiana | Jan 2005 | Oct 2014 |
| New York | Jul 2006 | Jun 2011 |
| Nebraska | Jul 2006 | Jun 2015 |
| Maryland | Sep 2007 | Jan 2013 |
| Texas | Aug 2010 | Jun 2015 |
| Hawaii | Aug 2012 | Dec 2013 |
| Nevada | Nov 2012 | Oct 2014 |
| Louisiana | Sep 2014 | Jun 2015 |
| Kentucky | Nov 2014 | Jun 2015 |
| Michigan | Nov 2014 | Jun 2015 |
| Ohio | Nov 2014 | Jun 2015 |
| Tennessee | Nov 2014 | Jun 2015 |

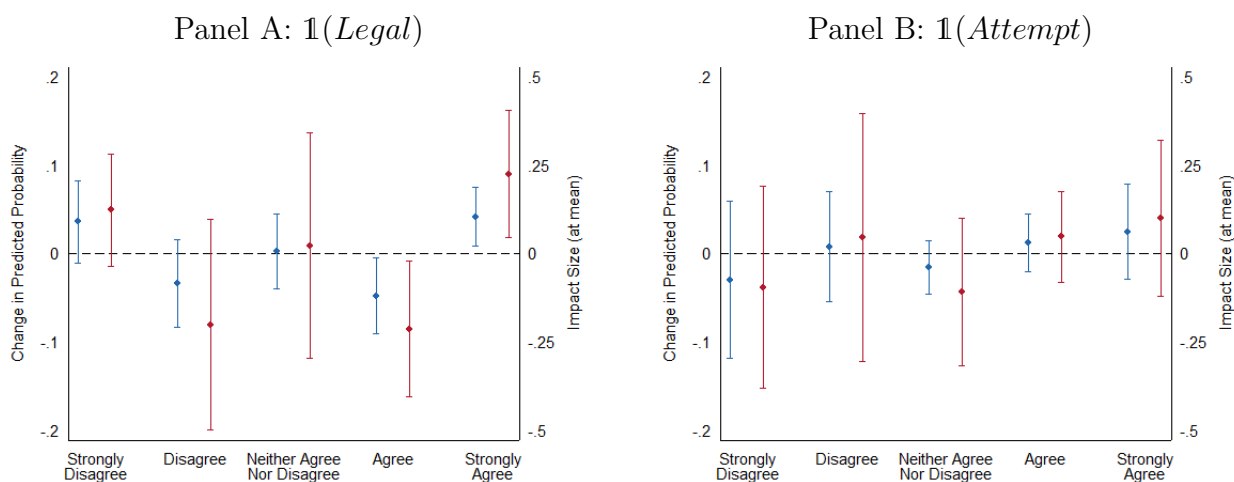
Notes: We present the month and state where there was a failed attempt to legalize same-sex marriage through the courts.

In Figure 16, we present the estimated coefficients of a multinomial logit model where we add in an indicator variable capturing failed attempts in the same spirit as our legal indicator which captured successful attempts.¹⁶ We report the predicted probabilities for both the original legal variable as well as our new attempt variable. In Panel B, we see a markedly different pattern to the estimated coefficients on attempts than that of legality in Panel A, which are virtually unchanged from the model in which the attempt variable is not included. The

¹⁶We define $\mathbb{1}(Attempt)$ as equal to one if there had been a failed attempt to legalize same-sex marriage previous to respondent i responding to the survey in state s in year y .

point estimates on $\mathbb{1}(Attempt)$ suggest a pattern of increasing mass at “strongly agree” and “agree” and decreasing mass in the middle category of “neither agree nor disagree” and in “strongly disagree.” The movement toward extreme beliefs seems to be tied to successful attempts, not simply because of the surrounding political and social environment. It is the legalization itself that matters.

Figure 16. Impact of failed attempts to legalize same-sex marriage



Notes: On the left of each category (in blue), we plot the change in predicted probability with a 95% confidence interval from the multinomial logit model of the five-fold categorical response to “Do you agree or disagree? Homosexual couples should have the right to marry one another.” Point estimates are interpreted as the average change in the predicted probability of a respondent selecting each of the five different opinion choices due to legalization. On the right of each category (in red), we plot the implied impact—the percent change in likelihood of occupying that category, evaluated at its (control) mean.

This pattern suggests that movement in beliefs comes from those respondents that “lost.” When the attempt is successful, there is a backlash among those who disagree, leading to more-strongly held beliefs. When the attempt is not successful, there is no corresponding strengthening of beliefs among those disagreeing. Thus, especially for those who strengthen their opposition, the outcome of the attempt matters.

Direct Measures of Polarization

In Table 11, we consider direct measures of polarization around changes in legalization. As a potentially intuitive approach, we first consider changes in state-level standard deviations of beliefs on same-sex marriage induced by legalization in Column (1). We aggregate the data to a state-year observation and run our preferred specification using the standard deviation of beliefs in a given state-year as the dependent variable. Although not statistically significant, we see a 15.1-percent increase in the within-state standard deviation of responses.

In Column (2) of Table 11, we ask directly how the Esteban-Ray index of polarization moves with legalization. This index incorporates both the ordered nature of beliefs and the mass of respondents at each belief choice. In this context, the index ranges between 0 and 1.¹⁷ It is here that we find a significant increase in polarization directly—with the advent of legalization, the ER polarization index increases by 0.065, corresponding with an increase in state-level polarization of 17 percent. This is a large effect, with legalization explaining roughly 65-percent of a standard deviation in state-level polarization. Having identified that this is derived from movement toward more-strongly held belief, we are inclined to suggest that the smoothness of any transition toward support over time may be slowed by this underlying empirical regularity.

¹⁷Esteban and Ray (1994) define polarization as

$$ER = \sum_{i=1}^5 \sum_{j=1}^5 p_i^2 \cdot p_j |b_i - b_j|$$

where p_i represents the proportion of the sample with belief b_i . In our case, i and j capture the five dimensions of categorical response (i.e., “strongly disagree” through “strongly agree”). This measure is maximized at 1 when 50 percent of the sample is at each extreme and the measure is minimized at 0 when 100 percent of the sample is at one belief choice.

Table 11. Polarization around same-sex-marriage legalization

| | State-year standard deviation (1) | Esteban-Ray polarization index ^a (2) |
|----------------------------|--|--|
| $\mathbb{1}(\text{Legal})$ | 0.031 (0.161) | 0.065*** (0.031) |
| Observations | 291 | 291 |
| R^2 | 0.413 | 0.405 |
| Control mean (dep var) | 1.43 | 0.38 |
| Control sd (dep var) | 0.23 | .10 |
| Impact size (at mean) | 0.022 | 0.171 |
| Effect size (at mean) | 0.135 | 0.648 |

Notes: The data is collapsed to a state-year observation. Specifications include state and year fixed effects, state-specific linear trends and individual-level controls. Standard errors in parentheses, allowing for clustering at the state level. *** significant at 1%; ** significant at 5%; * significant at 10%.

^a Following Esteban and Ray (1994), we define polarization for each state-year as $ER = \frac{\sum_{i=1}^5 \sum_{j=1}^5 p_i^2 \cdot p_j |b_i - b_j|}{\sum_{i=1}^5 p_i^2}$ where p_i represents the proportion of the sample with belief b_i . In our case, i and j capture the five categorical response (i.e., “strongly disagree” through “strongly agree”), and the index, ranges from the complete mass being at any single category (i.e., $ER = 0$) to half the mass at “Strongly disagree” and half the mass at “Strongly agree” (i.e., $ER = 1$).

Longevity of Changes in Beliefs and Polarization

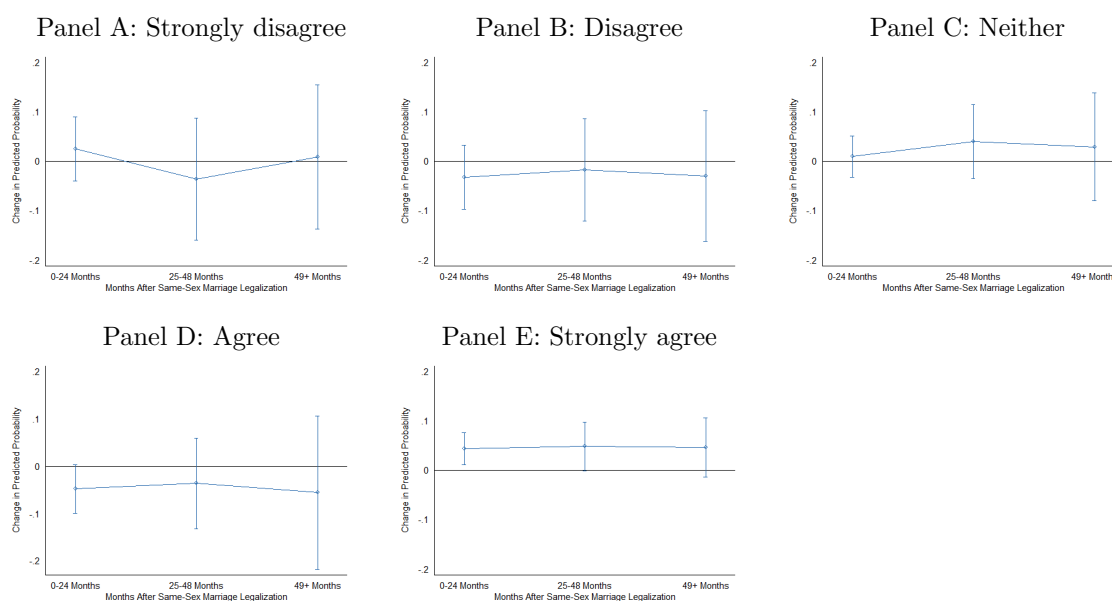
We have shown that beliefs and state-level polarization change systematically with the legalization of same-sex marriage. We now explore preliminary evidence of whether these changes are temporary or permanent. Unfortunately, the timing of the laws surrounding same-sex marriage limit our ability to perform this exercise as much of the policy variation occurs near the end of the time series available in our data.

However, we can provide preliminary evidence by relaxing the assumption that the effect of legalization is constant, no matter how much time has passed since the legalization. We stratify the independent variable of interest $\mathbb{1}(\textit{Legal})$ into three indicator variables based on the time between legalization and each respondent's survey date: 0-24 months, 25-48 months, and 49+ months after legalization.

We report the results of the multinomial logit specification with these three independent variables in Figure 17. These estimates generally lack power because of the limited number of states that legalized same-sex marriage early enough to contribute to all three time periods.¹⁸ In Panel A, the point estimates suggest that the initial increase in probability of selecting “strongly disagree” does not continue over time and is simply a one-time shock immediately after the law is passed. That does not seem to be the case for panels B, D, and E—the estimated coefficients remain relatively constant over time, representative of a consistent movement toward support induced by legalization.

¹⁸By March of 2016, the most-recent GSS survey, only 17 states had legalized same-sex marriage more than two years previously, and only seven had legalized same-sex marriage more than four years previously.

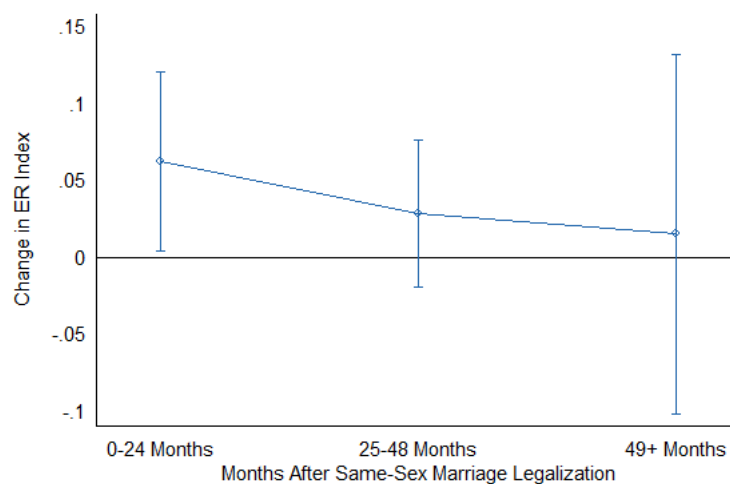
Figure 17. Permanence of changes in belief induced by the legalization of same-sex marriage



Notes: We plot the change in predicted probability over time since legalization with a 95% confidence interval. Estimates come from the multinomial logit model of the five-fold categorical response to “Do you agree or disagree? Homosexual couples should have the right to marry one another.” Point estimates are interpreted as the average change in the predicted probability of a respondent selecting each of the five different opinion choices due to legalization.

Turning to polarization, we again relax the assumption that the effect of legalization does not depend on the time since passage. In Figure 18, we report the changes in state-level polarization, measured by the Esteban-Ray polarization index. As time passes, the estimated coefficient attenuates to near 0, suggesting that the effect on state-level polarization is only temporary.

Figure 18. Permanence of changes in state-level polarization induced by same-sex marriage legalization



Notes: We plot the change in state-level polarization over time since legalization with a 95% confidence interval. Point estimates are interpreted as changes in the Esteban-Ray polarization index. The control mean for this index is 0.38.

Conclusion

Over some 25 years, efforts have been made to legalize or ban same-sex marriage, culminating in the 2015 U.S. Supreme Court decision that legalized same-sex marriage nationwide. We exploit the timing of legalization to retrieve an estimate of the effect of legalization on individual belief and state-level polarization.

With the exception of some suggestive evidence of weakening support associated with the imposition of a binding precedent, we find little evidence of legalization inducing people to switch between support and opposition of same-sex marriage. However, we do find movement toward strong support and strong opposition, seemingly coming from weaker support and weaker opposition. This effect is also somewhat concentrated demographically—white respondents are tending to move toward strong support while Hispanic respondents tend to move toward strong opposition. In addition, we see evidence of divergence between educational groups.

Using direct measures of how different within-state belief is at a moment in time, we find increases in state-level polarization that are on the order of roughly 65-percent of a standard deviation.

The impact of changing beliefs may extend well beyond the implications discussed in this paper. More extreme beliefs, even if temporary, could affect public policy, with fewer policies being enacted overall, or with increased turnover of policy as new parties are elected. It may also affect peer networks, with increased polarization leading to differential selection into friend groups, which in turn leads to decreased diversity of belief among networks. It may even affect crime, changing the incidence of hate crimes.

Although the policy prescription of these results is not to avoid making public policy altogether, it informs policy-makers on another potential cost of passing policy. Changing beliefs and increased polarization are a previously-unexamined consequence of policy, and are important to consider in future legislation.

CHAPTER IV
CONFIDENCE AND CONTRITION: IS CHEATING INTERNALIZED IN
PERFORMANCE ASSESSMENTS?

This chapter was co-authored with Glen Waddell and Michael Kuhn.

Introduction

In a summary of surveys of undergraduate academic dishonesty between 1962 and 2010, McCabe, Butterfield, and Trevino (2012) finds that academic dishonesty was, in essence, routine: in the 1999/2000 wave, cheating behavior ranged from 8 percent of the sample “turning in papers done entirely or in part by other students,” to 56 percent “getting questions or answers from someone who has already taken the same exam.” Of the nine cheating behaviors surveyed, 83 percent of the sample reported engaging in at least one. In a 2009 Harris poll, 28 percent of American workers also admitted that they would “act immorally” to keep their job (Park, 2009). Even with layers of safeguards in place, 300 public companies in the U.S. were found to have committed \$120 billion in financial misstatement and misappropriation between 1998-2007 (Beasley, Carcello, Hermanson, and Neal, 2010). Dishonest behavior is widespread.

Becker (1968) is among the earliest studies on the determinants of cheating, and forms the view that people will engage in dishonest behavior when the benefits of that behavior outweigh the costs. This precipitated literatures on the pecuniary and probabilistic determinants of cheating in the workplace: executive-compensation packages (Burns and Kedia, 2006; Efendi, Srivastava, and Swanson, 2007), competition and tournament incentives (Berentsen, 2002; Bunn,

Caudill, and Gropper, 1992; Gilpatric, 2011; Jacob and Levitt, 2003; Kräkel, 2007; Schwierien and Weichselbaumer, 2010), decreased deterrence (Curry and Mongrain, 2009), decreased monitoring (Kerkvliet and Sigmund, 1999; Nagin, Rebitzer, Sanders, and Taylor, 2002), team environments (Conrads, Irlenbusch, Rilke, and Walkowitz, 2013), and productivity (Gill, Prowse, and Vlassopoulos, 2013). Recent experimental work suggests that cheating may be a function of individual character—more responsive to across-individual variation in social- and self-image than to within-individual variation in motive, means, and opportunity (Fischbacher and Föllmi-Heusi, 2013; Gneezy, Kajackaite, and Sobel, 2018; Mazar, Amir, and Ariely, 2008; Weisel and Shalvi, 2015). (See Abeler, Nosenzo, and Raymond (2016) for a broad review.)

While the determinants of cheating are well-studied, the consequences of cheating have for the most part been overlooked. Chance, Norton, Gino, and Ariely (2011) finds that those who are given answers to an experimental test while taking the exam interpret their elevated performance as a sign of intelligence. Robert and Arnab (2013) exploits experimental variation in peer dishonesty to identify increased dishonesty, suggesting that dishonesty in one participant induces more dishonesty in others. Gneezy, Imas, and Madarasz (2014) finds that immoral behavior leads to feelings of guilt, and thereby to increased charitable donations. In this paper, we consider an additional consequence of dishonest behavior by examining the link between cheating and confidence.

The connection between cheating and confidence is natural, we believe. For example, consider two employees who report having achieved the same level of productivity—Employee A having done so honestly, but Employee B having inflated his productivity. While they might both receive performance pay for their

work, they subsequently face different signal-extraction problems in evaluating their own workplace ability. That is, while Employee A can attribute the performance pay to her productivity quite easily, Employee B must consider that his pay may have been induced at the margin by dishonesty. By extension, where such a signal extraction is not fully executed, one easily imagines that Employee B may subsequently impart bias to his self assessment—thus, we imagine, a role for dishonesty in the endogenous development of overconfidence.

There is also reason to believe that individuals are actually quite poor at tasks similar to that faced by Employee B. Haggag and Pope (2016) identifies “attribution biases,” in which consumers are unable to separate the state of the world in which consumption occurs from the state-independent utility of consumption. Enke and Zimmerman (2018) identifies “correlation neglect,” in which forecasters treat correlated signals about future economic growth as independent. Enke (2018) identifies “selection neglect,” in which individuals fail to recognize censoring in the news-generation process that informs their opinions. A large literature identifies “outcome bias,” in which outputs that are produced by both effort and luck are over-attributed to effort (Brownback and Kuhn, 2018; Charness and Levine, 2007; Cushman, Dreber, Wang, and Costa, 2009; de Oliveira, Smith, and Spraggon, 2017; Gurdal, Miller, and Rustichini, 2013; Rubin and Sheremeta, 2015; Sarsons, 2017).

In order for Employee B to accurately assess his ability, he must recall and accept the fact that he engaged in immoral behavior in the past. Yet, studies on motivated reasoning and “moral wiggle room” suggest that individuals are less responsive to negative information than to positive (Eil and Rao, 2011), and will

pay costs to avoid this information (Dana, Weber, and Kuang, 2007). Thus, we expect full attribution in this environment to be particularly difficult.

We also examine the role that gender might play in ability assessments. Beginning with Niederle and Vesterlund (2007), a sizable experimental literature has analyzed why there appears to be differential selection into competition and cooperation between men and women. Recent work suggests that overconfidence may play a larger role in that difference than previously thought. Veldhuizen (2017) finds that 48 percent of the competition gap can be explained by differences in confidence.¹ Kuhn and Villeval (2015) finds that under-confidence in women drives greater selection into team cooperation than is evident in men. Could male overconfidence come in part from differences in dishonest behavior? Dreber and Johannesson (2008), Erat and Gneezy (2012), and Kajackaite and Gneezy (2017) all find that men lie more than women, so long as they can personally benefit from the lie. Immoral behavior is often risky, of course, and Charness and Gneezy (2012) and Borghans, Golsteyn, Heckman, and Meijers (2009) find that women are more risk-averse than men. Thus, we anticipate that dishonesty may well be an important mechanism by which overconfidence propagates in men.

In this paper, we report on the results of two online experiments designed to measure the relationship between cheating and confidence. In both experiments, we first identify cheating at the individual level. Roughly half of subjects cheat, making the online implementation of our task more effective at drawing out cheating than previous laboratory efforts (Abeler et al., 2016). We then measure subsequent confidence with a Gneezy and Potters (1997) instrument, in which subjects invest in their own future performance. This measure of confidence is

¹An additional 37 percent of the gap can be explained by the interaction between confidence and risk preferences.

referred to as “revealed” confidence. We also measure confidence with an un-incentivized “cheap-talk” assessment. This measure is referred to as “stated” confidence.

Cheaters are more confident, by both measures. This means that a manager attempting to infer productivity based on stated and revealed signals of employee confidence may instead identify propensity to cheat. Additionally, we find that this assessment problem is different for men and women. For women, high confidence is indicative of marginal cheating—productive individuals nudge their performance report just over a threshold that earns them a bonus. For men, high confidence is indicative of maximum cheating—individuals report perfect performance without exerting effort.

We find important differences between our two methods of eliciting confidence. The association between cheating and confidence is much larger for stated confidence than for revealed confidence, and the expected gender gap exists *only in stated confidence*. In Experiment 1, we randomly vary the *ex-ante* incentive to cheat, but find little impact on cheating—this is consistent with the experimental literature noted above. In Experiment 2, we randomly vary the *ex-post* reward from cheating, and find that stated confidence is increased by the reward from cheating (driven by men), but that revealed confidence, if anything, is *reduced* by the reward from cheating (driven by women). These results suggest that cheating is not fully integrated out of self assessments, and that revealed confidence performance assessments may limit both unjust confidence from bad actors and the gender gap among them.

Experiment 1

Design

We used Amazon’s Mechanical Turk (MTurk) for our study—in doing so, we expect the social distance and monitoring difficulty inherent in interacting remotely and anonymously to be conducive to cheating. All participants received a \$1 participation fee, and the median time for subjects to complete the study is roughly 17 minutes.

The study is built around subjects’ performance on a hidden-object task. Participants are shown a picture (see Figure A2) and asked to find hidden items from a list of twelve potential objects. They are told that if they find nine of the twelve objects within the four minutes they are given, they will earn a reward. Likewise, they are told that if they find all twelve objects they will earn an even larger reward. However, only eight of the twelve objects on the list appear in the picture. Subjects self-report the number of objects they find and can thereby misrepresent their true performance on the task. Below the picture, subjects can tick off each object, allowing us to track which objects they reported finding. We can therefore distinguish cheaters who “barely” cheated by nudging themselves over the payment threshold (i.e., from eight to nine) from cheaters who report finding several missing objects. Subjects can advance from the picture at any time. 574 subjects completed this experiment, 350 (61.0 percent) of whom reported finding the correct eight objects in the picture.

The reward for ‘finding’ nine objects was randomly assigned (\$0.10, \$0.50, \$1.50, or \$2.75), but was always made known to the subject prior to performing the task (this will be a key difference between Experiment 1 and 2). We varied the

reward in order to create an instrumental variable for cheating, which we discuss in depth in Section IV. The marginal reward for “finding” twelve objects was also varied (\$0.40, \$0.50, or \$2.25).

After Task 1, subjects move on to a second object search task, Task 2. There are a number of key differences between Task 1 and Task 2. First, subjects are told that the new picture (see Figure A3) in Task 2 will be overlaid with a grid and they will have to report the grid location of each object that they find. In this way, we signal to participants that cheating is not possible on this task. Second, subjects learn that they will not be able to advance past Task 2 before time expires.

Third, subjects learn that all earnings from the second object-search task—both successful investments and uninvested endowments—are given to the Make-a-Wish[®] Foundation. Separating Task 2 from personal financial gain in Task 1 is designed to limit the impact of income effects on investment choices in the second.

Fourth, we describe a Gneezy and Potters (1997) instrument to subjects in which they can invest in their own performance on Task 2. Each participant is given an endowment of \$2 that they can invest in any cent-increment. If they are successful—which again means finding nine of twelve objects—their investment is tripled. If they are unsuccessful their investment is lost. Whatever they don’t invest is kept. Assuming that subjects would like to maximize their expected donation to the Make-a-Wish[®] Foundation, their investment is a revealed measure of their confidence. Mean investment is \$1.07 (S.D. = \$0.75), and we use standardized investment as a dependent variable. Importantly, the Gneezy and Potters (1997) investment decision is designed to elicit risk preference and, as such, our subjects’ investment should also be related to risk preference. Therefore, we elicit stated

risk preference and use it as a control variable.² All subjects have to answer three comprehension questions correctly before proceeding past the description of Task 2.

Following Task 2, subjects complete a brief survey.³ In the survey, we ask subjects to provide a statement of their confidence: “How well do you believe you would perform on similar hidden object tasks in the future? Please choose a value from 0 to 10.”

Results

We break the results from Experiment 1 into three sections. First, we examine the frequency and nature of cheating. Second, we examine the relationship between cheating and confidence. Third, we exploit the varying financial incentive to cheat in the first task as an instrumental variable for the impact of cheating on confidence.

Prevalence of Cheating

Unlike many experimental cheating paradigms, we find considerable cheating in our mTurk study. While it was only possible to find eight objects, 48.8 percent of subjects reported finding more than eight. Subjects also responded to the incentive to cheat fully, and claim that all objects were found—18.3 percent of subjects reported finding twelve objects. The distribution of reports among cheaters is bimodal. In an environment where there were rewards for nine and for twelve objects to be found, very few cheaters reported finding ten or eleven

²Subjects respond to, “In general, are you a person who likes to take risks or do you try and avoid taking risks? Please choose a value from 0 to 10.” In this query, zero corresponds to “I am not at-all willing to take risks,” and ten corresponds to “I am very willing to take risks.” The average response is 4.74 (S.D. = 2.48), and we use standardized risk as a control variable.

³Full survey is available upon request.

objects; there is a clear collapsing on the reporting of nine or twelve objects. This distinction, within the set of those who cheated, increases our confidence that our study exhibited its intended moral framework. Clearly, if the difference between cheaters and non-cheaters was purely a realization of the ability to cheat, we would expect to see no reports of nine objects. Conditional on finding all eight objects that were in the picture, cheating rates are much higher—of subjects who correctly indicated finding the eight feasible objects, 73.1 percent reported finding more than eight objects.

The partition of the sample into cheaters and non-cheaters is endogenous, as we would expect it to be in the workplace. However, the focus of this study is specifically on cheating, and barring exogenous variation in cheating, we look for observable determinants of cheating that can mitigate omitted variable bias when examining the relationship between cheating and confidence. We anticipated that risk preference, in particular, would be a strong predictor of cheating. In Column (1) of Table 12, we regress an indicator variable for cheating on our measure of risk tolerance. We find that a standard-deviation increase in risk tolerance is associated with a 3.8-percent increase in the likelihood of cheating. However, when we control for performance, or limit the sample to those who find the eight possible objects, this relationship is attenuated. In Column (2), we add performance, impulsivity, gender, employment status, and education as other personal characteristics that could correlate with cheating, as well as fixed effects for hour of day, day of week, week of month and month of year.⁴ In so doing, the coefficient on risk is reduced

⁴Performance is measured by the number of possible objects found. Impulsivity is measured using the standardized number of incorrect questions from the Cognitive Reflection Test (CRT, Frederick (2005)). Employed is an indicator variable equal to one for those employed either full or part time. College educated is an indicator variable equal to one for those who have obtained at least a bachelor's degree.

by nearly 50 percent. In columns (3) and (4), we limit the sample to subjects who found the eight possible objects—the risk coefficient remains small and insignificant. In Column (4), being employed and having a college degree are both predictive of cheating. We proceed by using these control variables, and showing results both with and without the limitation to subjects finding the eight real objects.

Table 12. Determinants of Cheating

| | Full Sample | | Restricted to those who found all possible objects | |
|--------------------|-------------|----------|--|---------|
| | (1) | (2) | (3) | (4) |
| Risk Tolerance | 0.038* | 0.021 | 0.022 | 0.011 |
| | (0.021) | (0.020) | (0.024) | (0.025) |
| Performance | | 0.160*** | | |
| | | (0.019) | | |
| Impulsivity | | -0.007 | | -0.005 |
| | | (0.020) | | (0.025) |
| 1(Female) | | 0.004 | | -0.002 |
| | | (0.042) | | (0.052) |
| 1(Employed) | | 0.113** | | 0.115* |
| | | (0.047) | | (0.066) |
| 1(CollegeEducated) | | 0.063 | | 0.106** |
| | | (0.039) | | (0.049) |
| Constant | 0.489 | | 0.731 | |
| | (0.021) | | (0.024) | |
| Time & Date FEs | N | Y | N | Y |
| Observations | 574 | 574 | 350 | 350 |

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Dependent variable is an indicator for cheating, $\mathbf{1}(\text{Objects found} \geq 9)$.

In Figure 19 we show the full distribution of the number of reported objects, separately by gender. When we consider the overall group of cheaters, we find no

difference in the likelihood of cheating between men and women—48.3 percent of women cheat in our environment, and 49.3 percent of men cheat (difference: $p = 0.83$, two-tail t -test). However, men are more likely than women to report that they have found twelve objects—22.8 percent of men, and only 14.2 percent of women report finding twelve (difference: $p = 0.01$, two-tail t -test). As individuals who report twelve have checked all of the available boxes, their stated performance need not inform their true performance on the task at all. In this sense, men are significantly more likely to *substitute* cheating for effort. On the other hand, among the individuals who report fewer than twelve objects, we can observe whether the eight correct objects are a subset of the report. In this case, women are more likely to *complement* effort with cheating, wherein they seemingly first exert effort and then nudge themselves over the threshold, although this difference is not statistically significant at conventional levels (28.8 percent of women versus 23.2 percent of men, $p = 0.12$, two-tail t -test). Indeed, individuals who cheat maximally spend an average of 50 seconds less time on the task than individuals who cheat marginally ($p < 0.01$, two-tail t -test).⁵ Marginal cheaters spend *significantly more* time on the task than non-cheaters.⁶


Cheating and Confidence

As described in Section IV, we elicited both stated measures of subjects' confidence, and revealed measures in Task 2. We first consider the problem faced by a manager trying to infer the productivity of an employee following an assessment of their confidence. In Table 13, we show the fraction of the sample

⁵300 seconds is the time limit, and the typical “marginal cheater” spends 280 seconds on the task.

⁶Marginal cheaters: 280 seconds, non-cheaters: 265 seconds, difference: $p = 0.01$, two-tail t -test.

Figure 19. Objects Reported by Gender, Experiment 1



Notes: We plot the histogram of respondents who self-reported finding a given number of objects in experiment 1, by gender.

who cheated, separately by gender and confidence-assessment tercile. In Column (1), we show the fraction reporting 9-to-11 objects, and in Column (2) we show the fraction reporting all-twelve objects for the maximum reward. We find that the manager's inference should depend on employee gender when the employee's confidence assessment is high. Specifically, 43 percent of female workers who state that they are highly confident are marginal cheaters, 35 percent are non-cheaters, and 23 percent are maximum cheaters. These are similar in revealed measures of confidence—42 percent, 37 percent, and 22 percent, respectively. Highly-confident men are more likely to be maximum cheaters than highly-confident women: 40 percent (31 percent) of male workers who either state or reveal that they are highly confident are also maximum cheaters. They are also less likely to be marginal cheaters: 31 percent of male workers both state and reveal that they are highly confident are also marginal cheaters. Especially when making inference based on cheap talk, managers should be aware that reporting high levels of confidence implies different things about the productivity and cheating behavior of male versus female employees.

Moving away from the manager's association problem, perhaps without true performance and other variables to condition on, we present regression estimates of the relationship between cheating and confidence. We first analyze stated confidence, based on how well subjects expect to do on a similar task in the future. In Table 14, we regress pooled and gender-stratified estimates of participants' stated levels of confidence in their future performance on a similar task on an indicator variable for whether the individual cheated on the first task. In columns (1) through (3), we present estimates for the full sample. We control for risk tolerance, true performance on the task, impulsivity, employment status,

Table 13. Confidence and Propensity to Cheat

| | | Cheated to 9-11 (1) | Cheated to 12 (2) |
|-------------------------------------|-------|------------------------|----------------------|
| <i>Panel A: Stated Confidence</i> | | | |
| Low Confidence | Men | 0.043 | 0.043 |
| | Women | 0.122 | 0.061 |
| | | $p = 0.16$ | $p = 0.68$ |
| Medium Confidence | Men | 0.301 | 0.086 |
| | Women | 0.316 | 0.094 |
| | | $p = 0.81$ | $p = 0.84$ |
| High Confidence | Men | 0.311 | 0.402 |
| | Women | 0.426 | 0.228 |
| | | $p = 0.05$ | $p < 0.01$ |
| <i>Panel B: Revealed Confidence</i> | | | |
| Low Confidence | Men | 0.154 | 0.209 |
| | Women | 0.225 | 0.135 |
| | | $p = 0.22$ | $p = 0.19$ |
| Medium Confidence | Men | 0.329 | 0.151 |
| | Women | 0.345 | 0.106 |
| | | $p = 0.82$ | $p = 0.37$ |
| High Confidence | Men | 0.306 | 0.306 |
| | Women | 0.420 | 0.210 |
| | | $p = 0.09$ | $p = 0.12$ |

Notes: We report the proportion of men and women within confidence terciles who cheated to 9-11 and who cheated to 12. Variables are standardized stated confidence in future performance and standardized investment in future performance.

and education, as well as time and date fixed effects. Being a cheater is associated with having three-fifths of a standard-deviation higher stated confidence. The effect is considerably larger for men than for women, although this difference is not significant at conventional levels ($p = 0.12$). In columns (4) through (6), we limit the sample to those who found all eight possible objects, and obtain similar estimates.

Table 14. Cheating and Stated Confidence

| | Full Sample | | | Restricted to those who found all possible objects | | |
|--|---------------------|----------------------------------|---------------------|--|----------------------------------|---------------------|
| | All (1) | Women (2) | Men (3) | All (4) | Women (5) | Men (6) |
| $\mathbb{1}(\text{Cheater})$ | 0.585*** (0.094) | 0.461*** (0.115) | 0.742*** (0.151) | 0.634*** (0.116) | 0.567*** (0.137) | 0.714*** (0.197) |
| Risk Tolerance | 0.195*** (0.040) | 0.209*** (0.059) | 0.195*** (0.063) | 0.178*** (0.047) | 0.157** (0.069) | 0.171** (0.079) |
| Performance | 0.125*** (0.039) | 0.116** (0.049) | 0.127** (0.056) | | | |
| Impulsivity | 0.041 (0.040) | 0.006 (0.053) | 0.065 (0.059) | 0.049 (0.046) | 0.063 (0.062) | 0.036 (0.072) |
| $\mathbb{1}(\text{Employed})$ | 0.087 (0.088) | 0.019 (0.114) | 0.235* (0.128) | 0.037 (0.107) | -0.063 (0.135) | 0.202 (0.174) |
| $\mathbb{1}(\text{College Educated})$ | -0.042 (0.076) | -0.023 (0.106) | -0.097 (0.114) | -0.006 (0.091) | 0.067 (0.125) | -0.117 (0.143) |
| $H_0 : \text{Female} = \text{Male}$ (Cheater) | | $\chi^2(1) = 2.41$ $p = 0.12$ | | | $\chi^2(1) = 0.44$ $p = 0.51$ | |
| Time & Date FEs | Y | Y | Y | Y | Y | Y |
| Observations | 574 | 302 | 272 | 350 | 183 | 167 |

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.
Dependent variable is standardized stated confidence in future performance.

Next, we consider the variation in revealed confidence by analyzing subjects' decisions on Task 2, the task in which subjects made an investment decision. In

Column (1) of Table 15, we find that cheaters are willing to invest significantly more money in their future *verified* performance—cheaters invest about one-third of a standard deviation more than do non-cheaters. Unlike in stated preference, in the full sample there is little evidence of a gender difference in the relationship between cheating and confidence in revealed preferences, as shown in columns (2) and (3). Limiting the sample to those who found all possible objects, in columns (4) through (6), reveals an effect that is larger for men than women, but the difference is not statistically significant.

Table 15. Cheating and Revealed Confidence

| | Full Sample | | | Restricted to those who found all possible objects | | |
|--|---------------------|----------------------------------|---------------------|--|----------------------------------|---------------------|
| | All (1) | Female (2) | Male (3) | All (4) | Female (5) | Male (6) |
| $\mathbb{1}(\text{Cheater})$ | 0.320*** (0.086) | 0.330*** (0.113) | 0.360** (0.138) | 0.328*** (0.109) | 0.249 (0.156) | 0.445*** (0.161) |
| Risk Tolerance | 0.312*** (0.040) | 0.293*** (0.060) | 0.351*** (0.058) | 0.299*** (0.050) | 0.311*** (0.076) | 0.331*** (0.073) |
| Performance | 0.005 (0.033) | 0.033 (0.050) | -0.024 (0.048) | | | |
| Impulsivity | -0.065* (0.039) | -0.076 (0.053) | -0.078 (0.061) | -0.051 (0.052) | -0.092 (0.076) | -0.071 (0.080) |
| $\mathbb{1}(\text{Employed})$ | -0.067 (0.093) | -0.108 (0.113) | 0.045 (0.165) | -0.046 (0.117) | -0.018 (0.141) | -0.033 (0.224) |
| $\mathbb{1}(\text{College Educated})$ | -0.086 (0.080) | -0.277** (0.108) | 0.122 (0.125) | -0.064 (0.102) | -0.291** (0.146) | 0.231 (0.158) |
| $H_0 : \text{Female} = \text{Male}$ (Cheater) | | $\chi^2(1) = 0.03$ $p = 0.63$ | | | $\chi^2(1) = 0.88$ $p = 0.35$ | |
| Time & Date FEs | Y | Y | Y | Y | Y | Y |
| Observations | 574 | 302 | 272 | 350 | 183 | 167 |

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Dependent variable is standardized investment in future performance.

We also test whether this relationship between cheating and confidence differs for standardized stated and revealed preference, in Table 16. We report the difference between the point estimates from the stated- and revealed-preference models and also report the statistical significance from a test of the null that the coefficients are equal across models. Notably, these differences are positive across all columns, suggesting a larger positive association between cheating and confidence in stated preference than exists in revealed preference. The gender-pooled difference and the difference for men (in columns 1 and 3) are statistically significant at the five-percent level. While the difference is larger for men than women, it is not statistically different, and in addition, this gap goes away when we restrict the sample to those finding all eight objects.

Table 16. Testing Stated vs. Revealed Preference (Table 14 vs. Table 15)

| | Full Sample | | | Restricted to those who found all possible objects | | |
|---|-------------|----------------------------------|------------|--|----------------------------------|------------|
| | All (1) | Women (2) | Men (3) | All (4) | Women (5) | Men (6) |
| Difference in $\mathbb{1}(\text{Cheater})$ Coefficients (Stated – Revealed) | 0.265** | 0.131 | 0.382** | 0.306** | 0.318 | 0.270 |
| $H_0 : \text{Female} = \text{Male}$ | | $\chi^2(1) = 1.21$ $p = 0.27$ | | | $\chi^2(1) = 0.03$ $p = 0.87$ | |
| Observations | 574 | 302 | 272 | 350 | 183 | 167 |

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Dependent variable is standardized investment in future performance.

In the case of revealed confidence through investment in Task 2, we find substantial bunching at no investment, investing exactly half of the endowment

and full investment. As such, we also examine the relationship between cheating and the full distribution of investment. Figure 20, shows the distribution of investment by the decision to cheat for the full sample. In Panel A, we show the unadjusted-investment variable, and in Panel B we adjust for risk preference by using the residual investment from a regression of investment on risk tolerance. The distributions for cheaters dominate the distributions for non-cheaters.⁷ The same is true when we treat women and men separately. Within both men and women, we reject the equality of distributions by cheating for both unadjusted and risk-adjusted investment, as well.⁸

Induced Cheating and Confidence


Within Experiment 1, we attempted to induce cheating by experimentally varying the marginal reward associated with cheating. In order for this incentive to serve as a valid instrumental variable (IV) for cheating, it must predict cheating and have no direct effect on stated or revealed confidence. The exclusion restriction we satisfy by experimental design.⁹ When we regress an indicator for cheating on the marginal incentive to cheat, we find a coefficient of 0.028 ($p = 0.17$, robust standard errors), suggesting that a \$1-increase in the marginal incentive to cheat

⁷A Kolmogorov-Smirnov (KS) test of distributions rejects the null hypothesis of equal distributions ($p < 0.01$ for both risk-adjusted and unadjusted investment).

⁸For women, $p < 0.01$ for both unadjusted and risk-adjusted investment. For men, $p = 0.01$ ($p = 0.02$) for unadjusted (risk-adjusted) investment.

⁹There is some nuance to this issue. Incentives determine wealth, and there could be an income effect (although we do not expect them over such a small range). However, the incentive size predicts wealth *only among cheaters*. An income effect among cheaters would imply that cheaters are inferring greater skill from greater wealth, failing to adjust for their past behavior. This is behavior that we want to capture in our estimate. There could be alternative impacts of wealth on confidence, but we test for these directly in Experiment 2, and find no evidence thereof. Moreover, revealed confidence elicited in Task 2 is on behalf of a charity, which is unaffected by the participants income in Task 1.

Figure 20. Distribution of Investment by Cheating, Experiment 1



Notes: We plot the cumulative density function for participants who cheated and participants who did not cheat in experiment 1. We plot these separately by with an unadjusted measure of investment and residual investment from a regression of investment on our measure for risk tolerance.

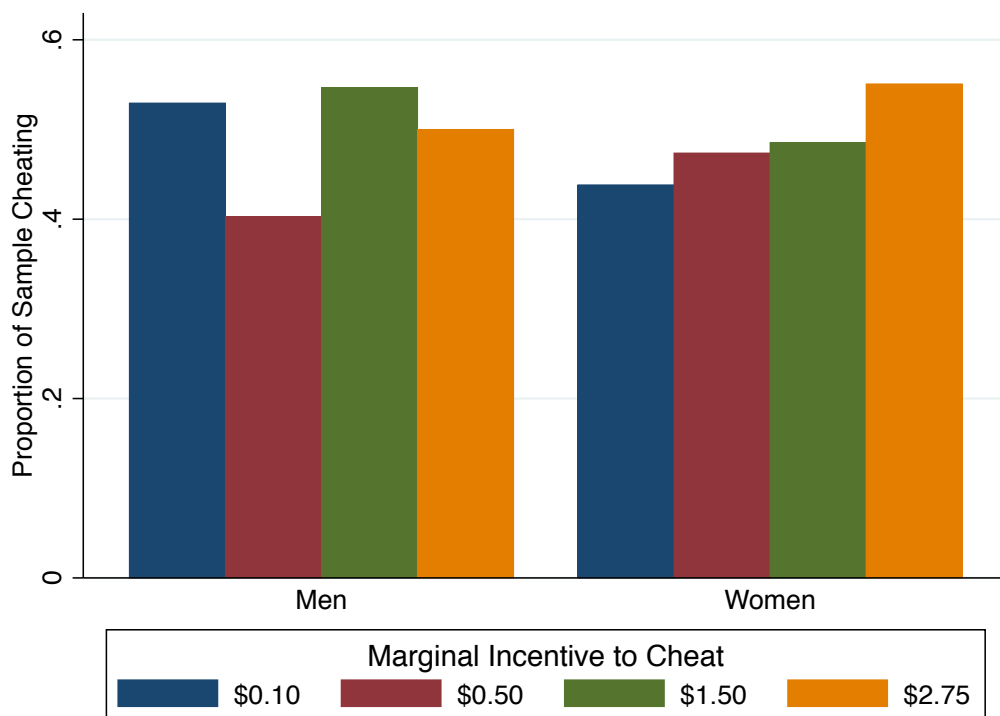
would only increase cheating by 2.8 percent. Thus, while the incentive is predictive of cheating, its relevance is not significant at conventional levels.

In Figure 21, we present the proportion of individuals cheating by gender and the marginal incentive to cheat. This reveals that the impact of the incentive is very different for men and women. For women, cheating is linear in the incentive, and of the expected sign, which suggests that it should work well as an instrument. When we regress cheating on the incentive for women only, the coefficient is substantially larger (0.039), although it is not more precise ($p = 0.17$). For men, behavior does not respond so clearly to incentives. If anything, cheating is distinctly lower when the reward for cheating is equal to \$0.50 than it is when the reward is \$0.10. We regress cheating on an indicator for whether the reward is \$0.50, and find that cheating at this reward is 12.2 percent lower than at other incentive levels ($p = 0.07$).

These somewhat-weak first stages motivate the alternative approach we take in Experiment 2. However, in Table 17 we present the instrumental-variable estimates of the impact of cheating on both stated and revealed confidence. We do this separately for men and women, using the incentive itself for women, and an indicator variable for an incentive equal to \$0.50 for men.¹⁰ For women, cheating increases stated preference by 2.1 standard deviations ($p = 0.21$) and revealed-preference investment by 1.7 standard deviations ($p = 0.26$). For men, cheating increases stated preference by 0.7 standard deviations ($p = 0.52$) and revealed-preference investment by 1.9 standard deviations ($p = 0.17$). While imprecise, in most cases our IV estimates are substantially larger than the OLS estimates. We

¹⁰Different instruments for men and women mean that the estimates are average treatment effects specific to different localities for men and women. As such, coefficients should not be compared across genders.

Figure 21. Cheating propensity, by incentive to cheat, Experiment 1



Notes: We plot the proportion of participants who cheated in each of four treatment groups in experiment 1, by gender.

interpret this as weak evidence that the significant relationship between cheating and confidence belies a causal impact.

Table 17. Impact of Cheating on Stated and Revealed Confidence

| | Women | | Men | |
|--------------|-------------------|-------------------|-------------------|-------------------|
| | Stated (1) | Revealed (2) | Stated (3) | Revealed (4) |
| 1(Cheater) | 2.076 (1.660) | 1.690 (1.505) | 0.696 (1.077) | 1.883 (1.361) |
| Constant | -1.005 (0.811) | -0.814 (0.735) | -0.297 (0.536) | -0.880 (0.673) |
| Observations | 302 | 302 | 272 | 272 |

Notes: Robust standard errors in parentheses. Dependent variables are standardized stated confidence in future performance and standardized investment in future performance.

Experiment 2

In Experiment 2, we modified our design to feature randomization of the *ex-post* rewards gained from cheating. We also added a new treatment in which succeeding at Task 1 is possible without cheating, to determine whether an income effect exists.

In this version of the study, only the expected value to cheating (\$1.50) is known prior to Task 1. Specifically, we told subjects that, “The average bonus payment is that people receive is \$1.50, but it can be more or less.” Subjects learn the realized reward after the first task is performed (and the opportunity to cheat

has passed). This experimentally tests whether, within the set of cheaters, higher rewards from cheating lead to higher confidence.

Design

As before, MTurk subjects are asked to find twelve objects in a picture, needing to find nine to earn the reward. In the ‘Impossible’ treatment, only eight of the listed objects are present, just as in Experiment 1. In the ‘Possible’ treatment, however, nine of the listed objects are actually present, and success is therefore possible without cheating. Subjects are randomized into one of these two treatments upon entry into the study.

We do not provide an incentive for finding all-twelve objects in this version of the experiment. If a participant in either treatment reports finding nine or more objects, their realized reward is drawn randomly from the set $\{\$0.25, \$1.50, \$2.75\}$, with equal probabilities. The Possible treatment serves to difference out any impact of the reward itself, apart from the interaction between cheating and the reward. Following the first task, and the revelation of the reward, subjects participate in a verified investment task similar to that in Experiment 1. We make one change to the investment task for Experiment 2: subjects participate for their own gain rather than for a charity. Because the Possible treatment serves as a control group, we do not need to rely on the charity task to avoid an income effect.

We identify the effect of rewards on confidence for cheaters using a difference-in-differences style design. Our empirical specification is

$$\text{Confidence}_i = \alpha + \beta_1 \mathbf{1}(\text{Impossible})_i + \beta_2 \text{Reward}_i + \beta_3 \mathbf{1}(\text{Impossible}) \times \text{Reward}_i + \epsilon_i \quad . \quad (4.1)$$

Because we use our treatment variation between Impossible and Possible to identify the impact of rewards to cheating separately from the impact of rewards, the coefficient of interest in these models comes from the interaction of an “Impossible” indicator variable and the realized reward for success.

Initially, we limit the sample to those who ‘found’ nine objects and received a reward—these participants were cheaters in the Impossible treatment and non-cheaters in the Possible treatment.¹¹ Of subjects in the Impossible treatment, 41.7 percent cheated to obtain a reward with expected value \$1.50. This is slightly less than the percent who cheated for a guaranteed reward of \$1.50 in Experiment 1 (48.3 percent). 71.6 percent of subjects in the Possible treatment earned a bonus, but a number of them did so by claiming to have found more than nine objects. Accounting for this, 60.9 percent of subjects in this treatment reported nine objects and received a bonus. Overall, we have 158 Impossible subjects and 110 Possible subjects in our sample. When we limit the Impossible treatment to include only those who found all eight objects, this excludes 19 individuals.

Results

We first estimate the impact of the rewards to cheating on stated confidence, which we present in Table 18. In columns (1) through (3), all cheaters in the Impossible treatment are included in the sample. We find no impact of the randomized reward in the Possible treatment, suggesting that there is no direct income effect on confidence in our studies. There is a dramatic level-difference of 0.57 standard deviations between the two treatments. This reflects the fact that those in the Possible treatment *actually succeeded* at the task. Their confidence

¹¹Technically, they could also be very lucky guessers in the Possible treatment, but any minor noise of that nature would only bias our estimate of the interaction term towards zero.

should be very high. There is a positive and significant effect of the reward on stated confidence in the Impossible treatment. Pooling women and men (in Column 1), we find that among cheaters in the Impossible treatment a \$1 (two-thirds of the expected reward) increase in the *ex-post* random reward to cheating increases stated confidence in future performance by about one-quarter of a standard deviation more than it does among successful individuals in the Possible treatment. Although this effect is slightly larger for men than women, the difference is not statistically significant. In columns (4) through (6), we limit the Impossible sample to include only cheaters who found all-eight objects. Results are very similar in these specifications, albeit with a slightly larger gender gap. In both selection-inclusive (Experiment 1) and selection-exclusive (Experiment 2) specifications, our findings are clear that being rewarded for cheating is associated with higher stated confidence.

Table 18. Impact of Rewards from Cheating on Stated Confidence

| | Full Sample | | | Restricted to those who found all possible objects | | |
|---|----------------------------------|----------------------|--------------------|--|--------------------|--------------------|
| | All (1) | Women (2) | Men (3) | All (4) | Women (5) | Men (6) |
| Reward | 0.011 (0.071) | 0.062 (0.095) | -0.048 (0.106) | 0.011 (0.071) | 0.062 (0.095) | -0.048 (0.106) |
| $\mathbb{1}(\text{Impossible})$ | -0.571*** (0.192) | -0.601*** (0.253) | -0.486* (0.285) | -0.467** (0.189) | -0.435* (0.254) | -0.444* (0.263) |
| Reward \times $\mathbb{1}(\text{Impossible})$ | 0.230** (0.100) | 0.178 (0.138) | 0.253* (0.141) | 0.177* (0.010) | 0.105 (0.140) | 0.224* (0.135) |
| Constant | 0.230 (0.133) | 0.107 (0.161) | 0.384 (0.214) | 0.230 (0.133) | 0.107 (0.161) | 0.384 (0.214) |
| H_0 : Female = Male (Reward \times Impossible) | $\chi^2(1) = 0.15$ $p = 0.70$ | | | $\chi^2(1) = 0.39$ $p = 0.53$ | | |
| Observations | 268 | 145 | 123 | 249 | 138 | 111 |

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Dependent variable is standardized stated confidence in future performance.

Next, in Table 19, we consider whether the rewards to cheating likewise influence revealed confidence, as exhibited by the subjects' costly investments in their future performance. We find very different empirical regularities in revealed preference: the ex-post-random rewards to cheating have little influence on investment in either treatment. In Column (5), we find that rewards have a negative, marginally significant effect on investment for women who cheat (relative to non-cheaters). On the other hand, in Column (6) we find that rewards have a positive, but insignificant effect for men who cheat (relative to non-cheaters). Switching from a cheap talk elicitation of confidence to a costly signal appears to mitigate the causal impact of the rewards from cheating on confidence, even though it did not fully mitigate the association between cheating and confidence.

Table 19. Impact of Rewards from Cheating on Revealed Confidence

| | Full Sample | | | Restricted to those who found all possible objects | | |
|---|-------------------|----------------------------------|-------------------|--|----------------------------------|-------------------|
| | All (1) | Women (2) | Men (3) | All (4) | Women (5) | Men (6) |
| Reward | -0.070 (0.076) | 0.061 (0.091) | -0.193 (0.119) | -0.070 (0.076) | 0.061 (0.091) | -0.193 (0.120) |
| $\mathbf{1}(\text{Impossible})$ | -0.197 (0.195) | -0.069 (0.244) | -0.317 (0.316) | -0.133 (0.202) | 0.068 (0.247) | -0.359 (0.340) |
| Reward \times $\mathbf{1}(\text{Impossible})$ | -0.030 (0.104) | -0.155 (0.133) | 0.085 (0.160) | -0.070 (0.108) | -0.240* (0.135) | 0.102 (0.173) |
| Constant | 0.322 (0.146) | 0.170 (0.177) | 0.477 (0.238) | 0.322 (0.146) | 0.170 (0.177) | 0.477 (0.239) |
| $H_0 : \text{Female} = \text{Male}$ (Reward \times $\mathbf{1}(\text{Impossible})$) | | $\chi^2(1) = 1.37$ $p = 0.24$ | | | $\chi^2(1) = 2.50$ $p = 0.11$ | |
| Observations | 268 | 145 | 123 | 249 | 138 | 111 |

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.
Dependent variable is standardized investment in future performance.

As before, we test whether the Reward/Impossible interaction is different for stated and revealed preference in Table 20. Again, the calculated differences

are all positive, suggesting that the impact of a large ex-post reward on confidence is higher for stated preference than it is for revealed preference. This difference is again larger for men than women, although the difference-in-differences is not statistically significant.

Table 20. Testing Stated vs. Revealed Preference (Table 18 vs. Table 19)

| | Full Sample | | | Restricted to those who found all possible objects | | |
|---|-------------|----------------------------------|------------|--|----------------------------------|------------|
| | All (1) | Women (2) | Men (3) | All (4) | Women (5) | Men (6) |
| Difference in Reward \times Impossible Coefficients (Stated – Revealed) | 0.259** | 0.169 | 0.333* | 0.246* | 0.122 | 0.344* |
| H_0 : Female = Male | | $\chi^2(1) = 0.43$ $p = 0.51$ | | | $\chi^2(1) = 0.72$ $p = 0.40$ | |
| Observations | 268 | 145 | 123 | 249 | 138 | 111 |

Notes: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Dependent variable is standardized investment in future performance.

Discussion

Despite the prevalence of cheating, lying, and other forms of dishonest behavior in the workplace, we still know little about the consequences of this behavior. In this paper, we examine the relationship between cheating and confidence. Using an experimental design that captures substantially more cheating than other work in this field, we link cheating to changes in both stated and revealed confidence elicited with a Gneezy and Potters (1997) investment task.

We find a positive relationship between cheating and levels of both stated beliefs of future ability and revealed beliefs through investment, suggesting that cheaters are more confident in their ability than non-cheaters. There are also important gender differences. First, the relationship between cheating and confidence is stronger for men than for women when confidence is only stated, but not when it is revealed. Moreover, managers faced with reports of high confidence should treat them differently for men and women. For men, such reports are more indicative of maximal cheating that substitutes for effort. For women, they are more indicative of marginal cheating that complements effort.

We also took two approaches to identifying a causal relationship between cheating and confidence. In Experiment 1, we varied the *ex-ante* rewards to cheating to create an instrumental variable. While we estimate a strong effect of cheating on confidence, our IV first stage is weak, not allowing strong inference in the resulting second stage. In Experiment 2, we varied the *ex-post* rewards to cheating to create experimental variation in the degree to which cheating paid dividends. We also introduced a new control group—the Possible treatment—to distinguish the effects of ill-gotten gains from properly-earned rewards. Cheaters who earned a higher reward have a higher stated confidence (an effect that is weakly greater for men), but we find no difference in subsequent investment decisions based on the reward from cheating. If anything, women who cheat display contrition in the revealed measure (lower confidence) in response to larger rewards to cheating, while men do not.

While employers seemingly continue to embrace methods of performance review that facilitate cheap talk, or the strategic revelation of information by the informed agent, there is little known about the systematic nature by which cheating

or dishonest revelation can influence either stated beliefs of ability or costly signals thereof. Given that the accumulated rewards from cheating do appear to influence confidence, we should worry that workplaces that feature the potential to cheat or opportunities for cheap-talk self promotion will feature excess advancement of men who are willing to substitute cheating for effort. Similarly confident women may not be completely honest, but they are more likely to be non-cheaters or cheaters who complemented their hard work with a slight nudge over the finish line.

Goldin (2014) attributes differences in gender composition and earnings gaps across industries to long-hours and/or low-flexibility premiums. Niederle and Vesterlund (2007) and Babcock, Recalde, Vesterlund, and Weingart (2017) single out gender differences in the participation in competitive and non-promotable tasks, respectively, as contributing to gender differences in occupation choices and in advancement trajectories. Similar industries can be categorized as conducive to gender gaps based on both theories; business, finance, and law are typical examples. These industries also feature ample opportunities to cut corners and strong upward pressure within firms. Our work suggests that the link between cheating and confidence may explain excess advancement of unproductive male cheaters within such industries, and that firms could select more equally on gender if they are careful to consider only costly signals of employee confidence.

CHAPTER V

CONCLUSION

This dissertation examines three different topics to better understand how people make decisions in a variety of environments. I econometrically and experimentally examine risk-taking behavior, normative beliefs, and self-deception, to determine how these behaviors and beliefs are affected by changes in the policy environment and by previous behavior.

In Chapter II, Glen Waddell and I examine a tournament environment where players are eliminated for insufficient performance half-way through the tournament. We find that in that environment, compared to themselves, when players are on the elimination side of the threshold, they hit fewer strokes, take more risks, and succeed at risks more often. In addition, we find that those lower-ability players are the most reactive to the threshold, having the largest estimated effects. Untangling the improvements in performance due to a particular type of risk-taking from improvements due to increased focus, the estimates suggest that 23 percent of the improvement in performance induced by potential elimination is due to productive increases in risk taking. In addition to the direct policy implications to professional sports, these results also may be applicable to workplace incentive structures, suggesting that structures with elimination structures may induce productive risk-taking behavior among those on the margin of elimination.

I examine the effect of public policy on normative beliefs in Chapter III, by exploiting the timing of the legalization of same-sex marriage. Although there is little evidence of legalization inducing switching between support and opposition, I find evidence that legalization strengthens beliefs—moving participant beliefs


toward the extremes. Examining heterogeneity, I find diverging beliefs among educational groups and racial groups. Measuring state-level polarization, I find substantial increases in polarization due the legalization of same-sex marriage. These results are clear in their implication that an additional cost of public policy is its effect on beliefs and polarization. Warnings of the consequences of increasing polarization are numerous in news reports as increased polarization may lead to more-extreme or to fewer public policies enacted, to divergence in peer groups and political-party beliefs, and even to increases into hate crime.

Finally, in Chapter IV, I, along with Glen Waddell and Michael Kuhn, experimentally consider the link between cheating and confidence. We use an experimental design that allows us to identify cheaters and then to link cheating to both stated and revealed confidence. This link by itself may be crucial to a manager, who often relies on self-reports of performance. As more glowing self-reports are linked to cheating, managers who reward high self-confidence may also be rewarding dishonest behavior. In addition, we vary the *ex-post* rewards to cheating in order to identify a causal effect of cheating, finding an increased level of stated-confidence induced by cheating. However, we find no evidence of a causal link between cheating and revealed-confidence. These results then may explain excess advancement of unproductive male cheaters, and that hiring differences in gender may be eliminated if employers only consider costly signals of employee confidence.

This dissertation includes unpublished co-authored material.

APPENDIX
 SUPPLEMENTAL FIGURES AND TABLES

Figure A1. Support for same-sex marriage, 1988 & 2004–2016



Notes: We plot the proportion of GSS respondents in each categorical response to the GSS question “Do you agree or disagree? Homosexual couples should have the right to marry one another,” in all GSS surveys in which the question was asked (i.e., 1988, 2004, 2006, 2008, 2010, 2012, 2014, and 2016).

Table A1. Strength of opinion on same-sex marriage around same-sex-marriage legalization, 2004-2016

| | Change in Probability (1) |
|----------------------------|---------------------------------|
| Strongly disagree | 0.034 (0.023) |
| Disagree | -0.033 (0.025) |
| Neither agree nor disagree | 0.0026 (0.022) |
| Agree | -0.047** (0.022) |
| Strongly agree | 0.044** (0.017) |
| Pseudo R^2 | 0.117 |
| Observations | 10539 |

Notes: We report the average marginal effects from a multinomial specification across the five categorical variables. The specification includes state and year fixed effects, state-specific linear trends, and individual demographic controls. Standard errors in parentheses, allowing for clustering at the state level. “Less than high school” is the omitted group. *** significant at 1%; ** significant at 5%; * significant at 10%.

Figure A2. Task 1




Please check each object as you find it. Remember, you must find at least 9 objects to earn a bonus payment.

- | | |
|---|--|
| <input type="checkbox"/> A horse | <input type="checkbox"/> An alligator |
| <input type="checkbox"/> An orange crayon | <input type="checkbox"/> A rolling pin |
| <input type="checkbox"/> A dinosaur | <input type="checkbox"/> A gun |
| <input type="checkbox"/> A striped candy cane | <input type="checkbox"/> A green car |
| <input type="checkbox"/> A yo-yo | <input type="checkbox"/> A sailboat |
| <input type="checkbox"/> A banana | <input type="checkbox"/> A black and yellow flag |

Notes: The hidden object task shown to participants in experiment 1. Participants were asked to find a banana, a dinosaur, an orange crayon, a black and yellow flag, an airplane, a sailboat, a basketball, a yo-yo, a gun, a rolling pin, a taco, and a bird.

Figure A3. Task 2



Write the location of the hidden object in the boxes below. For example, if you found a bird in row 4 and column D, write 4D in the box under bird.

| | | | | | | |
|----------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| | a duck | an elephant | a roller skate | a red apple | a black 8 ball | a helicopter |
| Location | <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="text"/> |

| | | | | | | |
|----------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| | a teddy bear | a clock | a red dice | a winged horse | a chair | a tennis racket |
| Location | <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="text"/> |

Notes: The hidden object task shown to participants in experiment 2. Participants were asked to find and report the location of a duck, an elephant, a roller skate, a red apple, a black 8 ball, a helicopter, a teddy bear, a clock, a red dice, a winged horse, a chair, and a tennis racket.

REFERENCES CITED

- Abeler, J., D. Nosenzo, and C. Raymond (2016). Preferences for truth-telling. CESifo Working Paper Series No. 6087.
- Alesina, A. and E. L. Ferrara (2005, September). Ethnic diversity and economic performance. *Journal of Economic Literature* 43(3), 762–800.
- Alesina, A. and P. Giuliano (2015, December). Culture and institutions. *Journal of Economic Literature* 53(4), 898–944.
- Babcock, L., M. P. Recalde, L. Vesterlund, and L. Weingart (2017). Gender differences in accepting and receiving requests for tasks with low promotability. *American Economic Review* 107(3), 714–747.
- Barabas, J. (2009). Not the next IRA: How health savings accounts shape public opinion. *Journal of Health Politics, Policy and Law* 34(2), 181–217.
- Barron, J. M. and G. R. Waddell (2008). Work hard, not smart: Stock options in executive compensation. *Journal of Economic Behavior & Organization* 66, 767–790.
- Beasley, M. S., J. V. Carcello, D. R. Hermanson, and T. L. Neal (2010). Fraudulent financial reporting 1998-2007. Technical report, Committee of Sponsoring Organizations of the Treadway Commission.
- Becker, B. E. and M. A. Huselid (1992). The incentive effects of tournament compensation systems. *Administrative Science Quarterly* 37, 336–350.
- Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy* 76(2), 169–217.
- Berentsen, A. (2002). The economics of doping. *European Journal of Political Economy* 18(1), 109 – 127.
- Bishin, B. G., T. J. Hayes, M. B. Incantalupo, and C. A. Smith (2015). Opinion backlash and public attitudes: Are political advances in gay rights counterproductive? *American Journal of Political Science*.
- Borghans, L., B. H. H. Golsteyn, J. J. Heckman, and H. Meijers (2009). Gender differences in risk aversion and ambiguity aversion. *Journal of the European Economic Association* 7(2/3), 649–658.

- Bothner, M. S., J. han Kang, and T. E. Stuart (2007). Competitive crowding and risk taking in a tournament: Evidence from nascar racing. *Administrative Science Quarterly* 52(2), 208–247.
- Brown, J. (2011). Quitters never win: The (adverse) incentive effects of competing with superstars. *Journal of Political Economy* 119(5), 982–1013.
- Brownback, A. and M. A. Kuhn (2018). Understanding outcome bias: Asymmetric sophistication and biased beliefs. Working Paper.
- Bunn, D. N., S. B. Caudill, and D. M. Gropper (1992). Crime in the classroom: An economic analysis of undergraduate student cheating behavior. *The Journal of Economic Education* 23(3), 197–207.
- Burns, N. and S. Kedia (2006). The impact of performance-based compensation on misreporting. *Journal of financial economics* 79(1), 35–67.
- Cabral, L. (2003). R&d competition when firms choose variance. *Journal of Economics & Management Strategy* 12(1), 139–150.
- Campbell, A. L. (2011). Policy feedbacks and the impact of policy designs on public opinion. *Journal of Health Politics, Policy and Law* 36(6), 961–973.
- Cantoni, D., Y. Chen, D. Y. Yang, N. Yuchtman, and Y. J. Zhang (2017). Curriculum and ideology. *Journal of Political Economy* 125(2), 338–392.
- Carrell, S. E. and M. L. Hoekstra (2010). Externalities in the classroom: How children exposed to domestic violence affect everyone’s kids. *American Economic Journal: Applied Economics* 2(1), 211–28.
- Chance, Z., M. I. Norton, F. Gino, and D. Ariely (2011). Temporal view of the costs and benefits of self-deception. *Proceedings of the National Academy of Sciences* 108(3), 15655–15659.
- Charness, G. and U. Gneezy (2012). Strong evidence for gender differences in risk taking. *Journal of Economic Behavior and Organization* 83, 50–58.
- Charness, G. and D. I. Levine (2007). Intention and stochastic outcomes: An experimental study. *The Economic Journal* 117(522), 1051–1072.
- Chevalier, J. and G. Ellison (1997). Risk taking by mutual funds as a response to incentives. *Journal of Political Economy* 105(6), 1167–1200.
- Clots-Figueras, I. and P. Masella (2013). Education, language and identity. *The Economic Journal* 123(570), F332–F357.
- Conrads, J., B. Irlenbusch, R. M. Rilke, and G. Walkowitz (2013). Lying and team incentives. *Journal of Economic Psychology* 34, 1–7.

- Curry, P. A. and S. Mongrain (2009). Deterrence in rank-order tournaments. *Review of Law and Economics* 5, 723.
- Cushman, F., A. Dreber, Y. Wang, and J. Costa (2009). Accidental outcomes guide punishment in a “trembling hand” game. *PloS one* 4(8), e6699.
- Dana, J., R. A. Weber, and J. X. Kuang (2007). Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness. *Economic Theory* 33, 67–80.
- de Oliveira, A. C., A. Smith, and J. Spraggon (2017). Reward the lucky? An experimental investigation of the impact of agency on luck and bonuses. *Journal of Economic Psychology* 62, 87–97.
- Dee, T. S. (2008). Forsaking all others? The effects of same-sex partnership laws on risky sex. *The Economic Journal* 118(530), 1055–1078.
- DellaVigna, S. and E. Kaplan (2007). The Fox News effect: Media bias and voting. *The Quarterly Journal of Economics* 122(3), 1187–1234.
- Di Tella, R., S. Galiani, and E. Schargrodsky (2007). The formation of beliefs: Evidence from the allocation of land titles to squatters. *The Quarterly Journal of Economics* 122(1), 209–241.
- Dreber, A. and M. Johannesson (2008). Gender differences in deception. *Economics Letters* 99, 197–199.
- Easterly, W. and R. Levine (1997). Africa’s growth tragedy: policies and ethnic divisions. *The Quarterly Journal of Economics* 112(4), 1203–1250.
- Efendi, J., A. Srivastava, and E. P. Swanson (2007). Why do corporate managers misstate financial statements? the role of option compensation and other factors. *Journal of financial economics* 85(3), 667–708.
- Ehrenberg, R. G. and M. L. Bognanno (1990a). Do tournaments have incentive effects? *Journal of Political Economy* 98(6), 1307–1324.
- Ehrenberg, R. G. and M. L. Bognanno (1990b). The incentive effects of tournaments revisited: Evidence from the european pga tour. *Industrial and Labor Relations Review* 43(3), 74S–88S.
- Eil, D. and J. M. Rao (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics* 3, 114–138.
- Enke, B. (2018). What you see is all there is. Working paper.

- Enke, B. and F. Zimmerman (2018). Correlation neglect in belief formation. *Review of Economic Studies Forthcoming*.
- Erat, S. and U. Gneezy (2012). White lies. *Management Science* 58(4), 723–733.
- Erikson, R. S. and L. Stoker (2011). Caught in the draft: The effects of vietnam draft lottery status on political attitudes. *The American Political Science Review* 105(2), 221–237.
- Esteban, J.-M. and D. Ray (1994). On the measurement of polarization. *Econometrica* 62(4), 819–851.
- Fischbacher, U. and F. Föllmi-Heusi (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association* 11(3), 525–547.
- Flores, A. R. and S. Barclay (2016). Backlash, consensus, legitimacy, or polarization: The effect of same-sex marriage policy on mass attitudes. *Political Research Quarterly* 69(1), 43–56.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives* 19(4), 25–42.
- Friedman, W., M. Kremer, E. Miguel, and R. Thornton (2016). Education as liberation? *Economica* 83(329), 1–30.
- Garcia, J. M. and E. F. Stephenson (2015). Does stableford scoring incentivize more aggressive golf? *Journal of Sports Economics* 16(6), 647–663.
- Genakos, C. and M. Pagliero (2012). Interim rank, risk taking, and performance in dynamic tournaments. *Journal of Political Economy* 120(4), 782–813.
- Gill, D., V. Prowse, and M. Vlassopoulos (2013). Cheating in the workplace: An experimental study of the impact of bonuses and productivity. *Journal of Economic Behavior and Organization* 96, 120 – 134.
- Gilpatric, S. M. (2011). Cheating in contests. *Economic Inquiry* 49(4), 1042–1053.
- Gneezy, U., A. Imas, and K. Madarasz (2014). Conscience accounting: Emotion dynamics and social behavior. *Management Science* 60(11), 2645–2658.
- Gneezy, U., A. Kajackaite, and J. Sobel (2018). Lying aversion and the size of the lie. *American Economic Review Forthcoming*.
- Gneezy, U. and J. Potters (1997). An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics* 112(2), 631–645.

- Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review* 104(4), 1091–1119.
- Grund, C., J. Hocker, and S. Zimmermann (2013). Incidence and consequences of risk-taking behavior in tournaments—evidence from the nba. *Economic Inquiry* 51(2), 1489–1501.
- Gurdal, M. Y., J. B. Miller, and A. Rustichini (2013). Why blame? *Journal of Political Economy* 121(6), 1205–1247.
- Guryan, J., K. Kroft, and M. J. Notowidigdo (2009). Peer effects in the workplace: Evidence from random groupings in professional golf tournaments. *American Economic Journal: Applied Economics* 1(4), 34–68.
- Gusmano, M. K., M. Schlesinger, and T. Thomas (2002). Policy feedback and public opinion: The role of employer responsibility in social policy. *Journal of Health Politics, Policy and Law* 27(5), 731–772.
- Haggag, K. and D. G. Pope (2016). Attribution bias in consumer choice. Working Paper.
- Hanushek, E. A., J. F. Kain, and S. G. Rivkin (2009). New evidence about brown v. board of education: The complex effects of school racial composition on achievement. *Journal of Labor Economics* 27(3), 349–383.
- Harris, A. P. (2015). Judicially activated: Electoral response to same-sex marriage in iowa. *University of Pennsylvania*.
- Hatzenbuehler, M. L., K. M. Keyes, and D. S. Hasin (2009). State-level policies and psychiatric morbidity in lesbian, gay, and bisexual populations. *American Journal of Public Health* 99(12), 2275–2281.
- Heffington, C., B. B. Park, and L. K. Williams (2017). The most important problem dataset (mipd): a new dataset on american issue importance. *Conflict Management and Peace Science*, 0738894217691463.
- Hetling, A. and M. L. McDermott (2008). Judging a book by its cover: Did perceptions of the 1996 us welfare reforms affect public support for spending on the poor? *Journal of Social Policy* 37(3), 471–487.
- Hetling, A., M. L. McDermott, and M. Mapps (2008). Symbolism versus policy learning. *American Politics Research* 36(3), 335–357.
- Hoxby, C. (2000). Peer effects in the classroom: Learning from gender and race variation. Technical report, National Bureau of Economic Research.

- Imas, A. (2016). The realization effect: Risk-taking after realized versus paper losses. *The American Economic Review* 106(8), 2086–2109.
- Jacob, B. A. and S. D. Levitt (2003). Rotten apples: An investigation of the prevalence and predictors of teacher cheating. *The Quarterly Journal of Economics* 118(3), 843–877.
- Kajackaite, A. and U. Gneezy (2017). Incentives and cheating. *Games and Economic Behavior* 102, 433–444.
- Kerkvliet, J. and C. L. Sigmund (1999). Can we control cheating in the classroom? *The Journal of Economic Education* 30(4), 331–343.
- Knoeber, C. R. and W. N. Thurman (1994). Testing the theory of tournaments: An empirical analysis of broiler production. *Journal of Labor Economics* 12(2), 155–179.
- Kräkel, M. (2007). Doping and cheating in contest-like situations. *European Journal of Political Economy* 23(4), 988 – 1006.
- Kreitzer, R. J., A. J. Hamilton, and C. J. Tolbert (2014). Does policy adoption change opinions on minority rights? the effects of legalizing same-sex marriage. *Political Research Quarterly* 67(4), 795–808.
- Kuhn, P. and M. C. Villeval (2015). Are women more attracted to cooperation than men? *Economic Journal* 125, 115–140.
- La Porta, R., F. Lopez-de Silanes, A. Shleifer, and R. Vishny (1999). The quality of government. *The Journal of Law, Economics, and Organization* 15(1), 222–279.
- Lazear, E. P. and S. Rosen (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy* 89(5), 841–864.
- Martin, G. J. and A. Yurukoglu (2017, September). Bias in cable news: Persuasion and polarization. *American Economic Review* 107(9), 2565–99.
- Mazar, N., O. Amir, and D. Ariely (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research* 45(6), 633–644.
- McCabe, D. L., K. D. Butterfield, and L. K. Trevino (2012). *Cheating in College: Why Students Do it and What Educators Can Do About It*. Johns Hopkins University Press.
- McFall, T. and K. W. Rotthoff (2016). Risk taking dynamics in tournaments: Evidence from professional golf. *Working Paper*.

- Montalvo, J. G. and M. Reynal-Querol (2005). Ethnic polarization, potential conflict, and civil wars. *The American Economic Review* 95(3), 796–816.
- Montalvo, J. G. and M. Reynal-Querol (2008). Discrete polarisation with an application to the determinants of genocides. *The Economic Journal* 118(533), 1835–1865.
- Mueller-Langer, F. and P. A. Versbach (2013). Leading-effect vs. risk-taking in dynamic tournaments: Evidence from a real-life randomized experiment. *Munich Discussion Paper No. 2013-6*.
- Nagin, D. S., J. B. Rebitzer, S. Sanders, and L. J. Taylor (2002). Monitoring, motivation, and management: The determinants of opportunistic behavior in a field experiment. *The American Economic Review* 92(4), 850–873.
- Niederle, M. and L. Vesterlund (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics* 122(3), 1067–1101.
- Nivola, P. S. and D. W. Brady (2008). *Red and Blue Nation?: Consequences and Correction of America's Polarized Politics*, Volume 2. Brookings Institution Press.
- Oprea, R. (2014). Survival versus profit maximization in a dynamic stochastic experiment. *Econometrica* 82(6), 2225–2255.
- Orszag, J. M. (1994). A new look at incentive effects and golf tournaments. *Economics Letters* 46(1), 77–88.
- Ozbeklik, S. and J. K. Smith (2017). Risk taking in competition: Evidence from match play golf tournaments. *Journal of Corporate Finance* 44, 506–523.
- Pacheco, J. (2013). Attitudinal policy feedback and public opinion: the impact of smoking bans on attitudes towards smokers, secondhand smoke, and antismoking policies. *Public Opinion Quarterly* 77(3), 714–734.
- Papyrakis, E. and P. H. Mo (2014). Fractionalization, polarization, and economic growth: Identifying the transmission channels. *Economic Inquiry* 52(3), 1204 – 1218.
- Park, A. (March 12 2009). Lie, cheat, flirt. what people will do to keep a job. *Time*.
- Pew Research Center (2014, June). Political polarization in the american public.
- Poole, K. T. (2005). *Spatial Models of Parliamentary Voting*. Cambridge University Press.

- Pope, D. G. and M. E. Schweitzer (2011). Is tiger woods loss averse? persistent bias in the face of experience, competition, and high stakes. *The American Economic Review* 101(1), 129–157.
- Prendergast, C. (1999). The provision of incentives in firms. *Journal of Economic Literature* 37(1), 7–63.
- Robert, I. and M. Arnab (2013). Is dishonesty contagious? *Economic Inquiry* 51(1), 722–734.
- Rubin, J. and R. Sheremeta (2015). Principal–agent settings with random shocks. *Management Science* 62(4), 985–999.
- Sarsons, H. (2017). Interpreting signals: Evidence from doctor referrals. Working Paper.
- Schwieren, C. and D. Weichselbaumer (2010). Does competition enhance performance or cheating? a laboratory experiment. *Journal of Economic Psychology* 31(3), 241–253.
- Soss, J. (1999). Lessons of welfare: Policy design, political learning, and political action. *The American Political Science Review* 93(2), 363–380.
- Soss, J. and S. F. Schram (2007). A public transformed? welfare reform as policy feedback. *American Political Science Review* 101(1), 111–127.
- Tatum, A. K. (2017). The interaction of same-sex marriage access with sexual minority identity on mental health and subjective wellbeing. *Journal of Homosexuality* 64(5), 638–653.
- Taylor, B. A. and J. G. Trogdon (2002). Losing to win: Tournament incentives in the national basketball association. *Journal of Labor Economics* 20(1), 23–41.
- Taylor, J. (2003). Risk-taking behavior in mutual fund tournaments. *Journal of Economic Behavior & Organization* 50(3), 373–383.
- Van Weelden, R. (2015). The welfare implications of electoral polarization. *Social Choice and Welfare* 45(4), 653–686.
- Veldhuizen, Roel, v. (2017). Gender differences in tournament choices: Risk preferences, overconfidence or competitiveness? Working Paper.
- Weisel, O. and S. Shalvi (2015). The collaborative roots of corruption. *Proceedings of the National Academy of Sciences* 112(34), 10651–10656.