Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Bounded Rationality in Strategic Decisions: Undershooting in a Resource Pool-Choice Dilemma

Christopher K. Hsee, Ying Zeng, Xilin Li, Alex Imas

informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–15
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Bounded Rationality in Strategic Decisions: Undershooting in a Resource Pool-Choice Dilemma

Christopher K. Hsee,[a] Ying Zeng,[b] Xilin Li,[a] Alex Imas[a]

[a] Booth School of Business, University of Chicago, Chicago, Illinois 60637; [b] Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada

Contact: christopher.hsee@chicagobooth.edu, https://orcid.org/0000-0002-8391-7847 (CKH); yying.zeng@rotman.utoronto.ca, https://orcid.org/0000-0001-9839-1182 (YZ); xilin.li@chicagobooth.edu, https://orcid.org/0000-0003-1343-5962 (XL); alex.imas@chicagobooth.edu, https://orcid.org/0000-0002-1651-7245 (AI)

Received: October 21, 2019
Revised: April 8, 2020
Accepted: June 9, 2020
Published Online in Articles in Advance: February 17, 2021

https://doi.org/10.1287/mnsc.2020.3814

Abstract. This research studies a resource *pool-choice dilemma*, in which a group of resource seekers independently choose between a larger pool containing more resources and a smaller pool containing fewer resources, knowing that the resources in each pool will be divided equally among its choosers, so that the more (fewer) people choose a certain pool, the fewer (more) resources each of them will get. This setting corresponds to many real-world situations, ranging from students choosing majors as a function of job opportunities to entrepreneurs choosing markets as a function of customer bases. Ten studies reveal a systematic *undershooting bias*: fewer people choose the larger pool relative to both the normative equilibrium benchmark and chance (random choice), thus advantaging those who choose the larger pool and disadvantaging those who choose the smaller pool. We present evidence showing that the undershooting bias is driven by bounded rationality in strategic thinking and discuss the relationship between our paradigm and other coordination games.

History: Accepted by Yuval Rottenstreich, decision analysis.
Supplemental Material: The appendix is available at https://doi.org/10.1287/mnsc.2020.3814.

Keywords: coordination • mixed strategy • level-k • cognitive hierarchy • resource competition • market entry game

Introduction

Imagine that you are one of many individual vendors who are allowed to sell ice cream at one of two events tomorrow and must decide independently which one to go to. Every vendor knows that one event will have more attendees than the other and that every attendee will likely buy one ice cream and is as likely to buy it from one vendor as from another. Thus, the more (fewer) vendors go to a certain event, the fewer (more) attendees will buy each vendor's ice cream, resulting in less (more) profit for each vendor. Which event will you go to?

Pool-Choice Dilemma

This is an example of a *pool-choice dilemma* in resource competition: A group of resource seekers choose between a larger pool containing more resources and a smaller pool containing fewer resources, knowing that the resources in each pool will be shared evenly by its choosers. Even though real-life decisions are inevitably more complicated than the stylized case above, many decisions share its characteristics; examples include firms deciding whether to enter a large market with more customers or a niche market with fewer customers, students deciding whether to major in a field with more expected jobs or one with fewer expected jobs, and residents in a virus-stricken area who need face masks deciding whether to rush to a store with more remaining masks or one with fewer. In each case, one pool contains more resources than the other, and the more (fewer) competitors choose to acquire resources from a certain pool, the fewer (more) resources will be available to each of its choosers. In order to maximize one's own outcome, one must think about the choices of others. Biases in such decisions carry potentially serious consequences for both individuals and the efficiency of the system as a whole, since deviations from optimal behavior can generate substantial inequality between people.

The pool-choice dilemma described above can be abstracted in the following account-choice game: N players choose between a larger account containing $\$l$ and a smaller account containing $\$s$ (where $s < l < N \times s$). The players are anonymous and cannot communicate with each other. Each player must make a choice between the two accounts and must choose independently. After everyone makes their choices, the money in each account will be divided equally among those who choose that account, so that the more (fewer) players choose a certain account, the less (more) money each of them will receive.

This game has a unique equilibrium *outcome*: $l/(l+s)$ of the players should choose the larger account in expectation (with the remaining players choosing the smaller account). For example, if the larger account contains $55 and the smaller account contains $45, the unique equilibrium outcome is that 55% of the players choose the larger account. The $l/(l+s)$ proportion is the only equilibrium outcome because, regardless of which account they choose, all players earn the same amount and no player has an incentive to deviate and choose a different account.

Rational players could reach the above outcome by adopting various equilibrium strategies. One such equilibrium strategy is a mixed strategy: each player chooses the larger account with a probability of $l/(l+s)$. In the example above, for instance, each player could draw a card from a deck containing 11 red cards and 9 black cards, and choose the larger account if the card drawn is red or the smaller account if the card drawn is black.[1]

In reality, however, people may not adopt such equilibrium strategies and may not reach the $l/(l+s)$ normative benchmark. Deviations from the normative equilibrium proportion fall under two categories. If the proportion of players choosing the larger account exceeds the equilibrium proportion, then those choosing the larger account will stand to earn less than those choosing the smaller account; we term this phenomenon *overshooting*. On the other hand, if the proportion of players choosing the larger account falls below the equilibrium proportion, then those choosing the smaller account will stand to earn less than those choosing the larger account; we term this phenomenon *undershooting*.

Both overshooting and undershooting can cause substantial inequality. For example, suppose that 100 players play the account-choice game in which the larger account contains $55 and the smaller account contains $45. If the players overshoot, say, 70% of them choose the larger account, then each of those choosing the larger account will end up with only $0.79, whereas each of those choosing the smaller account will end up with almost twice as much—$1.50. If the players undershoot, say, 40% of them choose the larger account, then each of those choosing the larger account will end up with $1.38, whereas each of those choosing the smaller account will end up with a much smaller share—only $0.75.

Undershooting

We predict that people in the pool-choice dilemma will systematically undershoot, beyond what mere noise or random choice could explain. Our prediction builds on models that incorporate bounded rationality into people's beliefs about the actions and beliefs of others (e.g., Stahl and Wilson 1994, Nagel 1995, Costa-Gomes and Crawford 2006, Melkonyan et al. 2018), particularly how *level-k* thinking affects decisions in coordination games (Crawford et al. 2008). According to this framework, people adopt different levels of strategic thinking, and a higher-level thinker uses a strategy that best responds to the behavior of lower-level thinkers. This bounded rationality in strategic thinking leads to predictable patterns in observed behavior.

In the context of the pool-choice dilemma, we define "K1 thinkers" as those who choose the larger pool simply because it contains more resources. We define "K2 thinkers" as those who believe that others are K1; in turn, they best respond by choosing the smaller pool.[2] We define "K3 thinkers" as those who believe that others are K2 thinkers, and therefore best respond by choosing the larger pool. Note that K1 thinkers are nonstrategic in the sense that they do not consider the behavior of others, whereas those with higher levels of reasoning are strategic in the sense that they best respond to what they perceive others are doing. These classifications are closely related to the setup of Crawford et al. (2008), who assume that K1 thinkers choose the payoff-salient option in coordination games with payoff asymmetries.[3] Theoretically, there could be even higher-level thinkers, but the literature shows that few people reason above the K3 classification (Costa-Gomes and Crawford 2006, Arad and Rubinstein 2012).

Extensive psychological research shows that people have a positive illusion about themselves, believing that they are better than others on a variety of dimensions, including abilities, future outcomes, and perceived control (e.g., Taylor and Brown 1988, Alicke et al. 1995, Alicke and Govorun 2005, Moore and Healy 2008, Scopelliti et al. 2015, Barasz et al. 2016). Likewise, research in game theory suggests that people interacting strategically believe that they are more sophisticated and think more deeply than other players in the same situation even though others are drawn from the same population, and are on average just as smart and sophisticated (Nagel 1995, Camerer 2003, Arad and Rubinstein 2012). For example, using a series of guessing games with various parameters to separate players' thinking levels, Costa-Gomes and Crawford (2006) find that the modal player in a game considers other players to be nonstrategic and then reacts strategically based on that belief, implying that the modal player is a K2 thinker.

Building on these findings, we predict that the modal player in our paradigm believes that most other players are nonstrategic and will choose the larger pool, and responds to this belief by choosing the smaller pool. In other words, we propose that the modal player in our setting is a K2 thinker. This leads to fewer people choosing the larger pool than the equilibrium proportion, generating an undershooting bias.

Relationship with Other Paradigms

Our pool-choice paradigm is related to classic coordination games such as the market entry game (Kahneman 1988, Sundali et al. 1995) and the minority game (Challet and Zhang 1997). Our prediction of a systematic undershooting bias contrasts with the findings from this existing literature which finds no systematic biases on the aggregate and leads Kahneman to conclude that behaviors in those games look "like magic" (Kahneman 1988).

However, there is a critical difference between our setting and the other paradigms: Our setting has an *apparently superior option*, namely, the larger pool. The presence of the apparently superior option serves as the starting point from which to reason strategically about others' behaviors (Crawford and Iriberri 2007, Crawford et al. 2008). Specifically, it leads people, who view others as nonstrategic to believe that they will choose the apparently superior option (the larger pool) and, to best respond to this belief by choosing its alternative (the smaller pool). This generates an undershooting bias.

By contrast, classic resource coordination games do not entail options that are ex-ante apparently superior than their alternatives. Take, for example, the market entry game (Kahneman 1988, Sundali et al. 1995), in which players choose between a riskless option (not entering the market) and a risky option (entering). The riskless option yields the same payoff regardless of others' decisions; the risky option has a lower payoff when more players decide to enter the market. Since risk is positively correlated with the expected payoff, neither the riskless option nor the risky option appears superior at first glance. The lack of an apparently superior option means that there is no salient starting point for players (assuming they are K2 thinkers) to predict the behaviors of others, and therefore, no systematic bias is expected in this context. To provide evidence for this argument, we conducted a study comparing our pool-choice dilemma with the market entry game directly and found evidence in support of it. We outline this study and present results in the General Discussion section.

In the minority game, an odd number of players choose between two identical options, and players who select the less popular option win a prize (Challet and Zhang 1997). Since the game involves two identical options, neither option is apparently superior, and no systematic bias is observed (Chmura and Güth 2011, Linde et al. 2014).

Our paradigm is also reminiscent of the route choice game (Selten et al. 2007), in which commuters choose between two different routes to a set destination with the goal of minimizing travel time. The routes vary in length and traffic capacity such that if the number of commuters is the same on each route, travel time will be shorter on one route than on the other. Commuters know the total number of other commuters but not the formula that determines the travel time. Because of the uncertainty of the formula, there is no salient feature that makes either the main route or the side route apparently superior. As a result, players' strategies do not systematically depart from the theoretical predictions, and outcomes quickly converge to equilibrium.

Our work is conceptually related to research on the *X-Y game* (Crawford et al. 2008), in which players are paired and choose between options with more salient labels versus options with higher payoffs; pairs who coordinate on the same option receive more money. Players often fail to coordinate when the label-salient option in a choice set is different from the payoff-superior option. The X-Y game is similar to our paradigm in the sense that both generate a systematic departure from the equilibrium benchmark that can be modeled using level-*k* thinking, but it differs from ours in important ways. In our paradigm, the larger pool is the only salient (apparently superior) option, and people "compete" for resources through their decisions. In the X-Y game, one option is more salient in label and the other more salient in payoffs, and there is no element of competition. As a result, these games are designed to model different types of real-world settings. Decisions between pools with asymmetric resources (as in our paradigm) are meant to model competitive situations such as the choice of where to conduct business or job search; decisions between options that differ in the saliency of labels or payoffs are meant to model classic coordination game settings such as the ones described in Schelling (1960). Departures from the equilibrium benchmark in the former generate substantial *inequality* in the system—creating winners and losers—whereas miscoordination in the latter decreases the payoffs for all players. The implications of the results of these two paradigms differ substantially for both the welfare of the players and potential interventions.[4]

The pool-choice dilemma we study is also related to foraging behavior, in which multiple foragers search for resources in multiple patches (e.g., Harper 1982, Goldstone and Ashpole 2004, Goldstone et al. 2005, Hills 2006, Hills et al. 2015). However, our research differs from the foraging work in both design and outcome. The typical foraging study does not tell foragers in advance which patch contains more resources (i.e., foragers are not aware which option is apparently superior) and lets foragers explore and learn through repeated trials. As a result, the foraging research either finds no systematic biases relative to the ideal free distribution (e.g. Harper 1982, Sokolowski and Tonneau 2004) or finds only mild biases relative to the rational predictions that cannot be distinguished from noise (Kennedy and Gray 1993, Goldstone and Gureckis 2009).[5]

Study Overview

We now proceed to report results from 10 studies that tested our predictions for the pool-choice dilemma. Study 1 to Study 4 demonstrated and replicated the predicted undershooting bias in different contexts and with different levels of asymmetry. Study 5 and Study 6 explored the psychological mechanism underlying the bias by asking players to explain their reasons for their decisions, providing support for our model of bounded rationality in strategic thinking. Study 7 and Study 8 further tested our theory by directly manipulating participants' thinking level. Finally, two studies reported in the appendix compare the pool-choice dilemma with the market entry game and explore a potential boundary condition of the undershooting bias, respectively. Table 1 provides an overview of the studies. The method section of each study reports how we determined the sample size, all data exclusions (if any), all manipulations, and all measures.

A note about our samples: Whereas most prior research on games used students as samples, most of our studies used workers from Amazon Mechanical Turk (MTurk). We did so not only because MTurkers are relatively easy to recruit, but also because they are more diverse in age, education, and experiences, and more representative of the general public than students (Buhrmester et al. 2011). However, to ensure the generality of our research, two of the studies (Study 2 and Study 7) used university students as participants.

All our studies involve only one-shot decisions. Studying one-shot decisions is important both theoretically and practically. On a theoretical level, one-shot decisions represent the cleanest approach for studying strategic reasoning (Camerer et al. 2004, Crawford and Iriberri 2007, Crawford et al. 2008). Unless learning dynamics of the self and others are modeled explicitly, repeated decisions will not identify the types of strategic considerations we aim to study in our paradigm. Practically, many real-world decisions are made infrequently or only once in a lifetime (Thaler 2015). Even if they can repeat a decision, real-world decision makers are often in what Hogarth et al. (2015) refer to as "wicked learning environments" and rarely receive appropriate and useful feedback for them to learn.

Study 1

Method

Study 1 tested the predicted undershooting effect in an incentive-compatible account-choice game. We aimed to recruit 150 participants from MTurk to guarantee relatively high statistical power, and we received completed responses from 151 participants ($N_{\text{female}} = 81$; $M_{\text{age}} = 32.05$).

Table 1. Overview and Main Results of All Studies

Study	Purpose(s) and/or feature(s)	Condition	Equilibrium benchmark	% choosing the larger pool
1	Tests undershooting in an incentive-compatible account-choice game		55.0	35.1* [27.5, 43.3]
2	Replicates undershooting in a realistic mask-seeking context, using a student sample		50.03	35.6* [27.0, 44.9]
3	Replicates undershooting in an incentive-compatible and naturally occurring survey choice context		52.0	33.7* [24.6, 43.8]
4	Replicates undershooting with different levels of asymmetry in an ice-cream vendor context	Low asymmetry	52.0	36.2* [28.6, 44.4]
		High asymmetry	92.0	64.7* [56.5, 72.3]
5	Replicates undershooting in an incentive-compatible account-choice game and explores thinking level		52.0	40.0* [30.3, 50.3]
6	Replicates undershooting in an incentive-compatible account-choice game and explores both others' and one's own thinking levels		55.0	42.0* [32.7, 51.7]
7	Tests the effect of thinking-level manipulation in a mask-seeking context, using a student sample	Control	50.5	35.9* [28.4, 44.1]
		Hint	50.5	62.6 [54.2, 70.4]
8	Tests the effect of thinking-level manipulation on prediction of others' choices and one's own choices in a gold-seeking context	Control	50.5	24.0* [17.4, 31.7]
		Prediction	50.5	27.3* [20.4, 35.2]
		Hint+prediction	50.5	47.3 [39.1, 55.6]
Additional studies reported in an appendix				
A1	Compares our pool-choice dilemma with the market entry game in a water-seller context	Pool choice	50.5	33.3* [25.9, 41.5]
		Market entry	50.5	49.7 [41.4, 58.0]
A2	Tests the effect of vividness in an incentive-compatible account-choice game, using a student sample	Control	55.0	33.3* [24.4, 43.2]
		Treatment	55.0	57.9 [48.0, 67.4]

Notes. The results in Study 7 are from round 2. Numbers in brackets are 95% confidence intervals.
*Significant undershooting.

All participants read the following instructions:

By participating in this study, you are guaranteed to receive the payment advertised for the study. In addition, everyone in this study has a chance to get more money. Here is how (everyone has the following information):

At the end of the study, we will randomly pick 20 winners from all the participants, and let them get some extra money.

Each winner will get money from one of two accounts, one large and one small. The larger account contains $55, and the smaller account contains $45.

Everyone must decide in advance which account to get money from if they win.

After the winners are picked, the money in each account will be divided equally among the winners who choose that account. Therefore, the more winners choose a certain account, the less money each winner will receive.

Note: This is not hypothetical, but real. You actually have a chance to win, and, if you win, which account you choose will influence how much money you can get.

After they made their choices, participants were asked the following comprehension questions:

Which of the following is true?

- The more winners choose a certain account, the more money each of them will get.

- The more winners choose a certain account, the less money each of them will get.

- Regardless of how many winners choose a certain account, each of them will get the same amount of money.

If half of the winners choose the larger account and half the winners choose the smaller account, then which of the following would be true?

- Those who choose the larger account will get more money.

- Those who choose the smaller account will get more money.

- The two groups of winners will get the same amount of money.

After the study, we picked 20 winners and paid them according to their choices, as promised, with each winner receiving $5 on average.

Results and Discussion

According to the equilibrium prediction, the proportion of participants choosing the larger account should be 55/(55+45) = 55%. In support of our behavioral prediction, we found a significant undershooting bias: only 35.1% of participants chose the larger account, which was not only below the equilibrium benchmark of 55%, but also below 50% ($\chi^2(1, n = 151) = 13.41$, $\varphi = 0.30$, $p < 0.001$).[6] Because it was below 50%, this result cannot be ascribed to mere noise or mere random choice.

The result was essentially the same if we included only the participants who answered both comprehension questions correctly ($n = 114$): Among them, 34.2% chose the larger account, significantly below both the equilibrium benchmark and 50% ($\chi^2(1, n = 114) = 11.37$, $p = 0.001$, $\varphi = 0.32$).

The undershooting bias greatly affected participants' earnings. Had there been no bias (i.e., had the proportion of participants choosing the larger account matched the equilibrium benchmark of 55%), then each winner would have received $5.00. But because of undershooting, each of those who chose the larger account received $7.86 while each of those who chose the smaller account received less than half—only $3.46. This represents substantial inequality in the system.

Study 2

Study 2 was a replication of Study 1 using an enriched mask-seeking scenario inspired by real events surrounding the concurrent coronavirus pandemic. The study used a different participant sample than Study 1. Furthermore, it involved only a slight asymmetry in resources (i.e., 2001 vs. 1999) so that the equilibrium benchmark (i.e., 50.03%) was essentially 50%; we were curious whether such a minimal asymmetry could still trigger an undershooting bias. The study was preregistered at http://aspredicted.org/blind.php?x=2ci8ek.

Method

Participants in the study were students from a large university in China who did the study for a payment. We hoped to recruit 100 participants and received completed responses from 118 participants due to a higher than expected response rate ($N_{\text{female}} = 56$; $M_{\text{age}} = 21.39$). All participants read the following (originally in Chinese):

Some students at your university are still living on campus. There are two convenience stores on campus—A and B. To prevent the spread of the coronavirus, both stores decide to give away the face masks they have free of charge to students at 9 a.m. tomorrow morning. Face masks are in great shortage now, and many students—including you—plan to go to the stores to claim masks tomorrow morning, and want to get as many as possible.

Every student can go to only store to claim masks, must decide independently which store to go to, and cannot discuss with one another. Everyone has the following information: Store A has 2001 masks and Store B has 1999 masks. Each store will distribute the masks they have equally to the students who wait there to claim masks. Thus, the fewer the students who go to a certain store to claim masks, the more the masks each student who goes there will get; conversely, the more the students who go to a certain store to claim masks, the fewer the masks each student who goes there will get.

Given the information above, which store would you go to to claim masks?

Results and Discussion

Again, the study revealed a significant undershooting bias: only 35.6% of the participants chose to go to the larger store, which was both below the equilibrium benchmark (50.03%) and chance (50%) ($\chi^2(1, n = 118) = 9.80$, $p < 0.005$, $\varphi = 0.29$). Replicating the result of Study 1, this finding shows that the undershooting bias occurs not only in the abstract account-choice game but also in an enriched and realistic context. Additionally, the bias can be generated by even a slight asymmetry in resources between the two options.

Study 3

Whereas Study 1 tested the undershooting bias in an incentive-compatible but abstract game and Study 2 tested the bias in a realistic but hypothetical scenario, Study 3 tested the bias in a realistic (naturally occurring) and incentive-compatible setting. The study was preregistered at https://aspredicted.org/blind.php?x=4k86ep.

Method

The study asked MTurkers to choose between two surveys. To be consistent with the instructions to the participants (see below), we set a target sample size of 100 and received completed responses from 101 participants ($N_{\text{female}} = 53$; $M_{\text{age}} = 35.66$).

We first asked the participants to complete an unrelated short consumer survey for $0.15. After that, we told them they could do another survey and earn extra money. We gave the following instructions:

We are recruiting about 100 workers (including you) to do an additional survey and we will pay you additional money. Everyone has the following information:

There are two surveys each worker can choose from—Survey A or Survey B. They are similar in content and length (about 5 minutes). Each worker can do only one survey, not both. Each survey has its own budget. Survey A has a budget of $26 and Survey B has a budget of $24. The budget for each survey will be divided equally among the workers who choose to do that survey. Therefore, the fewer workers choose to do a certain survey, the more money each of them will get; conversely, the more workers choose to do a certain survey, the less money each of them will get.

Which survey will you do?

Even though doing the extra survey was optional, everyone made a choice between Survey A and Survey B. They then completed the extra survey (which, again, was unrelated to the present study) and received the promised payment.

Results and Discussion

Only 33.7% of the participants chose Survey A (the survey with a higher budget), which was significantly below both the equilibrium benchmark of 52% and 50% ($\chi^2(1, n = 101) = 10.78$, $p = 0.001$, $\varphi = 0.33$). This choice bias engendered severe inequality in earnings. Although the two surveys were equally long, those who did Survey A earned $0.76, while those who did Survey B earned only $0.36, which is not a trivial difference for workers on this platform. This study replicated the undershooting bias in a naturally occurring survey choice setting among workers who make a living by performing similar tasks.

Study 4

Whereas the studies reported so far demonstrated the undershooting bias when the asymmetry in resources was small, Study 4 varied the level of asymmetry to see if people would still undershoot when the asymmetry was high. The study adopted an ice-cream vendor scenario similar to the one introduced at the beginning of the article. The study was preregistered at https://aspredicted.org/blind.php?x=vn492d.

Method

We aimed to recruit 300 participants from MTurk to ensure that each condition would have a sample size of 150, and we received completed responses from 302 participants ($N_{\text{female}} = 139$; $M_{\text{age}} = 37.59$). Participants were randomly assigned to either a low-asymmetry condition (52%) or a high-asymmetry condition (92%). They all read the following; the number before a "/" was for the low-asymmetry participants and the number after it was for the high-asymmetry participants:

You are one of 20 ice cream vendors who plan to go to one of two outdoor events tomorrow to sell ice cream. Of those events, one is larger than the other. The larger event is expected to have 2600/4600 customers and the smaller event is expected to have 2400/400 customers.

Each of you can go to only one event, and must decide independently which one to go to. You cannot change your mind afterward and cannot discuss your decision with each other.

How much profit a vendor will make tomorrow depends on how many customers will buy their ice cream. Every customer will buy exactly one serving of ice cream, and every vendor will make a $1 profit by selling one serving of ice cream. Thus, the fewer vendors go to a certain event, the more customers at the event will buy each vendor's ice cream, and the more profit each vendor will get. Specifically, if N vendors go to a certain event, then 1/Nth of the customers at that event will buy each vendor's ice cream, and each vendor will get a 1/Nth share of the total profit from that event.

Again, the large event is expected to have 2600/4600 customers and the smaller event is expected to have 2400/400 customers. Each vendor can go to only one event. Given all the information above, which event will you go to?

Results and Discussion

We observed undershooting in both conditions. In the low-asymmetry condition, the proportion of participants choosing the larger event—36.2%—was significantly below both the equilibrium benchmark of 52% and 50% ($\chi^2(1, n = 152) = 11.61$, $p = 0.001$, $\varphi = 0.28$). In the high-asymmetry condition, the proportion of participants choosing the larger event—64.7%—was significantly below the equilibrium benchmark of 92% ($\chi^2(1, n = 150) = 152.26$, $p < 0.001$, $\varphi = 1.00$), but significantly above 50% ($\chi^2(1, n = 150) = 12.91$, $p < 0.001$, $\varphi = 0.29$). This latter result indicates that, when the resources in the two pools are highly asymmetric, most people would choose the larger pool instead of the smaller pool, but the proportion of people choosing the larger pool is still below the equilibrium benchmark.

Because the choice proportion in the high-asymmetry condition falls between the equilibrium benchmark and chance (50%), we do not know whether it is a result of a systematic undershooting bias or a combination of unbiased choice and random choice. This ambiguity is inevitable if the choice options involve highly asymmetric resources, as in the high-asymmetry condition of this study. To avoid this issue, all the remaining studies adopt low asymmetries.

Study 5

Study 5 explored the mechanism underlying the undershooting bias by asking participants to explain the reasons behind their choices. These data were used to begin our investigation into whether bounded rationality in strategic thinking was driving the undershooting bias. Like Study 1, Study 5 used an account-choice game, but unlike Study 1, which incentivized only some randomly chosen participants, Study 5 paid everyone according to their choice.

Method

We aimed to recruit 100 participants from MTurk and received completed responses from 100 participants ($N_{\text{female}} = 61$; $M_{\text{age}} = 35.35$). All participants received the following instructions:

In this study, you will play an "account choice game" with 11 other MTurkers. By playing this game, you can get some extra money. The rules of the game are as follows:

We have two accounts: a larger account containing 520 cents, and a smaller account containing 480 cents. Every participant can get money from one of these accounts, and everyone must decide in advance which account to get money from.

At the end of the study, we will count how many participants choose the larger account and how many choose the smaller account, and we will divide the money in each account equally among the participants who choose that account. Thus, the more participants who choose a certain account, the less money each of them will get. For example, if only 2 participants choose a certain account, then each of them will get 1/2 of the money in that account; if 10 participants choose a certain account, then each of them will get only 1/10 of the money in that account.

Given the above rules, which account would you choose?

After participants made their choice, the next screen asked for their rationale. Specifically, those who chose the larger account were asked:

Why did you choose the larger account instead of the smaller account?

(LA) Because it contains more money and I wanted to get more money.

(LB) Because I figured that most other participants would choose the smaller account and I wanted to choose a different account from what they would choose.

(LC) Because I figured that most other participants would choose the larger account and I wanted to choose the same account as they would.

(LD) Other reason; please specify: ________

Those who chose the smaller account were asked:

Why did you choose the smaller account instead of the larger account?

(SA) Because it contains less money and I wanted to get less money.

(SB) Because I figured that most other participants would choose the larger account and I wanted to choose a different account from what they would choose.

(SC) Because I figured that most other participants would choose the smaller account and I wanted to choose the same account as they would.

(SD) Other reason; please specify: ________

The labels before the options are added here for ease of reference. Note that option LA represented K1 thinking, option SB represented K2 thinking, option LB represented K3 thinking, and the remaining options represented other thinking styles.[7]

At the end of the study, we paid everyone according to their earnings in the game.

Results and Discussion

Choice. We replicated the undershooting effect: the proportion of participants choosing the larger account was 40.0%, significantly below both the equilibrium benchmark of 52% and chance ($\chi^2(1, n = 100) = 4.00$, $p < 0.05$, $\varphi = 0.20$).

Figure 1. Study 5 Thinking Level Results

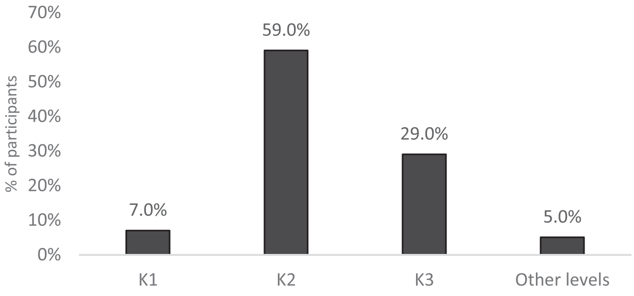

Thinking Level. Figure 1 shows the percentages of participants associated with each thinking level, as inferred from their answers to the questionnaire. As the figure shows, 59.0% of the participants can be classified as K2 thinkers, and this proportion was higher than any other types of thinkers (goodness-of-fit $\chi^2(2, n = 100) = 10.23$, $p < 0.01$, compared with the next most frequent type, which was only 29.0%). This result supports our proposition that the modal participant was a K2 thinker.

To probe further, we conducted separate analyses of the reasons given by those choosing the larger account (the minority) and those choosing the smaller account (the majority). Among those choosing the larger account (the minority), 17.5% can be classified as K1 thinkers, 72.5% as K3 thinkers, and the remaining 10.0% as other types. Among those choosing the smaller account, almost everyone (98.3%) can be classified as a K2 thinker. These results suggest that the undershooting bias was largely driven by K2 thinking.

Study 6

Whereas Study 5 used the reasons participants gave for their choices to infer their levels of thinking, Study 6 provided participants with information about the different levels of potential strategic reasoning and asked them to evaluate both others' level of thinking and their own.

Method

To have at least 100 players (as specified in the instructions), we aimed to recruit 110 participants from MTurk and ended up receiving completed responses from 112 participants ($N_{\text{female}} = 49$; $M_{\text{age}} = 38.5$).

All participants read the following instructions:

By participating in this study, everyone is guaranteed to receive the minimal payment advertised for this study. In addition, every participant will get some extra money. The rules are as follows:

We have two accounts, one larger and one smaller: The larger account contains $55, and smaller account contains $45. We will recruit about 100 participants for this study, and we will give all the money ($55+$45 = $100) to these 100 participants, including you.

Every participant will get money from only one of these accounts, and everyone must decide in advance which account to choose, the larger one or the smaller one.

At the end of the study, we will divide the money in each account equally among the participants who choose that account, and send the money to them. Thus, the more participants choose a certain account, the less money each participant who chooses that account will get.

This is real, not hypothetical. Everyone will receive the extra money as a bonus at the end of the experiment.

Given the above rules, which account would you choose?

After they made their decision, participants were led to the next page and read the following, adapted from our earlier definitions of the different levels of strategic thinking:

Let us define a few terms. Please read carefully.

A Level Zero Thinker is someone who makes a random choice.

A Level One Thinker is someone who chooses the larger account without thinking about what account most other participants would choose.

A Level Two Thinker is someone who assumes that most other participants would choose the larger account and therefore chooses the smaller account him/herself.

A Level Three Thinker is someone who assumes that most other participants would choose the smaller account and therefore chooses the larger account him/herself.

And so on and so forth.

In your opinion, most other participants are _____ (choose one below):

- *Level Zero Thinkers*
- *Level One Thinkers*
- *Level Two Thinkers*
- *Level Three Thinkers*
- *Higher than Level Three Thinkers*

In your opinion, you are _____ (choose one below):

- *a Level Zero Thinker*
- *a Level One Thinker*
- *a Level Two Thinker*
- *a Level Three Thinker*
- *a Higher than Level Three Thinker*

After they answered the above question, participants continued to another page and answered the following comprehension question:

To make sure you understand the instructions of this study, please answer the following question(s):

- *How much money I will get depends both on what account I choose and what account most other participants choose.*

- Regardless of what account most other participants choose, I will always get more money by choosing the larger account.
- Regardless of what account most other participants choose, I will always get more money by choosing the smaller account.
- Regardless of which account most other participants choose, I will always get the same amount of money.

After the study, all participants were paid according to their earnings in the game.

Results and Discussion

Choice. Participants again undershot: 42% chose the larger account, significantly below the equilibrium benchmark of 55% ($\chi^2(1, n = 112) = 7.69$, $p < 0.01$, $\varphi = 0.26$) and marginally significantly below 50% ($\chi^2(1, n = 112) = 2.89$, $p = 0.089$, $\varphi = 0.16$). If we only included participants who answered the comprehension question correctly ($n = 98$), the proportion choosing the larger account was even lower—37.8%, significantly below both the equilibrium benchmark and 50% ($\chi^2(1, n = 98) = 5.88$, $p < 0.05$, $\varphi = 0.24$).

Thinking Level. Figure 2 shows people's judgements about others' thinking level as well as their own thinking level. The results are stark: the most common answer for others' thinking level was K1 (46.4%, marginally significantly above the next most common prediction, 28.6%, goodness-of-fit $\chi^2(2, n = 112) = 4.76$, $p = 0.09$), and the most common answer for one's own thinking level was K2 (49.1%, significantly above the next most common response, 16.1%, goodness-of-fit $\chi^2(2, n = 112) = 18.75$, $p < 0.001$). Furthermore, the combination of K1 for others and K2 for the self (31.3%) was significantly higher than any other combinations (goodness-of-fit $\chi^2(2, n = 112) = 12.52$, $p < 0.005$, compared with the second most common combination, which was only 9.8%). Consistent with our theory, these results indicate that the modal participant considers others to be K1 and themselves to be K2. (Very few participants considered others to be making random choices—K0 thinkers—and even fewer identified themselves as such.)

Figure 2. Study 6 Thinking Level Results

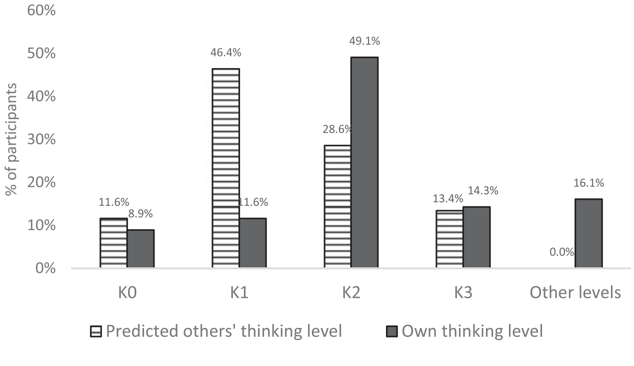

To further test our theory, we compared the responses of those who chose the smaller account (the majority) with the responses of those who chose the larger account (the minority), and we found that the thinking-level pattern we theorized was more prevalent among those who chose the smaller account than those who chose the larger account. Among those choosing the smaller account, most (56.9%) considered others to be K1 and most (76.9%) considered themselves to be K2; each of these percentages was significantly higher than the corresponding percentage among those choosing the larger account (31.9% and 10.6%, respectively, both p's < 0.01).

These results corroborate the results of Study 5 by further suggesting that the undershooting bias reflects the modal participant's belief that others have a lower level of strategic reasoning, choosing the larger account without considering the actions of others, while they themselves are more sophisticated and respond by choosing the smaller account instead.

Study 7

Whereas Study 5 and Study 6 tested our theory without a manipulation and using a postchoice questionnaire, Study 7 manipulated participants' thinking level and tested their choices and reasoning simultaneously. The study used participants from yet another population—students in Canada—and was preregistered at https://aspredicted.org/blind.php?x=cp37ep.

The study comprised two between-participant conditions: control and hint. All participants read a mask-seeking scenario like the one in Study 2 and indicated their choice and reasoning behind this choice. After this, participants were asked to indicate their choice and reasoning behind this choice again. The only difference between the two conditions was that, between the first and second set of questions, participants in the hint condition received a thinking-level manipulating hint asking them to assume that the others were as smart as they were and would reason in a similar way. Participants in the control condition did not receive this information.

We predicted (a) that participants in both conditions would undershoot in the first round and, more importantly, (b) that participants in the control condition would still undershoot in the second round but participants in the hint condition would not.

Method

Participants in this study were students from a large university in Canada who did this and other unrelated studies for course credits. We aimed to recruit 300 participants to ensure that we had at least 150 participants per condition, and we received completed responses from 314 participants ($N_{\text{female}} = 198$; $M_{\text{age}} = 19.63$).

All participants first read the following:

A potentially fatal virus is spreading on a remote island. The virus affects only visitors to the island, and does not affect the islanders because they are already immune to the virus.

You are one of 20 visitors who are stuck on the island and do not know when you can be evacuated. These visitors are strangers to each other and don't talk to each other.

An effective way to prevent one from catching the virus is to wear a mask. Every visitor is selfish, wants to get as many masks as possible and will not share their masks with others. Everyone has learned the following information:

There are only two stores on the island that have masks. The two stores are identical except that one is slightly larger than the other. The larger store has a total of 202 masks and the smaller store has a total of 198 masks. The two stores plan to give all the masks they have free of charge to the visitors.

Each visitor can go to only one store to claim masks, and must decide independently which store to go to.

Each store will distribute all its masks equally to the visitors who go to that store. Therefore, the fewer visitors go to a certain store, the more masks each of them will get. Conversely, the more visitors go to a certain store, the fewer masks each of them will get.

Which store would you go to and why?

(A) *I would go to the larger store, because it has more masks.*

(B) *I would go to the smaller store, because I think most other visitors would go to the larger store.*

(C) *I would go to the larger store, because I think most other visitors would go to the smaller store.*

(D) *Other; please specify*________

The labels A, B, C, and D are added here for ease of reference. Note that answer A corresponded to K1 thinking, answer B to K2 thinking, and answer C to K3 thinking.

After answering, participants in both conditions proceeded to the next page. Those in the hint condition then received the following instruction: "Assume that most other visitors are as smart as you are and would think the same way as you would." They were then asked to answer the above question again. Those in the control condition did not receive the information, but were also asked to answer the above question again. Participants in both conditions were given the same four possible answers as before.

Results and Discussion

Choice. We coded a participant's choice as going to the larger store if they chose answer A or C or if they chose answer D and specified that they would go to the larger store. We coded a participant's choice as going to the smaller store if they chose answer B or if they chose answer D and specified that they would go to the smaller store. If a participant did not specify his/her choice, we considered his/her response as missing. The percentage of such responses was relatively low and not significantly different across conditions: 2.5% and 2.5% during the first and the second rounds in the control condition and 3.8% and 6.4% during the first and the second rounds in the hint condition.

Table 2 presents the choice results. When answering the first time, only 35.9% of the participants (35.3% in the control condition and 36.4% in the hint condition) chose the larger store, significantly below both the equilibrium benchmark of 50.5% and chance ($\chi^2(1, n = 304) = 24.33$, $p < 0.001$, $\varphi = 0.28$), replicating the undershooting effect. When answering the second time, choices diverged in line with our predictions: In the control condition, still only 35.9% chose the larger store, not significantly different from their first choices ($\chi^2(1, n = 153) = 0.03$, $p > 0.5$). This suggests that the mere opportunity to make a choice a second time was not sufficient to change their original choice or overcome the bias. In the hint condition, however, the percentage of participants choosing the larger store rose to 62.6%, significantly above both the equilibrium benchmark and chance ($\chi^2(1, n = 147) = 9.31$, $p < 0.01$, $\varphi = 0.25$), revealing an *overshooting* bias. Although our theory did not predict an overshooting bias, this result supports our prediction that the thinking-level manipulating hint would turn off undershooting bias. Below, we provide evidence that this change in behavior was driven by a change in strategic reasoning.

Thinking Level. We classified a participant's thinking level to be K1, K2, K3, or other levels if they chose answer A, B, C, or D, respectively. Figure 3(a) shows the percentages of participants fitting each of these levels based on their initial answer. Because the hint manipulation had not been introduced, the results of the two conditions were naturally similar. Replicating the pattern from prior studies, the majority of the participants could be classified as K2.

Figure 3(b) shows the percentages of participants fitting each level the second time they answered, after

Table 2. Study 7 Choice Results

Condition	% choosing the larger store in Round 1	% choosing the larger store in Round 2
Control	35.3%	35.9%
Hint	36.4%	62.6%

Figure 3. (a) Study 7 Round 1 Thinking Level Results; (b) Study 7 Round 2 Thinking Level Results

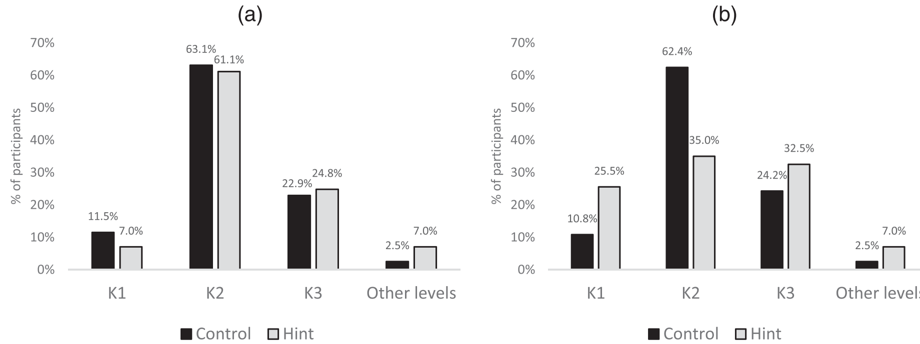

the manipulation was introduced. As the figure reveals, the pattern in the control condition did not change much from the first to the second round, but the pattern in the hint condition changed considerably: the proportion of K2 thinkers in the second round (35.0%) was significantly lower than that in the first round (61.1%, $\chi^2(1, n = 157) = 47.04$, $p < 0.001$, $\varphi = 0.56$), and the overall thinking-level distribution became flatter.[8] These results suggest that the hint was effective in changing people's level of strategic reasoning, which eliminated the undershooting bias.

Study 8

Study 8 was a replication and extension of Study 7. Like Study 7, it manipulated thinking level, but unlike Study 7, it asked (some) participants to first predict others' choices before making their own. We preregistered the study at https://aspredicted.org/blind.php?x=qb5df6.

The study consisted of three between-participants conditions: control (without prediction), prediction, and hint + prediction. All participants first read a pool-choice dilemma involving gold seeking and a choice between a larger mine and a smaller mine. Participants in the control condition were asked to make their choice without being asked to predict others' choices first. Participants in the prediction condition were first asked to predict others' choices and then to make their own decision. Participants in the hint + prediction condition were first given a thinking-level manipulating prompt stating that others were similarly sophisticated as themselves. They then predicted the choices of others and made their own.

We predicted that most participants in the control condition would choose the smaller mine, replicating the undershooting effect. More importantly, based on our theory that the modal player is a K2 thinker, we predicted that most participants in the prediction condition would predict that others would choose the larger mine and choose the smaller mine themselves—just like those in the control condition. In contrast, we predicted that the thinking-level manipulation would change both participants' predictions of others' choices and their own; that is, we predicted that participants in the hint + prediction condition would be less likely to predict that others would choose the larger mine and less likely to choose the smaller mine themselves.

Method

We expected to recruit 450 participants from MTurk so that each condition would have a sample size of around 150. We received completed responses from 450 participants ($N_{\text{female}} = 255$; $M_{\text{age}} = 40.39$). All participants received the following instructions:

> *Imagine the following: You are one of 20 independent gold seekers who have just learned of two newly discovered gold mines. It is estimated that the larger mine contains 51 kilograms (kg) of gold and the smaller mine contains 50 kilograms (kg) of gold.*
>
> *Each of you can get gold from only one of these two mines, and must decide independently which mine to go to. You cannot change your mind afterward and cannot discuss your decision with the other gold seekers.*
>
> *Once each of you have made your choice, you will go to your chosen mine to excavate gold. All the gold in each mine will belong to those who go to that mine, and will be shared equally among them. Thus, the fewer gold seekers go to a certain mine, the more gold each of them will get, and the more gold seekers go to a certain mine, the less gold each of them will get.*

Participants in the control condition were directly asked "Which mine would you go to?"

Participants in the prediction condition were first asked, "First, make a prediction: Which mine do you think most of the other gold seekers would go to?" After they had made their prediction, the participants were asked, "Now that you have predicted which mine most of the other gold seekers would go to, which mine would you go to?"

Table 3. Study 8 Prediction and Choice Results

Condition	% predicting most others would choose the larger mine	% choosing the larger mine themselves
Control	N/A	24.0%
Prediction	74.0%	27.3%
Hint + prediction	49.3%	47.3%

Participants in the hint + prediction condition first received a hint: "Please be aware that like you, the other gold seekers are smart and sophisticated." They were then asked to predict which mine most of the other gold seekers would go to given the hint and then to decide which mine they would go to.

Results and Discussion

The results, shown in Table 3, supported our predictions. In the control condition, only 24.0% of the participants chose the larger mine, significantly below both the equilibrium benchmark of 50.5% and chance ($\chi^2(1, n = 150) = 40.56$, $p < 0.001$, $\varphi = 0.52$).

In the prediction condition, most participants (74.0%) predicted that most others would choose the larger mine ($\chi^2(1, n = 150) = 34.56$, $p < 0.001$, $\varphi = 0.48$), and most (72.7%) chose the smaller mine themselves. Only 27.3% chose the larger mine, significantly below both the equilibrium benchmark and chance ($\chi^2(1, n = 150) = 30.83$, $p < 0.001$, $\varphi = 0.45$). The majority of the participants in this condition (67.3%) both predicted that others would choose the larger mine and chose the smaller mine themselves. These results support our proposition that the modal player is a K2 thinker who expects others to choose the larger pool and responds by choosing the smaller pool themselves. Note that the choice result in the prediction condition (27.3%) was quite similar to that in the control (without prediction) condition (24.0%), which suggests that, even if people are not explicitly asked to predict others' choices before making their own, they perhaps implicitly do so anyway.

Finally, we examine the results in the hint + prediction condition in which the participants received a thinking-level manipulating hint before predicting others' choices and making their own choices. Compared with the participants in the prediction condition who did not receive the hint, significantly fewer participants in the hint + prediction condition predicted that others would choose the larger mine (49.3% vs. 74.0%, $\chi^2(1, n = 300) = 19.30$, $p < 0.001$, $\varphi = 0.25$) and significantly more chose the larger mine themselves (47.3% vs. 27.3%, $\chi^2(1, n = 300) = 12.82$, $p < 0.001$, $\varphi = 0.21$). In fact, the participants in the hint + prediction condition no longer exhibited systematic bias ($p > 0.1$ when comparing 47.3% with the equilibrium benchmark or chance). Although we could not tell whether the choice result in this condition reflected the use of equilibrium strategies, random choice, or a combination of the two, it at least demonstrated that the thinking-level manipulating hint is effective in changing people's beliefs about others' responses and overcoming their own undershooting bias.

General Discussion

This research studies how resource-seekers choose between a larger pool and a smaller pool, knowing that the more (fewer) people who choose a certain pool, the fewer (more) resources each chooser will acquire. Using both an abstract incentive-compatible game and enriched hypothetical scenarios, we document a robust undershooting bias that creates severe resource-allocation inequality, leaving those choosing the larger pool with significantly more resources than those choosing the smaller pool. Drawing on nonequilibrium models of strategic thinking (e.g., Costa-Gomes and Crawford 2006), we attribute the undershooting bias to people's misprediction of the strategies of others. Specifically, we show that the modal player can be classified as a K2 thinker—believing that most others are nonstrategic and would choose the larger pool, and best respond by choosing the smaller pool themselves. This proposition is not only consistent with our main finding (undershooting), but is also supported by the studies that explored the underlying mechanism and aimed to manipulate thinking levels. We devote the remainder of this section to open questions and tentative answers.

How Is the Pool-Choice Dilemma Different from Other Resource Coordination Problems?

As we argued in the Introduction, a key difference is that our paradigm has an apparently superior option, whereas the other games do not. To test this proposition, we conducted a study to compare our pool-choice dilemma with the market entry game, which, like our setting, involves competition over potential resources. We adopted a scenario like the vendor case introduced at the beginning of the article and designed the stimuli so that the two conditions were identical except for the features unique to each setting. In the pool-choice condition, participants chose between going to a larger event or a smaller event; in the market-entry condition, participants chose between going to an event (a risky option) or not going to the event (a riskless option). As predicted, in the pool-choice condition, the larger event was perceived by a significant majority of the participants as the apparently superior option, whereas in the market-entry condition, neither the riskless option (not entering the market) nor the risky option (entering the market) was perceived by a significant majority as apparently superior. Importantly, in line with our prediction that the existence of a salient superior option is key for generating a systematic bias, the pool-choice condition again showed the systematic undershooting bias whereas no systematic bias was observed in the market-entry condition. See Study A1 for details.

Will the Undershooting Bias Disappear If Players Are Visible to Each Other?

According to our proposition that the bias is driven by bounded rationality in strategic reasoning, people undershoot because they mistakenly believe that others do not think as deeply as they do. Prior research argues that people are less likely to have such incorrect beliefs if others are vivid than if they are abstract, because it is easier to identify others with oneself as vividness increases (e.g., Alicke et al. 1995, Hsee and Weber 1997, Small and Loewenstein 2003, Sah and Loewenstein 2012, Steffel and Le Boeuf 2014). Building on this literature, we conjecture that people facing a pool-choice dilemma are less likely to undershoot if they are "vivid" and identifiable to each other. Study A2, reported in the appendix, tested this conjecture. A group of students played a version of the incentive-compatible account-choice game on a social media platform. They were assigned to either a control condition, in which the players were completely anonymous, or a treatment condition, in which the players could see each other's profile and nickname. Participants in the control condition again showed a significant undershooting bias (just like participants in the other studies), but participants in the treatment condition did not. See Study A2 for details.

We expect the undershooting bias to be more prevalent in situations where the resource seekers are anonymous and "abstract" to each other than in situations where the resource seekers are identified and "vivid" to each other. In real life, both types of situations exist and abound. For example, in some situations, ice cream vendors are from different regions and do not know each other; in other situations, they are from the same region and know each other well. In some situations, mask-seeking customers are strangers, and in other situations, they are neighbors. In each example, the undershooting bias is more likely to occur in the former type of situations than in the latter. This discussion highlights the need for future research to identify and explore the boundary conditions for the undershooting bias.

Are There Individual Differences?

We have replicated the undershooting bias among both online workers in the United States and university students from Canada and China. This gives us some confidence that the effect applies to different types of individuals. However, we do not claim that it applies to everyone. For example, we doubt that the effect applies to people who have received game-theory training or are otherwise highly sophisticated. It is possible that they know how to reach the equilibrium solution and do not exhibit a systematic bias. It is also possible that they show an *overshooting* bias—they expect others to be K2 thinkers, choosing the smaller pool, and best respond through K3 thinking, choosing the larger pool instead.

What Will Happen If There Are Only Two Resource Seekers?

The pool-choice dilemma studied in this research always involves many resource seekers (e.g., 20). As such, it is practically impossible that all the resource seekers choose the same pool, leaving the resources in the other pool unclaimed. If there are only a small number of resource seekers, say, only two, there will be a good chance that both resource seekers choose the same pool, leaving the resources in the other pool unclaimed. Because of this difference and other possible differences, people may behave differently in the two-player setting; we are currently investigating this topic in a separate project (Hsee et al. 2020).

Will People Learn Over Repeated Trials?

Suppose that a group of people have played the account-choice game once, learned how many players chose the larger account and how many chose the smaller account, and realized that those choosing the larger account earned more money. They are now given an opportunity to play the game again. Will they still undershoot? Likely not. Based on the results of Nagel (1995) and subsequent research in this area, people learn from feedback by conditioning their strategy on the modal response in the previous period. In our setting, this process predicts that players will learn that most others chose the smaller account in the previous round, and best respond by choosing the larger account in the subsequent round—leading to an overshooting bias. Iterating this process many times will likely lead to behavior conforming to the equilibrium benchmark (see Nagel (1995) and Camerer (2003), for evidence on such convergence). However, this does not undermine the undershooting bias documented in this research, because, as noted earlier, many real-life decisions cannot be repeated, and even if a decision can be repeated, decision makers in real life rarely receive the kind of clear and valid feedback provided in game theory experiments.

Will the Undershooting Bias Disappear If There Were More than Two Pools?

We have focused on pool-choice dilemmas with only two unequal pools (e.g., a larger account and a smaller account). But suppose that players are choosing among four unequal accounts, containing $28, $26, $24, and $22. According to equilibrium predictions, 28% of the players should choose the $28 account, followed by 26%, 24%, and 22% for the respective remaining accounts. However, we predict that a disproportionally small percentage of the players would choose either the largest or the smallest account, and a disproportionally high percentage of the players would choose the third largest account ($24). These predictions are based on a combination of prior research showing extremity aversion (Simonson 1989, Simonson and Tversky 1992, Neumann et al. 2016) and our framework of limited strategic reasoning. If the predicted biases indeed occur, they would advantage the players choosing the largest and the smallest accounts and disadvantage those choosing the third largest account.

Although this work is far from conclusive, it makes at least two contributions: drawing attention to an understudied problem—the pool-choice dilemma—and documenting a systematic bias—undershooting. We hope that this work will spur future work to further probe the psychology behind the effect, explore boundary conditions, and examine the generalizability and robustness of our findings in real-world settings.

Acknowledgments

The authors thank Colin Camerer, Emir Kamenica, Devin Pope, Shane Frederick, Frank Yu, Zheng Gong, and participants in George Wu's laboratory at Chicago Booth for their helpful comments and suggestions during early stages of this project.

Footnotes

[1] There are other equilibrium strategies. For example, if 100 players (labeled 1 through 100) are choosing between a $55 account and a $45 account, then another equilibrium strategy is for players 1–55 to choose the larger account and players 56–100 to choose the smaller account. However, regardless of which equilibrium strategy players adopt, the only equilibrium outcome is that 55% of the players end up choosing the larger account.

[2] Unless otherwise specified, we focus on situations in which the asymmetry between the larger and the smaller accounts is not very high. If the asymmetry is great—for example, if the larger account contains 90% of all the money—then players may choose the larger account even if they believe most others also will choose the larger account.

[3] Our numbering of thinking level is consistent with that by Costa-Gomes and Crawford (2006), Ho et al. (1998), and Nagel (1995), but different from that by Crawford, Gneezy, and Rottenstreich (2008), who define level 0 thinkers (instead of K1) as non-strategic players who would choose the payoff-superior option.

[4] For example, although undershooting in our paradigm is a departure from the equilibrium prediction, it is still Pareto efficient. In contrast, miscoordination in Crawford et al. (2008) is not Pareto efficient.

[5] For example, if one patch contains 60% of the resources and another contains 40%, the proportion of foragers going to the larger patch falls below 60% but not below 50%. Such mild biases could be attributed to a combination of unbiased choices and noise (random choice).

[6] If the choice proportion in a study is significantly below both the equilibrium benchmark (always above 50%) and 50%, we report only the statistics from the test comparing the choice proportion with 50%, because that test is more stringent.

[7] Theoretically, these options could represent even higher thinking levels; for example, option SB could represent K4 thinking. In practice, however, people have rarely been observed to adopt such higher levels of strategic reasoning (Costa-Gomes and Crawford 2006; Nagel 1995).

[8] Curiously, the hint manipulation increased the proportion of K1 thinkers. A possible reason is that the manipulation led some people to feel that it was impossible to predict others' choices, and thereby simply to go to the larger store, since it had more masks. Strictly speaking, these individuals were not K1 thinkers who never thought about others' choices; rather, they thought about others' choices, found them too hard to predict, and ignored them.

References

Alicke MD, Govorun O (2005) The better-than-average effect. Alicke MD, Dunning D, Krueger J, eds. *The Self in Social Judgment* (Psychology Press, New York), 85–106.

Alicke MD, Klotz ML, Breitenbecher DL, Yurak TJ, Vredenburg DS (1995) Personal contact, individuation, and the better-than-average effect. *J. Personality Soc. Psych.* 68(5):804–825.

Arad A, Rubinstein A (2012) The 11–20 money request game: A level-k reasoning study. *Amer. Econom. Rev.* 102(7):3561–3573.

Barasz K, Kim T, John LK (2016) The role of (dis)similarity in (mis) predicting others' preferences. *J. Marketing Res.* 53(4):597–607.

Buhrmester M, Kwang T, Gosling SD (2011) Amazon's mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspect. Psych. Sci.* 6(1):3–5.

Camerer CF (2003) *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton University Press, Princeton, NJ).

Camerer CF, Ho TH, Chong JK (2004) A cognitive hierarchy model of games. *Quart. J. Econom.* 119(3):861–898.

Challet D, Zhang YC (1997) Emergence of cooperation and organization in an evolutionary game. *Phys. A Statist. Mech. Appl.* 246 (3-4):407–418.

Chmura T, Güth W (2011) The minority of three-game: An experimental and theoretical analysis. *Games* 2(3):333–354.

Costa-Gomes MA, Crawford VP (2006) Cognition and behavior in two-person guessing games: An experimental study. *Amer. Econom. Rev.* 96(5):1737–1768.

Crawford VP, Iriberri N (2007) Fatal attraction: Salience, naivete, and sophistication in experimental" hide-and-seek" games. *Amer. Econom. Rev.* 97(5):1731–1750.

Crawford VP, Gneezy U, Rottenstreich Y (2008) The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures. *Amer. Econom. Rev.* 98(4):1443–1458.

Goldstone RL, Ashpole BC (2004) Human foraging behavior in a virtual environment. *Psychonomic Bull. Rev.* 11(3):508–514.

Goldstone RL, Ashpole BC, Roberts ME (2005) Knowledge of resources and competitors in human foraging. *Psychon. Bull. Rev.* 12(1):81–87.

Goldstone RL, Gureckis TM (2009) Collective behavior. *Topics Cognitve Sci.* 1(3):412–438.

Harper D (1982) Competitive foraging in mallards: "Ideal free" ducks. *Anim. Behav.* 30(2):575–584.

Hills TT (2006) Animal foraging and the evolution of goal-directed cognition. *Cognitive Sci.* 30(1):3–41.

Hills TT, Todd PM, Lazer D, Redish AD, Couzin ID, Bateson M, Cools R, et al. (2015) Exploration vs. exploitation in space, mind, and society. *Trends Cognitive Sci.* 19(1):46–54.

Ho TH, Camerer CF, Weigelt K (1998) Iterated dominance and iterated best response in experimental "p-beauty contests." *Am. Econ. Rev.* 88(4):947–969.

Hogarth RM, Lejarraga T, Soyer E (2015) The two settings of kind and wicked learning environments. *Current Directions Psych. Sci.* 24(5):379–385.

Hsee CK, Weber EU (1997) A fundamental prediction error: Self–others discrepancies in risk preference. *J. Experiment. Psych. General* 126(1):45–53.

Hsee CK, Li X, Imas A, Zeng Y (2020) Biases in a Two-Player Pool-Choice Dilemma. Working paper, University of Chicago, Chicago.

Kahneman D (1988) Experimental economics: A psychological perspective. Tietz R, Albers W, Selten R, ed. *Bounded Rational Behavior in Experimental Games and Markets* (Springer, Berlin), 11–18.

Kennedy M, Gray RD (1993) Can ecological theory predict the distribution of foraging animals? A critical analysis of experiments on the ideal free distribution. *Oikos* 68(1):158–166.

Linde J, Sonnemans J, Tuinstra J (2014) Strategies and evolution in the minority game: A multi-round strategy experiment. *Games Econom. Behav.* 86:77–95.

Melkonyan T, Zeitoun H, Chater N (2018) Collusion in Bertrand vs. Cournot competition: A virtual bargaining approach. *Management Sci.* 64(12):5599–5609.

Moore DA, Healy PJ (2008) The trouble with overconfidence. *Psych. Rev.* 115(2):502–517.

Nagel R (1995) Unraveling in guessing games: An experimental study. *Amer. Econom. Rev.* 85(5):1313–1326.

Neumann N, Böckenholt U, Sinha A (2016) A meta-analysis of extremeness aversion. *J. Consumer Psych.* 26(2):193–212.

Sah S, Loewenstein G (2012) —ore affected = more neglected: Amplification of bias in advice to the unidentified and many. *Soc. Psych. Personality Sci.* 3(3):365–372.

Schelling TC (1960) *The Strategy of Conflict* (Harvard University Press, Cambridge, MA).

Scopelliti I, Morewedge CK, Mccormick E, Min HL, Lebrecht S, Kassam KS (2015) Bias blind spot: structure, measurement, and consequences. *Management Sci.* 61(10):2468–2486.

Selten R, Chmura T, Pitz T, Kube S, Schreckenberg M (2007) Commuters route choice behaviour. *Games Econom. Behav.* 58(2):394–406.

Simonson I (1989) Choice based on reasons: The case of attraction and compromise effects. *J. Consumer Res.* 16(2):158–174.

Simonson I, Tversky A (1992) Choice in context: Tradeoff contrast and extremeness aversion. *J. Marketing Res.* 29(3):281–295.

Small DA, Loewenstein G (2003) Helping a victim or helping the victim: Altruism and identifiability. *J. Risk Uncertainty* 26(1):5–16.

Sokolowski MBC, Tonneau F (2004) Human group behavior: The ideal free distribution in a three-patch situation. *Behav. Processes* 65(3):269–272.

Stahl DO, Wilson PW (1994) Experimental evidence on players' models of other players. *J. Econom. Behav. Organ.* 25(3):309–327.

Steffel M, Le Boeuf RA (2014) Overindividuation in gift giving: Shopping for multiple recipients leads givers to choose unique but less preferred gifts. *J. Consumer Res.* 40(6):1167–1180.

Sundali JA, Rapoport A, Seale DA (1995) Coordination in market entry games with symmetric players. *Organ. Behav. Hum. Decision Processes* 64(2):203–218.

Taylor SE, Brown JD (1988) Illusion and well-being: a social psychological perspective on mental health. *Psych. Bull.* 103(2):193–210.

Thaler RH (2015) *Misbehaving: The Making of Behavioral Economics* (WW Norton, New York).