

Rationalizing the Past: A Taste for Consistency

Erik Eyster
London School of Economics
`e.eyster@lse.ac.uk`

30 December 2002

Abstract

People dislike knowing that they would be better off having made different past choices. This paper models the behavior of someone with a taste for rationalizing past actions by taking current actions for which those past actions were optimal. When past and present actions are strategic complements, then a past action higher than optimal leads to a present action also higher than optimal—the sunk-cost effect—and a past action lower than optimal leads to a present action also lower than optimal—an “unsunk-cost effect.” A taste for consistency gives rise to the declining-price anomaly in sequential auctions. A rational monopolist may best respond to consumers who rationalize the past by limiting the menu of tariffs it offers them. People who underestimate their future taste for consistency overbid in wars of attrition and procrastinate in search models.

Acknowledgments For helpful comments and discussions, I thank Bob Anderson, Mark Armstrong, Jimmy Chan, Aaron Edlin, Rui de Figueiredo, Stefano della Vigna, Roman Inderst, Thomas Kittsteiner, Paul Klemperer, Botond Köszegi, David Laibson, Ulrike Malmendier, Meg Meyer, Mike Peters, Dražen Prelec, Kevin Roberts, Eldar Shafir, Chris Shannon, Rebecca Stone, and especially Matthew Rabin, as well as seminar participants at the BRC Conference on Behavioral Economics, Berkeley, Bristol, Cambridge, Oxford, and UCL.

1 Introduction

People dislike knowing that they would have been better off had they made different past choices. Rather than simply regret the past, people may do things in the present to make their past choices look better. For example, someone who has paid a lot of money to join a health club may exercise when injured to justify the money spent. A labor union that has conducted a long strike may reject a contract that it would have accepted before the strike began to warrant the lost wages. A country that has lost many soldiers may continue to fight a war that it would never otherwise enter in order to validate the soldiers' deaths.

This paper models the behavior of a decision maker with a taste for rationalizing her past choices through current choice. People who are uncertain about the consequences of their actions often make choices that later they wish they could change. A choice that was optimal *ex ante* (before learning its payoff consequences), often turns out to be suboptimal *ex post*. Someone who learns that a past choice was suboptimal should recognize that it is “water under the bridge” and not allow regret to affect her current choice. Yet a large literature in psychology indicates that people do not always behave in this way. Not only do people regret past behavior, but they also seek to rationalize it through current behavior. A past choice that was suboptimal given one current action may not have been suboptimal given another current action. If so, then a person can rationalize her past choice by changing her current action; often someone can choose a current action consistent with her past choice having been optimal.

In Section 2, I review the psychology evidence for this type of behavior. One example of this phenomenon is the widely-recognized “sunk-cost” effect: costs that have been paid (or “sunk”) in the past affect people’s behavior in ways that economic theory predicts that they do not. Thaler (1980, p. 11) provides the following example:

“A family pays \$40 for tickets to a basketball game to be played 60 miles from their home. On the day of the game there is a snowstorm. They decide to go anyway, but note in passing that had the tickets been given to them, they would have stayed home.”

Standard economic theory predicts that whether the family bought the tickets itself or received them as a gift would not affect its decision to attend the game. But if the family skips the game, then even though buying the tickets may have been the right thing to do at the time, it later turns out to have been a “mistake”: a cheaper way to skip the game would have been by not buying the tickets in the first place. On the other hand, if the family attends the game, then buying the tickets was optimal: there was no better way to attend the game than by buying tickets. Thus, the family can rationalize its past choice to buy tickets by attending a

game that otherwise it would not want to attend; attending the game is consistent with the choice to buy tickets having been optimal, whereas not attending the game is not. Had the family been given the tickets, it would have no incentive to attend the game to rationalize past choice.¹

Section 3 presents a formal model that captures a concern for rationalizing past choice. To illustrate the model in the context of Thaler’s example, suppose that tickets cost \$40, that attending the game is worth \$80 but costs \$100 in the snow (and nothing otherwise). Suppose in addition that the family cannot attend the game without buying tickets in advance, that the family bought tickets, and that on gameday it snows.² In my model, by skipping the game the family receives a utility that is a linear combination of (i) its actual payoff, -40 (the cost of the tickets) and (ii) the difference between its actual payoff (-40) and the best payoff it could get *given that it skips the game* (0 , by not buying tickets). By attending the game the family receives a utility that is a linear combination of (i) its actual payoff, -60 (the value of attending the game minus the combined costs of tickets and attending the game) and (ii) the difference between its actual payoff (-60) and the best payoff it could get *given that it attends the game* (-60). In this way, the family trades the “material” loss from attending a game that it would rather not attend off against the “psychological” loss from buying tickets that it does not use. If the family puts weight α on (i) and β on (ii), then not attending yields the family $(\alpha + \beta)(-40)$ while attending yields $\alpha(-60)$. If $\alpha < 2\beta$, then the family attends the game; namely, it exhibits the sunk-cost effect.

This type of behavior incorporates two departures from standard decision theory. Firstly, why should someone regret a decision that was optimal at the time it was made but later turns out to be suboptimal? Secondly, why should such regret affect current decisions? A person may care about whether past decisions were optimal purely for consumption value—she dislikes thinking that she would be better off if she could change the past—or because she wants to believe that when making her past decision she had good information about its payoff consequences, in which case she would not have done something that turns out to be suboptimal. (She also may wish to convince someone else that she had good information.) Reducing regret by manipulating current choice involves a form of self deception. Thaler’s family deceives itself into believing that it wants to attend the game, which then rationalizes its choice to buy tickets. In a sense, it refuses to accept that it has made a “mistake” by acting

¹Throughout, I use the term “sunk-cost effect” (Thaler (1980)) rather than “sunk-cost fallacy” in order to distinguish the behavior I describe from people’s tendency to confuse average with marginal effects. I believe that this is a distinct phenomenon, a good example of which is Gottfries and Hylton’s (1987) survey showing that many MIT students on a meal plan with declining average prices would switch to another with higher marginal prices but lower average prices under the mistaken belief that doing so saves money.

²Assuming that tickets cost more on the day of the game than beforehand works equally well.

in such a way that it did not. Such self deception seems especially likely when preferences over current actions are more malleable than those over past actions. Attending a basketball game in the snow has no obvious value, whereas not using a \$40-dollar ticket has an obvious cost.

This paper extends my interpretation of Thaler’s story to model the behavior of someone who attempts to rationalize her past choice through current choice. In the basic version of my model, a decision maker chooses an action in each of two periods. In the first period, she faces uncertainty about a “state of nature” that affects her payoffs. After choosing her first-period action, the decision maker learns the state of nature. In the second period, she chooses another action, at which point she receives a payoff that depends upon both of her actions and the state. Standard decision theory calls for the decision maker in the second period to maximize her payoff given her first-period action and the state. By contrast, a decision maker who cares only about rationalizing her first-period choice chooses the second-period action for which her first-period action was as close as possible to optimal. In other words, she chooses her second-period action to minimize her regret from her first-period action, namely the difference between her actual payoff and the best payoff she could obtain by re-choosing her first-period action. A decision maker who cares both about payoff maximization and rationalizing past choice trades one off against the other.

Section 4 describes several features that a choice environment must exhibit in order for a preference for consistency to affect choice. Firstly, there must be meaningful choice in the first period, for otherwise the decision maker has nothing to rationalize in the second period. Secondly, the decision maker must face meaningful uncertainty, for otherwise she can choose *ex post* optimal actions in both periods. Thirdly, the decision maker’s payoff function must include some complementarities (positive or negative) between first- and second-period actions (or the set of choices available to the decision maker in the second period must depend upon first-period choice) so that by manipulating second-period choice the decision maker can rationalize first-period choice.

The decision maker’s first-period behavior depends crucially upon her beliefs about her future preference for consistency. Suppose, for instance, that $\alpha = \beta = 1$ in Thaler’s story, so the family attends the game in the snow. Let μ be the probability of snow. Then a “sophisticated” decision maker who correctly predicts her future preference for consistency recognizes that if she purchases the tickets her expected utility is $(1 - \mu)40 + \mu(-60)$, which exceeds her payoff from not purchasing tickets, 0, if $\mu \leq \frac{2}{5}$. A “fully naive” decision maker who incorrectly believes that $\beta = 0$ and therefore that she will not care about consistency in the future thinks that she will attend the game if and only if it does not snow. She perceives

her expected utility from buying the tickets to be $(1 - \mu)40 + \mu(-40)$, which exceeds her utility from not buying tickets if and only if $\mu \leq \frac{1}{2}$. Thus, for $\mu \in (\frac{2}{5}, \frac{1}{2})$ a fully naive decision maker purchases tickets when she should not because she incorrectly predicts that she will neither attend the game in the snow nor regret not going.

More generally, a decision maker may recognize that she will care about consistency in the future but naively underestimate the extent to which that is the case. A decision maker who is anything other than fully-naive in the first-period recognizes that her second-period choice may depend upon her first-period choice set because how much she regrets her first-period choice depends upon her first-period choice set. As a result, her first-period choice may depend upon so-called irrelevant alternatives in her choice set. For example, a consumer choosing from a menu of tariffs may have preferences that depend on her choice set because how much she ultimately regrets choosing one tariff depends on how much cheaper she could consume the same amount under another tariff. Indeed, anytime a decision maker who cares about consistency chooses x over y from some first-period choice set S , while a decision maker who does not care about consistency chooses y over x from S , there exists some choice set $T \supset \{x, y\}$ from which both decision makers choose y . Thus, someone whose taste for consistency influences her choice from some choice set violates the Weak Axiom of Revealed Preference.

Section 5 shows how a preference for consistency captures the famous sunk-cost effect. To do this, I define a two-period decision problem to have strategic complements if in each possible state of the world, raising the first-period action increases the decision maker's incentive to raise her second-period action. One example is an investor choosing how much to invest in some project in each of two periods, where higher first-period investment increases her incentive to invest in the second period. In this setting, if the investor turns out to have invested “too much” in the first period—more than would be optimal given that in the second period she maximizes her payoff given her actual first period investment—then a preference for consistency leads her to increase her investment in the second period; the investor rationalizes her overly high first-period investment by investing more in the second period, the sunk-cost effect. If the investor turns out to have invested too little in the first period, then she rationalizes her overly-low first-period investment by investing less in the second period, a kind of “unsunk-cost effect.”

One application of these concepts is to a consumption setting where someone first chooses one of two possible two-part tariffs, then learns how much she values consumption, and finally chooses how much to consume. Someone who has chosen a tariff with a high fixed price and low marginal price has incentive to raise her demand to rationalize her choice of tariff, the

sunk-cost effect. Someone who has chosen a tariff with a low fixed price and high marginal price has incentive to lower her demand to rationalize her choice of tariff, the unsunk-cost effect. In a simple example, I show that when the unsunk-cost effect dominates, a monopolist may want to limit the number of tariffs it offers to consumers, even when consumers are fully naive and fail to account for the fact that more tariffs lead to more regret. A similar effect may operate in sequential auctions. Someone who has passed up a cheap opportunity to purchase some good may be reluctant to pay more for the same good in some subsequent auction. This accords with Ashenfelter's (1989) "declining-price anomaly": prices tend to decline in sequential auction, which conflicts with Bayesian Nash analysis.

Section 6 extends the model to many periods and applies it to a war of attrition as well as a model of search. In the many-period version of the model, the decision maker prefers profiles of actions in which each action is optimal given all other actions. Specifically, in the last period, the decision maker maximizes a linear combination of her payoff and the total regret from all of her past actions, namely how much better she could do by changing her first-period action keeping all others fixed, how much better she could do by changing her second-period action keeping all others fixed, etc. In a war of attrition (or Shubik's 1971 dollar auction), each period two bidders choose whether they want to remain in the auction or drop out. The auction ends when the first bidder drops out; if the other bidder remains in, then he wins a prize. Remaining in the auction costs each bidder a constant amount every period, so that a bidder who remains in the auction for a while before dropping out regrets her past decisions to stay in. Sophisticated bidders drop out faster the more they wish to avoid regret; intuitively, regret makes them risk averse. Naive bidders, however, may either drop out faster or slower than bidders who do not care about consistency, but when the auction goes on for long enough they eventually drop out slower. When bidders are naive enough, the seller's expected revenue exceeds the value of the prize, a result consistent with classroom and laboratory evidence.

An intuition similar to the unsunk effect explains procrastination in a multi-period search model. Consider someone searching for a job who receives a new offer every period; offers differ in lifetime wages, where each period's wages are drawn from a constant and known distribution. Each period, the worker can either accept that period's offer or continue to search. The job seeker's problem ends when he accepts a job, and his (material) payoff is just the discounted value of the lifetime wages he accepts. A worker who receives several bad early wage offers may find himself in a position where unlikely to find a job whose discounted lifetime wages exceed what he could have received had he accepted his first offer. A preference for consistency leads sophisticated workers to act as if they are risk averse, accepting lower

wages than would an expected-utility maximizer. But naive workers who receive bad early wage offers tend to search too long; each period the naive worker thinks that next period she will care less about regret than she does today, which provides her incentive to keep searching. The longer she searches, the more the naive worker has to regret, which further decreases the likelihood that she will accept a job. Section 7 concludes with some problems and extensions with the model.

This paper is related to several distinct literatures in economics. The first literature shows that the sunk-cost effect may be rational behavior in a signaling context, or may emerge through evolutionary dynamics. Prendergast and Stole (1996) show that a decision maker who wishes to convince an audience that she learns quickly may appear to exhibit a sunk-cost effect. In their model, an investor seeks to maximize profits from a project with uncertain profitability while simultaneously convincing the market that she has precise information about that profitability. After many periods, the investor downplays new information because substantial changes in her investment strategy reveal that she has imprecise information. In this sense, the investor becomes committed to a course of action, something akin to the sunk-cost effect. In certain settings, my model may be viewed as a reduced-form approach to a problem that includes Prendergast-Stole-style signaling motives. Carmichael and MacLeod (2000) show that the sunk-cost effect may help solve the hold-up problem: someone who sinks a cost in making a relationship-specific investment and is known to be unwilling to agree to any trade that does not allow her to recoup the cost of her investment is invulnerable to the hold-up problem.

The premise that people dislike knowing that they would have been better off having made different past choices corresponds to an economics literature on “regret theory” introduced simultaneously by Bell (1982) and Loomes and Sugden (1982, 1983, 1986). In their models, a decision maker choosing among lotteries may not choose the lottery that maximizes her expected utility if it causes her more *ex post* regret than some other lottery. Both papers argue that this can explain Allais’s (1953) famous anomaly: people prefer \$1 million for sure to a lottery giving a ninety-percent chance of \$5 million and a ten-percent chance of nothing, but they prefer a nine-percent chance of \$5 million and a ninety-one percent chance of nothing to a ten-percent chance of \$1 and a ninety-percent chance of nothing. These preferences are inconsistent with expected-utility theory (they are not linear in probability), but may be explained by regret: someone who foregoes the sure \$1 million to take the risky lottery and comes away with nothing would regret having passed up \$1 million, while someone who takes the \$1 million may never know whether she would have received the \$5 million or nothing.

In a many-period setting, how much someone regrets choice in one period may depend

upon what she does in another. The crucial issue in extending regret theory to multiple periods is whether in the second period people treat regret as “sunk” and choose their action to maximize their payoff, or whether regret affects their choice. This paper takes the latter approach: one argument is that someone who initially is willing to sacrifice future expected payoff to reduce expected regret ought to be willing later to sacrifice current payoff to reduce regret. However, a decision maker who thinks she will only care about regret a small amount in the future thinks that a taste for consistency will have minimal effect on future actions. In this case, the only effect that a taste for consistency has on her preferences over first-period actions is to push her towards favoring those first-period actions that lead to the least expected regret.

Several papers have incorporated Festinger’s (1957) seminal concept of cognitive dissonance into economics.³ People experience unpleasurable cognitive dissonance when they hold two opposing beliefs concurrently; they can reduce dissonance by changing one of those beliefs. Akerlof and Dickens (1982) present a model of cognitive dissonance where workers in a dangerous industry choose whether to use safety equipment.⁴ Because believing that her job is safe conflicts with wearing a safety helmet, a worker may choose not to wear a helmet to reassure herself of her safety. Rabin (1994) models people who choose their moral standards as well as how much of an immoral activity to engage in; people may choose to convince themselves that an immoral activity is moral in order to engage in it, but this involves some cost. Rabin and Schrag (1999) model “confirmatory bias,” where people who believe that A is more likely than B sometimes misread evidence in favor B as supporting A instead. While Rabin and Schrag do not explain why people do this, surely part of the explanation is that people dislike seeing previously maintained hypotheses discredited. Yariv (2002) addresses this by modeling someone forming beliefs and taking actions over time. A person who believes that A is more likely than B today and acts accordingly would like to believe A more likely than B tomorrow as well. All of these models share the feature that someone’s desire to hold positive beliefs about herself—be it that she is knowledgeable, prudent, or moral—conflicts with information that she learns. As a result, people may not learn from new information in order to maintain positive self image. This paper takes a very different tack: rather than reduce dissonance by maintaining beliefs under which her previous choices were optimal, a decision maker avoids recognizing that a past action was suboptimal by acting in such a way that it was not: in essence, she changes her current preferences in such a way as to take an action for which the past action was optimal. Whatever the psychological merits of this approach—some of which are discussed in the next section—because my model operates

³Aronson (1995) provides a nice introduction to dissonance and self justification.

⁴In the conclusion, I discuss why this behavior does not fit my model.

through preferences in a simple way, it can be applied to a wide range of decision problems, unlike the models of cognitive dissonance.

2 Evidence

The psychology literature contains a wealth of examples of the sunk-cost effect (often called “escalation of commitment” or “entrapment”) and other instances of rationalizing past behavior.⁵ The most convincing single experiment comes from Arkes and Blumer (1985), who provided funds for the Ohio University Theatre to give discounts on season tickets randomly to the first sixty people to buy season tickets. One-third of patrons paid the full price of \$15; one-third received a \$2 discount; and one-third received a \$7 discount. Patrons without discounts attended an average of 4.11 of the first five plays, while those with \$2 and \$7 discounts attended an average of 3.32 and 3.29 plays, respectively. (Both differences are significant at the five-percent level.)⁶ As in Thaler’s story, the more someone pays for a ticket, the more incentive that person has to use the ticket.

Arkes and Blumer (1985) obtain the same qualitative result using survey data. They presented Ohio and Oregon college students with the following hypothetical situation.

“Assume that you have spent \$100 on a ticket for a weekend ski trip to Michigan. Several weeks later you buy a \$50 ticket for a weekend ski trip to Wisconsin. You think that you will enjoy the Wisconsin ski trip more than the Michigan ski trip. As you are just putting your just-purchased Wisconsin ski trip ticket in your wallet, you notice that the Michigan ski trip and the Wisconsin ski trip are on the same weekend! It’s too late to sell either ticket, and you cannot return either one. You must use one ticket and not the other. Which ski trip will you go on?” (Arkes and Blumer, 1985, p. 126.)

Thirty-three of the sixty-one subjects chose to go to Michigan, despite the fact that they think that they would enjoy the Wisconsin trip more. People regret wasting a \$50 ticket more than a \$100 ticket.

Empirically testing for the sunk-cost fallacy is complicated by the fact that information-based explanations for behavior are often difficult to rule out. The most convincing study comes from Camerer and Weber (1999), who corroborate Staw and Hoang’s (1995) earlier finding that National Basketball Association (NBA) players play more minutes the earlier

⁵In animal behavior, the sunk-cost effect goes by the term “Concorde Effect,” and its existence has been actively debated for thirty years (see Trivers, 1972; Dawkins and Borckmann, 1980; and Arkes and Ayton, 1999).

⁶Attendance at the last five plays averaged only 2.09 and hardly varied across treatments. One explanation for this might be that people care more about being consistent with actions taken in the recent past than in the distant past.

they are drafted, controlling for performance.⁷ A team that used a valuable draft pick to acquire a player has more incentive to play the player to rationalize its pick. Using data from 1986 to 1991, Camerer and Weber control for players' quality using past performance, a pre-draft scouting report, and predicted performance as a function of draft number; the scouting report proxies for priors on the players' abilities, and predicted performance controls for the informational value of draft order. Camerer and Weber find that being drafted ten places earlier (out of approximately 50 players drafted) increases playing time by approximately 200 minutes over the course of the season, which exceeds ten percent of the average rookie's playing time. This effect diminishes significantly for players traded after being drafted, where the team had no incentive to rationalize the draft decision. Nevertheless, it is hard to rule out the hypothesis that the team that drafts the player starts with beliefs about the player's ability that are too optimistic and only learns very gradually.

Dick and Lord (1998) provide strong evidence of the sunk-cost effect in a marketing experiment on the effects of membership fees on shopping. Eighty business-school students started with 300 "utility" points and were told that the ten subjects finishing the experiment with the most utility would win \$25. Shopping at Store A gave an average utility of 40 points, Store B 20 points, and Store C 15 points, though the actual utility was stochastic.⁸ Half of subjects were informed that they could shop at Store A if and only if they paid a one-off membership fee of 100 points. The other half did not have to pay this membership fee. Virtually all subjects in both treatments started shopping at Store A. At the beginning of either period 3 or period 7, subjects were told that "a change in governmental regulations concerning dealing with Japanese distributors of videotapes" had lowered the average utility from shopping at Store A—the only store using Japanese distributors—to the average utility from shopping at Store B.⁹ Of those subjects informed of the utility change in period 3, subjects who paid the membership fee shopped at Store A an average of 6.25 times in the final eight rounds, while those who did not pay the fee shopped at Store A an average of only 4.65 times: subjects who paid the membership fee continued to shop at Store A to rationalize paying the membership fee. Of those subjects informed of the utility change in period 7, subjects who paid the membership fee shopped at Store A an average of 2.5 times in the last four rounds, while those who did not pay the fee shopped at Store A 2.55 times. In this treatment, there was no need to rationalize paying the membership fee, which had already

⁷Players entering the NBA are chosen by teams in a draft where teams pick a limited number of players in a preassigned order.

⁸The three stores also varied in utility-irrelevant ways.

⁹Given the experiment's tournament structure, equilibrium payoff-maximizing behavior surely calls for some subjects to shop at both Stores A and B (and probably Store C), an intuition familiar from the racing literature.

provided a net benefit from it of $6 \cdot (40 - 20) - 100 = 20$, the difference in utility relative to the best competitor, Store B, less the membership fee.

A large literature in organizational behavior discusses sunk-cost effects (see Bazerman 2002 for an excellent survey). Bazerman *et al* (1982) report an experiment where one half of subjects were asked to promote one of three managers to regional director of a hypothetical corporation on the basis of various performance data; the other half of subjects did not participate in this stage of the experiment. All subjects were then presented with evidence that the new regional director performed poorly in the two years following promotion. The subjects were then asked whether they wanted to raise the director's salary, increase her vacation time, promote her again, etc. On all of these counts, subjects who initially promoted the director were more likely to take actions consistent with their promotion decision having been the right one. For example, subjects who promoted the employee offered her an average raise of 9.72% as compared to an average of 8.87% by those who had no role in her hire (subjects were told that the firm averaged a 10% raise).¹⁰ Barron, Chulkov, and Waddell (2001) use Compustat's ExecuComp's database of 2,226 publicly-traded firms from 1992 to 1999 to examine the relationship between executive turnover and project termination. In the year following the departure of a CEO or other top executive, firms are significantly more likely terminate an operation; that is, they may end projects that exist only because of the sunk-cost effect. The effect is even stronger two years after the CEOs departure, which suggests that the relationship is not due to firms first deciding to terminate unsuccessful projects and then firing the executives responsible.

Bastardi and Shafir (1998) show that people tend to pursue unnecessary information, which they the use in order to rationalize having sought it.¹¹ In one experiment, they ask 261 subjects to

“Imagine that you are on the admissions committee of Princeton University. You are reviewing the file of an applicant who plays varsity soccer, has supportive letters of recommendation, and is editor of the school newspaper. The applicant has a combined SAT score of 1250 and a high school average of a B.” (Bastardi and Shafir, 1998, pp. 21-22.)

When asked whether they would like to accept or reject the candidate, 57% choose accept and 43% reject. Bastardi and Shafir told a second group of 278 subjects the same story,

¹⁰Schoorman (1988) analyzes supervisors' evaluations of approximately 400 secretaries, finding that evaluations are correlated with whether the supervisors were personally involved in hiring decisions. Supervisors who previously opposed to hiring a worker gave that worker lower average performance evaluations than supervisors uninvolved in hiring the worker; supervisors who previously favored hiring the worker gave higher performance evaluations. While these findings are suggestive, they can also be explained by the fact that supervisors involved in hiring have better information about the workers.

¹¹See also Bastardi and Shafir (2000).

except that now “you have two conflicting reports on the applicant’s high school average grade. The guidance counselor’s report indicates a B average, while the school office reported an A average.” These subjects were asked whether they wanted to accept or reject the applicant, or instead wait a few days for the school to check its records for clarification. In this case, 21% chose to accept, 5% to reject, and 74% to wait for more information. Those who waited were told that the average grade was a B and asked whether they wanted to accept or reject. 34% chose accept and 66% reject. Bastardi and Shafir note that this behavior conflicts with Savage’s sure-thing principle: someone willing to accept based on a B average (and presumably an A average as well) should not gather costly information about the candidate’s grade. More germane to this paper, despite having the same final information as those subjects in the control group, subjects who sought information were much more likely to reject the applicant. One interpretation of this experiment is that people seek information to satisfy curiosity, but then feel compelled to rationalize acquiring it by using it: employing the admissions rule “accept iff A” rationalizes acquiring the information because it is implementable only after learning the applicant’s grade.

Laboratory evidence suggests not only that people display a sunk-cost effect but also an unsunk-cost effect. Tykocinski, Pittman, and Tuttle (1995) report an experiment with 120 student participants where

“[T]he participant was considering joining a frequent flyer program before taking a trip for the holidays. If the participant decided to join, he or she would accumulate 5500 miles (toward a free ticket, on accumulating 20,000 miles). In the two experimental conditions, the participant was also said to have considered joining this program once before, near the beginning of the year, but had not done so. The number of miles that would have been accumulated had the participant joined then was manipulated. In the large-difference condition, if the participant had joined at the beginning of the year, after this current trip he or she would have accumulated 15,500 miles, versus 7,500 miles in the small-difference condition.” (Tykocinski *et al*, pp. 795-796.)

On a scale from 0 (“not at all likely”) to 10 (“extremely likely”) participants assigned an average likelihood of joining the program of 5.63 in the large-difference treatment, and 7.45 and 7.36 in the small-difference and control treatments, respectively; assessments in the large-difference treatment were significantly lower than in the other two. A past decision not to join a frequent-flyer program is only a mistake if one joins now; hence, the more miles one passed up, the less inclined one is to join today.¹²

A natural question is whether people make choices that are consistent with their previous actions being optimal to rationalize those past choices or because they misread new information to conform to their previous beliefs. That is, do past actions determine current

¹²See Tykocinski and Pittman (1998) for related findings.

beliefs, and current actions then follow; or do past actions determine current preferences, and beliefs then follow? To address this issue, Arkes and Hutzel (2000) presented subjects with the following problem:

“As the director of research in an airline company, you have invested 90% (10%) of the company’s research money into a research project...to build a plane that would not be detected by conventional radar... Just recently, another firm began marketing a plane that cannot be detected by radar. Also, it is apparent that their plane is much faster and more economical than the plane your company is building.”

Subjects were asked to assess the likelihood of the project’s success with and without investing the rest of the research budget, as well as whether they would invest. Arkes and Hutzel varied the order in which they asked these questions. Subjects first asked whether they wanted to invest provided an average estimate of the likelihood of success of 0.57 while subjects first asked to estimate the likelihood of success on provided an average estimate of 0.50 ($p < 0.02$). This difference suggests that beliefs follow actions and not the other way around: having sunk a cost changes subjects’ preferences, which in turn change their beliefs.

A recent literature on projection bias or empathy gaps relates to this paper in two ways. Firstly, experiments show that people predict that their future preferences correspond more closely to their current preferences than those future preferences do. Loewenstein and Adler (1995) explore whether people correctly predict the extent to which their preferences change with their endowments. In one treatment, subjects without mugs predicted on average that once endowed with mugs they would name a selling price of \$3.27 for the mugs, whereas in fact they asked an average of \$4.56. This accords with other evidence that people underestimate the extent to which their future preferences differ from their current preferences. While it does not provide direct evidence that people underestimate the extent to which they care about future regret, this literature does suggest the direction of any possible bias in beliefs: people who do not currently regret an action seem more likely to underestimate than overestimate the extent to which they will regret it in the future. Secondly, the very act of asking people to predict their future preferences affects their future choice. Loewenstein and Adler include a control group that simply received mugs and then was asked to name a selling price—they never predicted their future selling price. These subjects demanded an average of \$4.98. Hence, it appears that subjects who are first asked to predict their future selling prices later shade their selling prices in order to reduce the magnitude of their past prediction error.

Not all psychology experiments support the sunk-cost effect. Conland and Garland (1993, 1998) and Boehne and Paese (2000) argue that people may appear to attend to sunk-cost when really they dislike starting projects that they do not complete. They present experiments

designed to show that such a model fits behavior better than either an economic model or a model where people only care about sunk-costs, although a model where people trade the two off such as mine does better than all three. Heath (1995) presents evidence of a countervailing heuristic: people who have previously invested resources in a project are especially unhappy when the total amount invested exceeds the value of the project. Heath shows how this can lead to the very opposite of the sunk-cost fallacy. Consider someone who invested \$6.58 in a project, only to later learn that another investment of \$1.55 is needed for a return of \$6.66. Fifty-eight percent of subjects refused, while no subjects in a control group offered the same marginal opportunity without investing the \$6.58 refused. Heath’s results indicate the presence of a framing effect that may operate on top of the sunk-cost effect.

3 A Two-Period Model

A decision maker (DM throughout) makes a choice in each of two periods, 1 and 2. Ω is a finite set of payoff-relevant “states of the world.” Let μ be a probability measure that represents the DM’s priors over Ω and puts positive weight on every element of Ω . In the first period, the DM takes an action $a_1 \in A_1$. The DM’s action a_1 does not affect the probability distribution over Ω . In the second period, the DM learns the true state of the world, $\omega \in \Omega$, before choosing an action $a_2 \in A_2(a_1, \omega)$. The set of actions available to the DM in the second period may depend both upon the state and the DM’s first-period action. Let $g : \{(a_1, a_2, \omega) : a_1 \in A_1, \omega \in \Omega, a_2 \in A_2(a_1, \omega)\} \rightarrow \mathbb{R}$ be the DM’s “material” payoff function, which depends on her actions in both periods as well as on the state of the world.

A DM with a taste for consistency chooses an action in period 2 to rationalize her choice in period 1. In particular, she dislikes choosing actions in period 2 that cause her to regret her action in period 1. To make this precise requires a definition of what it means for a DM to regret her previous actions. Let $A_1(a_2, \omega) = \{a_1 \in A_1 : a_2 \in A_2(a_1, \omega)\}$ denote the set of first-period choices that are consistent with a_2 being in the DM’s second-period choice set, given the state ω . Throughout, let $a_1^*(a_2, \omega) \equiv \arg \max_{a_1 \in A_1(a_2, \omega)} \{g(a_1, a_2, \omega)\}$.

Definition 1 *The DM’s regret from action $a_1 \in A_1$ in state $\omega \in \Omega$ given that she takes action $a_2 \in A_2(a_1, \omega)$ is $r(a_1|a_2, \omega) \equiv g(a_1^*(a_2, \omega), a_2, \omega) - g(a_1, a_2, \omega)$.*

The DM’s regret from action a_1 in state ω given action a_2 is the difference between the maximum payoff that she could receive if she could choose a_1 again, fixing a_2 and ω , and her actual payoff. In Thaler’s basketball-game example, the family’s regret from buying tickets given that it skips the game in the snow is simply the maximum payoff it could receive by skipping the game in the snow—the payoff it would have received had it not bought the

tickets and skipped the game in the snow—minus its actual payoff. Its regret from buying tickets given that it attends the game is zero, for buying tickets was the best way that it could attend the game. The DM experiences regret only when she could improve her payoff by re-choosing her first-period action.¹³ The most important feature of regret in my model is not its magnitude, but rather how it depends on a_2 .

A DM who cares about consistency may choose a second-period action that does not maximize her payoff given a_1 and ω if that action produces little regret.

Definition 2 *A decision maker is a ρ -rationalizer if for each $a_1 \in A_1$ and $\omega \in \Omega$,*

$$a_2 \in \arg \max_{a_2 \in A_2(a_1, \omega)} \{g(a_1, a_2, \omega) - \rho r(a_1 | a_2, \omega)\},$$

for $\rho \in \mathbb{R}_+$.

When $\rho = 0$, the DM simply maximizes her payoff given her first-period action and the state. When $\rho \rightarrow \infty$, the decision maker chooses a_2 to minimize her regret; that is, she chooses her second-period action such that her *first-period action* comes as close as possible to maximizing her payoff given her *second-period action* and the state of the world.¹⁴ Choosing a_2 in this way “rationalizes” a_1 in the sense that given a_2 , no other action in A_1 yields as high a payoff as a_1 . For $\rho > 0$, the DM trades off maximizing her payoff and rationalizing her first-period action. Throughout, I use the term “utility” to refer to $g - \rho r$, and “payoff” to refer to g . When either $\rho = 0$ or $r(a_1 | a_2, \omega) = 0$ —either the DM does not care about regret or her first-period action was optimal—her utility is normalized to her payoff g .

The ρ -rationalizer’s problem is equivalent to choosing $a_2 \in A_2(a_1, \omega)$ to minimize the loss function

$$-\left(\max_{a_2 \in A_2(a_1, \omega)} \{g(a_1, a_2, \omega)\} - g(a_1, a_2, \omega)\right) - \rho \left(\max_{a_1 \in A_1(a_2, \omega)} \{g(a_1, a_2, \omega)\} - g(a_1, a_2, \omega)\right).$$

A ρ -rationalizer’s loss function consists of two components. The first is that a_2 may be suboptimal given a_1 and ω ; this is the standard part of the loss function that a DM maximizing g would set to zero. The second is the non-standard part of the ρ -rationalizer’s loss function: she wants her first period action, a_1 , to be optimal given a_2 and ω . But because she DM has already chosen a_1 , the only way that she can reduce the loss from a suboptimal a_1 is by manipulating a_2 .

As in most dynamic choice, the DM’s first-period behavior depends crucially upon both her beliefs about her second-period behavior and her current preferences over second-period payoffs.

¹³When there are only two actions in A_1 , Definition 1 coincides with a special case of Loomes and Sugden’s (1982, 1983, 1986) regret theory.

¹⁴Recall that g is only defined for (a_1, a_2, ω) such that $a_1 \in A_1(a_2, \omega)$.

Definition 3 A ρ -rationalizer is ν -naive if

$$a_1 \in \arg \max_{a_1 \in A_1} \left\{ \sum_{\omega \in \Omega} \max_{a_2 \in A_2(a_1, \omega)} \{g(a_1, a_2, \omega) - \rho(1 - \nu)r(a_1|a_2, \omega)\} \mu(\omega) \right\},$$

for $\nu \in [0, 1]$.

When $\nu = 0$, a ρ -rationalizer chooses her first-period action to maximize her expected payoff given correct beliefs about her behavior in the second period; I refer to such a DM as “sophisticated.” The only way that a sophisticated DM departs from expected-payoff maximization is by caring about regret. When $\nu = 1$, the DM chooses her first period action under the incorrect belief that she is a payoff maximizer in the second period; I refer to such a DM as “fully naive.” When $\nu \in (0, 1)$, the DM recognizes that she is a ρ -rationalizer in the second period, but underestimates the extent to which she cares about regret, namely ρ .¹⁵ Clearly, a ν -naive ρ -rationalizer and a fully sophisticated $(1 - \nu)\rho$ -rationalizer behave the same way in the first period.

Whatever ν , Definition 3 imposes the restriction that the DM’s first-period preferences over her second-period payoffs coincide with her beliefs about her second-period preferences. For example, as $\rho \rightarrow \infty$ a ρ -rationalizer cares only about consistency in the second period. If such a ρ -rationalizer is fully naive, then in the first period she thinks that she will not care at all about payoff in the second period; also, in the first period she does not care about second-period payoff. That is, unlike the literature on time inconsistency where people today disapprove of the present-biased preferences they will have tomorrow, a DM with a taste for consistency always approves of her predicted future behavior. She simply mispredicts her future preferences.

Definition 2 characterizes the behavior of a ρ -rationalizer in a setting without payoff uncertainty in the second period. But in many applications, the DM may still face uncertainty in the second period. To model this situation, let \mathcal{P}_1 be a partition of Ω .¹⁶ After making her first period choice, the DM learns which element P_1 of \mathcal{P}_1 contains the state ω ; that is, the DM learns only that $\omega \in P_1(\omega) = \{\omega' \in \Omega : \omega' \in P_1\}$, namely she can rule out some states of the world, but possibly not all. Assume that for each $a_1 \in A_1, P_1 \in \mathcal{P}_1, \omega, \omega' \in P_1$, $A_2(a_1, \omega) = A_2(a_1, \omega')$; within each of the DM’s possible information sets, the set of actions available to her must be constant. Otherwise, knowing which actions are available would enable the DM to refine her information. The natural extension of ρ -rationalization to this

¹⁵The model could easily be extended to include cases where the DM overestimates ρ .

¹⁶ \mathcal{P}_1 is a partition of Ω if $\cup_{P_1 \in \mathcal{P}_1} P_1 = \Omega$ and for each $P_1, P'_1 \in \mathcal{P}_1$, if $P_1 \cap P'_1 \neq \emptyset$, then $P_1 = P'_1$.

context is for each $a_1 \in A_1$, $P_1 \in \mathcal{P}_1$, and ρ , a ρ -rationalizer chooses

$$a_2 \in \arg \max_{a_2 \in A_2(a_1, \omega \in P_1(\omega))} \left\{ \sum_{\omega \in P_1(\omega)} (g(a_1, a_2, \omega) - \rho r(a_1 | a_2, \omega)) \mu(\omega | P_1(\omega)) \right\},$$

where $\mu(\omega | P_1(\omega)) \equiv \frac{\mu(\omega)}{\sum_{\omega' \in P_1(\omega)} \mu(\omega')}$. In the second period, a ρ -rationalizer simply maximizes her expected utility, with utility defined as before. In particular, note that the DM does not consider regret from choosing an a_2 that turns out to be suboptimal.¹⁷

4 General Results

In many settings, a taste for consistency does not affect the DM's decisions. Let $a_1^*(\rho, \nu)$ denote a ν -naive ρ -rationalizer's optimal first-period action. Let $a_2^*(a_1, \omega; \rho)$ denote a ρ -rationalizer's optimal second-period action in state ω following the first-period action a_1 . To minimize notation I assume henceforth that the DM's preferred action in each period exists and is unique (unless otherwise noted).

Proposition 1 *Suppose that A_1 , A_2 , or Ω is a singleton set, or for each $\omega \in \Omega$, $a_1, a'_1 \in A_1$, $a_2 \in A_2$, $g(a_1, a_2, \omega) = g_1(a_1, \omega) + g_2(a_2, \omega)$ for some g_1 and g_2 , and $A_2(a_1, \omega) = A_2(a'_1, \omega)$. Then for each $\rho, \rho' \in \mathbb{R}_+$, $\nu, \nu' \in [0, 1]$, $\omega \in \Omega$, $a_1^*(\rho, \nu) = a_1^*(\rho', \nu')$ and $a_2^*(a_1^*(\rho, \nu), \omega; \rho) = a_2^*(a_1^*(\rho, \nu), \omega; \rho')$.*

Proposition 1 states that a preference for consistency only affects choice when payoffs are uncertain and exhibit complementarities between actions, and there is choice in both periods.¹⁸ Without choice in the first period, there can be no issue of rationalizing first-period actions through second-period actions. Without uncertainty, the DM trades the cost of choosing a_1 that is suboptimal given a_2 off against the cost of choosing a_2 that is suboptimal given a_1 . A naive DM misperceives the way that she trades off these two losses in the second period and for some a_1 may mispredict $a_2^*(a_1, \omega; \rho)$. However, by choosing the optimal first-period action a_1^* she need not trade off these losses, for a_2^* maximizes both $g(a_1^*, a_2, \omega)$ and $g(a_1^*, a_2, \omega) - \max_{a_1 \in A_1} g(a_1, a_2, \omega)$. Without choice in the second period, a preference for consistency clearly cannot affect actions in the second period. Likewise, if the payoff function is separable in a_1 and a_2 , then the extent to which a_1 is suboptimal in any given state does not depend upon a_2 , and therefore the DM cannot use a_2 to rationalize a_1 (as long as A_2

¹⁷Proposition 1 in the next section implies that adding a third period with a singleton action space would not lead to regret concerns that affect second-period choice. That is, only regret that varies with future choice affects current choice.

¹⁸Unless in the text, all proofs are in the appendix.

does not depend on a_1).¹⁹ However, in both of these cases the DM may experience regret. The fact that the DM's utility function is linear in regret implies that first-period choice does not depend upon either ρ or ν .²⁰ To summarize, a taste for consistency affects choice only if there is meaningful choice in both periods, meaningful uncertainty, and complementarities between actions in the payoff function (or action-dependent choice sets).²¹

In the second period, the DM wants to rationalize a_1 by choosing a_2 for which a_1 was optimal. How close a_1 was to optimal given a_2 depends of course upon what other first-period actions the DM had available. A DM who recognizes this in the first period has preferences over actions in A_1 that also depend upon A_1 . Let $a_1^*(A_1; \rho, \nu)$ be a ν -naive ρ -rationalizer's first-period choice as a function of her choice set.

Definition 4 *A ν -naive ρ -rationalizer's first-period choice satisfies the Weak Axiom of Revealed Preference (WARP) if whenever $a_1, a'_1 \in A_1 \cap A'_1$, if $a_1 = a_1^*(A_1; \rho, \nu)$, then $a'_1 \neq a_1^*(A'_1; \rho, \nu)$.*

Choices that satisfy WARP do not depend on choice sets. The next proposition shows that a preference for consistency affects first-period choice if and only if WARP is violated.

Proposition 2 *A ν -naive ρ -rationalizer's first-period choice satisfies WARP if and only if for each A_1 , $a_1^*(A_1; \rho, \nu) = a_1^*(A_1; 0, 0)$.²²*

The “if” direction of the Proposition 2 is trivial: a DM who does not care about consistency satisfies the Weak Axiom, and so too does a ν -naive ρ -rationalizer making the same choices. The “only if” direction shows that anytime a preference for consistency matters, choice violates WARP.

The following example illustrates the proof of this direction. By Proposition 1, any such example must have at least two first-period choices, two states of nature, and two second-

¹⁹The condition that g be additively separable in a_1 and a_2 is equivalent to the condition that g be both sub- and super-modular in a_1 and a_2 , or that for twice-differentiable g the cross-partial between a_1 and a_2 be zero. These conditions play an important role in the next section.

²⁰Without choice in the second period, the sophisticated version of my model closely resembles Loomes and Sugden's (1982) regret theory with a linear regret-rejoice function. Loomes and Sugden note that in this case, the DM prefers a_1 to a'_1 if and only if a_1 yields a higher expected payoff, which is equivalent to the statement in Proposition 1 that ρ does not affect first-period choice when A_2 is a singleton set. While I believe that without second-period choice people would choose first-period actions to avoid regret, I use linear preferences to simplify the analysis as well as to distinguish between a taste for rationalizing past actions and a taste for avoiding regret.

²¹A similar result is that a taste for consistency does not affect choice when for each $a_1, a'_1 \in A_1$ and $\omega \in \Omega$, $A_2(a_1, \omega) \cap A_2(a'_1, \omega) = \emptyset$. In that case, the only a_1 consistent with choosing $a_2 \in A_2(a'_1, \omega)$ is a'_1 , and therefore the DM cannot experience regret whatever her second-period choice.

This certainly complicates application of my model, for which actions take the same label is of course a modeling choice.

²²Recall that for each ρ and ν , $a_1^*(A_1; 0, 0) = a_1^*(A_1; \rho, 0) = a_1^*(A_1; 0, \nu)$.

period choices. Such $2 \times 2 \times 2$ decision problems are represented by two 2×2 matrices, where entry (a_1, a_2) in matrix ω gives $g(a_1, a_2, \omega)$.

	ω	a_2	a'_2	ω'	a_2	a'_2
Example 1	a_1	$2 + \varepsilon$	0	a_1	0	1
	a'_1	0	1	a'_1	0	2

Suppose that each state occurs with probability $\frac{1}{2}$, and $\varepsilon > 0$. When $\rho = 0$, the action a_1 yields an expected payoff of $\frac{3}{2} + \frac{\varepsilon}{2}$, and action a'_1 yields an expected payoff of $\frac{3}{2}$. Consider now a fully sophisticated ρ -rationalizer with $\rho > 0$. The DM's expected payoff from a_1 is $\frac{1}{2}(2 + \varepsilon + 1 + \rho(1 - 2)) = \frac{3}{2} + \frac{\varepsilon - \rho}{2}$, so long as $1 + \rho(1 - 2) > 0$ or $\rho < 1$, which I assume for simplicity. The DM's expected utility from a'_1 does not depend on ρ since in each state her planned second-period action causes her no regret. Thus for $\rho > \varepsilon$ a ρ -rationalizer chooses a'_1 over a_1 , in contrast to a DM who does not care about consistency. Illustrating the proof of Proposition 2 necessitates finding some action a''_1 that when added to the DM's choice set reverses her choice from a'_1 to a_1 . Consider Example 2.

	ω	a_2	a'_2	ω'	a_2	a'_2
Example 2	a_1	$2 + \varepsilon$	0	a_1	0	1
	a'_1	0	1	a'_1	0	2
	a''_1	$2 + \varepsilon$	$2 + \varepsilon$	a''_1	-1	-1

The presence of a''_1 neither affects second-period following a_1 nor a'_1 : in state ω , a'_1 was optimal given either second-period choice, and so it does not affect the DM's choice; in state ω' , a'_1 was suboptimal given both second-period actions and again it does not affect the DM's choice. What a''_1 does is to make the DM regret action a'_1 in state ω , something absent under the old choice set. As a result, her expected utility from a'_1 is now $\frac{1}{2}(1 + \rho(1 - (2 + \varepsilon)) + 2) = \frac{3}{2} - \frac{\rho}{2} - \frac{\varepsilon\rho}{2}$, less than her payoff from choosing a_1 . Finally, note that the DM's expected payoff from a_1 is higher than that from a''_1 .

4.1 Characterizing First-Period Choice

Let $\hat{\rho} \equiv \rho(1 - v)$, and define

$$V(a_1, \hat{\rho}) \equiv \sum_{\omega \in \Omega} [g(a_1, a_2^*(a_1, \omega; \hat{\rho}), \omega) - \hat{\rho}r(a_1 | a_2^*(a_1, \omega; \hat{\rho}), \omega)] \mu(\omega),$$

the DM's perceived expected utility from choosing the first-period action a_1 given her perceived second-period behavior $a_2^*(a_1, \omega; \hat{\rho})$. Let $V(\hat{\rho}) \equiv \max_{a_1 \in A_1} V(a_1, \hat{\rho})$, the DM's perceived expected utility, and $a_1^*(\hat{\rho}) \equiv \max_{a_1 \in A_1} V(a_1, \hat{\rho})$.

Proposition 3 *For each $a_1 \in A_1$, $V(a_1, \hat{\rho})$ is decreasing, continuous, convex, and for almost all $\hat{\rho} > 0$, $\frac{\partial V(a_1, \hat{\rho})}{\partial \hat{\rho}} = -\sum_{\omega \in \Omega} r(a_1 | a_2^*(a_1, \omega; \hat{\rho}), \omega) \mu(\omega)$. Likewise, $V(\hat{\rho})$ is decreasing, continuous, convex, and for almost all $\hat{\rho} > 0$, $\frac{\partial V(\hat{\rho})}{\partial \hat{\rho}} = -\sum_{\omega \in \Omega} r(a_1^*(\hat{\rho}) | a_2^*(a_1^*(\hat{\rho}), \omega), \omega) \mu(\omega)$.*

For each $a_1 \in A_1$, $\omega \in \Omega$, $a_2 \in A_2(a_1, \omega)$, $g(a_1, a_2, \omega) - \hat{\rho}r(a_1|a_2, \omega)$ is a decreasing, linear function of $\hat{\rho}$, and $\max_{a_2 \in A_2(a_1, \omega)} g(a_1, a_2, \omega) - \hat{\rho}r(a_1|a_2, \omega)$ is the upper envelope over $a_2 \in A_2(a_1, \omega)$. As such, it is decreasing, continuous, and convex. $V(a_1, \hat{\rho})$ is simply a weighted sum of these upper envelopes, and so it too is decreasing, continuous, and convex. Because it is decreasing, it is almost-everywhere differentiable; combining this with Theorem 1 of Milgrom and Segal (2002) yields its derivative.²³

Proposition 3 expresses how an increase in the DM's taste for consistency leads her to favor first-period actions that lead to low regret. Convexity of $V(\hat{\rho})$ and the formula for its derivative imply that for $\hat{\rho}' > \hat{\rho}$,

$$\sum_{\omega \in \Omega} r(a_1^*(\hat{\rho}')|a_2^*(a_1^*(\hat{\rho}'), \omega)), \omega) \mu(\omega) < \sum_{\omega \in \Omega} r(a_1^*(\hat{\rho})|a_2^*(a_1^*(\hat{\rho}), \omega)), \omega) \mu(\omega).$$

Quite intuitively, the more the DM wishes to avoid regret, the lower her expected regret. When the DM cares little about consistency, she takes actions that lead to high expected payoffs and also high regret; when she cares more about consistency, she takes actions with low expected payoffs and low regret. Increasing how much the DM cares about consistency reduces her expected utility more in the first instance because her regret is higher.

4.2 Risk Attitudes

One class of problems that plays a particularly important role in applications is where some first-period action \bar{a}_1 guarantees a payoff that is independent of both the state of nature and the DM's second-period action. Suppose that rather than choose \bar{a}_1 , the DM chooses some other action that gives payoffs that depend both upon the state and her second-period action. Suppose also that all uncertainty is not resolved by the second period: when the state is ω , the DM learns only that $\omega \in P_1(\omega)$. In this case, the DM's payoff from the profile of actions (a_1, a_2) is uncertain; let $(g(a_1, a_2, \omega))_{\omega \in P_1(\omega)}$ represent the payoff lottery induced by the actions a_1 and a_2 when $\omega \in P_1(\omega)$.

²³ Another method for computing its derivative comes from using first-order conditions in the case where g and r are continuously differentiable, so

$$\begin{aligned} \frac{\partial V(a_1, \hat{\rho})}{\partial \hat{\rho}} &= - \sum_{\omega \in \Omega} r(a_1|a_2^*(a_1, \omega; \hat{\rho}), \omega) \mu(\omega) \\ &\quad + \sum_{\omega \in \Omega} \left[\frac{\partial g(a_1, a_2^*(a_1, \omega; \hat{\rho}), \omega)}{\partial a_2^*(a_1, \omega; \hat{\rho})} - \hat{\rho} \frac{\partial r(a_1|a_2^*(a_1, \omega; \hat{\rho}), \omega)}{\partial a_2^*(a_1, \omega; \hat{\rho})} \right] \left(\frac{\partial a_2^*(a_1, \omega; \hat{\rho})}{\partial \hat{\rho}} \right) \mu(\omega) \\ &= - \sum_{\omega \in \Omega} r(a_1|a_2^*(a_1, \omega; \hat{\rho}), \omega) \mu(\omega), \end{aligned}$$

since the second term in the first equality vanishes from the first-order condition for $a_2^*(a_1, \omega; \hat{\rho})$, so long as it is an interior maximizer.

Proposition 4 Suppose that $A_1 = \{a_1, \bar{a}_1\}$, and for each $\omega \in \Omega$, $A_2(a_1, \omega) = A_2(\bar{a}_1, \omega)$ and for each $a_2 \in A_2(\bar{a}_1, \omega)$, $g(\bar{a}_1, a_2, \omega) = k \in \mathbb{R}$.

1. If for some $P_1(\omega)$, $(g(a_1, a_2, \omega))_{\omega \in P_1(\omega)}$ second-order stochastically dominates $(g(a_1, a'_2, \omega))_{\omega \in P_1(\omega)}$, then

$$\sum_{\omega \in P_1(\omega)} (g(a_1, a_2, \omega) - \rho r(a_1|a_2, \omega)) \mu(\omega|P_1(\omega)) \geq \sum_{\omega \in P_1(\omega)} (g(a_1, a'_2, \omega) - \rho r(a_1|a'_2, \omega)) \mu(\omega|P_1(\omega)),$$

where equality implies

$$k \notin \left(\min_{\omega \in P_1(\omega)} \{g(a_1, a_2, \omega), g(a_1, a'_2, \omega)\}, \max_{\omega \in P_1(\omega)} \{g(a_1, a_2, \omega), g(a_1, a'_2, \omega)\} \right).^{24}$$

2. If for some $P_1(\omega)$ and $\rho > 0$,

$$\sum_{\omega \in P_1(\omega)} (g(a_1, a_2, \omega; k) - \rho r(a_1|a_2, \omega; k)) \mu(\omega|P_1(\omega)) = (1 + \rho)l - \rho \max\{k, l\},$$

then if $l < k$ and $k' > k$ (or $l > k$ and $k' < k$),

$$\sum_{\omega \in P_1(\omega)} (g(a_1, a_2, \omega; k') - \rho r(a_1|a_2, \omega; k')) \mu(\omega|P_1(\omega)) \geq (1 + \rho)l - \rho \max\{k', l\}.$$

The first statement expresses that a taste for consistency makes a ρ -rationalizer “risk averse”: it makes her utility function concave. For payoff outcomes below k , the DM’s utility function has a slope $1 + \rho$, while for outcomes above k it has a slope of 1. Thus, her marginal utility of payoffs is decreasing, and she is risk averse. The second statement expresses how the DM’s risk aversion depends upon her past foregone payoff. If her certainty equivalent for a given lottery is higher than her foregone payoff ($l > k$), then decreasing her foregone payoff increases the value of the lottery but not that of the certainty equivalent; thus it increases the lottery’s certainty equivalent. If her certainty equivalent is less than her foregone payoff, then increasing the value of that foregone payoff decreases the value of her certainty equivalent more than that of the lottery; thus it increases the lottery’s certainty equivalent. In both cases, the DM becomes less risk averse.

5 The Sunk- and Unsunk-Cost Effects

In this section, I show how a taste for rationalizing past actions generates the sunk- and unsunk-cost effects, and explore some of their implications. Consider an investor choosing how much to invest in a risky project. In the first period, she invests $a_1 \in \mathbb{R}_+$. She then learns

²⁴The distribution F second-order stochastically dominates G if for each x , $\int_{-\infty}^x G(t)dt \geq \int_{-\infty}^x F(t)dt$.

the project's rate of return $\omega \in \Omega$ and invests $a_2 \in \mathbb{R}_+$. Her final payoff is $\omega(a_1 + a_2) - C(a_1, a_2)$, where C is the cost of investment. The cost of investment is known, but the return is unknown when making first-period investment. Suppose that $\frac{\partial^2 C(a_1, a_2)}{\partial a_1 \partial a_2} < 0$, so that the marginal cost of second-period investment decreases with first-period investment: more investment in the first period makes investment in the second-period cheaper. How does a preference for consistency affect the investor's behavior?

Definition 5 ((Bulow, Geanakoplos, and Klemperer 1985)) *The DM's problem g has strategic complements if for each $\omega \in \Omega$, $a'_1 \geq a_1$, and $a'_2 \geq a_2$, $g(a'_1, a'_2, \omega) - g(a'_1, a_2, \omega) \geq g(a_1, a'_2, \omega) - g(a_1, a_2, \omega)$.²⁵*

A decision problem that has strategic complements has the feature that high first-period actions increase the DM's incentives to increase second-period actions. The investment example has strategic complements because early investment, by reducing the marginal cost of later investment, increases the incentive for later investment. Recall that $a_1^*(a_2, \omega) \equiv \arg \max_{a_1 \in A_1(a_2, \omega)} \{g(a_1, a_2, \omega)\}$, and define $a_2^*(a_1, \omega) \in \arg \max_{a_2 \in A_2(a_1, \omega)} \{g(a_1, a_2, \omega)\}$.

Proposition 5 *Suppose that g has strategic complements. If $a_1 \geq a_1^*(a_2^*(a_1, \omega), \omega)$, then for each $\rho > 0$, $a_2(a_1, \omega; \rho) \geq a_2(a_1, \omega; 0)$. If $a_1 \leq a_1^*(a_2^*(a_1, \omega), \omega)$, then for each $\rho > 0$, $a_2(a_1, \omega; \rho) \leq a_2(a_1, \omega; 0)$. Suppose furthermore that for each $a_1 \in A_1$, $\omega \in \Omega$, $g(a_1, a_2, \omega)$ is strictly quasi-concave in a_2 .²⁶ Then if $a_1 \geq a_1^*(a_2^*(a_1, \omega), \omega)$, for each $\rho' > \rho$, $a_2(a_1, \omega; \rho') \geq a_2(a_1, \omega; \rho)$. If $a_1 \leq a_1^*(a_2^*(a_1, \omega), \omega)$, then for each $\rho' > \rho$, $a_2(a_1, \omega; \rho') \leq a_2(a_1, \omega; \rho)$.*

In a situation with strategic complements, a ρ -rationalizer rationalizes having chosen a_1 too high by increasing a_2 . For $a_1 \geq a_1^*(a_2)$, the DM's regret from choosing a_1 decreases in a_2 because of strategic complements; raising a_2 increases raises $a_1^*(a_2)$, justifying high first-period choice.²⁷ The second part of the proposition provides the stronger result that when $a_1 \geq a_1^*(a_2)$ the DM's second-period action increases in ρ , so long as the payoff function is quasi-concave. Without quasi-concavity, the result would not necessarily hold. Starting from $a_2^*(a_1)$, consistency motives first may lead the DM to increase a_2 some; but then a payoff function that is locally increasing in a_2 might lead the DM to raise ρ still further. In this case, a DM with a lower ρ may choose a higher a_2 .

Proposition 5 captures the essence of the sunk-cost effect. A ρ -rationalizer who has invested "too much" in some project where past investment and current investment are strategic

²⁵If g is twice differentiable, then the DM's problem has strategic complements if $\frac{\partial^2 g}{\partial a_1 \partial a_2} \geq 0$.

²⁶The function $g(a_1, a_2, \omega)$ is strictly quasi-concave in a_2 if whenever $a_2 \neq a'_2$ and $g(a_1, a'_2, \omega) \geq g(a_1, a_2, \omega)$, for each $\alpha \in (0, 1)$, $g(a_1, \alpha a'_2 + (1 - \alpha)a_2, \omega) > g(a_1, a_2, \omega)$. It implies that g has a single peak in a_2 .

²⁷Of course, analogous results hold for strategic substitutes (after first-period actions that are too high, ρ -rationalizers choose lower second-period actions) and follow from reversing the order on a_2 .

complements invests too much in current investment; she throws good money after the bad. The DM chooses first-period investment as a function of her beliefs about the rate of return on the project. When the true rate of return turns out to be low relative to the DM's forecasts, then $a_1 > a_1^*(a_2^*(a_1))$, and a DM who cares about consistency has incentive to choose second period investment higher than the level that would maximize payoffs, the sunk-cost effect. When the true rate of return turns out to be high relative to the DM's forecasts, then $a_1 < a_1^*(a_2^*(a_1))$, and a DM who cares about consistency has incentive to choose second period investment lower than the level that would maximize payoffs, an unsunk-cost effect.

Example 3 $A_1 = A_2 = \{0, 1\}$ and $g(a_1, a_2, \omega) = a_1 a_2 r - a_1 - \omega a_2$.

Example 3 describes a situation where a project pays off r if and only if an investor invests in each of two periods; the investor does not know the second-period cost of investment, $\omega \in [\underline{\omega}, \bar{\omega}]$, when choosing first-period investment.²⁸ The investor's problem has strategic complements, for investing in the first period increases the investor's incentive to invest in the second period. Clearly if the investor does not invest in the first period, then she does not in the second. If she does invest in the first, then by investing in the second she receives a payoff of $r - \omega - 1$ from investing and does not regret her initial decision to invest. If she invests in the first period but not in the second, then her utility is $-1 - \rho$; she regrets wasting her first-period investment. The higher ρ , the more likely the investor is to complete the project; this is the conclusion of Proposition 5. In the first period, a ν -naive ρ -rationalizer perceives her expected utility following first-period investment to be

$$V(1; \rho, \nu) = -1 + \int_{\underline{\omega}}^{r+\rho(1-\nu)} (r - \omega) d\mu(\omega) - \rho(1 - \nu) \int_{r+\rho(1-\nu)}^{\bar{\omega}} d\mu(\omega).$$

The DM thinks that she completes the project if and only if $r - \omega \geq -\rho(1 - \nu)$, or $\omega \leq r + \rho(1 - \nu)$; the third term in $V(1; \rho, \nu)$ is her perceived regret from first-period investment when she does not invest in the second period. For $\nu < 1$, $\frac{\partial V}{\partial \rho(1-\nu)} < 0$; the more the investor thinks she will care about regret, the less likely she will invest. In this setting where the sunk-cost effect dominates, a taste for consistency lowers the investor's first period investment.

Example 4 $A_1 = A_2 = \{0, 1\}$ and $g(a_1, a_2, \omega) = a_2 r - a_1 - a_2 - a_2(1 - a_1)\omega$, for $\omega \in [\underline{\omega}, \bar{\omega}]$ with $\underline{\omega} \in (0, 1)$ and $\bar{\omega} > 1$.

Example 4 describes a situation where the investor completes a project worth $r > 2$ if and only if she invests in the second period. The cost of investing in the second period is

²⁸Example 3 has the same structure as Thaler's basketball example.

1 if she invested in the first period and $1 + \omega$ if she did not invest in the first period, while investing in the first period always costs 1. Whereas Example 3 depicts a project that is only completed by investing in both periods, Example 4 depicts one that can be completed either by investing in both periods or by investing only in the second period. This captures the idea that by working very hard in the second period the investor can complete the project without investing in the first period, but that completing the project in the second period costs less when it was started in the first period than when it was not. Uncertainty in this example revolves around whether the cheapest way to complete the project is by investing in both periods or only the first. (The restrictions on $\underline{\omega}$ and $\bar{\omega}$ mean that in certain states of the world investing in period 1 maximizes profits, while in others investing in period 2 maximizes profits.) As in Example 3, the DM's payoff function has strategic complements, for $\frac{\partial^2 g}{\partial a_1 \partial a_2} = \omega > 0$; it is always cheaper to complete the project in the second period after investing in the first period than after not investing in the first period. An investor who invests in the first period clearly invests in the second period and, hence, never regrets first-period investment. An investor who does not invest in the first period or in the second period receives a payoff of zero and never regrets not investing in the first period—investment in the first period is worthless without investment in the second period. An investor who invests only in the second period receives a payoff of $r - 1 - \omega$. When $\omega > 1$, she regrets not investing in the first, and her utility is $r - 1 - \omega + \rho(1 - \omega)$. The higher is ρ , the less likely is the investor to complete the project; this is the unsunk-cost effect described in Proposition 5—the investor takes an action lower than that maximizes her payoff to rationalize her overly-low first period action.

In the first period, a ν -naive ρ -rationalizer perceives her expected utility following not investing to be

$$V(0; \rho, \nu) = \int_{\underline{\omega}}^1 r - 1 - \omega d\mu(\omega) + \int_1^{\frac{r+\rho(1-\nu)-1}{1+\rho(1-\nu)}} r - 1 - \omega + \rho(1 - \nu)(1 - \omega) d\mu(\omega).$$

The DM thinks that she will complete the project if and only if $r - 1 - \omega + \rho(1 - \omega) \geq 0$ or $\omega \leq \frac{r+\rho(1-\nu)-1}{1+\rho(1-\nu)}$. When $\omega < 1$, then it is cheapest to complete the entire project in the second period, and the DM does not regret not investing in the first period; this is the first term of V . When $\omega > 1$, then the DM completes the project but regrets not investing in the first period. When $\omega > \frac{r+\rho(1-\nu)-1}{1+\rho(1-\nu)}$, then the DM does not complete the project and therefore experiences no regret. For $\nu < 1$, $\frac{\partial V}{\partial \rho(1-\nu)} < 0$; the more the investor thinks she will care about regret, the less attractive waiting until the second period to begin the project, and the more likely the investor is to invest in the first period. In this setting where the unsunk-cost effect dominates, a taste for consistency increases the investor's first period investment.

In many situations, a consumer learns over time how much she values consumption of a good, and a firm may ask the consumer to pay for consumption or choose a tariff that specifies payment for future consumption before the consumer knows her valuation. For example, a telephone company may charge a monthly hook-up fee then some (possibly nonlinear) price per call. If all consumers have the same beliefs about their future value of telephone calls, then a monopoly with constant marginal costs follows a simple optimal pricing policy: it maximizes profits plus consumer surplus by pricing calls at marginal cost, then extracts all of the consumer's expected surplus through a fixed fee. But if consumers are heterogeneous in the first period, the monopolist optimally offers a menu of tariffs. Intuitively, consumers with high expected demand pay a high marginal price for early units of demand and then a low marginal price for later units of demand, while those with low expected demand pay a low marginal price for early units and a high marginal price for later units. Courty and Li (2000) analyze this kind of a sequential-screening model where consumers' demand is either zero or one.

When consumers care about consistency, they care both about tariffs offered but not chosen and outside options. In the phone example, a consumer who opted for the high-demand plan who turns out to have low demand may regret her tariff choice either because she should have opted for the low-demand tariff or because she should not have chosen either tariff. In both cases, the "sunk-cost effect" may encourage the consumer to consume more than she normally would. In this section I first analyze how a preference for consistency affects choice of tariffs, and then give a simple example where it leads a profit-maximizing monopolist to reduce the number of tariffs it offers.

A consumer with type θ values consumption of q units of a good at $v(q, \theta) = \theta q - q^2/2$. The consumer's payoff is linear in the value of consumption and her payment $t(q)$, $\theta q - q^2/2 - t(q)$. Initially the consumer does not know her type; she has beliefs given by the distribution function $F(\theta)$. Before learning her type, the consumer chooses from two different two-part tariffs, $(0, 1)$ and (α, β) , where $\alpha > 0$ and $\beta < 1$. The first tariff charges the consumer 1 for every unit she consumes. The second tariff charges the consumer a fixed price of α followed by a marginal price of β . After learning her type, the consumer chooses consumption q and pays a the total cost of consumption, either q or $\alpha + \beta q$, depending on her choice of tariffs.

Ordering $(0, 1) < (\alpha, \beta)$ makes the consumer's problem one of strategic complements, for choosing the higher tariff leads to a lower second-period marginal price, which increases the consumer's incentive to raise q . From Proposition 5, the sunk-cost effect leads high types of θ who chose (α, β) to increase her consumption to rationalize having paid the fixed price α , and the unsunk-effect leads low types of θ who chose $(0, 1)$ to decrease her consumption to

rationalize having chosen not to pay the fixed price (to pay the higher marginal price). How does a preference for consistency affect the consumer's choice of tariffs?

If the consumer chooses the (α, β) scheme, then her demand is

$$q(\theta | (\alpha, \beta)) = \begin{cases} \theta - \beta + \rho(1 - \beta) & \theta < \frac{\alpha}{1-\beta} + \beta - \rho(1 - \beta) \\ \frac{\alpha}{1-\beta} & \theta \in \left(\frac{\alpha}{1-\beta} + \beta - \rho(1 - \beta), \beta + \frac{\alpha}{1-\beta} \right) \\ \theta - \beta & \theta > \beta + \frac{\alpha}{1-\beta}. \end{cases}$$

When her type is high, she consumes as usual, setting price equal to marginal utility. For intermediate values of θ , she consumes in such a way that her regret is zero—the price of $\frac{\alpha}{1-\beta}$ is the same under both tariffs. When θ is low, the consumer increases demand to rationalize her choice of tariff, but she does not raise demand high enough that (α, β) was the cheapest tariff. Likewise, if the consumer chooses the $(0, 1)$ scheme, then her demand is

$$q(\theta | (\alpha, \beta)) = \begin{cases} \theta - 1 & \theta < \frac{\alpha}{1-\beta} + 1 \\ \frac{\alpha}{1-\beta} & \theta \in \left(\frac{\alpha}{1-\beta} + \beta - \rho(1 - \beta), \frac{\alpha}{1-\beta} + 1 \right) \\ \theta - 1 - \rho(1 - \beta) & \theta > \frac{\alpha}{1-\beta} + \rho(1 - \beta) + 1. \end{cases}$$

In this case, the unsunk cost leads the consumer to underconsume when her valuation is high.

When the unsunk-cost effect dominates, a monopolist may find it optimal to limit the number of tariffs it offers consumers for two reasons. Firstly, consumers anticipate future regret from choosing the wrong tariff, which lowers their expected utility from consuming from the monopolist in the first place. Secondly, even when consumers naively believe that they will not care about regret and therefore overestimate their utility from purchasing from the monopolist, the unsunk-cost effect may dominate the sunk-cost effect, which reduces the monopolist's incentive to offer multiple tariffs to extract consumer's information rents. Consider a simple example, where a consumer has demand for two goods; she values the first at 1 and the second at either 1 or 2. In a first period, the consumer believes either that her valuation for the second good will be 2 with probability p_l or p_h , where $p_l < p_h$. Fraction μ_l of consumers hold beliefs p_l . For simplicity, suppose that there are no costs of production.

As a benchmark, suppose that the consumer does not care about consistency. The firm can offer two types of tariffs: one that sells both goods to the consumer at a fixed price, and one that sells the goods independently. Suppose that the firm offers to sell both goods for $2 + p_h$ and to sell the two goods for prices $p_1 = 1$ and $p_2 = 2$. High-belief consumers are indifferent between the two tariffs and willing to accept the first. Low-belief consumers strictly prefer the second tariff. Suppose that $(2 + p_h)(1 - \mu_l) + 3\mu_l p_l > 2$ so that the monopolist prefers this scheme to offering a single tariff that sells each good at a price of 1. The monopolist can do no better than this: it rations by not selling to low-beliefs consumers who ultimately have

high valuations for the second object, and then extracts all the rents from the consumers it does serve.

Now suppose that the consumer cares about rationalizing her choice of tariff. If the firm offered the same tariffs, $2 + p_h$ and $(1, 2)$, then sophisticated low-belief consumers who cared enough about regret would not accept either tariff. A sophisticated consumer knows that if her valuation turns out to be high and she paid the marginal price of 2 then she would regret not having paid $2 + p_h$ for both goods initially; if she wishes strongly enough to avoid regret, then she will not choose the tariff $(1, 2)$. Nor would she pay $2 + p_h$ to consume both goods when her expected valuation is only $2 + p_l$. A fully naive low-belief consumer would accept $(0, 2)$, expecting zero surplus, but then neither high nor low types would consume in the second period: the marginal price of 2 leaves high types with no surplus and positive regret. Indeed, for ρ sufficiently high, even if consumers are fully naive, the firm will offer a single tariff, either $(1, 2)$ —serving all the high types—or serving either all high-belief consumers with $2 + p_h$, or all consumers with $2 + p_l$.

6 A Many-Period Model

This section extends the basic two-period model to many periods and provides three applications: Ashenfelter’s (1989) declining-price anomaly, Shubik’s (1971) dollar auction, and a model of search over time.

Again let Ω be a finite set of payoff relevant states of the world. In each of T periods the DM takes an action. In period t , her information is represented by the partition \mathcal{P}_{t-1} . When the state is ω , in period t the DM learns only that $\omega \in P_{t-1}(\omega) = \{\omega' \in \Omega : \omega' \in P_t\}$. For each t , \mathcal{P}_t refines \mathcal{P}_{t-1} ; that is for each ω , $P_t(\omega) \subset P_{t-1}(\omega)$ —over time the DM can rule out more and more states. Again let $g(a_1, \dots, a_T, \omega)$ represent the DM’s payoffs as a function of all of her actions and the state. Let $a \equiv (a_1, \dots, a_T)$ be the profile of all of her actions, and let $a_{-t} \equiv (a_1, \dots, a_{t-1}, a_{t+1}, \dots, a_T)$, the profile of all actions other than a_t . As before, all payoffs occur at the end of the game. A ρ -rationalizer chooses a_T to maximize

$$g(a, \omega) + \rho \left(\sum_{t=1}^{T-1} g(a, \omega) - \max_{a'_t} g(a_{-t}, a'_t, \omega) \right).$$

For $\rho > 0$, the DM wants each of her actions to be optimal given all other actions—past and future. The more decisions the DM took in the past, the more weight she puts on her consistency, namely total regret. A feature of these preferences that plays a prominent role in both applications is that when it’s too late to rationalize choices in the distant past, it may not be too late to rationalize actions in the recent past. The natural generalization of

ν -naivety is that the DM thinks in periods $t < T$ that she will be a $\rho(1 - \nu)$ rationalizer in period T .

The many-period model is considerably more complicated to work with than the two-period model. When choosing a_2 , the DM knows that how much she regrets a_1 depends not only on her current choice, but also on her future choice. The two applications in this section share a common feature that considerably simplifies analysis. In each case, the DM's action space each period is binary: she can either "end her decision problem," guaranteeing herself a sure payoff, or continue the problem. Thus, when choosing whether to end the problem in period t , the DM's regret only depends upon whether it would have been better to have ended in each period $s < t$.

6.1 The Declining-Price Anomaly

Ashenfelter (1989) describes prices in sequential English auctions of identical goods. When all bidders have unit demand and independent and identically-distributed valuations, Bayesian Nash analysis predicts that the expected prices are constant across auctions.²⁹ By contrast, Ashenfelter finds that in wine auctions prices tend to decline. Many authors have analyzed how departures from the assumptions of risk neutrality, independent valuations, and unit demand affect the expected price path in sequential auctions.³⁰ In this section, I show that the unsunk-cost fallacy leads to declining prices.

The intuition for this result is directly related to the finding that people who once passed up an opportunity to buy something on sale later decline to buy it for full price, even when doing so is in their material interest. Many experimental papers provide evidence of this phenomenon. For example, Tykocinski, Pittman, and Tuttle (1995) ask 108 student subjects to imagine that:

“Your friend called you at the beginning of October and told you that the intended to buy a special pass to Ski Liberty (where you both like to ski). He said that the deal was that if you bought the pass before the 15th of October the pass would cost you only \$40 (\$80) instead of the \$100 regular price. Although it seemed like a good idea, you forgot to do it by the 15th. The next time your friend called he

²⁹One way to see this comes from the Revenue Equivalence Theorem. Consider auctioning 2 units to n bidders. Auctioning both units simultaneously in a sealed-bid, uniform-price auction yields twice the third-highest valuation. In sequential English auctions, bidders clearly bid their valuations in the last auction, so the price is the third-highest valuation. Since both the sequential auction and the uniform-price auction leave no surplus to a bidder with the lowest possible valuation and allocate the objects in the same way, the Revenue Equivalence Theorem implies that they generate the same expected revenue. Thus, the expected revenue in the first of the two sequential English auctions is also the expected third-highest valuation, and the expected prices are constant across the sequential auctions.

³⁰See McAfee and Vincent (1993) for an analysis of risk aversion, Black and de Meza (1992) for the effect of allowing a winning bidder to purchase all subsequent lots at the same price (a common feature of wine auctions, although Ashenfelter (1989) finds the declining-price anomaly even in its absence), and de Frutos and Rosenthal (1998) for a characterization of when prices fall on average if bidders have common values.

told you that although you missed the deadline you could still get a pass for \$90 if you pay this week. How likely are you to spend the \$90?" (Tykocinski *et al*, 1995, p. 794)

A control group told nothing about the previous sale reported a mean likelihood of 6.36, and subjects who were told that the previous price was \$80 reported an average likelihood of 7.25; the difference is insignificant. By contrast, those told an initial price of \$40 gave an average likelihood of 4.94, significantly lower than the other two.

Two identical units of a good are to be auctioned in sequential English auctions. To simplify analysis, assume that after the first auction the winning bidder's bid is announced.³¹ Each Bidder i has unit demand with valuation v_i drawn from the common distribution F . Let $v_1 \equiv \max_{j \neq i} \{v_j\}$, the highest valuation of a bidder other than Bidder i , and $v_2 \equiv \max_{j \neq i} \{v_j\} \setminus \{v_1\}$ be the second-highest of these other bidders' valuations. Let b^1 and b^2 denote the bidding rules of bidders $j \neq i$ in the first and second auctions, respectively.

To fix ideas, think of Bidder i as choosing in the first period whether to pay $b^1(v_1)$ for the right to consume the object. In the second auction, she chooses whether to pay $b^2(v_2)$ for the good if she did not win the first auction. In the third period, if she owns the good, she chooses whether to "consume" it. A bidder experiences regret when she consumes the good at a higher price than she could have consumed it at had she won the first auction; she does not experience regret if she does not win the auction. This corresponds to Tykocinski, Pittman, and Tuttle's (1995) example, where taking a ski holiday leads to regret when one had a past opportunity to purchase a cheap ticket, but not otherwise.

More formally, consider a three-period model where bidders first decide how much to bid in the first auction, b^1 , then how much to bid in the second auction, b^2 , and finally whether to consume the good, $c \in \{0, 1\}$. For simplicity, suppose that if the bidder does not win one of the two auctions, $c = 0$, namely she cannot consume the object. Let

$$g_i(b_1^1, \dots, b_n^1, b_1^2, \dots, b_n^2, c) = \begin{cases} cv_i - \max_{j \neq i} b_j^1 & b_i^1 > \max_{j \neq i} b_j^1 \\ cv_i - \max_{j \notin \{i, \arg \max_j \{b_j^1\}\}} b_j^2 & b_i^1 < \max_{j \neq i} b_j^1 \text{ and } b_i^2 > \max_{j \notin \{i, \arg \max_j \{b_j^1\}\}} b_j^2 \\ 0 & \text{otherwise} \end{cases}$$

represent Bidder i 's payoffs. If she wins the first auction and consumes the object, then her payoff is her valuation less the price she pays in the first auction. If she wins the second auction and consumes the object, then again she gets her valuation less the price she pays, which is the second-highest bid from any bidder who did not win the first auction. In this way, the rules of the auction prevent bidders from winning more than one object: a bidder

³¹Alternatively, this could describe a sequential sealed-bid, second-price auction with all first-auction bids announced before the second auction.

who wins the first and places the highest bid in the second wins only the first auction. Of course, the most natural way to model sequential auctions would not impose this restriction on bidding. This modeling approach would be unnecessary if regret in the many-period model were included not only the sum of regret from each past action, but also the difference between current payoff and the best payoff that could be obtained re-optimizing over all past actions simultaneously. In this case, a bidder would regret winning the second auction at a price higher than that at which she could have won the first.

First note that a bidder who has won one of the two auctions always consumes the good. As a result, a bidder who does not consume the good in the third period does not regret either of her bids; given that she does not consume the object, it was optimal to lose both auctions. A bidder who does consume the object never regrets her bid in the second auction. If she won the first auction, then her bid in the second auction does not affect payoffs. If she loses both auctions, then she does not consume the good, and her second-auction bid is optimal. If she loses the first auction and wins the second, then her bid in the second auction was optimal, for it was the only way that she could consume the object. Thus, it suffices to consider how a bidder regrets her bid in the first auction.

If Bidder i loses the first auction and $b^1(v_1) > v_i$, then in the second auction Bidder i bids her valuation v_i , for win or lose the second auction she does not regret losing the first. If $b^1(v_1) < v_i$ and Bidder i wins the good in the second auction and pays $b^2(v_2) > b^1(v_1)$, then she regrets not having won the first auction. In this case, her utility is

$$v_i - b^2(v_2) + \rho(b^1(v_1) - b^2(v_2)).$$

Let $F_2(v_2|v_1)$ be the distribution of the second-highest valuation of Bidders $-i$ given that the highest valuation is v_1 . Then Bidder i chooses her second-auction bid \hat{b}^2 to maximize

$$\int_0^{(b^2)^{-1}(\hat{b}^2)} v_i - b^2(v_2) + \rho \max \{ (b^1(v_1) - b^2(v_2)), 0 \} dF(v_2|v_1),$$

which is done when $\hat{b}^2 = \frac{1}{1+\rho}v_i + \frac{\rho}{1+\rho}b^1(v_1)$. As usual in private-values English auctions, this bidding function is weakly dominant; in particular, it does not depend upon all the bidders' having the same ρ . Thus, if b^1 is the winning bid in the first auction, then in a symmetric equilibrium of the second auction a ρ -rationalizer with valuation v bids

$$b^2(v) = \begin{cases} v & v < b^1 \\ \frac{1}{1+\rho}v + \frac{\rho}{1+\rho}b^1 & v > b^1. \end{cases}$$

A ρ -rationalizer with a high valuation shades her bid in the second auction to avoid regret.

Now consider the first auction, and suppose that all bidders are fully sophisticated (and share the same ρ). Suppose again that Bidder i loses the first auction. If $b^1(v_1) > v_i$, then ρ

does not affect Bidder i 's bid in the second-auction. Furthermore, it does not affect Bidder j 's bid when Bidder i wins the second auction, namely when $v_j < v_i$, and therefore ρ does not affect Bidder i 's perceived utility from losing the first auction. If $b^1(v_1) < v_i$, then Bidder i 's expected utility from the second auction is

$$\begin{aligned} & \int_0^{b^1(v_1)} v_i - v_2 dF(v_2|v_1) \\ & + \int_{b^1(v_1)}^{v_i} v_i - \frac{1}{1+\rho}v_2 - \frac{\rho}{1+\rho}b^1(v_1) + \rho \left(b^1(v_1) - \frac{1}{1+\rho}v_2 - \frac{\rho}{1+\rho}b^1(v_1) \right) dF(v_2|v_1) \\ = & \int_0^{v_i} v_i - v_2 dF(v_2|v_1), \end{aligned}$$

which does not depend on ρ either. Intuitively, ρ lowers the bids of bidders $j \neq i$, but also causes Bidder i regret when she wins. The fact that each bidder cares about regret in the same way means that these two effects exactly offset each other, and so Bidder i 's expected utility from losing the auction does not depend on ρ , whatever $b^1(v_1)$. Since Bidder i 's expected utility from losing the first auction does not depend on ρ , nor does her utility from winning the first auction, her first-auction bid does not depend on ρ either.

Sophisticated ρ -rationalizers recognize that the price in the second auction is on average lower than that in the first but nevertheless bid high in the first in order to avoid regret in the second. But what about naive bidders who recognize that the price tends to fall in the second auction but underestimate the extent to which they will care about regret? To analyze the naif bidding, suppose that each Bidder i correctly predicts the distribution of the other bidders' second-auction bids, but misperceives her own second-auction bid. Proposition 7 summarizes equilibrium bidding in both the sophisticated and naive cases.

Proposition 6 *Suppose that $\rho, \rho' > 0$ and $\nu = 0$. Then $b^1(v_i; \rho) = b^1(v_i; \rho')$ for each ρ, ρ' , and*

$$b^2(v_i) = \begin{cases} v_i & v_i < b^1 \\ \frac{1}{1+\rho}v_i + \frac{\rho}{1+\rho}b^1 & v_i > b^1. \end{cases}$$

The expected price in the first auction is strictly greater than that in the second auction, and the difference increases in ρ . If $\nu < 1$ ($\nu = 1$) then the expected price in the first auction is strictly greater than (equal to) the expected price in the second auction.

The bidding function for sophisticates is derived in the text, and the fact that $\frac{db^2}{d\rho} < 0$ when $v_i < b^1$ and is equal to zero otherwise implies that the difference in expected revenue increases in ρ . If bidders are fully naive, then the pivotal bidder drops out of the first auction when $v_i - b(v_i) = v_i - Ep$, where Ep is her expectation of the price in the second auction. She is indifferent between winning the auction at her bid and moving on to the second auction,

where she knows that she will win for sure. Thus, $b(v_i) = Ep$, namely the bidder drops out at her expectation of the price in the second auction. The result that with fully naive bidders the expected price is constant across auctions follows by noting that a fully naive bidder correctly predicts the price in the second auction when she is pivotal, despite the fact that she mispredicts her own bidding rule. This follows from the fact that the pivotal bidder knows that she has the highest valuation in the second auction and therefore she is simply taking her expectation of the highest bid of a bidder with a valuation lower than her own, which does not depend on her perception of her own second-period bid as long as she bids at least as high for every valuation as the other bidders. When $\nu < 1$, then $v_i - b(v_i) = v_i - Ep - Er$, where $Er > 0$ is expected regret. Thus $b(v_i) > Ep$, and bidders in the first auction expect the price to decline. Once more, the pivotal bidder correctly predicts the price in the second auction, and therefore the expected price does decline.

6.2 The Dollar Auction (War of Attrition)

Shubik (1971) describes how auctioning \$1 in an ascending auction where the highest bidder wins the dollar but both the highest and second-highest bidders pay their bids usually yields significantly more than \$1 in seller revenue. Of course, in no Nash equilibrium can the seller's expected revenue exceed \$1, for that would imply a negative expected surplus for at least one bidder, which cannot occur in equilibrium—that bidder could always guarantee herself a payoff of zero by not participating. Since Shubik, much research has been done on the dollar auction, which with two bidders is equivalent to a war of attrition. Bazerman (2002) reports that auctioning \$20 bills in this way typically generates revenue between \$40 and \$140, and that in sixteen years of running these auctions he has made in excess of \$25,000! What accounts for such marked overbidding?³²

To explore this issue, I consider a simplified version of the dollar auction where a prize worth $w > 2$ (where w is an integer) is to be awarded to one of two bidders. In each period $t \in \{1, 2, \dots\}$ each of the two bidders decides simultaneously whether to remain in the auction or to drop out. If both bidders remain in the auction through period $t-1$, and Bidder i remains in the auction in period t while Bidder j drops out, then Bidder i wins w and pays a total of t , and Bidder j wins nothing and pays a total of $t-1$. If both bidders drop out in period $t-1$, then neither wins the object, and both pay $t-1$. (This feature considerably simplifies analysis.) If each bidder remains in the auction in period t , then the auction moves to period $t+1$.

As a benchmark, consider first the symmetric, stationary Nash equilibrium.³³ Suppose

³²Teger (1980) provides an extensive analysis of experimental dollar auctions.

³³This game also has many asymmetric equilibria (e.g. Bidder 1 always stays in and Bidder 2 never enters).

that the auction reaches period t , where Bidder 1's strategy calls for her to drop out with probability p . By dropping out, Bidder 2 receives a sure payoff of $-(t-1)$. By staying in for period t and dropping out in $t+1$, he receives an expected payoff of $-t + pw$, his total costs plus the probability that he wins the object in period t or $t+1$ times w . In order for Bidder 2 to follow a strategy that calls for him to drop out with positive probability each period, he must be indifferent between dropping out in period t and staying in for period t and dropping out in period $t+1$, which implies that $p = 1/w$. Hence, it is a Nash equilibrium for each Bidder to drop out each period with probability $1/w$. Intuitively, a stationary equilibrium exists because when bidders treat their past bids as sunk, the subgame starting in period t is the same at that starting at $t+1$.

Now consider the behavior of a sophisticated ρ -rationalizer who has stayed in the auction through period $t-1$. In equilibrium, a sophisticated ρ -rationalizer correctly predicts her own strategy and payoffs, as well as the other bidder's strategy. Again, if Bidder 1's strategy calls for her to drop out in period t with probability p , then by dropping out in period t Bidder 2 receives a sure utility of

$$-(t-1) + \rho \left(\sum_{s=1}^{t-1} -(t-1) - [-(s-1)] \right) = -(t-1) - \rho \frac{(t-1)t}{2}.$$

The second term on the left-hand side expresses Bidder 2's regret from not dropping out in periods $s < t$ when he drops out in t : given that Bidder 2 drops out in t , it was a mistake to have remained in the auction for each $s < t$. Bidder 2's regret from continuing to bid in period s given that the auction lasts until period t when he drops out is the payoff that he could have received by dropping out in period s , $-(s-1)$, less his actual payoff, $-(t-1)$. Bidder 2's regret from dropping out in period t is convex in t ; this follows from the feature that as the auction progresses, bidders have more and more past decisions to rationalize. By remaining in the auction for period t and dropping out in period $t+1$, Bidder 2 receives an expected utility of

$$p(w-t) + (1-p) \left(-t + \rho \left(\sum_{s=1}^t -t - [-(s-1)] \right) \right) = -t + pw - \rho(1-p) \frac{t(t+1)}{2},$$

so long as $w > t$, namely the prize is sufficiently large that winning it gives a better payoff than dropping out in the first period. In order that Bidder 2 be indifferent between dropping out in period t and remaining in for period t and dropping out in period $t+1$, $p(t; \rho) = \frac{2+2\rho t}{2w+\rho t(t+1)}$. When $t = 1$, a ρ -rationalizer does not regret dropping out in the first period, but she does regret dropping out in the second; this means that Bidder 1 must drop out at $t = 1$ with a

Given the formulation of the game, it is also an equilibrium for both bidders to stay in forever.

higher probability for Bidder 2 to be indifferent between dropping out and staying in. For small t , increasing t increases Bidder 2's regret from dropping out in $t+1$, which is proportional to $t(t+1)$, more than her regret from dropping out in t , which is proportional to $t(t-1)$, and again Bidder 1 must drop out with higher probability for Bidder 2 to be indifferent between staying in and dropping out. But as t increases, Bidder 2 starts to regret dropping out in period t more than $t+1$. Regret is convex in t , and so avoiding regret with probability p by winning the auction becomes increasingly important, meaning that p must fall for Bidder 2 to be indifferent. When $w < t$, Bidder 2 experiences regret even when he wins the prize. But he does not regret all his past decisions: although staying in at $s = 1$ was mistake, staying in periods $s = t - 1$ was not. Bidder 2 only regrets staying in at s such that $-(s - 1) > w - t$, or $s < t - w + 1$. Thus, his expected utility from staying in is

$$\begin{aligned} & p \left(w - t + \rho \left(\sum_{s=1}^{t-w} w - t - [-(s - 1)] \right) \right) + (1 - p) \left(-t + \rho \left(\sum_{s=1}^t -t - [-(s - 1)] \right) \right) \\ &= -t + pw - \rho p \frac{(t - w + 1)(t - w)}{2} - \rho(1 - p) \frac{t(t + 1)}{2}. \end{aligned}$$

In order that Bidder 2 be indifferent between dropping out in period t and remaining in for period t and dropping out in period $t + 1$, $p(t; \rho) = \frac{2+2\rho t}{2w+\rho w(2t+1-w)}$.

It is straightforward to verify that for each t and ρ , $\frac{\partial p(t; \rho)}{\partial \rho} > 0$: sophisticated ρ -rationalizers drop out faster than do bidders who do not care about rationalizing their past bids. This of course implies that the seller's expected revenue is declining in ρ . As $t \rightarrow \infty$, $p(t; \rho) \rightarrow 1/w$, the stationary Nash equilibrium; when $p = 1/w$, for large t the bidders' regret from dropping out in period t is the same as their expected regret from dropping out period $t + 1$. Of course, the low seller revenue in auctions with sophisticated ρ -rationalizers does not accord with the experimental evidence reported above.

The behavior of naifs is more complicated. First, assume that in equilibrium naifs correctly predict each other's strategies, although they naively mispredict their own future preferences and strategies.³⁴ Second, because payoffs are unspecified if neither bidder ever drops out of the auction, it helps to look at the limit of a slightly perturbed version of the game. In the ε -perturbed version of the dollar auction, if both bidders remain in the auction through period t , then after period t the auction ends with probability $\varepsilon \left(\frac{1}{2}\right)^t$, and both players pay $(1 + \rho)t$ and get nothing. This perturbation ensures that the auction will not run forever; it rules out the equilibrium where each bidder remains in the auction forever, knowing that the other does the same, because that way the auction never ends so that neither bidder pays

³⁴Naive bidders who underestimated each other's concern for consistency would overestimate each other's probability of dropping out each period, thereby increasing their own incentive to stay in the auction. This would only strengthen the results on overbidding.

her bid. This perturbation would not affect the above analysis of sophisticates who drop out every period with positive probability. Proposition 7 summarizes equilibrium.

Proposition 7 *A symmetric Nash equilibrium in ε -perturbed dollar auction when both bidders are ν -naive ρ -rationalizers is for each bidder to drop out in period t with probability $p(t; \rho, \nu, \varepsilon)$, where*

$$\lim_{\varepsilon \rightarrow 0} p(t; \rho, \nu, \varepsilon) = \begin{cases} \frac{2+2\rho t - \rho \nu t(t+1)}{2w + \rho(1-\nu)t(t+1)} & t \leq \max \left\{ w, \frac{\rho(2-\nu) + \sqrt{\rho^2(2-\nu)^2 + 8\rho\nu}}{2\rho\nu} \right\} \\ \frac{2+2\rho t - \rho \nu t(t+1)}{2w + \rho(1-\nu)w(2t+1-w)} & w < t < \frac{\rho(2-\nu) + \sqrt{\rho^2(2-\nu)^2 + 8\rho\nu}}{2\rho\nu} \\ 0 & t > \frac{\rho(2-\nu) + \sqrt{\rho^2(2-\nu)^2 + 8\rho\nu}}{2\rho\nu}. \end{cases}$$

In the first period, a ν -naive ρ -rationalizer is less likely to enter the auction than a bidder who does not care about regret. Like a sophisticate he has no past bids to rationalize, and to the extent that he recognizes that he will care about regret in the next period he favors dropping out immediately. Once the auction has gone on long enough, naive bidders become less and less likely to drop out—dropping out means regretting past decisions to stay in, while naivety means that the bidder underestimates his expected regret from staying in. In other words, naive bidders continue to bid not because they expect to win in the future, but because they believe that they will not regret dropping out in the next round. As the proof of Proposition 7 establishes, once the auction has progressed long enough, in each period naive bidders believe that they will drop out with positive probability in the next period, but then when that period arrives they do not. Essentially, bidders are procrastinating in dropping out of the auction, each period thinking that next period it will be less painful to drop out than it actually is.³⁵

Figure 1 illustrates equilibrium bidding for $w = 10$, $\rho = 1/2$, $\nu \in \{0, 1/10\}$. The symmetric Nash equilibrium is for each bidder to drop out with constant probability $1/w = 1/10$. Sophisticated ρ -rationalizers are always more likely to drop out of the auction than bidders who do not care about consistency. In early periods, naive bidders also drop out with probability higher than $1/10$. But as the auction progresses, they escalate their commitment, until eventually they never quit; after period 20, neither bidder ever leaves the auction.

Corollary 1 *For each $\rho, \nu > 0$, there exists some \bar{t} such that for each $t > \bar{t}$, $p(t; \rho, \nu) = 0$. For each $\rho > 0$, there exists $\bar{\nu}$ such that for each $\nu > \bar{\nu}$, the seller's expected revenue is higher than in the stationary Nash equilibrium with $\rho = 0$.³⁶*

³⁵This is qualitatively similar to O'Donoghue and Rabin's (1999, 2001) results on procrastination by people who underestimate the extent to which their future preferences are present-biased.

³⁶When $\rho = 0$, since each bidder is indifferent between dropping out for a sure payoff of zero and staying in the auction, her expected payoff must be 0. The seller's expected revenue conditional on selling the object is w , but with positive probability both bidders drop out in the same period and the seller retains the object.

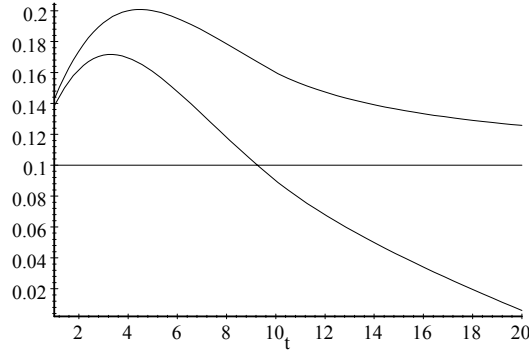


Figure 1: $w = 10$, $\rho = 1/2$, $\nu = 0$ above, $v = 1/10$ below

Both parts of the Corollary follow immediately from the bidding function described in Proposition 7. The first part says that any ρ -rationalizer who is even slightly naive will eventually stay in the dollar auction forever. The second part notes that as long as bidders are sufficiently naive, a taste for consistency leads to higher expected seller revenue than the Nash equilibrium described above. When $\nu = 1$, the bidding functions in Proposition 7 imply that naive bidders always bid more than in Nash equilibrium, which establishes the claim. A taste for consistency when combined with naivety can account for the overbidding in laboratory dollar auctions.

6.3 Naive Procrastination in Search

This section uses an example to show how naivety combines with a taste for consistency to produce procrastination in search models. In the dollar auction, bidders who sink higher and higher bids have incentive to win the auction to rationalize their past bids; here discounting plays a similar role, as people are loathe today to accept jobs, write papers, etc. with values lower than the discounted value of past opportunities. Proposition 4 shows how regret leads to risk aversion in choice over lotteries, which explains why sophisticated ρ -rationalizers drop out quickly in the dollar auction. Likewise, in a search model regret leads a sophisticated ρ -rationalizer to lower her reservation wage, meaning that she accepts an offer too soon. To illustrate how a taste for consistency can lead to procrastination as well, I consider the behavior of a decision maker who is fully naive.

Every period t a worker receives a job offer with wage $w_t \in \{1, 2, 3\}$, where each occurs with probability $1/3$, and offers are independent across periods. In each period t , the worker chooses $a_t \in \{0, 1\}$, namely whether to accept that period's wage offer if she has not already

accepted a previous offer.³⁷ Let $\tau = \min\{t : a_t = 1\}$, the first period that the worker accepts an offer. In period t , the worker knows her current offer as well as any past offers but does not know future options. The worker is an exponential discounter with discount factor δ so that her material payoff is $g(a_1, a_2, \dots, w_1, w_2, \dots) = a_\tau \delta^\tau w_\tau$.

The worker's regret from accepting the wage w_τ depends upon her history of her rejected wage offers. If $w_s \delta^s > w_\tau \delta^\tau$, then she regrets not having accepted w_s when accepting w_τ . Let $h_\tau(w_\tau) \equiv \{w_s : s < \tau, w_s \delta^s > w_\tau \delta^\tau\}$, those wage offers that had they been accepted when offered would have led to higher payoffs than accepting w_τ in period τ . A ρ -rationalizer attempts to maximize $\delta^\tau w_\tau - \rho \left(\sum_{w_s \in h_\tau} \delta^\tau w_\tau - \delta^s w_s \right)$.

Suppose that the worker is fully naive, so that in each period she thinks that she will never experience future regret. Consider some period t such that for each $s < t$, $a_s = 0$. A naive worker believes that in the future she will not care about regret and will simply maximize her expected payoff. This is done by using the threshold rule r that calls for her to accept w_{t+k} iff $\delta^{t+k} w_{t+k} \geq \sum_{s=t+k+1}^{\infty} \delta^s (1 - F(r)) F^{s-(t+k+1)}(r) E[w_s | w_s \geq r]$. As a result, the worker accepts option w_t if

$$\delta^t w_t + \rho \left(\sum_{w_s \in h_t} \delta^t w_t - \delta^s w_s \right) \geq \sum_{s=t+1}^{\infty} \delta^s (1 - F(r)) F^{s-(t+1)}(r) E[w_s | w_s \geq r],$$

namely if her utility from accepting her current option is higher than her expected payoff from following the threshold rule r in all future periods. Hence, the worker accepts offers w_t such that

$$w_t + \rho \left(\sum_{w_s \in h_t} w_t - \delta^{s-t} w_s \right) \geq \frac{\delta \int_{w \geq r} w dF(w)}{1 - \delta F(r)}.$$

For $\rho = 0$, the worker simply exercises w_t whenever it exceeds the option value of waiting. As $\rho \rightarrow \infty$, the worker only accepts wages that rationalize all previous decisions to keep searching. If $r = 1$, the worker accepts all wages next period, for an expected discounted payoff of $2\delta^{t+1}$. If $r = 2$, the worker's expected payoff is $\frac{5\delta^{t+1}}{3-\delta}$, and for $r = 3$ it is $\frac{3\delta^{t+1}}{3-2\delta}$.

In the first period, the worker has nothing to regret and no expectation of future regret, and therefore she behaves just as she expects to behave in the future, so she uses the threshold r . For $\delta < \frac{1}{2}$, the worker accepts $w_1 = 1$. For $\delta > \frac{6}{7}$, accepts only $w_1 = 3$. For intermediate values of $\delta \in [\frac{1}{2}, \frac{6}{7}]$, the worker is selective but not too selective, accepting $w_1 \geq 2$ and expecting to use $r = 2$ in the future. I focus on this case. If $w_1 \in \{2, 3\}$, then the worker accepts her first job offer, and so consider instead the case where $w_1 = 1$. In the second period, because $\delta^{-1} 1 \leq r = 2$, $w_1 = 1 \notin h_2(w_2 = 2)$, namely the worker does not regret

³⁷ A naive worker in period t would never want to go back and accept an offer that she rejected in $t - k$.

having passed up the offer of $w_1 = 1$ when she accepts $w_2 = 2$. Thus, the worker accepts $w_2 = 2$. Suppose instead that $w_2 = 1$. If $\delta^{-2} > 2$ or $(\delta \lesssim \frac{5}{7})$, then $w_1 = 1 \in h_2(w_3 = 2)$, meaning that the worker regrets not having accepted the first-period offer of 1 when accepting a third-period offer of 2. (By the same logic as before, the worker does not regret rejecting $w_2 = 1$ when accepting $w_3 = 2$.) Furthermore, if $(1 + \rho)2 - \rho\delta^{-2} < \frac{5\delta}{3-\delta}$ (or $\rho > \frac{7\delta-6}{(3-\delta)(2-\delta^{-2})}$), then the worker's regret is strong enough that she rejects $w_3 = 2$. If $\delta^{-2} > 3$ or $\delta < \frac{1}{3^{\frac{1}{2}}} \approx \frac{4}{7}$, then the worker regrets rejecting $w_1 = 1$ even when accepting the highest possible wage of $w_3 = 3$. If $(1 + \rho)3 - \rho\delta^{-2} < \frac{3\delta}{3-2\delta}$ (if $\rho > \frac{9\delta-9}{(3-2\delta)(3-\delta^{-2})}$), then the worker rejects even the highest possible wage of 3. For example, if $\delta = \frac{1}{2}$, then if $\rho > \frac{9}{4}$, the worker rejects $w_3 = 3$. Hence, for $\delta = \frac{1}{2}$ and $\rho > \frac{9}{4}$, a naive worker searches forever with probability no less than $\frac{1}{9}$. Again, a naive worker procrastinates because accepting an offer forces her to accept that she has made mistakes in the past by searching for so long, whereas she mistakenly believes that in the future she will find accepting a wage easier than she does now.

7 Conclusion

People dislike regretting their past choices. One way that they avoid regretting past choices is by taking current choices that make their past choices look better: people have a taste for current actions for which their past actions were optimal. A taste for rationalizing the past affects choice when material payoffs have strategic complementarities, namely when current actions affect the material incentives to take past actions. When actions are strategic complements, a preference for rationalizing the past leads to a sunk-cost effect—past actions that were too high increase the incentive for high current actions—as well as an “unsunk-effect”—past actions that were too low increase the incentive for low current actions. The sunk-cost effect may lead a sophisticated firm to be cautious about beginning an investment project because it knows that once started, it will have a tendency to continue the project. The unsunk-cost effect may lead that same firm to begin an investment project because it knows that if it does not start now, it may never start. Which of these effect dominates depends of course on the environment. An intuition similar to the unsunk-cost effect leads to the declining price anomaly in sequential auctions; bidders do not want to win an auction at a higher price than they could have won some previous auction. People who naively believe that in the future they will care less about rationalizing their past behavior than they actually do may overbid in wars of attrition (the dollar auction) and procrastinate in search for jobs, etc.

In individual decision problems with strategic complements, someone can rationalize an overly-high first-period action by raising her second-period action. This may have interesting

interactions with time-inconsistent, present-biased preferences.³⁸ DellaVigna and Malmendier (2002) examine consumers' choice of tariffs at health clubs, finding that people too often choose tariffs with zero marginal prices; they pay an average of \$17 per visit despite having the option of a \$10 pay-as-you-go tariff. Suppose that attending the gym entails paying immediate costs for future rewards. A sophisticated ρ -rationalizer who would like to attend the gym in the future may sink a membership fee to give her future selves incentive to attend the gym; in this way, the sunk-cost effect may help to counteract a future taste for immediate gratification.³⁹

All of the interpersonal, strategic settings in the paper share the feature that a player can “end the game” by taking some action. Extending the model to strategic settings without this feature presents the complication that what one agent does in one period may depend upon what another did in some earlier period. As a result, asking whether one player's past action was optimal given her current action makes little sense. The same problem arises in a single-person setting where the person's first-period action affects the distribution over states of the world. An extension of the paper would be model how people conceptualize regret in these settings.

In my model, regret makes people risk averse over payoffs. Many authors have suggested that someone who admits to herself that she has made a past mistake may willingly accept risk in the hopes of getting back to some reference point, or undoing her mistake. Kahneman and Tversky's (1979) prospect theory proposes that people have a reference-dependent value function for money: it is convex for money outcomes below the reference point, concave for outcomes above the reference point, and kinked at the reference point such that the slope in the losses domain is higher than that in the gains domain. Section 4 illustrates a setting in which my model generates a natural “reference point”—the best foregone, safe payoff—at which preferences are kinked, but people are risk neutral over lotteries entirely above or below that reference point. Extending the model to incorporate Kahneman and Tversky's risk attitudes may help capture escalation in the dollar auction with less naivety. Nevertheless,

³⁸Laibson (1997) introduces the now standard model of (β, δ) discounting, where people are exponential discounters except that they discount all future periods by the additional factor β .

³⁹My model has implications in other settings where present-biased preferences have been applied. For example, Gruber and Köszegi (2001) and O'Donoghue and Rabin (2001) apply these preferences to Becker and Murphy's (1988) model of addiction, where a good is addictive if current consumption lowers future utility at the same time that it increases future marginal utility of consumption. Because smoking yesterday and today are strategic complements, in my model a smoker regrets not having quit yesterday more if he quits today than if he continues to smoke. In other words, people are reluctant to quit smoking (start exercising, etc.) because they regret not having quit in the past. A paternalistic government might tell people that “it is never too late to give up smoking ” (as in Britain) not so much to inform them that the health benefits of quitting are large—people already know this—but rather to convince them that they do not have “too much invested to quit.”

explaining the very severe overbidding in most classroom dollar auctions through risk attitudes would require an implausibly high degree of risk-lovingness, which suggest that naivety plays a central role in these auctions.

Another way to extend the model would be by adding “rejoicing” to regret. Not only may people experience more displeasure the worse some past choice performed relative to the best alternative, but they may experience more pleasure the better some past choice performed relative to the next best alternative. Indeed, Loomes and Sugden (1982, 1983, 1986) make this a central feature of their regret theory. In some applications, rejoicing would counteract regret. For example, if bidders in an auction rejoiced at paying less for a good than they would have paid by winning an earlier auction for an identical good, then they would have incentive to bid higher in later auctions, undermining the declining price anomaly. Nevertheless, the psychology literature suggests that regret dominates rejoicing. Introducing a subordinate taste for rejoicing should not affect any of my results.

The extent to which regret influences preferences in any particular state of the world may depend on the *ex ante* likelihood of that state. For example, Thaler’s (1980) family may be less likely to attend the basketball game in the snow if it lives in Miami than if it lives in Boston. Because no one blames herself for failing to foresee snow in Miami, a family in Miami may feel little need to rationalize its choice to purchase tickets. A ρ -rationalizer’s preferences could easily be extended to allow for a state-dependent ρ , where $\rho(\omega)$ increases in $\mu(\omega)$: regret plays a more prominent role in predictable states.⁴⁰ A related idea is that people may regret some types of actions more than others. Gilovich and Medvec (1995) argue that people regret things that they actively do more than those that they passively do not do.

Because it operates naturally through preferences, my model can be applied to a wide variety of settings. However, framing effects impede its application. Consider, for instance, a firm that chooses some investment level, learns something about the quality of the project, and then chooses another investment level. In the first period, the firm invests q_1 . In the second period, the firm could be thought to choose either q_2 or $q_1 + q_2$. But fixing $q_1 + q_2$ and reoptimizing over q_1 is different than fixing q_2 and reoptimizing over q_1 . These framing issues underscore how important it is to understand how people conceptualize their choices. But they also serve as a reminder that while my model captures the idea that people manipulate choice to avoid regret, it is only a crude approximation of their decision making.

This paper has argued that people avoid regret by choosing actions for which past actions were optimal. People who do this engage in a form of self deception, essentially convincing themselves that their current preferences are different than what they would be if they had

⁴⁰This extension seems more appealing in Thaler’s example than in the dollar auction, where bidders likely would hold themselves responsible for an improbably large escalation in bids.

nothing to regret in the past. The cognitive dissonance literature argues that people fail to learn new information that conflicts with their previous beliefs and use other methods of dissonance reduction. In some settings, my model makes very different predictions from other types of dissonance reduction. For example, Akerlof and Dickens (1982) discuss why construction workers may eschew safety equipment in order to convince themselves that their jobs are safe and that they are therefore prudent decision makers. Suppose that a worker first chooses which industry to work in, then learns how dangerous the industry is, and finally chooses whether to use safety equipment. Working in a dangerous industry increases the worker's material incentive to wear safety equipment. In my model, this increases the worker's incentive to wear it—the better the safety equipment, the more prudent the decision to work in construction. Someone may be overly cautious in order to convince herself that she has not mistakenly chosen a job that is too dangerous.

8 Appendix

Proof of Proposition 1. If A_1 is a singleton set, then for each $\omega \in \Omega$ and $a_2 \in A_2$, $r(a_1|a_2, \omega) = 0$, and therefore ρ does not affect a_2 . Now suppose that $\Omega = \{\omega\}$, and let (a_1^*, a_2^*) maximize $g(a_1, a_2)$. If she chooses a_1^* , the DM believes that she will choose a_2 to maximize $g(a_1^*, a_2, \omega) + \rho(1 - \nu)(g(a_1^*, a_2, \omega) - \max_{a_1 \in A_1} g(a_1, a_2, \omega))$, which is done when $a_2 = a_2^*$ (no matter what ρ and ν). Since for each (a_1, a_2)

$$g(a_1, a_2, \omega) + \rho(1 - \nu) \left(g(a_1, a_2, \omega) - \max_{a_1 \in A_1} g(a_1, a_2, \omega) \right) \leq g(a_1^*, a_2^*, \omega),$$

the DM knows she can do no better than by choosing a_1^* and a_2^* . If $A_2 = \{a_2\}$, the DM's first-period objective function is

$$\sum_{\omega \in \Omega} \max_{a_2 \in A_2} \left\{ (1 + \rho(1 - \nu))g(a_1, a_2, \omega) - \rho(1 - \nu) \max_{a_1 \in A_1} g(a_1, a_2, \omega) \right\} \mu(\omega),$$

which simplifies to

$$\left(\sum_{\omega \in \Omega} (1 + \rho(1 - \nu))g(a_1, a_2, \omega)\mu(\omega) \right) - \rho(1 - \nu) \sum_{\omega \in \Omega} \max_{a_1 \in A_1} g(a_1, a_2, \omega)\mu(\omega).$$

The second term does not depend upon a_1 , and therefore the DM chooses a_1 to maximize $\sum_{\omega \in \Omega} g(a_1, a_2, \omega)\mu(\omega)$, and so a_1^* does not depend on ρ or ν . Finally, note that if $g(a_1, a_2, \omega) = g_1(a_1, \omega) + g_2(a_2, \omega)$,

$$\begin{aligned} r(a_1|a_2, \omega) &= g(a_1, \omega) + g(a_2, \omega) - \max_{a_1 \in A_1(a_2, \omega)} \{g_1(a_1, \omega) + g_2(a_2, \omega)\} \\ &= g(a_1, \omega) - \max_{a_1 \in A_1(a_2, \omega)} g_1(a_1, \omega). \end{aligned}$$

Suppose by way of contradiction that for some $a_2, a'_2 \in A_2(a_1, \omega)$, $A_1(a_2, \omega) \neq A_1(a'_2, \omega)$. Then there exists some $a'_1 \in A_1$ such that $a_2 \in A_2(a'_1, \omega)$ but $a'_2 \notin A_2(a'_1, \omega)$, and therefore $A_2(a_1, \omega) \neq A_2(a'_1, \omega)$, which contradicts the assumption. Thus $r(a_1|a_2, \omega)$ does not depend on a_2 , and hence a_2^* does not depend upon ρ or a_1 . But then the DM's first-period objective function is

$$\sum_{\omega \in \Omega} \max_{a_2 \in A_2} g_2(a_2, \omega)\mu(\omega) + (1 + \rho(1 - \nu)) \sum_{\omega \in \Omega} g(a_1, \omega)\mu(\omega) - \rho(1 - \nu) \sum_{\omega \in \Omega} \max_{a_1 \in A_1(\omega)} g_1(a_1, \omega)\mu(\omega).$$

Only the second term depends on a_1 , which implies that a_1^* does not depend on ρ or ν . ■

Proof of Proposition 2. (\Leftarrow) Trivial and therefore omitted.

(\Rightarrow) Suppose that for some A_1 , ρ , and ν , $a_1^*(A_1; \rho, \nu) \neq a_1^*(A_1; 0, 0)$; wlog suppose $\nu = 0$. The goal is to find another choice set $A'_1 \ni a_1^*(A_1; \rho, 0)$ such that $a_1^*(A'_1; \rho, 0) \neq a_1^*(A_1; \rho, 0)$. Let

$B \equiv \{a_1^*(A_1; \rho, 0), a_1^*(A_1; 0, 0)\}$. If $a_1^*(B; \rho, 0) = a_1^*(B; 0, 0) = a_1^*(A_1; 0, 0) \neq a_1^*(A_1; \rho, \nu)$, then the proof is complete, so assume that $a_1^*(B; \rho, 0) \neq a_1^*(B; 0, 0)$. Define

$$\begin{aligned}\bar{k} &\equiv \max_{\{(a_1, a_2, \omega): a_1 \in B, \omega \in \Omega, a_2 \in A_2(a_1, \omega)\}} g(a_1, a_2, \omega), \\ \underline{k} &\equiv \min_{\{(a_1, a_2, \omega): a_1 \in B, \omega \in \Omega, a_2 \in A_2(a_1, \omega)\}} g(a_1, a_2, \omega), \text{ and} \\ \Omega^* &\equiv \left\{ \omega : \max_{a_2 \in A_2(a_1^*(B; 0, 0), \omega)} g(a_1^*(B; 0, 0), a_2, \omega) \geq \max_{a_2 \in A_2(a_1^*(B; \rho, 0))} g(a_1^*(B; \rho, 0), a_2, \omega) \right\}.\end{aligned}$$

Both Ω^* and its complement must be non-empty: if Ω^* were empty, then clearly the DM with $\rho = 0$ would change her choice; if Ω/Ω^* were empty, then choosing $a_1^*(B; 0, 0)$ would lead to no regret and hence the ρ -rationalizer would choose it. Define b such that for each $\omega \in \Omega$, $A_2(b, \omega) = A_2(a_1^*(B; 0, 0, \omega) \cup A_2(a_1^*(B; \rho, 0, \omega)$, and for each $a_2 \in A_2(b, \omega)$,

$$g(b, a_2, \omega) = \begin{cases} \max_{a_2 \in A_2(a_1^*(A_1; 0, 0), \omega)} g(a_1^*(B; 0, 0), a_2, \omega) & \text{for } \omega \in \Omega^* \\ \frac{\underline{k} + \rho(\underline{k} - \bar{k}) - \mu(\Omega^*)\bar{k}}{1 - \mu(\Omega^*)} - \varepsilon & \text{for } \omega \notin \Omega^*, \end{cases}$$

where $\varepsilon > 0$. Consider the choice set $B \cup \{b\}$. First note that the DM does not choose b , for the highest utility it could yield is

$$\mu(\Omega^*)\bar{k} + (1 - \mu(\Omega^*)) \left(\frac{\underline{k} + \rho(\underline{k} - \bar{k}) - \mu(\Omega^*)\bar{k}}{1 - \mu(\Omega^*)} - \varepsilon \right) < \underline{k} + \rho(\underline{k} - \bar{k}),$$

the lowest utility that either $a_1^*(B; \rho, \nu)$ or $a_1^*(B; 0, 0)$ could yield. Next note that $\omega \in \Omega^*$, b was an optimal first period action and therefore whatever her first-period choice, a ρ -rationalizer chooses a_2 to maximize $g(\cdot, a_2, \omega)$. For $\omega \notin \Omega^*$, b does not affect a ρ -rationalizer's second-period action or utility.

A ρ -rationalizer's expected utility from choosing $a_1^*(B; \rho, 0)$ from $B \cup \{b\}$ is $V(a_1^*(B; \rho, 0), \rho)$

$$\begin{aligned}&= \sum_{\omega \in \Omega^*} \left[(1 + \rho) \max_{a_2 \in A_2(a_1^*(B; \rho, 0), \omega)} g(a_1^*(B; \rho, 0), a_2, \omega) - \rho \max_{a_2 \in A_2(a_1^*(B; 0, 0), \omega)} g(a_1^*(B; 0, 0), a_2, \omega) \right] \mu(\omega) \\ &\quad + \sum_{\omega \notin \Omega^*} \max_{a_2 \in A_2(a_1^*(B; \rho, 0), \omega)} g(a_1^*(B; \rho, 0), a_2, \omega) \mu(\omega),\end{aligned}$$

since for each $\omega \in \Omega^*$ and $a_2 \in A_2(a_1^*(B; 0, 0), \omega)$,

$$\max_{a_1 \in A_1(a_2, \omega)} g(a_1, a_2, \omega) = g(a_1^*(B; 0, 0), a_2, \omega) = \max_{a_2 \in A_2(a_1, \omega)} g(a_1^*(B; 0, 0), a_2, \omega);$$

for $\omega \notin \Omega^*$, a_1' leads to no regret. The DM's expected utility from $a_1^*(B; 0, 0)$ is no smaller than

$$\begin{aligned}\underline{V}(a_1^*(B; 0, 0), \rho) &\equiv \sum_{\omega \in \Omega^*} \max_{a_2 \in A_2(a_1, \omega)} g(a_1, a_2, \omega) \mu(\omega) \\ &\quad + \sum_{\omega \notin \Omega^*} \left[(1 + \rho) \max_{a_2 \in A_2(a_1, \omega)} g(a_1, a_2, \omega) - \left(\max_{a_2 \in A_2(a_1', \omega)} g(a_1, a_2, \omega) \right) \right] \mu(\omega),\end{aligned}$$

since for $\omega \notin \Omega^*$,

$$r(a_1^*(B; 0, 0)|a_1, \omega) \leq \max_{a_2 \in A_2(a_1^*(B; 0, 0), \omega)} g(a_1^*(B; 0, 0), a_2, \omega) - \max_{a_2 \in A_2(a_1^*(B; \rho, 0), \omega)} g(a_1^*(B; \rho, 0), a_2, \omega).$$

Then

$$\begin{aligned} & \underline{V}(a_1^*(B; 0, 0), \rho) - V(a_1^*(B; \rho, 0), \rho) \\ &= (1 + \rho) \left(\sum_{\omega \in \Omega} \max_{a_2 \in A_2(a_1^*(B; 0, 0), \omega)} g(a_1^*(B; 0, 0), a_2, \omega) - \sum_{\omega \in \Omega} \max_{a_2 \in A_2(a_1^*(B; \rho, 0), \omega)} g(a_1^*(B; \rho, 0), a_2, \omega) \right) \\ &= (1 + \rho) (V(a_1^*(B; 0, 0), 0) - V(a_1^*(B; \rho, 0), 0)) > 0, \end{aligned}$$

since a DM who does not care about consistency chooses $a_1^*(B; 0, 0)$ from B iff $V(a_1^*(B; 0, 0), 0) > V(a_1^*(B; \rho, 0), 0)$. Thus a ρ -rationalizer chooses $a_1^*(B; \rho, 0)$ from B and $a_1^*(B; 0, 0) (\neq a_1^*(B; \rho, 0))$ from $B \cup \{b\}$, violating WARP. ■

Proof of Proposition 4. The proofs of both statements are straightforward. Define

$$h_{\rho, k}(x) \equiv \begin{cases} x & x > k \\ x - \rho(k - x) & x < k. \end{cases}$$

Note $h_{\rho, k}(\lambda x + (1 - \lambda)x') \leq \lambda h_{\rho, k}(x) + (1 - \lambda)h_{\rho, k}(x')$; $h_{\rho, k}$ is concave. A ρ -rationalizer chooses a_2 to maximize $E[h_{\rho, k} \circ g(a_1, a_2, \omega) | \omega \in P_1(\omega)]$, the expectation of a concave function, and the result is well-known. Note equality holds only when preferences are linear, which happens only for $k \notin (\min_{\omega \in P_1(\omega)} g(a_1, a_2, \omega), \max_{\omega \in P_1(\omega)} g(a_1, a_2, \omega))$. The second result follows from noting that if $l < k$, then raising k to $k' > k$ decreases the right hand-side of the equality by $\rho(k' - k)$ and can decrease the left-hand side no more. ■

Proof of Proposition 5. Since ω plays no role in the proof, it is dropped to conserve notation. From the statement in the proof of Proposition 3 (in text) that $\max_{a_2 \in A_2(a_1, \omega)} g(a_1, a_2) - \rho r(a_1|a_2)$ is decreasing and convex in ρ , $r(a_1|a_2^*(a_1; \rho)) \leq r(a_1|a_2^*(a_1))$, which holds iff

$$g(a_1, a_2^*(a_1; \rho)) - g(a_1^*(a_2^*(a_1; \rho)), a_2^*(a_1; \rho)) \geq g(a_1, a_2^*(a_1)) - g(a_1^*(a_2^*(a_1)), a_2^*(a_1)).$$

Towards a contradiction, suppose that $a_2^*(a_1) > a_2^*(a_1; \rho)$. Then by strategic complements, since $a_1 > a_1^*(a_2^*(a_1))$ by assumption,

$$g(a_1, a_2^*(a_1)) - g(a_1^*(a_2^*(a_1)), a_2^*(a_1)) > g(a_1, a_2^*(a_1; \rho)) - g(a_1^*(a_2^*(a_1)), a_2^*(a_1; \rho)).$$

Combining the two inequalities gives $g(a_1^*(a_2^*(a_1)), a_2^*(a_1; \rho)) > g(a_1^*(a_2^*(a_1; \rho)), a_2^*(a_1; \rho))$, which contradicts the definition $a_1^*(a_2^*(a_1; \rho)) \equiv a_1 \in \arg \max_{a_1 \in A_1(a_2, \omega)} g(a_1, a_2^*(a_1; \rho))$. Thus, $a_2^*(a_1; \rho) \geq a_2^*(a_1)$.

Now suppose that $g(a_1, a_2)$ is strictly quasi-concave in a_2 and $\rho' > \rho > 0$. Suppose by way of contradiction that $a_2^*(a_1; \rho') < a_2^*(a_1; \rho)$. From the first statement of the Proposition, $a_2^*(a_1) \leq a_2^*(a_1; \rho') < a_2^*(a_1; \rho)$. Since g is strictly quasi-concave in a_2 , $g(a_1, a_2^*(a_1; \rho)) < g(a_1, a_2^*(a_1; \rho'))$. Thus $r(a_1|a_2^*(a_1; \rho)) < r(a_1|a_2^*(a_1; \rho'))$, otherwise $a_2^*(a_1; \rho)$ would not be chosen by a ρ -rationalizer. But this contradicts the result in the proof of Proposition 3 that $r(a_1|a_2^*(a_1; \rho')) \leq r(a_1|a_2^*(a_1; \rho))$. ■

Proof of Proposition 7. Suppose that $w > t$. Again Bidder 2's payoff from dropping out in period t is $-(t-1) - \rho \frac{(t-1)t}{2}$. His perceived payoff from dropping out in period $t+1$ is

$$\frac{\varepsilon}{2}(1+\rho)(-t) + \left(1 - \frac{\varepsilon}{2}\right) \left(-t + pw - \rho(1-\nu)(1-p)\frac{t(t+1)}{2}\right).$$

If Bidder 2 is indifferent between dropping out in period t and waiting until period $t+1$ to drop out, then

$$p = \frac{1 - t - \rho \frac{(t-1)t}{2} + \frac{\varepsilon}{2}(1+\rho)t + \left(1 - \frac{\varepsilon}{2}\right) \left(t + \rho(1-\nu)\frac{t(t+1)}{2}\right)}{\left(1 - \frac{\varepsilon}{2}\right) \left(w + \rho(1-\nu)\frac{t(t+1)}{2}\right)},$$

the limit of which as $\varepsilon \rightarrow 0$ is as in the Proposition. The argument for $w < t$ is similar and therefore omitted. For these indifference conditions to hold, Bidder 2 must believe in period t that she drops out in period $t+1$ with positive probability. Suppose by way of contradiction that in some period \hat{t} , some Bidder 2 believes that he will stay in with probability one in $\hat{t}+1$. Since naivety leads Bidder 2 to overestimate the probability that he will drop out, and the equilibrium is symmetric, Bidder 2 believes that Bidder 1 will also remain in period $\hat{t}+1$ with probability one. Either Bidder 2 believes that there exists some period $t' > \hat{t}+1$ in which Bidder 2 drops out with positive probability but where neither drops out in $t'-1$ with positive probability, or he believes that both will stay in the auction forever. The introduction of ε rules out the latter, for by dropping out in period $\hat{t}+1$ he gets a perceived payoff of $(1+\rho(1-\nu))(-\hat{t})$, while by staying in forever he gets an expected payoff less than $(1+\rho)(-(\hat{t}+1))$. The former cannot occur either, for given that Bidder 1's strategy calls for her to stay in for sure in period $t'-1$, Bidder 2 believes that he would strictly prefer dropping out in $t'-1$ to dropping out in t' . Thus, the equilibrium condition that each bidder thinks that she will drop out in the next period with positive probability holds. ■

References

- [1] Akerlof, George and William T. Dickens. "The Economic Consequences of Cognitive Dissonance," *American Economic Review* 72, June 1982, 307-319.
- [2] Allais, Maurice. "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine," *Econometrica* 21, 1953: 503-546.
- [3] Arkes, Hal R. and Peter Ayton. "The Sunk Cost and Concorde Effects: Are Humans Less Rational than Lower Animals?" *Psychological Bulletin* 125(5), 1999: 591-600.
- [4] Arkes, Hal R. and Catherine Blumer, "The Psychology of Sunk Cost," *Organizational Behavior and Human Decision Processes* 87(2), March 2001: 371-385.
- [5] Arkes, Hal R. and Laura Hutzel. "The Role of Probability of Success Estimates in the Sunk Cost Effect." *Journal of Behavioral Decision Making* 13, 2000: 295-306.
- [6] Aronson, Elliot. *The Social Animal, Seventh Edition*. New York: W.H. Freeman and Company, 1996.
- [7] Ashenfelter, Orley. "How Auctions Work for Wine and Art." *Journal of Economic Perspectives*, 3(3), Summer 1989: 23-36
- [8] Barron, John M.; Chulkov, Dmitriy V; and Glen R. Waddell. "The Escalation Phenomenon and Executive Turnover: Theory and Evidence." Mimeo, Purdue University, July 2001.
- [9] Bastardi, Anthony and Eldar Shafir. "On the Pursuit and Misuse of Useless Information." *Journal of Personality and Social Psychology* 75(1), 1998, 19-32.
- [10] Bastardi, Anthony and Eldar Shafir. "Nonconsequential Reasoning and Its Consequences." *Current Directions in Psychological Science*. 9(6), December 2000, 216-219.
- [11] Bazerman, Max H. *Judgment in Managerial Decision Making, Fifth Edition*. New York: John Wiley & Sons, Ltd., 2002.
- [12] Bazerman, Max H.; Beekun, Rafik I.; and F. David Schoorman. "Performance Evaluation in a Dynamic Context: A Laboratory Study of the Impact of a Prior Commitment to the Ratee," *Journal of Applied Psychology* 67(6), 1982: 873-876.
- [13] Becker, Gary and Kevin M. Murphy. "A Theory of Rational Addiction." *Journal of Political Economy* 96(4), August 1988: 675-700
- [14] Bell, David E. "Regret in Decision Making Under Uncertainty," *Operations Research* 30(5), September-October 1982: 961-981.
- [15] Black, Jane and David De Meza. "Systematic Price Differences between Successive Auctions Are No Anomaly." *Journal of Economics and Management Strategy* 1(4), Winter 1992: 607-28
- [16] Boehne, Donna M. and Paul W. Paese. "Deciding Whether to Complete or Terminate an Unfinished Project: A Strong Test of the Project Completion Hypothesis." *Organizational Behavior and Human Decision Processes*, 81(2), March 2000: 178-194.
- [17] Bornstein, Brian H. and Gretchen B. Chapman. "Learning Lessons from Sunk Costs." *Journal of Experimental Psychology: Applied* 4(1), 1995: 251-269.

- [18] Bulow, Jeremy I.; Geanakoplos, John D.; and Paul D. Klemperer. "Multimarket Oligopoly: Strategic Substitutes and Complements." *Journal of Political Economy*; 93(3), June 1985: 488-511.
- [19] Camerer, Colin and Roberto Weber. "The Econometrics and Behavioral Economics of Escalation of Commitment: A Re-Examination of Staw and Hoang's NBA Date," *Journal of Economic Behavior and Organization* 39(1), May 1999, 59-82.
- [20] Carmichael, Lorne and W. Bentley MacLeod. "Caring about Sunk Costs: A Behavioral Solution to Hold-Up Problems with Small Stakes." Olin Working Paper 99-19, USC Law School, January 2000.
- [21] Courty, Pascal and Hao Li. "Sequential Screening," *Review of Economic Studies* 67, 2000, 697-717.
- [22] Dawkins, Richard and H. Jane Brockmann. "Do Digger Wasps Commit the Concorde Fallacy?" *Animal Behavior* 28, 1980: 892-896.
- [23] De Frutos, Maria Angeles and Robert W. Rosenthal. "On Some Myths about Sequenced Common-Value Auctions." *Games and Economic Behavior* 23(2) May 1998: 201-21
- [24] DellaVigna, Stefano and Ulrike Malmendier. "Self Control in the Market: Evidence from the Health Club Industry." Mimeo, University of California at Berkeley, July 2002.
- [25] Dick, Alan S. and Kenneth R. Lord. "The Impact of Membership Fees on Consumer Attitude and Choice," *Psychology & Marketing* 15(1), January 1998: 41-58.
- [26] Festinger, Leon. *A Theory of Cognitive Dissonance*. Evanston, IL. Row, Peterson: 1957.
- [27] Gilovich, Thomas and Victoria Husted Medvec. "The Experience of Regret: What, When, and Why." *Psychological Review* 102(2), 1995: 379-395.
- [28] Gottfries, Nils and Keith Hylton. "Are M.I.T. Students Rational? Report on a Survey." *Journal of Economic Behavior and Organization* 8(1), March 1987: 113-120.
- [29] Gruber, Jonathan and Botond Köszegi. "Is Addiction Rational? Theory and Evidence." *Quarterly Journal of Economics* 116(4), 2001: 1261-1305
- [30] Heath, Chip. "Escalation and De-Escalation of Commitment in Response to Sunk Costs: The Role of Budgeting in Mental Accounting." *Organizational Behavior and Human Decision Processes*, 62(1), April 1995: 38-54.
- [31] Kahneman, Daniel and Amos Tversky. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47, 1979: 263-291.
- [32] Laibson, David. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112(2), May 1997: 443-477.
- [33] Loewenstein, George and David Adler. "A Bias in the Prediction of Tastes," *Economic Journal* 105 (431), 929-937, 1995.
- [34] Loomes, Graham and Robert Sugden. "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *The Economic Journal* 92, December 1982: 805-824.
- [35] Loomes, Graham and Robert Sugden. "A Rationale for Preference Reversal," *American Economic Review* 73(3), 1983: 428-432.

- [36] Loomes, Graham and Robert Sugden. "Disappointment and Dynamic Consistency in Choice under Uncertainty," *Review of Economic Studies* 53, 1986: 271-282.
- [37] McAfee, R. Preston and Daniel Vincent. "The Declining Price Anomaly." *Journal of Economic Theory* 60(1), June 1993: 191-212
- [38] Milgrom, Paul and Ilya Segal. "Envelope Theorems for Arbitrary Choice Sets." *Econometrica* 70(2), March 2002: 583-601.
- [39] O'Donoghue, Ted and Matthew Rabin. "Doing It Now or Later." *American Economic Review* 89(1), March 1999: 103-124.
- [40] O'Donoghue, Ted and Matthew Rabin. "Choice and Procrastination." *Quarterly Journal of Economics*, February 2001: 121-160.
- [41] O'Donoghue, Ted and Matthew Rabin. "Addiction and Present-Biased Preferences," Mimeo, University of California at Berkeley, December 2001.
- [42] Prendergast, Canice and Lars Stole. "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning." *Journal of Political Economy* 104(6), 1105-1134.
- [43] Rabin, Matthew. "Cognitive Dissonance and Social Change," *Journal of Economic Behavior and Organization*; 23(2), March 1994, pp. 177-94
- [44] Rabin, Matthew and Joel Schrag. "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics* 114(1), February 1999: 37-82.
- [45] Schoorman, F. David. "Escalation Bias in Performance Appraisals: An Unintended Consequence of Supervisor Participation in Hiring Decisions," *Journal of Applied Psychology* 73(1), 1988: 58-62.
- [46] Shubik, Martin. "The Dollar Auction Game: A Paradox in Noncooperative Behavior and Escalation," *Journal of Conflict Resolution* 15(1), 1971: 109-111.
- [47] Sugden, Robert. "An Axiomatic Foundation for Regret Theory," *Journal of Economic Theory* 60, 1993: 159-180.
- [48] Teger, Alan I. *Too Much Invested to Quit*. New York, NY. Pergamon Press: 1980.
- [49] Thaler, Richard H. "Toward a Positive Theory of Consumer Choice," *Journal of Economic Behavior and Organization* 1, 1980: 39-60, reprinted in Thaler, Richard H., *Quasi-Rational Economics*, Russell Sage Foundation, New York, 1991.
- [50] Tykocinski, Orit E.; and Thane S. Pittman. "The Consequences of Doing Nothing: Inaction Inertia as Avoidance of Anticipated Counterfactual Regret." *Journal of Personality and Social Psychology* 75(3), 1998, 607-616.
- [51] Tykocinski, Orit E.; Pittman, Thane S.; and Erin E. Tuttle. "Inaction Inertia: Foregoing Future Benefits as a Result of an Initial Failure to Act." *Journal of Personality and Social Psychology* 68(5), 1995, 793-803.
- [52] Yariv, Leeat. "I'll See it When I Believe It—A Simple Model of Cognitive Consistency." Cowles Foundation Discussion Paper 1352, February 2002.