

ESSAYS ON APPLIED INDUSTRIAL ORGANIZATION

Ph.D. Dissertation
Department of Economics
University of Siena

July 2007

Ph.D. Candidate
Dora Kadar

Supervisor Prof. Roberto Renó

Acknowledgements

To Gizike

I would like to dedicate this work to my Mom, whose support in the past years has been invaluable.

I am equally indebted to Francesco Drago, with whom I co-authored one chapter of this dissertation and shared exciting trips and countless hours of discussions and learning. Thanks for his constant encouragement, patience and help throughout the program.

Many other people have contributed to this dissertation in different ways. In particular, I wish to thank Ugo Pagano and Sam Bowles for a great PhD program in Siena and for my visiting positions in highly acknowledged universities in the world: Berkeley, California and IDEI, Toulouse; Marianna Belloc, for an unforgettable academic year in Berkeley and for guiding my first steps in econometrics, Roberto Reno, for being my supervisor; Bronwyn Hall and Paul Seabright, for sharing their data with me and for valuable comments and suggestions during my visits in Berkeley and Toulouse; Monica Schnitzer, for discussions and guidance during my visit in Munich; my professors in Siena and Berkeley, who have thought me during these years; and Laszlo Kadar, Zoltan Kadar, Pietro Vertova, Roberto Galbiati, Simone D'Alessandro, Martina Pignatti Morano and Luigi Marattin, for the laughs and support.

Dora Kadar
New York, July 2007

Table of Contents

Abstract	p. 1
Chapter 1. Behavioral Consumer Choice	p. 3
Chapter 2. Are Acquirers Rational?	p. 28
Chapter 3. Agency Costs, Synergy Gains and Bargaining in Corporate vs. Financial Takeovers	p. 60
Chapter 4. Rebate or Bait? A Model on Time Inconsistent Consumer Behavior	p. 91

Abstract

The thesis investigates the question of rationality in consumer and managerial choices in the area of industrial organization. The first essay (Chapter 1) reviews the literature on behavioral consumer choice, its policy implications and its limitations. In the past few decades a great deal of evidence has been collected by psychologists and economists confirming that consumers' choices often do not reflect their normative preferences. These studies demonstrate that consumers' choices are affected by, among others, analytic errors, misperceptions, passivity, inattention and impatience. The findings raise important issues such as how to identify consumers' true preferences and what the responsibility of the policy-maker is when mistakes by consumers during the decision making process can be expected.

In the context of managerial choice the second essay (Chapter 2) investigates the seemingly irrational behavior of acquirers. A major finding of the literature examining the stock price changes driven by merger announcements is that the combined firm gains significantly, the target firm captures the whole stock market surplus and acquirers on average hardly break even. If all empirical evidence is correct the following question arises: Why do firms volunteer as buyers if they cannot expect to gain? The paper argues that previous empirical investigations only partially detected the value created by mergers. An empirical model applied to a large dataset from Compustat with information on mergers and acquisitions finds that using observable corporate characteristics as explanatory variables the model predicts with high probability the value-creating mergers and gives low probability to the value-preserving and -destroying ones. This result suggests that investors should be able to predict profitable mergers already before the announcement is made and in efficient markets this information should be incorporated in the stock prices by the time of the merger announcement. The short window event study performed to test this hypothesis confirms significant abnormal stock price returns preceding the announcement day for the acquirers classified into the higher acquisition probability classes. Hence, acquirers' behavior may be explained without resorting to some sort of irrationality on part of the acquirers.

Departing from the question of rationality Chapter 3 empirically investigates the claim that private equity takeovers, an increasingly important share of the global M&A market, "cheat target shareholders out of a fair price". The paper analyzes the differences in takeover premia paid by corporate and financial investors (predominantly private equity firms) for 226 quoted companies in the UK between 2004 and 2007. Financial investors are documented

to pay only 70 per cent of the premia given by corporate investors, 16 per cent and 23 per cent, respectively. The gap is significant and of approximately the same magnitude when several characteristics of the targets and differences in the hypothesized sources of takeover premia are controlled for: agency costs, bargaining power and synergy gains. Agency costs and bargaining considerations motivated by differences in the ownership structures of targets and bidders are an important source of takeover premia as suggested by the principal-agent literature. Moreover, financial synergies, but not real synergies, account for a significant part of the premia as well. When splitting the sample along the type of investors, most of the model's coefficients change signs between the two regressions suggesting important differences in the merger pricing behavior of industrial corporations and private equity firms. The results of the analysis support the argument that private equity firms pay significantly less for their targets and this does not seem to be a result of lower value generated by this type of takeovers. However, it can not be safely rejected that the observed pattern is not exclusive to private equity firms but holds for private firms in general.

The last essay (Chapter 4) develops a behavioral model capturing the interactions of time inconsistent preferences and deadlines and applies it to explain consumers' behavior regarding rebates in the US. The model combines partial naivete about future self-control problems and the sunk-cost effect (regret) assuming that agents deviating from their past choices suffer a certain emotional disutility from having brought a bad decision in the past. In the context of the application the model explains why in a multi period setting a large number of consumers respond to rebate offers intending to redeem the rebate and then fail to provide the necessary effort when it comes to collect their money. Moreover, the model explains the empirical finding that consumer failure to accomplish a task planned in the past (e.g. redeeming the rebate) is more likely when the deadline of completion is longer. This prediction is supported by experimental studies on various forms of procrastination and by field and experimental evidence on mail-in-rebates. The paper reviews a number of areas for which the theory may have important implications.

Chapter 1

Behavioral consumer choice

Abstract

This paper reviews the literature on behavioral consumer choice, its policy implications and its limitations. In the past few decades a great deal of evidence has been collected by psychologists and economists confirming that consumers' choices often do not reflect their normative preferences. These studies demonstrate that consumers' choices are affected by, among others, analytic errors, misperceptions, passivity, inattention and impatience. The findings raise important issues such as how to identify consumers' true preferences and what the responsibility of the policy-maker is when mistakes by consumers during the decision making process can be expected. From collecting evidence and building theoretical models economists working on this literature have now moved to the analysis of these normative questions. Of course, such issues are not likely to be settled easily and indeed have divided the economic arena and inspired an exciting intellectual debate.

1 Introduction

”An article can have no value unless it has utility. No one will give anything for an article unless it yield him satisfaction. Doubtless, people are sometimes foolish, and buy things, as children do, to please a moment’s fancy; but at least they think at the moment that there is a wish to be gratified. Doubtless, too, people often buy things which, though yielding pleasure for the moment, or postponing pain, are in the end harmful. But here ... we must accept the consumer as the final judge. The fact that he is willing to give up something in order to procure an article proves once for all that for him it has utility, it fills a want”.

Frank Taussig, 1912

In his presidential address opening the session of the American Economics Association (2006), Daniel McFadden has challenged the long held view that we have to accept the observed consumer behavior as the ultimate revelation of consumer preferences. Why is this question so important? Revealed preference theorists emphasized that consumers’ preferences can be inferred from demand behavior (Samuelson, 1938, 1948; Little, 1949; Houthakker, 1950; Afriat, 1967). In other words, observing which choice among the feasible alternatives a person makes reveals her preferences. Understanding preferences means that we know what an individual wants to achieve and how she would behave under different circumstances. But if observed consumer choices deviate from those implied by consumers’ normative preferences than our understanding of consumers’ behavior may be mistaken as well as our predictions for her future actions. Moreover, the policy recommendations that we make based upon the revealed preferences may not serve the consumers’ best interests. Clearly the question of whether we understand consumers and are able to promote their well-being is a highly important one for economics.

According to Parkin (1990), a rational choice is the best possible action from the point of view of the person making the choice, given that person’s preferences and given the information available at the time the decision is made. If there is a discrepancy between revealed and normative preferences then the choices made by the consumers are suboptimal, i.e. irrational

in the sense of the above definition.¹ An increasing number of studies document that human behavior is jointly determined by both normative preferences and other factors such as analytic errors, myopic impulses, inattention and passivity. Individuals do make inferior choices, choices that they would change if they had complete information, unlimited capabilities and no lack of willpower. Studies surveying related work in different fields include Rabin (1998) summarizing the psychological findings, Beshears et al. (2007) reviewing the literature on savings and investment decisions, and Bernheim and Rangel (2006) collecting evidence in the domain of public economics.

This paper gives an account of the development and the current state of the behavioral consumer choice literature. From collecting evidence and building theoretical models economists working on this literature have now moved to the analysis of important normative questions such as the responsibility of the policy-maker when mistakes by consumers during the decision making process can be expected. Unsurprisingly, these issues have got into the focus of a heated debate between behavioral economists suggesting some sort of paternalism on the one side, and the representatives of strict liberalism on the other.

The remainder of the paper is organized as follows. Section 2 reviews studies that suggest deviations from what the revealed preference theory assumes, i.e. individual choice behavior reflects underlying stable preferences. Section 3 brings examples on how profit-maximizing firms react to such consumer behavior. The next section reviews the remedies offered by behavioral economists to the issue of how to identify consumers' normative objectives if they are hidden beneath mistakes and misperception. Section 5 is concerned with the policy implications of this literature, while Section 6 summarizes the theory's limitations and relevance bringing some counter-examples and opposing opinions. The last section concludes.

2 Consumer choice deviating from standard preferences

One of the foundational assumptions of neoclassical economics is that individual choice behavior reflects underlying stable preferences. However, a number of field studies and experiments

¹The statement is not quite correct if revealed preferences deviate from normative preferences due to limited information, but this issue will be ignored throughout this paper.

demonstrate that consumers' preferences are imprecise or "fuzzy", are context specific and subject to framing, depend on the individuals' reference points, choice sets, default options, willpower and a number of other factors. But even if preferences were clear and complete, they would often not be reflected in consumer choices because consumers make analytic errors and judgmental mistakes, their estimates of the products that they purchase are imperfect, and the frustration due to their choice overload may result in their sticking to the status quo.

2.1 Imprecise, unstable preferences and context-specificity

Except for a few field studies the evidence on imprecise and unstable preferences come from experiments. For example, Ariely et al. (2005) show that willingness to pay for a product (in their experiments, mugs and chocolate) depend on the prices that consumers face. Specifically, consumers' valuation of a product is much higher if the price distribution is skewed to the right than if it is skewed to the left. Although they emphasize that this is not the result of participants' attempt to infer the quality of the goods from the price distribution (they are valuing a familiar chocolate bar), there is ample evidence for the latter as well. Indeed, consumers often judge lower priced items to be of lower quality (e.g. Rao and Monroe, 1989), and their beliefs may even alter the actual efficacy of the product that they purchase. Shiv et al. (2005) demonstrate in a series of experiments, that participants who purchase a product (energy drink) at a discounted price may derive less benefit from it (solve less puzzles) than those who buy the product at the regular price.

Anchoring is another example of how consumers may be influenced by context and framing that should be irrelevant to choice. A well-known experiment by Ariely, Loewenstein, and Prelec (2003) shows that subjects' valuations of products are affected by arbitrary "anchors" such as a person's social security number. In this study students were asked whether they would buy a product (e.g. a bottle of wine) for a price obtained from the last two digits of their social security number. After this response they were also asked to state their maximum willingness to pay for the same product. The results showed that students with high social security number had much higher valuation of the same product than those with lower social security numbers.

In a recent field study Ahlee and Malmendier (2005) show using data on E-bay auctions that in 51 per cent of all purchases consumers pay a higher price as the result of the auction than the buy-it-now price for the same product posted on the same web-site. Although this may not be evidence for imprecise preferences as it may only reflect systematic mistakes on part of the consumers, the authors also show that consumers' willingness to pay is affected by irrelevant information. For example, the final price bid for the products is significantly higher if the retailer mentions that the manufacturer's price is higher than the retailer's suggesting that the consumer is concluding a good deal.

2.2 Reference dependence and loss aversion

Another robust finding of experimental studies related to preferences is reference dependence. People have been repeatedly shown to value opportunities in relation to certain reference points, and not only consider absolute measures in their decision making. Moreover, changes of reference points often lead to preference reversals (Tversky and Kahneman, 1991). A behavioral pattern frequently observed in the context of reference dependence is loss aversion. Decision-makers dislike losses more than appreciate equal sized gains relative to their reference points. A well-known study demonstrating the latter behavior on the housing market of Boston is that of Genesove and Mayer (2001). Using data on almost 6 thousand property listings between 1990 and 1997, the authors show that property owners, whose original purchase price exceeds that of their expected selling price, set much higher asking prices than other sellers in the area, and as a result, they typically remain longer on and eventually withdraw from the market.

In an interesting field experiment Camerer et al. (1997) study the behavior of New York City cabdrivers to test intertemporal substitution on the labor market. Contrary to what dynamic models of labor supply would predict, i.e. workers work longer hours in high-wage and shorter hours in low wage periods, they find that NYC cabdrivers quit early on busy days (due to for example bad weather or subway breakdown) and drive longer hours on low wage days. The authors conclude that the evidence is consistent with the hypothesis that drivers are averse of falling below a target income.

Although field experiments on reference dependence in the consumer choice literature have

been limited, theoretical models examining the implications of consumers' loss aversion in the marketplace have been developed by a number of authors (e.g. Koszegi and Rabin, 2006, and Heidhues and Koszegi, 2005).

2.3 Choice overload, default options and status-quo effect

If consumers trust their own tastes and aim at maximizing their utilities they should favor more choices to less and welcome market opportunities. Neoclassical economics relies on this concept of the consumer, and argues for decentralized, competitive markets. In contrast, an increasing number of economists claim that more choice is not always better. Loewenstein (1999) argues that contrarily to what standard economic theory would suggest, consumers may not benefit from too many choices but become intimidated by the overwhelming options that they face day by day. Expanded choice not only provides the opportunity to improve one's situation but entails certain costs: time costs that the choice overload results in, the cost of making a bad decision, and psychic costs, such as anxiety about making decisions, especially when consumers lack the expertise required to evaluate the different options. This argument is supported by the experiment by Sethi-Iyengar and Lepper (2000) which found that consumers consider a progressively shrinking number of options when the choice set is significantly expanded.

Another indication for how choice can be intimidating, i.e. consumers may not trust their own capabilities to bring good decisions, is that default options have an enormous impact on consumers' decisions. This behavioral bias is also called status quo bias and have a number of sources (Camerer, 2003). Loss aversion, as explained above, can easily accommodate this behavior. A further source is the omission/commission bias explained by Ritoy and Baron (2000). That is, people tend to care much more about errors that are a result of their own actions, than about errors that are a result of their inaction. Finally, procrastination has also been mentioned as a source of status quo bias. People repeatedly delay beneficial actions due to a mistaken belief that they will take them in the future.

Madrian and Shea (2001) and Choi et al. (2004, 2003) show how defaults affect 401k participation, saving rates and asset allocation. They find that when employees are enrolled automatically in a 401k plan, very few of them opt out; while if they are not enrolled automat-

ically less than half of them enroll on their own during their first year of employment.

Further example to this passive behavior is the endowment effect in which consumers show a reluctance to give up any position that they hold. An excellent demonstration of this effect is the well-known cup experiment of Knetsch (1989) in which a random assignment of coffee cups in a class, followed by an opportunity to trade, produced a large gap between willingness to pay and willingness to accept. While the subjects who possess the cups are reluctant to part with them, the ones without the cups are unwilling to give up their money to acquire them. Several other studies confirmed this behavioral pattern (e.g. Kahneman, Knetsch, Thaler, 1990, and McFadden, 2006).

2.4 Regret, sunk cost effect and taste for consistency

People dislike losing, as discussed above, but they have been argued to feel even worse about it when they think that they could have done better if they had made a different decision (Sugden, 1985). Moreover, regret does not need to arise ex post. Kraheimer and Stone (2005) show in a dynamic context, that in light with the empirical evidence, regret may lead to excess conservatism or a tendency to make up for missed opportunities.

The feelings of regret are exacerbated by what Fischhoff (1975) refers to as "hindsight bias" - the tendency to view outcomes, after the fact, as having been more predictable than they actually were when the decision was made. "While people tend to avoid decisions with a high potential for regret, some such decisions are unavoidable, and feelings of regret about decisions that turned out badly remain an important source of personal misery and thus an added potential cost of expanding choice."

The other side of the same token is that people tend to stick to past decisions in order to avoid regret. Various researchers have found that costs that have been paid (or "sunk") in the past affect people's behavior in ways that are not consistent with economic theory (Thaler 1980; Thaler and Johnson 1990; Shefrin 2000). An excellent description of this behavior is Thaler's well-known family example:

"A family pays 40 dollars 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.” Thaler (1980, p. 11).

If they were thinking rationally, at the point of deciding whether to attend the game the family should ignore how much they paid for the tickets and should only take into account the future costs and benefits of attending the game. In a series of experiments, Arkes and Blumer (1985) demonstrated that monetary sunk costs matter to individuals when making decisions. In hypothetical situations and in situations involving actual money, subjects in these experiments made decisions not on the basis of expected marginal costs and benefits, but on the basis of expenses already incurred.

Eyster (2002) develops a model in which he interprets this evidence as a ”taste for consistency”, i.e. consumers may bring suboptimal choices in the present in order to justify their past decisions. The experiments of Charness and Levin (2004) support this interpretation by showing that people tend to repeat their past choices significantly more often if those choices were the result of their own decisions than if they were required to make that choice.

In a related study Prendergast and Stole (1996) develop a signaling model in which agents send a signal to their future selves (or some third party) about the quality of their decisions. In this model, suboptimal choices may arise from the need to signal one’s talent not only to third parties but future selves as well. A feature of the model is that agents underreact new information if it prompts a deviation from past decisions. Similarly, Yariv (2005) builds on psychological experiments documenting that individuals tend to interpret new evidence in ways that confirm their prior beliefs. When the new evidence strongly contradicts prior beliefs individuals may even change their beliefs to accommodate their past actions. In Yariv’s model the agents can choose their beliefs directly in order to make themselves feel better about past decisions. Current actions are then taken in accordance with the resulting beliefs and so reflect past judgments even if they are mistaken. The argument developed here resembles closely the cognitive dissonance model of Akerlof and Dickens (1982).

2.5 Time-inconsistent choices

Another mechanism, somewhat distinct from those above, leading to suboptimal choices is time inconsistency. Standard economic models assume that individuals discount the future exponentially, i.e. their choices are consistent and their discount rate does not change over time. For example, if today an agent prefers 1050 dollars in two months to 1000 dollars in one month, after 30 days she should still prefer the 1050 dollars due 1 month later to the immediate 1000 dollars. In the past few decades, behavioral economists have challenged the time consistency hypothesis implied by exponential discounting. Several experiments have been designed to test whether individuals behave according to the exponential model. In the pioneering study of Thaler (1981) students were asked to compare immediate and delayed payoffs to elicit annual discount rates. Thaler (1981) finds that the annualized discount rate for a 3 month delay is two to five times larger than the annualized discount rate for a 3 year horizon. This form of discounting indicates that agents prefer a larger, later reward over a smaller, earlier one as long as the rewards are sufficiently distant in time; however, as both rewards get closer in time, the agent may choose the smaller, earlier reward.

A functional form that captures this feature is the quasi-hyperbolic one (Strotz, 1956; Phelps and Pollack, 1968), which is equal to 1 for $t = 0$ and to $\beta\delta$ for $t > 0$, with $\beta \leq 1$. The present value of a flow of future utilities $u(c_t)_{t \geq 0}$ is given by:

$$u_0 + \beta \sum_{t=1}^T \delta^t u_t$$

This functional form captures the regularities found by experimental studies (e.g. Kirby and Herrnstein, 1995; Bension et al., 1989) namely time inconsistency with larger short run than long run discount rate. Indeed, the discount factor between today and tomorrow is $\beta\delta$, while the discount factor in any two periods in the future is $\delta \geq \beta\delta$. The short-run discount factor, β , reflects the degree of self-control problems. The agent may know about his future self-control problem depending on his level of sophistication (O'Donoghue and Rabin, 1999). Sophisticated agents know their degree of self-control problems in the future, while the naives (or partially naives) overestimate their future self-control (naivete is a form of bounded rationality about future individual preferences).

Studies in diverse fields such as investments and savings (Madrian and Shea, 2001) and labor market (DellaVigna and Paserman, 2005 and Drago, 2006) highlight behavioral regularities that support the idea that people have hyperbolic preferences and are naive about their future self-control problems. In the field of consumer choice, DellaVigna and Malmendier (2006) analyze a large dataset on gym-users who choose between three types of contract. Gym-users can choose to pay on a per visit basis, or they can sign up to monthly or annual membership contracts. It is assumed that members incur a short-run disutility with each exercise visit and obtain a long-term reward in the form of better health. The hyperbolic discounting model implies that consumers will prefer a higher visit rate *ex ante* than they will actually choose *ex post*. The empirical finding is that the usage rates of members on the monthly and annual plans are so low that these consumers would be better off using the per visit plan. The evidence provided shows that this is the case for the majority of members in the sample. The authors argue that the delays in contract cancellation can only be explained, within the hyperbolic framework, by naivete about future self-control problems that lead gym-users to systematically overestimate their usage rate in the future.

Another good example of suboptimal choices is the observed consumer behavior in the context of rebate promotions. Ample empirical and anecdotal evidence show that a large number of consumers respond to mail-in-rebate promotions fully intending to get back part of the purchase price but fail to exert the necessary effort to collect their money subsequently (e.g. Silk, 2005). On average 60 percent of consumers fail to redeem the rebate. As explained in Drago and Kadar (2006) this evidence can hardly be reconciled with a standard rational model of consumer behavior. It is very likely that also in this case consumer behavior involves time inconsistency and naivete in future self-control problems.

3 Profit-maximizing firms reaction to deviations from standard preferences

A few studies examined how rational firms respond to deviations from standard preferences and biases in decision making. DellaVigna and Malmendier (2004) analyze the optimal contracts

designed by rational firms when consumers have time-inconsistent preferences and are partially nave. They distinguish between two types of products: investment goods, where first a costly effort needs to be made to gain benefits at a later point in time; and leisure goods, where the benefits precede the effort. They demonstrate that the profit-maximizing contract, by targeting consumers' misperception about future consumption, will price investment goods below and leisure goods above marginal cost. Sophisticated agents benefit from the commitment device that this price structure implicitly offers. The misperceptions of the nave agents, on the other hand, are exploited by these contracts. The authors claim that the predictions of their model are matched by empirical evidence in a number of industries: e.g. health clubs, credit cards and mobile phones. The findings of their model are also supported by the empirical study of magazine subscriptions by Oster and Scott Morton (2005). Magazines are classified into investment and leisure types and it is shown that for the former type subscription prices are higher than newsstand prices, *ceteris paribus*. The authors interpret this as evidence that publishers react to the present biased preferences of consumers when setting their prices.

Drago and Kadar (2006) suggest that rebates - previously held to be a simple price-discrimination device - may be a profit-maximizing tool for firms facing nave consumers with time-inconsistent preferences. The argument is based on the evidence that rebates are a promotional technique much favored by consumers but redemption rates are unexplainably low. The authors show that based on their characteristics rebates belong to the investment group category and their utilization rates by naives will be suboptimal.

In a different domain, Gabaix and Laibson (2006) show that if consumers are boundedly rational, i.e. they have imperfect estimates of the value of the products they purchase, firms in equilibrium take advantage of this fact by increasing the complexity of their products as well as their mark-ups. That is, firms have an incentive to mask information by inefficiently increasing the complexity of their products. The model has important implications to some results of industrial organization previously perceived to be unquestionable. For example, stronger competition increases complexity and often even increase mark-ups. This is exactly the opposite of what theories of perfect competition conclude. The model's predictions in case of myopic (nave) consumers - i.e. consumers who do not anticipate the future consequences of their current purchase - is that firms will price add-ons, such as a mini-bar in the hotel

room or a tank refill in case of rental cars, at monopoly prices. In the same area Ellison and Ellison (2005a, 2005b) similarly show that increased competition induces firms to engage in obfuscation - practices that hinder consumers to observe the true value of their products - and retain their mark-ups. The example they use to demonstrate this process is the observed changes in firms' behavior as a result of the diffusion of the Internet and the appearance of the price search engines. They argue that the easy price search made demand for certain products extremely price sensitive to which rational firms responded by putting back some friction into the market. Obfuscation can take different forms: from complicating the description of the product to more refined strategies, such as offering inefficiently low quality products at very low prices to attract consumers and then talk them into more pricey, higher quality versions. The discussion gives some ideas about why we should not expect bounded rationality to be transitory: firms make more profits on boundedly rational consumers.

Finally, Heidhues and Koszegi (2005) develop a model that analyzes the impact of consumer loss aversion on the pricing decision of a profit-maximizing monopoly. One of their findings is that firms can increase their revenues by not exposing the consumers to small price changes, i.e. their model gives an explanation for sticky prices.

4 Remedies to the problem of identifying consumers' preferences

The findings of the behavioral consumer choice literature raise an important question, i.e. how to conduct welfare analysis when the objective function inferred from consumers' choices do not represent a preference ordering or incorporate misperceptions. Beshears et al. (2007) discuss six methods that can help policy makers to infer normative preferences when revealed preferences likely deviate from them. Structural estimation would start from a behavioral model that incorporates the reasons for the deviation between revealed and normative preferences. When defaults are expected to have excessive effect on consumer behavior, an analysis of active decisions can be used to infer the underlying preferences. Similarly, long-run choices are argued to reflect better the consumers' interests than short-run choices. Aggregate revealed preferences

are argued to be closer to normative optimality than noisy, error-prone individual decisions. Finally, reported preferences and informed opinions can be suggestive of the deviation between actual behavior and the underlying preferences. Except for the first one, these suggestions are based on observations that identify situations in which consumers are less likely to make errors.

Ariely et al. (2005) also point out that their findings, i.e. consumers' willingness to pay depends on the price distribution, pose a problem for measurement of preferences and welfare implications. Indeed, procedures to measure 'true preferences' strongly rely on the assumption that preferences are independent of the price distribution. Ariely et al. (2005) discuss two variants of the basic method to elicit preferences and confirm that they decrease the preference dependence on the price distribution.

Koszegi and Rabin (2007) argue that revealed mistakes should be incorporated into economic models and policy prescriptions. Economists can identify the intentions of consumers and reconstruct the underlying preferences from consumer behavior and the knowledge of mistakes that they systematically commit. One major problem with this approach is that often mistakes are identified but the source of these mistakes are not. Policy recommendations however may strongly depend on the source of the mistake.

Benhabib and Bisin (2007) emphasize the need for structural models which have predictions for both processes and choices, to predict consumer behavior. Although they argue that standard economic theory has successfully rationalized behavioral findings by weakening its axioms and extending its choice sets (e.g. Gul and Pesendorfer, 2001) this has come at the expense of predictive power. The information gained from structural models that focus on process besides choice could help the theory regain its explanatory power.

5 Policy implications of behavioral consumer choice

The above methods help us better understand consumer behavior and the underlying preferences, but do not guide economists whether and how to interfere in situations where consumers are likely not able to make choices that serve their best interests.

On this issue, Thaler and Sustain (2003) advocate what they call the "libertarian paternalism". Under paternalism they mean "policies with the goal of influencing the choice of affected

parties in a way that will make those parties better off". The libertarian approach requires that no choice be forbidden and policy-makers do not impose their own values on consumers. Thaler and Sustain (2003) argue that in certain cases paternalism is inevitable. For example, as discussed above, default options have enormous impact on consumer choice, and thus planners have important responsibility in designing these options. Choi et al. (2003) analyze optimal defaults in the area of savings (e.g. 401K participation) where defaults have been found to have an important role in consumer choice. They show that if people have heterogeneous optimal savings rates than it may be optimal to set defaults that are far from the mean optimal savings rates. One reason for this is that people may overcome procrastination if the default option is farther from the savings rate that they try to achieve.

Going beyond the inevitable paternalism Thaler and Sustain (2003) also suggest that decision-makers who control factors that may affect choices - even if rationally they should not - should act in accordance with the maximization of the overall welfare of the population affected (i.e. based on a cost-benefit analysis). Assessing a body of negative evidence on consumer rationality, Kahneman (1994) arrives at a similar conclusion: "the observed deficiencies suggest the outline of a case in favor of some paternalistic interventions, when it is plausible that the state knows more about an individual's future tastes than the individual knows presently."

Camerer et al. (2003) have introduced the term "asymmetric paternalism" which is used for the type of regulation that creates large benefits for irrational agents and cause little or no harm for the fully rational ones. The approach admittedly relies on the assumption that bounds on rationality can be empirically identified and measured, and asymmetrically paternalistic policies can be assessed. The authors list a number of existing regulations that can be interpreted to be of this type, e.g. defaults, provision or reframing of information, cooling-off periods and limitation of consumer choices. Defaults have been discussed extensively in the literature. To provision of information belong, for example, regulations that legally require customers to read and sign a document that gives complete information about the product that they purchase. Cooling-off periods give chance to the customers to review and change their decisions in a limited period of time. All door-to-door sales are subject to this rule in the United States. Examples for consumer choice limitation are deadlines that are supposed (and have been

empirically verified) to help consumers to overcome procrastination. For example, Individual Retirement Account (IRA) contribution for a year must be made by 15 April of the next year.

An analysis similar to the study on optimal defaults is that of O'Donoghue and Rabin (2003) who illustrate optimal paternalism through modeling sin taxes for unhealthy goods. The framework that they suggest starts from assumptions on the distribution of rational and boundedly rational agents and analyzes which of the available policy instruments achieve the most efficient outcomes. The challenge of this approach - expressed by other authors as well - is to identify all types of mistakes that consumers potentially commit.

Koszegi and Rabin (2007) point at economists' reluctance to accept that consumers make mistakes, because that would mean they lose the revealed preference approach to infer welfare measures from behavior and they would have to (paternalistically) impose a welfare measure themselves. They argue that the normative conclusions of their approach, which derives preferences from the consumers' own behavior and intentions (but not necessarily from their choices), respect consumers' values and there is no need for the policymaker to impose her own welfare criterion.

Bernheim and Rangel (2006) discuss a number of possibilities to replace revealed preference with other normative principles. The approach represented by Sugden (2004) is to separate positive theories of consumer choice from normative welfare analysis, where the notion of welfare is based on opportunities. Instead of analyzing the choices that consumers make, these theories are concerned with the opportunity set from which individuals are free to choose. The underlying ideology is that justice requires the equality of opportunities, and whether individuals satisfy their preferences should be left for them to decide. The advantage of this approach is that it is viable even without information on consumer choice. However, as Bernheim and Rangel (2006) suggests, it is probable that many economists would resist such a radical deviation from the standard welfare analysis. A less radical version of the approach that separates positive and normative theories can be applied when consumer choices are believed to systematically deviate from normative preferences. In this case, consumers can be endowed with well-behaved lifetime preferences, and based on the observed anomalies a choice process - that may lead to suboptimal decisions - can be formulated. Description of consumer behavior would rely on the preferences and the choice process, while welfare analysis would consider preferences

alone. Another approach, motivated by time-inconsistency, allows for different consumer preferences at different points in time and analogously to the aggregation over many individuals aggregates over multiple selves. The analysis is then completed using some appropriate welfare criterion (e.g. Pareto optimality) or an aggregate welfare function to derive optimal policies.

Examples where the responsibility of policy-makers may be especially important are health care, social security and education, where consumers may not be able to judge properly the quality of the services provided and the comparison of complex programs may prove intimidating. McFadden (2006) summarizes the key learning points from a study on the Medicare Part D program, which provides substantial insurance against drug costs. Despite the complexity of the program, the study predicts that consumers will handle satisfactorily the choices offered by the different prescription plans. However, a substantial number of people are likely to opt out of the program by default, contrarily to their own interests. Moreover, a significant fraction of the consumers are expected to pick suboptimal plans and to procrastinate past the deadlines despite their intention to enroll.

In a similar vein Loewenstein (1999) explores the implications of adding more choice to the Social Security system. Under most arrangements workers are covered automatically and they do not have any influence on how their funds are managed. Using the evidence on behavioral consumer choice Loewenstein (1999) argues that substituting individually managed savings accounts for Social Security's existing benefits would result in some cases too conservative (loss aversion), while in others too risky (overconfidence) behavior, and the default option will have a significant effect on the choice outcomes. He concludes that when consumers lack expertise the consequences of new choices can be suboptimal decisions and frustration instead of welfare improvement.

6 Relevance of the literature, limitations and counter-examples

Early critics of the behavioral consumer choice literature argued that laboratory experiments can not replicate "real world" situations (Smith, 1991). Consumers may be more attentive when their choices involve higher stakes. Since then, a number of field experiments have shown that the predictions of the behavioral models hold in the real life as well. Moreover, numerous studies

find that consumers make mistakes in the area of very important, life-time decisions, such as retirement savings and investments (Beshears et al., 2007). The claim against these experiments seems to have been reversed now, as in his review of the book "Advances in behavioral economics (Camerer et al, 2004)" Harrison (2005) warns that field experiments should be complemented by laboratory experiments with the necessary controls that are "typically absent in the wild".

Other economists have argued that market forces, i.e. competition and arbitrage, coupled with evolution and learning will wipe out irrational consumer behavior from the market. Hendricks (2006) reviewing the survey of Ellison (2006) on bounded rationality gives voice to the skepticism over models that rely on consumers with persistent perception biases to explain behavior. Mullainathan and Thaler (2000) summarizes why the four forces mentioned above may not be sufficient to terminate irrational behavior. They argue that markets can only provide incentives to switch from irrational behavior to a rational one, but can not force consumers to do so. Arbitrage - being limited already in financial markets, let alone in other segments of consumer choice - is not strong enough to drive irrational behavior out of the market. Evolutionary arguments do not unambiguously predict the elimination of irrational behavior. In fact, evolutionary models - depending on the initial environment - can easily suggest that it is rational behavior that will be driven out of the market. As for learning, the experimental literature has shown that there can be complete lack of learning even in infinite horizons. Some decisions do not get repeated so often that by learning from their past mistakes consumers could move toward the optimal decision. In summary, Mullainathan and Thaler argue that unbounded rationality can not be defended on a purely theoretical ground and empirical studies are necessary to investigate this issue.

As consumer sovereignty has been a founding pillar of economic theory, direct empirical evidence justifying such regard toward individual choice is very limited. The fact that some have recently appeared suggests that some authors believe that the defense of the rationality assumption is necessary. In an interesting paper Miravete (2003) analyzes choices among optional calling plans in the telephone industry. He shows that individual expectation of future telephone usage plays a critical role in the subscription decisions and that the ex-post utilization support the choices made by consumers. Moreover, consumers whose ex-ante expectation of their future utilization was wrong, quickly change to plans that better suite their interests.

Waldfogel (2002) argues that even if consumers make mistakes their choices are still better than what others would make for them. Based on a survey of about 1,500 gifts and own purchases he concludes that consumers value their own purchases at about 18 per cent more than the gifts received from family and friends. He emphasizes that if close relatives who know the individuals well can not make better choices than the individuals themselves, there is no reason to expect that the state or any planner would do a better job. Therefore, he argues, that paternalism may not be warranted even if consumers make observable mistakes.

A further concern with the behavioral literature is that even if the different types of irrational behavior may persistently exist, building a new model every time a new source of bias awaits explanation is impractical. That is, this literature will not lead to a unified theory of decision making. The fear is that mistakes become a free parameter that can be used to fit any situation. Rabin (2004) points out that simultaneously to this claim, behavioral economists has been criticized for the opposite reason as well: they try to explain a wide range of phenomena repeating the same themes over and over again.

Sen (1990) predicts that it will be difficult to find replacements for the standard assumptions of rational behavior even though their inadequacies are hard to deny. There is little hope of finding an alternative structure that is as simple and usable as the traditional self-interest maximization and consistency of choice. According to Rubinstein (2005), the lack of elegance and generality is the major drawback of behavioral economics models. On the issue of elegance, Rabin (2002) argues that the models trying to explain the behavioral anomalies while retaining the rationality assumption are often significantly more complex than the models presented by behavioral economists. Also, it is reasonable to expect that the early behavioral models will be replaced by clearer and simpler versions as the field evolves. On the generality of these models, behavioral economists intend to generalize the standard models by incorporating features that allow them to explain a wider range of behavior than that consistent with the standard assumptions. Rabin (2002) argues that it is appropriate "to stop treating the classical assumptions as the maintained hypothesis in our analysis, and start treating them as special cases, corresponding to particular parameter values of a generalized model, and then investigating what are the best fits for those parameter values".

On a more theoretical point, Gul and Pesendorfer (2005) point out that the behavioral

patterns attacked by behavioral economics are consistent with many utility functions therefore welfare considerations can not be derived from an agent's behavior. Rubinstein (2005) disagrees with them saying that "behavioral economics makes it clear that one must make assumptions about the relation between the preferences that explain behavior and the preferences used in the welfare criterion". Indeed, the point made by Gul and Pesendorfer (2005) would pose difficulties to economics in general, as all economic theories have to make assumptions to connect preferences to behavior (Koszegi and Rabin, 2007).

The policy implications derived by behavioral economists from the observed evidence, i.e. some form of paternalism is inevitable and required to help consumers achieve their objectives, has also been strongly criticized. Rizzo (2007) attacks "libertarian paternalism" on a number of fronts. First, he finds the standard for better decisions, identified as "what people would do if they had "complete information, unlimited cognitive abilities, and no lack of will-power", overly ambitious. Clearly, behind this claim lies the assertion that we can not elicit the true preferences of consumers once we accept that they make mistakes. Second, he questions whether boundedly rational policy-makers should be trusted to bring better decisions than consumers whose welfare ultimately depend on them. He alleges that mistakes made at a governmental level are less likely to be corrected by individual or social processes than errors made by the consumers. In his view, libertarian paternalism is an oxymoron: libertarian means less government, while paternalism means more government.

7 Conclusion

There is ample evidence on the deviation between preferences revealed by consumer choice and consumers' normative preferences. This survey has reviewed a number of studies documenting that, among others, consumers' preferences are imprecise and unstable, their decisions are affected by irrelevant factors such as past decisions and sunk costs, they have bounded will-power and they increasingly prefer default options as the choice overload intensifies. As a consequence, consumers bring sub-optimal choices that make inferences about their true preferences complicated, if at all possible. Consumers' preferences are hardly dispensable for welfare analysis. Although behavioral economists have commenced providing various methods

to identify consumers' normative preferences even if their choices reflect mistakes and misperceptions, the question whether the policy-maker should interfere to direct consumers toward optimal choices remains.

References

Afriat, S. N. (1967): The Construction of Utility Functions from Expenditure Data, *International Economic Review* 8, 67-77.

Ahlee, H. and U. Malmendier (2005): Do Consumers Know Their Willingness to Pay? Evidence from eBay Auctions, mimeo, University of Stanford.

Akerlof, G. A. and W. T. Dickens (1982): The Economic Consequences of Cognitive Dissonance, *The American Economic Review*, 72, 307-319.

Ariely, D., Loewenstein, G., and D. Prelec (2003): Coherent Arbitrariness : Stable Demand Curves without Stable Preferences, *Quarterly Journal of Economics*, 118, 73-105.

Ariely, D., Koszegi, B., Mazaar, N. and K. Shampan'er (2005): Price-Sensitive Preferences, mimeo.

Arkes, H. R. and C. Blumer (1985): The psychology of sunk cost, *Organizational Behavior and Human Decision Processes*, 35, 124-140.

Benhabib, J. and A. Bisin (2007): Choice and Process: Theory Ahead of Measurement, mimeo.

Benzion, U., Rapoport, A. and J. Yagil (1989): Discount Rates Inferred from Decisions: an Experimental Study, *Management Science*, 35, 270-284.

Bernheim, D and A. Rangel (2006): Behavioral Welfare Economics, mimeo.

Beshears, J., Choi, J. J., Laibson, D. and B. C. Madrian (2007): How are Preferences Revealed, mimeo.

Camerer, C., Babcock, L., Loewenstein, G. and R. Thaler (1997): Labor Supply of New York City Cabdrivers: One Day at a Time, *Quarterly Journal of Economics*, 112, 407-441.

Camerer, C. F., Loewenstein, G. and M. Rabin (2004): *Advances in Behavioral Economics*, Princeton University Press, Princeton.

Camerer, C., Issacharoff, S., Loewenstein, G., O'Donoghue, T. and M. Rabin (2003): Reg-

ulation for Conservatives: Behavioral Economics and the Case for 'Asymmetric Paternalism', University of Pennsylvania Law Review, 151, 1211-1254.

Charness, G. and D. Levin (2004): When Optimal Choices Feel Wrong: A Study of Bayesian Updating, Complexity and Affect, Available at SSRN: <http://ssrn.com/abstract=437866>

Choi, J., Laibson, D., Madrian, B. and A. Metrick (2004): For Better or For Worse: Default Effects and 401(k) Savings Behavior, The Economics of Aging, University of Chicago Press, 81-121.

Choi, J., Laibson, D., Madrian, B. and A. Metrick (2003): Optimal Defaults, American Economic Review Papers and Proceedings, 180-185.

DellaVigna, S. and U. Malmendier (2004): "Contract Design and Self-Control: Theory and Evidence," Quarterly Journal of Economics, 119 (2), 353-402.

DellaVigna, S. and U. Malmendier (2006): "Paying Not to Go to the Gym," American Economic Review, 96, 694-719.

DellaVigna, S. and D. Paserman (2005): Job Search and Impatience, Journal of Labor Economics, 23, 527-588.

Drago, F. (2006): Career Consequences of Hyperbolic Time Preferences, IZA Discussion Papers, No. 2113.

Drago, F. and D.Kadar (2006): Rebate or Bait? A Model of Regret and Time Inconsistency in Consumer Behavior, CEPR Discussion Paper, No. 5757.

Ellison, G. (2006): Bounded Rationality in Industrial Organization, Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Cambridge University Press.

Ellison, G. and S. F. Ellison (2005a): Lessons About Markets from the Internet, Journal of Economic Perspectives, 19, 139-158, 2005

Ellison, G. and S. F. Ellison (2005b): Search, Obfuscation, and Price Elasticities on the Internet, mimeo.

Eyster, E. (2002): Rationalizing the Past: A Taste for Consistency, Mimeo, Nuffield College, Oxford.

Fischhoff, B. (1975): Hindsight Foresight: The Effect of Outcome Knowledge on Judgement Under Uncertainty, Journal of Experimental Psychology: Human Perception and Performance,

1, 288-299.

Gabaix, X. and David Laibson (2006): Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets, *Quarterly Journal of Economics*, 121, 505-540.

Genesove, D. and C. Mayer (2001): Loss Aversion and Seller Behaviour: Evidence from the Housing Market, *CEPR Discussion Papers*, No. 2813.

Gul, F. and W. Pesendorfer (2001): Temptation and Self-Control, *Econometrica*, 69, 1403-1436.

Gul, F. and W. Pesendorfer (2005): The Idea of Neuroeconomics, a lecture in the Nemmers conference.

Harrison, G.W. (2005): Book Review: *Advances in Behavioral Economics*, *Journal of Economic Psychology*, 25, 793-795.

Heidhues, P. and B. Koszegi (2005): The Impact of Consumer Loss Aversion on Pricing, *CEPR Discussion Paper*, No. 4849.

Hendricks, K. (2006): Price Discrimination and Irrational Consumers, mimeo.

Houthakker, H. S. (1950): Revealed Preference and the Utility Function, *Economica*, 17, 159-174.

Iyengar, S. S. and M. R. Lepper (2000): When Choice is Demotivating: Can One Desire Too Much of a Good Thing?, *Journal of Personality and Social Psychology*, 79, 995-1006.

Kahneman, D. (1994): New Challenges to the Rationality Assumption, *Journal of Institutional and Theoretical Economics*, 150, 18-36.

Kahneman, D. Knetch, J. and R. Thaler (1990): Experimental Tests of the Endowment Effect and the Coase Theorem, *Journal of Political Economy*, 98, 1325-1348.

Kahneman, D. and A. Tversky (2000): *Choices, Values and Frames*, Cambridge University Press.

Kirby, K. N. and R. J. Herrnstein (1995): Preference Reversals Due to Myopic Discounting of Delayed Reward, *Psychological Science*, 6, 83-89.

Knetsch, J. (1989): The Endowment Effect and Evidence of Nonreversible Indifference Curves, *American Economic Review*, 79, 1277-1284.

Koszegi, B. and M. Rabin (2007): Revealed Mistakes and Revealed Preferences, *Methodologies of Modern Economics*, forthcoming.

Koszegi, B. and M. Rabin (2006): A Model of Reference-Dependent Preferences, *Quarterly Journal of Economics*, 121, 1133-1166.

Krahmer, D. and R. Stone (2005): Regret in Dynamic Decision Problems, mimeo.

Little, I. M. D. (1949): A Reformulation of the Theory of Consumer's Behaviour, *Oxford Economic Papers*, 1, 90-99.

Loewenstein, G. (1999): Is More Choice Always Better?, *National Academy of Social Insurance, Social Security Brief*, 7.

Madrian, B. and D. Shea (2001): The Power of Suggestion: Inertia in 401(k) participation and Savings Behavior, *Quarterly Journal of Economics*, 116, 1149-1187.

McFadden, D. (2006): Free Markets and Fettered Consumers, AEA Presidential Address.

Miravete, E. (2003): Choosing the Wrong Calling Plan? Ignorance and Learning, *American Economic Review*, 93, 297-310.

Mullainathan, S. and R. H. Thaler (2000): Behavioral Economics, NBER Working Paper, No. 7948.

O'Donoghue, T. D. and M. Rabin (1999): Doing It Now or Later, *American Economic Review*, 89, 103-124.

O'Donoghue, T. and M. Rabin (2003): Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes, *American Economic Review*, 93, 186-191.

Oster, S. M. and F. M. Scott Morton (2005): Behavioral Biases Meet the Market: The Case of Magazine Subscription Prices, *Advances in Economic Analysis and Policy*, 5. Parkin, M. (1990): *Economics*. Addison-Wesley, Reading. Phelps, E. S. and R. A. Pollak (1968): On Second-Best National Saving and Game-Equilibrium Growth, *Review of Economic Studies*, 35, 185-199.

Prendergast, C. and L. Stole (1996). "Impetuous youngsters and jaded old timers: acquiring a reputation for learning." *Journal of Political Economy*, 104(6): 1105-1134.

Rabin, M. (1988): Psychology and Economics, *Journal of Economic Literature*, 36, 11-46.

Rabin, M. (2004): A Perspective on Psychology and Economics, UC Berkeley Working Paper, No. E02-313.

Rao, A. R. and K. B. Monroe (1989): The Effect of Price, Brand Name, and Store Name on Buyers' Perceptions of Product Quality: An Integrative Review, *Journal of Marketing*

Research, 26, 351-57.

Ritoy, I. and J. Baron (2000): Reluctance to Vaccinate: Omission Bias and Ambiguity, Behavioral Law and Economics, 168, 168-172.

Rizzo, M. (2007): Should Policies Nudge People to Make Certain Choices?, http://econlog.econlib.org/archives/2007/05/rizzo_vs_thaler

Rubinstein, A. (2006): Comments on Behavioural Economics, Advances in Economic Theory, Cambridge University Press, 246-254.

Samuelson, P. A. (1938): A Note on the Pure Theory of Consumer's Behavior, *Economica*, 5, 61-71.

Samuelson, P. A. (1948): Consumption Theory in Terms of Revealed Preference, *Economica*, 15, 243-253.

Shefrin, H. (2000). Recent Development in Behavioral Finance, *Journal of Private Portfolio Management*, 107-17.

Sen, A. (1990): Rational Behavior, in Eatwell, John, Milgate, Murray, and Newman Peter, *Utility and Probability* (New York: W. W. Norton & Company), 1998-2016.

Shiv, B., Carmon, Z. and D. Ariely (2005): Placebo Effects of Marketing Actions: Consumers May Get What They Pay For, *Journal of Marketing Research*, 42, 383-393.

Silk, T. G. (2005): Why Do We Buy but Fail to Redeem? The Influence of Pre-Purchase, Post-Purchase, and Repeat-Purchase Factors On Consumers' Subjective and Objective Probability of Redeeming Mail-in Rebates, Mimeo, University of Florida.

Smith, V. L. (1991): Review: Rational Choice: The Contrast Between Economics and Psychology, *Journal of Political Economy*, 99, 877-897.

Strotz, R. H. (1956): Myopia and Inconsistency in Dynamic Utility Maximization, *Review of Economic Studies*, 23, 165-180.

Sugden, R. (1985): Regret, Recrimination and Rationality, *Theory and Decision*, 19, 77-99.

Sugden, R. (2004): The Opportunity Criterion: Consumer Sovereignty Without the Assumption of Coherent Preferences, *American Economic Review*, 94, 1014-1033.

Taussig, F. (1912): *Principles of Economics*, New York: Macmillan.

Thaler, R. H. (1980): Toward a Positive Theory of Consumer Choice, *Journal of Economic Behavior and Organization* 1, 39-60.

Thaler, R. H. (1981): Some Empirical Evidence on Dynamic Inconsistency, *Economics Letters*, 8, 201-207.

Thaler, R. and Johnson, E. J. (1990): Gambling With the House Money and Trying To Break Even: The Effects of Prior Outcomes on Risky Choice, *Management Science*, 36, 643-660.

Thaler, R. and C. Sustain (2003): Libertarian Paternalism, *American Economic Review*, 93, 175-179.

Tversky, A. and D. Kahneman (1991): Loss Aversion in Riskless Choice: A Reference-Dependent Model, *Quarterly Journal of Economics*, 106, 1039-1061.

Waldfogel, (2002): Does Consumer Irrationality Trump Consumer Sovereignty? Evidence from Gifts and Own Purchases, mimeo.

Yariv, L. (2005): I'll see it when I believe it: A simple model of cognitive consistency, mimeo.

Chapter 2

Are Acquirers Rational?

Abstract

A major finding of the literature examining the stock price changes driven by merger announcements is that the combined firm gains significantly, the target firm captures the whole stock market surplus and acquirers on average hardly break even. If all empirical evidence is correct the following questions arise: Why do firms volunteer as buyers if they cannot expect to gain? Why do acquirers systematically overpay the target firms? Are acquirers rational? This paper argues that previous empirical investigations have only partially detected the effects of mergers. Following Hall (1999) I apply the propensity score method to a large dataset from Compustat with information on mergers and acquisitions. I find that using observable corporate characteristics as explanatory variables the model predicts with high probability the value-creating mergers (i.e. those that increase total factor productivity and R&D investment growth) and gives low probability to the value-preserving and -destroying ones. This result suggests that investors should be able to predict profitable mergers already before the announcement is made and in efficient markets this information should be incorporated in the stock prices by the time of the merger announcement. The event study performed to test this hypothesis confirms significant abnormal stock price returns preceding the announcement day for the acquirers classified into the higher acquisition probability classes.

1 Introduction

Mergers and acquisitions periodically reallocate massive resources of the economies and reshape the structure of industries.¹ The value that changes ownership during peak activity periods may account for 5% of world-wide GDP. The three main merger waves since the early 1960s have revealed an ever increasing corporate interest in takeover activities measured both by the number and the value of the deals. During the 1980s nearly half of all major US corporations received a takeover offer and the M&A activity in the 1990s is even more dramatic and widespread, as reported by Mitchell and Mulherin (1996) and Andrade et al (2001). Albeit this clear tendency there is still little agreement among empirical researchers in explaining why mergers take place, whether they create shareholder value and if they do, what the sources of the gains are.

Theoretical contributions have provided many possible reasons for why mergers might occur: improving efficiency by exploiting economies of scale or other synergies; attempts to create market power; gains from removing the incompetent management of the target firms; managerial hubris and empire building; diversification and reaction to unexpected shocks to industry structure. In order to identify empirically which of these reasons are the main drivers of takeover activities, a useful step would be to establish whether mergers and acquisitions create or destroy shareholder value. Despite the latter question has caught the attention of numerous scholars, univocal agreement has not been reached concerning the issue.

There have been two major approaches to the quantitative assessment of merger outcomes: event studies and operating performance studies. Event studies build on the theory that in efficient markets stock prices quickly adjust following any public information incorporating the news' entire expected future wealth effects (for reviews see Ravenscraft and Scherer, 1987 and Sirower, 1997). Accordingly, standard short window event studies measure the announcement period abnormal returns for both acquirers and targets, as well as for the acquirer and the target combined. Operating performance studies (e.g. Scherer and Ross, 1990; Healey et al, 1992; Bild, 1998), on the other hand, analyze the changes in reported or estimated performance

¹The differences in the two forms of corporate takeovers, mergers and acquisitions, are not relevant for the purposes of this study, in the followings I use the terms as synonyms.

indicators, such as return on equity or assets and growth rates of capital expenditure or total factor productivity.²

This paper addresses the puzzling pattern of abnormal stock price returns found by short window event studies in the 3-5 days around merger announcement dates. The stylized facts of this strand of research are that the combined firm gains significantly, the target firm captures all the stock market surplus and acquirers lose or break even (Jensen and Ruback, 1983; Bradley et al, 1988; Kaplan and Weisbach, 1992; Mulherin and Boone, 2000; DeLong, 2001). As for the magnitude of the announcement period stock price changes the following numbers are reported by surveys: the average announcement period abnormal returns over the three-five-day event window for the combined firms range from 1% to 3% suggesting that mergers do create shareholder value on average. While these returns may not seem extraordinary at first sight, the fact that they materialize in a very short period (3-5 days) clearly places corporate takeovers among the most profitable investment opportunities. A separate analysis of the same measures for the target firms and the acquirers reveals an unexpected pattern. The average announcement period abnormal return for target firms is in the range of 15-25%, while the corresponding number for the acquirers is between (-3)-1%. From these results, short window event studies generally conclude that mergers are on average profitable but the announcement period gains accrue entirely to the target firms' shareholders leaving acquirers hardly break even. These results are puzzling especially because takeover decisions are in most cases brought by acquirers. Why do some firms volunteer as buyers if they cannot expect to gain from their investment? Why do acquirers systematically overpay the target firms?³

The hypothesis that I propose, based on the findings of an operating performance study conducted by Hall (Hall, 1999), is the following: the value-creating mergers are predicted and expected by the stock market participants, and thus their stock prices incorporate the gains from the merger already before the announcement is made. The value-preserving and -destroying mergers, on the other hand, are unexpected, and it is their impact that traditional short window event studies measure. I test this hypothesis by analyzing the performance and the stock market returns of about 550 U.S. based companies that acquired another firm

²See section 3 for a discussion on operating performance studies.

³Previous attempts to answer these questions are discussed in section 2.

during the period 1975-1995. Applying the propensity score methodology of Rosenbaum and Rubin (1984) to construct appropriate control groups for the merging companies I estimate the probability that a firm will engage in corporate takeovers by running a logit regression on observable company characteristics. Based on the predicted probabilities of acquisition I classify all firms in the sample into 6 coarse groups and I compare the performance of the merging and not merging firms within each group. I find that firms with a high (low) a priori probability of making an acquisition have significantly higher (similar) total factor productivity and R&D growth after the merger than those firms that have the same probability but do not engage in corporate takeovers. Moreover, the short window event study confirms significant abnormal stock price returns preceding the announcement day for the acquirers classified into the highest acquisition probability classes.

The contribution of this paper to the literature is twofold. First, it gives a new explanation to the puzzling pattern of announcement date stock price returns observed by short window event studies. By doing so it casts doubt on previous findings, namely, acquirers hardly break even in terms of stock price returns and targets capture the whole stock market surplus arising from mergers. Furthermore, it offers a meaningful separation of mergers that sheds light on the differences obscured by using aggregate data. Classifying mergers based on the likeliness that they will take place is a natural starting point if we consider the extent of effort investors exert in predicting such movements. This paper joins a growing number of studies that bring together the two major methodologies (operating performance and event studies) applied in the literature (Healey et al, 1992; Dickerson et al, 1997), which allows one to conclude on the basis of a largely extended information set.

The rest of the paper is organized as follows. Section 2 details the hypothesis and explains the reasoning behind it. Section 3 contains the first step of the empirical strategy, the operational performance study. It describes the dataset and the methodology and summarizes the results. Section 4 briefs on the dataset and the methodology used for the event study and analyses the results. The last section concludes.

2 Hypothesis

This paper focuses on the findings of the short term event studies, i.e. the combined firms gain significantly, the target firms capture the whole surplus brought about by the merger and the acquirers hardly break even. Before detailing the hypothesis that I propose I briefly review some previous attempts to rationalize this phenomenon.

Some authors have suggested that the presence of competing bidders could allow targets to extract full value from the eventual acquirer (Molnar, 2002). However, Andrade et al. (2001) mention two reasons why this may not be the case. First, the majority of acquisitions only feature one bidder. Second, if gains were to stem from synergy effects, then it should imply that there is a unique match between the target and bidder that would allow for such gains to arise.

One may claim that the arguments of Andrade et al. (2001) are not entirely acceptable as we do not need to observe the presence of competing bidders in order to have them affect the outcome of the bidding process. In other words, it is enough that the company that bids for a target takes into account the possibility that a relatively low offer could trigger higher counteroffers by other interested parties. However, this kind of reasoning is somewhat partial as it ignores the possibility that there may be a number of potential targets for the acquirer. If competing bidders were to decrease the bargaining power of the acquirers, other potential targets should similarly decrease the bargaining power of the sellers.

Another explanation for the targets' success may be conjectured from studies tracing insider trading (Meulbroek and Hart, 1997). Information leaking before the announcement day allows investors to bid up the share prices of target firms and so capture the entire gain from the acquisition. However, insider trading is not so widely observed phenomenon that would allow for such generalization.

An important point concerning the distribution of gains between the merging parties is the size difference between targets and acquirers. As several authors have pointed out (Asquith et al (1983), Andrade et al (2001)), acquiring firms are on average significantly larger than their targets, thus one should expect a significantly lower percentage gain for the acquirers even in case of equal distribution of the gains. This, however, still does not explain why event studies

find negative or zero abnormal return for the acquiring companies.

The hypothesis advanced here builds on Hall (1999), who seeks to identify the effects of M&A activities on R&D and productivity growth. Using the propensity score method in order to construct an appropriate control sample for the merging companies, she estimates the probability of acquisition for each firm based on some corporate characteristics (size, Tobin's Q, cash flow etc.). Classifying the firms by their acquisition probability and comparing their outcomes (R&D investment, TFP growth) to those of the control groups, she finds the following:

1. Firms with a high predicted probability of making an acquisition, which actually make an acquisition, have a significantly higher increase in their R&D and a significantly higher total factor productivity (TFP) growth when compared to the control firms. I will term this type of acquisitions "value-creating mergers".
2. Firms with a low a priori acquisition probability that actually make an acquisition have a significantly lower increase in their R&D when compared to the control firms that have not merged, and their TFP growth around the merger is about the same as that of the non-merging firms. These acquisitions are called "value-preserving and -destroying mergers" in the followings.

The hypothesis that I propose based on the above findings is the following: mergers that eventually bring about significant corporate performance improvements are predictable, and therefore expected by investors well before the announcement date. Consequently, event studies with a short window around the announcement date will be unable to trace down the impacts of the value creating mergers, these being already incorporated in the share prices. As for the value-preserving and -destroying mergers, these mergers are unexpected and it is their impact that short window event studies measure. Thus, the average change in the acquirers' share prices at the announcement date will be negative or zero, but this finding only partially reflects the true merger effects. The abnormal stock price returns preceding the announcement date should be included in the analysis in order to account for the gains arising from value-creating mergers.

The hypothesis that acquirers do gain from corporate takeovers also has important implications for the distribution of gains between the merging parties. In particular, it casts doubt

on the finding that targets capture the whole stock market surplus, and, as discussed above, even a small percentage gain shown to be extracted by the acquirers could testify an equal distribution of the value created if size differences between the merging firms are large.

An important point should be addressed before I turn to the arguments that support the above hypothesis. One could claim that good mergers should be predictable for the target firms as well, and thus part of their gain might be incorporated in their share prices before the announcement is made. Although this would not affect our conclusion about whether acquirers fare better than it is currently held, it would have important implications about the distribution of gains. However, several studies (e.g. Palepu, 1986) find that predicting acquisitions is much less powerful for target firms than for acquirers. In Section 3 I confirm that this is the case in my dataset as well. For this reason it seems justified to expect that the announcement that a company is a takeover target is a total surprise for the investors, and the full effect of the merger on the targets' value can be measured at the announcement day.

The following considerations support the view that profitable takeovers may be expected by the market preceding the announcement date. First, the M&A departments of large consultancy firms, funds and investment banks spend significant resources on developing models that are capable to predict M&A activities. Market participants' effort to anticipate certain takeover activities are even documented in the Wall Street Journal Index by the comment "bid foreseen by analysts" that often accompany merger announcements. Moreover, it is not the takeover per se that needs to be predicted in order for the acquirers' stock prices to rise already before the announcement of the merger. It is enough, for example, if investors see promising growth opportunities in a firm. How managers eventually materialize these opportunities, be it through mergers or building new plants etc., is unimportant concerning the stock market's behavior.

Studies on insider trading build on a similar idea as the one proposed here, i.e. investors may have information about mergers already before the announcement date. They, however, mainly focus on the movements of the target firms' stock prices preceding the announcement date and rarely on those of the acquiring firms. This paper is closest in spirit to Asquith et al. (1983) and Schipper and Thompson (1983), two studies that analyze the abnormal stock price returns of acquirers preceding the announcement date. Both test the hypothesis that

when firms announce merger programs, i.e. are planning to undertake a series of takeovers, the gains of later acquisitions are capitalized in the beginning of the programs. They both find significant positive excess returns for different lengths of periods before the announcement of the mergers, which accords well with the hypothesis advanced in this paper. Contrarily to this study, however, their datasets include only firms announcing merger programs and in both cases the reason for examining the stock prices before the announcement date is the possible leakage of information, in other words: insider trading.

3 Operating performance after mergers⁴

Operating performance studies attempt to measure the impact of mergers and acquisitions by comparing various corporate performance indicators of the merged firms to those of the control firms. These indicators may vary from profitability measures, such as return on assets or operating margins to growth measures such as rates of R&D investment and capital expenditure. In this study the effects of corporate takeovers on total factor productivity (TFP) growth and different measures of R&D investment are investigated.

3.1 Dataset

The dataset, constructed using various files of Compustat, contains approximately 6,000 United States-based publicly traded manufacturing firms that existed for some time between 1975 and 1995. The variables used in this analysis, such as employment, sales, capital to labor ratio, R&D investment, are described in Table 1, their means are given in Table 2 and their distribution among industries in Table 3. The dataset includes 623 mergers with enough data for the analysis, their distribution in time is shown in Table 4.

⁴This section closely follows the discussion in Hall (1999) as I am using both her methodology and her dataset. The dataset has been cleared since she conducted her study, this accounts for the differences between the results reported in the two papers.

3.2 Methodology

The main methodological challenge of operating performance studies is "the lack of observable counter-factual", i.e. the lack of a control sample of non merging firms that are otherwise identical to the merged sample. In order to isolate the true effects of mergers one would ideally compare the change in the outcomes of a firm around the merger to what it would have experienced in case it did not merge. Given that merging and not merging of the same company cannot be observed at the same time an appropriate control group has to be constructed. Several studies use the performance of the industry average or simply that of the entire non-merging sample as a benchmark when assessing the mergers' impact. This procedure is unsatisfactory if the merged firms differ from the benchmark companies in relevant respects that may affect their performance. One solution to this problem offered by Rosenbaum and Rubin (1984) is the "propensity score" methodology. In this approach, one computes the probability of an event (the "treatment") for a sample of observations and compares the outcome after the event date for firms that have the same probability of the event, but did or did not experience it.

Consider the two treatments, merging and not merging (labeled 1 and 0, respectively) and let z indicate the treatment assignment. The propensity score is the conditional probability that a firm with vector x of observed covariates will be assigned to merger, $e(x) = Pr(z = 1|x)$. Rosenbaum and Rubin (1984) show that in subclasses of the entire sample that are homogenous in the propensity score, $e(x)$, the distribution of the observed covariates is the same for treated and control firms; formally x and z are conditionally independent given $e(x)$. This means that the outcomes of firms within the same subclasses are directly comparable if we are willing to assume that the selection mechanism for merger is ignorable conditional on the predictors of $e(x)$. In other words, one has to accept, that subclassification on the estimated propensity score cannot balance unobserved covariates, except to the extent that they are correlated with the observed variables.

I estimate the probabilities of the mergers using a logit regression with explanatory variables such as the size of the firm measured by the logarithm of employment, the R&D intensity measured by the R&D to sales ratio, a dummy for zero or insignificant amounts of R&D, the capital-labor ratio, the logarithm of Tobin's Q, a dummy for missing or implausibly small (below

0.1) Tobin's Q, the earnings-sales ratio, and a dummy for very negative earnings. All covariates are lagged one year with respect to the year of acquisition in order to avoid endogeneity problems. The regression is run on the data of approximately 6,000 firms with nearly 560 acquisitions.⁵ The results from this regression are shown in Table 5.

The first column in Table 5 shows the coefficients from a logit regression for the probability that a firm will acquire another public firm. Such firms are much larger than the others, have lower R&D intensities and higher Tobin's Q. They are also less likely to have small cash flow - either more negative or more positive cash flow is associated with acquiring a public firm. The finding that higher probabilities of acquisition are assigned to larger firms further strengthens the hypothesis that stock market participants may expect some mergers. Larger firms are under constant monitoring by analysts and investors alike who spend significant effort to anticipate the movements of the industry leaders.

The predicted probabilities derived from this regression are shown in Box Plots for the two groups of firms in Figure 1 (the acquirers versus the rest). The boxes indicate that there are important differences between the firms that acquire and the ones that do not. This result is important for two reasons: first, it confirms that the effort to construct an appropriate control group is justified, the results would probably be biased if I simply compared the outcomes of the merging firms to those of their industry average (which is a typical procedure in operating performance studies). Second, the mean probability of acquirers is significantly higher than that of the non-acquiring firms, thus the regression performs relatively well in terms of predicting the probabilities of acquisitions.

The same logit regression was run to analyze the probabilities that a company will be acquired. Table 6 confirms what I anticipated in the Introduction, most of the coefficients are insignificant and the power of the regression is mainly attributed to the year dummies (not shown). The significant overlapping in the boxes in Figure 2 testifies that there is little difference between the two groups (sellers and the rest), therefore it seems justified not to expect that investors foresee which companies will be acquired. In other words, we should expect that the announcement that a company is a takeover target is a total surprise for the investors,

⁵There are no satisfactory data one year preceding the acquisition for 63 mergers out of the 623 in the original sample.

and the full effect of the merger on the targets' value can be measured at the announcement day. In the next step I compare the performance (TFP and R&D growth) of the merging and non-merging firms in the different probability classes.

3.2.1 Total factor productivity growth

In order to calculate TFP growth I estimate a Cobb-Douglas production function using the panel data described above. As it has often been found in previous studies, simple OLS regressions produce coefficients in line with constant returns to scale (CRS) and with evidence from factor shares. However, these estimates are likely biased due to problems of simultaneity and unobserved, firm-specific effects. The first differenced GMM estimator proposed by Arellano and Bond (1991) yields unsatisfactory results, the capital coefficient is implausibly low and insignificant and tests easily reject the CRS assumption. Thus, I use the extended GMM estimator of Blundell and Bond (2000), that specifically addresses these problems, to estimate the following equation:

$$y_{it} = \beta_n n_{it} + \beta_k k_{it} + \gamma_t + (\eta_i + v_{it} + \omega_{it}) \quad (1)$$

$$v_{it} = \rho v_{i,t-1} + \epsilon_{it}$$

where y_{it} is log sales of the firm in year t , n_{it} is log employment, k_{it} is log capital stock and γ_t is a year-specific intercept. The error term is composed of an unobserved, firm-specific effect, η_i , a possibly autoregressive productivity shock, v_{it} , and measurement error, ω_{it} . This equation can be conveniently estimated by using the `xtabond2` command in Stata developed by Roodman (2004). Coefficients are estimated separately for the different industries and total factor productivity growth is gained as the residual from the first differenced version of equation 1. The TFP growth rates are then compared for two groups of firms that merge and do not merge, but that have the same probability of merger. The comparisons are done using medians and using a distribution free Kruskal-Wallis or rank sum test.

Changes in TFP growth are measured between the average of up to three years' worth of data before the merger to up to three years' data after, adjusted for the average length of the gap so that they correspond to annual rates. The pre-merger values used are the combined

TFP growth rates for the two merging firms (weighted by relative sales to adjust for the size differences between buyers and sellers). Overall there is very little difference between the two groups of firms, and this is confirmed in Table 7a. The Kruskal-Wallis test confirms that the hypothesis that there is no difference between the two distributions is not rejected at the 10% level.

However, when the firms are ranked by their probability to acquire, significant differences emerge. The observations are classified into 6 groups based on their estimated propensity to acquire other firms; the groupings are designed so that the number of acquired firms is roughly the same across them. Table 7a shows that firms with a lower propensity to merge (between 1 to 3%) that actually make an acquisition have a somewhat lower increase (in some cases this is a decrease) in their TFP growth when compared to firms that have not merged. This result is significant at the 10% level. On the contrary, firms with an acquisition probability higher than 10% have significantly higher TFP growth (with about 1% per annum) if they actually make an acquisition than if they do not.

3.2.2 R&D growth

In the next step changes in real R&D and in the R&D-to-sales ratio are investigated. I use the same procedure as for TFP growth to measure changes in these indicators. That is, changes are calculated from the 6 years' worth of data around the merger and the pre-merger values are the combined rates for the two merging firms. The results in Table 7b show similar pattern to that discussed above. Merging and not merging firms on average do not significantly differ concerning their R&D activity. However, Table 7b confirms that firms that have lower acquisition probability and do make an acquisition have significantly lower increase in R&D than that of the control group. The opposite relation is true between acquirers and control firms in the higher acquisition probability classes.

The implications of the overall results are the following: observing certain characteristics of the firms, such as size, Tobin's Q, capital/labor ratio and cash flow, one can predict the probability that a firm will acquire another one and can expect that the acquisition will create value in case of high predicted probabilities. In the followings I examine whether the stock

price movements of the different groups of merging firms accord with the hypothesis that stock market participants use such predictions.

4 Abnormal returns at and before announcement

Standard event studies measure the impact of an economic event on the value of firms using security prices observed over a relatively short time period. They usually restrict their attention to the stock price returns in the 3-5 days around the merger announcement date, building on the theory that in efficient markets stock prices quickly adjust following any public information incorporating the news' entire expected future wealth effects. The hypothesis in this paper suggests that abnormal stock price returns for some acquirers accumulate in a longer period preceding the announcement date of the mergers. But how long should be the period investigated? In the operation performance study I used company characteristics lagged one year to predict the acquisition probabilities. As such, a natural choice would be to analyze stock price movements throughout the whole year preceding the announcement of the mergers. However, several authors have raised concerns about long-run event studies for methodological deficiencies (Fama, 1998; Mitchell and Strafford, 2000; Brav, 2000). A major problem is that modeling expected returns for a longer horizon (often multi-year periods) can only roughly be estimated, therefore estimates of long-run abnormal returns will necessarily be imprecise. For this reason I limit the analysis to a relatively short period, 30 days, preceding the announcement date. I recognize that in this way I may not fully capture the abnormal returns accruing to acquirers, but I believe that the results are in any case informative about the hypothesis. As a robustness check I also calculate the abnormal returns for a longer period, 100 days, preceding the announcement of mergers.

An alternative way to test the hypothesis would be to compare the average cumulative returns of the acquirers in the whole year preceding the announcement date to those of the control groups. This could be done by matching the acquirers to a number of firms in the same probability class that were not engaged in corporate takeovers in the period in question. This procedure may get around the problem of estimating expected returns for a long horizon, as the normal return for the acquirer would simply be calculated as the actual return of non-merging

firms with the same acquisition probability. Due to limitations of the data this exercise will be left to future work.

4.1 Dataset

For the same firms analyzed in the previous section the announcement dates of acquisitions have been collected from the Wall Street Journal Index. Additionally, the announcement dates were cross-checked and information about the values of the transactions were collected using the Security Data Corp (SDC) data file. Unfortunately, the SDC contains only about half of the mergers in the original dataset, and does not have information about the largest mergers. The dataset includes daily stock price returns for the acquiring firms 300-300 days before and after the announcement date, and the same returns for the market index for the entire period between 1976-1995. The source of the information on stock prices was the Center for Research in Security Prices (CRSP).

4.2 Methodology

In order to gauge the merger's impact a measure of the abnormal return needs to be constructed. The abnormal return is the actual ex post return of the security over the event window minus the normal return of the firm over the event window. The normal return is defined as the expected return without conditioning on the event taking place. The expected return is calculated on the basis of a market model, a statistical model which relates the return of any security to the return of the market portfolio. The model's linear specification follows from the assumed joint normality of asset returns (MacKinlay, 1993). For any security i the market model is:

$$R_{it} = \alpha_i + \beta_i R_{mt} + \epsilon_{it} . \quad (2)$$

The term R_{it} is the realized logarithmic return at time t for firm i ($i = 1, \dots, N$); R_{mt} is the CRSP value-weighted market index logarithmic return at time t ; ϵ_{it} is the residual which is assumed to be serially independent and normally distributed with a zero mean and standard deviation of σ_i .

To obtain the parameters of the market model OLS regressions were run for each firm separately using daily-returns data starting 200 days (115-117 trading days) before the an-

nouncement date and ending 30 days (18-21 trading days) before the announcement date. The abnormal return for firm i on date t is given by:

$$AR_{it} = R_{it} + (\hat{\alpha}_i + \hat{\beta}_i * R_{mt}), \quad (3)$$

where $\hat{\alpha}_i$ and $\hat{\beta}_i$ are the predicted parameters of the market model. The cumulative abnormal returns are given by a summation of the abnormal returns over the 30 days preceding the announcement date:

$$CAR_i = \sum_{t_0}^T AR_{it} \quad (4)$$

The cumulative abnormal returns are calculated similarly over the 3 days around the announcement date. The average CAR for a group of N firms is given by:

$$\overline{CAR}_i = \frac{1}{N} \sum_{i=1}^N CAR_i \quad (5)$$

The statistical significance of abnormal returns can be tested by averaging abnormal returns across all stocks on each trading day and calculating the standard deviation of the daily averages, σ , from the time period used to estimate the market model (which is likely unaffected by merger news). The ratio of the individual daily average abnormal return to σ follows a Student t -distribution with $N-1$ degrees of freedom, N representing the number of firms in the analysis, assuming daily abnormal returns are normal, independent, and identically distributed across time. For cumulative returns, I divide the average cumulative daily abnormal return by $\sigma\sqrt{M}$, where M equals the cross company average number of trading days in the period. The ratio is approximately normally distributed assuming a zero mean average cumulative abnormal return and statistical independence of daily average abnormal returns (Franks and Harris, 1989). I classify the firms into different groups based on their acquisition probabilities calculated in the previous section and I calculate the average cumulative abnormal returns both preceding and on the announcement date.

4.3 Results

Table 8a and 8b report the means of CARs over the 3 days around the announcement date and the 30 days preceding the announcement date. Overall, there were 427 firms included in the

analysis from the original dataset (623 mergers). About 60 observations have been discarded because the exact announcement date of these mergers could not be found either in the Wall Street Journal Index or in the SDC data file. Also, about half of the merging firms engage in multiple acquisitions during the period analyzed with 4-5 acquisitions/firm not uncommon. For those firms that announce two corporate takeovers within a year (around 70), the second observation was excluded from the analysis because in these cases the normal return estimations were not reliable. Finally, observations where the relative size of targets to acquirers were lower than 5% (nearly 70 mergers) were not included in the analysis.

The last entry in the first column of Table 8a shows that the 427 acquirers in the sample had insignificant, 0.08% negative abnormal returns over the 3 days around the announcement date. The results remain qualitatively the same, i.e. not significantly different from zero, when firms in the different probability classes are separately analyzed. The findings accord well with the conclusion of previous studies; acquirers do not gain significant abnormal returns in a short period around the announcement of the mergers. However, the second set of columns in Table 8a confirm that the same firm gained 1.22% cumulative abnormal return preceding the announcement date. This result is significant at the 10% level. The firms in the lower probability classes had cumulative abnormal returns in the same period that are not significantly different from zero. Those in the higher three probability classes gained 2.84%, 1.95% and 2.13% CARs, that are significant at least at the 5% level. The percentage gains in these classes decrease somewhat with the probability of acquisition which might reflect that the value of the targets relative to that of the acquirers is decreasing. As I mentioned earlier, the dataset includes the relative value of the merging firms for only about half of the mergers. From the above analysis I excluded the mergers where the relative value of the two companies is less than 5% and I kept all the mergers for which I did not have information about the transaction value. The employment of such threshold for the relative size of the merging parties is a standard approach in the literature (Savor, 2006) since the inclusion of bids for very small firms would just add noise to the results. The lack of data on the relative values primarily affects the larger companies both because the SDC data file did not have information about the largest mergers and because the expected relative value in these cases are lower. For this reason the highest probability classes may be contaminated with mergers where the relative value of the merging firms are very low

and the above estimates of the abnormal returns might be seriously downward biased.

As an attempt to infer something about the robustness of these results I perform the same analysis on the 100-day-period preceding the announcement dates. Recalling the earlier discussion about long run event studies insignificant results are not surprising. It is noteworthy, however, that the point estimates accord well with the hypothesis. The results in Table 8b show somewhat higher CARs than those measured during the 30-day-interval for the acquirers in the higher probability classes. Results are qualitatively the same as before for the acquirers that have a lower propensity to merge. Note that the announcement date returns are slightly different from the ones previously reported because the normal return model is estimated using stock market returns from an earlier period. However, the conclusion about these results are the same as before, acquirers do not gain significant abnormal returns at the announcement of the mergers.

The overall results suggest that event studies analyzing the abnormal stock market returns in a very short period around the announcement dates of the mergers might have missed part of the effects of mergers. Note that the above results are only indicative about the magnitude of the abnormal stock price returns preceding the announcement dates. First, the results are likely downward biased as discussed earlier. Second, due to methodological weaknesses in estimating long-run expected returns, I preferred to concentrate on a relatively short period (30 days) preceding the announcement dates.

5 Discussion and Conclusion

In this paper I have argued that in order to fully capture the impact of mergers on acquirers stock prices the investigation needs to be extended to a longer period preceding the announcement dates. This claim has been formulated based on the results of an operational performance study, i.e. merged companies that have a priori high (low) probability to make an acquisition fare better (similarly) in terms of TFP and R&D growth than firms that have the same acquisition probability but do not engage in corporate takeovers. With a relatively short window event study I tested the hypothesis that investors are able to utilize this information and predict profitable mergers before the announcement of the mergers. The empirical analysis confirmed

a significant 1.2% overall gain for the acquirers during a 30-day-period preceding the announcement dates. Firms that had an a priori probability of acquisition lower than 5% had abnormal returns not significantly different from zero in the same period. Acquirers with higher acquisition probabilities gained significant cumulative abnormal returns in the magnitude of 2-3%. Note, that the 2-3 percent abnormal stock returns are brought about by the largest companies, which suggest that the absolute gains might be economically very relevant. Note also, that the CARs reported may be seriously downward biased. A major shortcoming of the current state of the empirical analysis is the lack of data on the relative size of target and acquirer for the majority of the mergers. I suggest that acquisitions in which the deal value is less than 5% of the acquirers value should not be included in an event study analysis. Failure to do so might have accounted for lots of observations having insignificant returns also in the high probability classes (where this problem is even more serious since the largest acquirers are in the highest probability classes), which in turn might have biased the CARs downward. The overall results cast doubt on previous findings according to which acquirers do not gain, at least not in terms of stock prices, from corporate takeovers. Moreover, it necessitates the reexamination of the claim that they on average overpay target firms.

The dataset used to derive these conclusions is uniquely extensive, and the combination of an operational performance and an event study allowed me to support my hypothesis exploiting as many information about these companies as possible. However, more could be done with a further enrichment of the information set, which will be the aim of future research. First, there are several ways to further test my hypothesis. The findings so far accord well with the ones of previous studies, i.e. the announcement date returns for the overall group of acquirers as well as for each group stratified by the propensity to merge are insignificant. However, previous studies find significant announcement date abnormal returns for the combined firm. If the hypothesis in this paper was correct we should find that targets of high probability mergers (value-creating mergers) earn very significant CARs in the short period (2-3 days) around the announcement date. We would expect that targets of value-preserving/destroying mergers do not earn significant returns in the announcement period. In other words, the gains for the combined firms found in previous studies should be entirely brought about by mergers that have a priori high probability to take place. This could be easily tested with data on the deal

value and the targets' value around announcement time.

A second way to check the robustness of the results of the event study would be to compare the average cumulative returns of the acquirers in the whole year preceding the announcement date to those of the control groups. This could be done by matching the acquirers to a number of firms in the same probability class - possibly by industry - that were not engaged in corporate takeovers in the period in question. This procedure would be more in line with the hypothesis in this paper, as I show that investors should be able to predict the profitable mergers already a year before the announcement dates. Moreover, this methodology would get around the problem of how to estimate long run expected returns, as the normal returns would be defined as the average returns of the control group.

References

Andrade, G., Mitchell, M., and E. Stafford (2001): New Evidence and Perspectives on Acquisitions, *Journal of Economic Perspectives*, 15, 103-120.

Arellano, M. and S. Bond (1991): Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations, *Review of Economic Studies*, 58, 277-297.

Asquith, P., Bruner, R. F., and D. W. Mullins, Jr. (1983): The Gains to Bidding Firms From Merger, *Journal of Financial Economics*, 11, 121-139.

Bild, M. (1998): Valuation of Takeovers, EFI: Stockholm.

Blundell, R. and S. Bond (2000): GMM Estimation with Persistent Panel Data: An Application to Production Functions, *Institute for Fiscal Studies*, W.P. No W99/4.

Bradley, M., Desai, A. and E. Kim (1988): Synergistic Gains from Corporate Acquisitions and Their Division between the Stockholders of Target and Acquiring firms, *Journal of Financial Economics*, 21, 3-40.

Brav, A (2000): Inference in Long-Horizon Event Studies: A Parametric Bootstrap Approach with Application to Initial Public Offerings, *Journal of Finance*, 55, 1979-2016.

DeLong, G. L. (2001): Stockholder Gains from Focusing versus Diversifying Bank Mergers, *Journal of Financial Economics*, 59, 221-252.

Dickerson, A., Gibson, H. and E. Tsakalotos (1997): The Impact of Acquisitions on Company Performance: Evidence from a Large Panel of U.K. Firms, *Oxford Economic Papers*, 49, 344-361.

Fama, E. F. (1998): Market Efficiency, Long-Term Returns and Behavioral Finance, *Journal of Financial Economics*, 49, 283-306.

Franks, J. R. and R. S. Harris (1989): Shareholder Wealth Effects of Corporate Takeovers: The U.K. Experience 1955-1985, *Journal of Financial Economics*, 23, 225-249.

Hall, Bronwyn H., 1999. 'Mergers and R&D Revisited', Berkeley, University of California, Mimeo.

Healey, Paul M., Palepu, Khrisna G., and Richard S. Ruback, 1992. 'Does Corporate Performance Improve After Mergers?', *Journal of Financial Economics*, vol. 31, pp. 135-175.

Jensen, M. and R. Ruback (1983): The Market for Corporate Control: the Scientific Evidence, *Journal of Financial Economics*, 11, 5-50.

Kaplan, S. and M. Weisbach (1992): The Success of Acquisitions: Evidence from Divestitures, *Journal of Finance*, 47, 107-138.

MacKinlay, Craig A., 1997. 'Event Studies in Economics and Finance', *Journal of Economic Literature*, vol. 35, pp. 13-39.

Meulbroek, Lisa K. and Carolyn Hart, 1997. 'The Effect of Illegal Insider Trading on Takeover Premia', *European Finance Review*, vol. 1, pp. 51-80.

Mitchell, M. and E. Stafford (2000): Managerial Decisions and Long-Term Stock Price Performance, *Journal of Business*, 73, 287-329.

Mitchell, M. L. and J. H. Mulherin (1996): The Impact of Industry Shocks on Takeover and Restructuring Activity, *Journal of Financial Economics*, 41, 193-229.

Molnar, J. (2002): Preemptive Horizontal Mergers: Theory and Evidence, Hungarian Academy of Sciences, Institute of Economics, DP. 2002/13.

Mulherin, J. H. and A. L. Boone (2000): Comparing Acquisitions and Divestitures, *Journal of Corporate Finance*, 6, 117-139.

Palepu, K. (1986): Predicting Takeover Targets: a Methodological and Empirical Analysis, *Journal of Accounting and Economics*, 8, 3-35.

Ravenscraft, David J., and F. M. Scherer, 1989. 'The Profitability of Mergers', *Journal of Industrial Economics*, 7, pp. 101-116.

Roodman, D. M. (2004): XTABOND2: STATA Module to Extend XTABOND Dynamic Panel Data Estimator, Boston College.

Rosenbaum, P. and D. B. Rubin (1984): Reducing Bias in Observational Studies Using Subclassification on the Propensity Score, *Journal of the American Statistical Association*, 79, 516-524.

Savor, P. (2006): Do Stock Mergers Create Value for Acquirers?, SSRN No. 881513.

Scherer, F. M. and D. Ross (1990): Industrial Market Structure and Economic Performance, Houghlon Niffin Company, Boston.

Schipper, K. and R. Thompson (1983): Evidence on the Capitalized Value of Merger Activity for Acquiring Firms, *Journal of Financial Economics*, 11, 437-467.

Sirower, M. L. (1997): *The Synergy Trap: How Companies Lose the Acquisition Game*, The Free Press, New York.

Table 1
U.S. Manufacturing Sector Variables

Variable	Source	Description
Sales	CS # 12	Sales (net) (\$M)
Employment	CS # 29	Number of employees (1000s)
R&D Spending	CS # 46	R&D expense (\$M)
Capital Expend.	CS # 30	Property, Plant & Equipment – Cap. exp. (\$M)
Net Plant	constructed	Net plant adjusted for inflation (\$M)
Net Capital Stock	constructed	Plant, inventories + other assets adjusted for inflation (\$M)
Market Value	constructed	Common equity + market value of long term debt + short term assets less short term liabilities + preferred

Source: Hall (1999)

CS = Standard and Poor's Compustat variable number

Table 2
Variable means and medians

Year		1976	1980	1985	1990	1995
Number of firms		2383	2415	2478	2472	2433
Employment	Mean	1120	1073	708	566	641
(numbers)	Median	1024	990	625	584	674
R&D to sales ratio	Mean	1.49%	1.89%	4.73%	5.66%	7.10%
(per cent)	Median	0.33%	0.47%	1.31%	1.40%	1.76%
Capital-labor ratio	Mean	20.2	22.4	24.1	24.0	26.5
(87\$K per worker)	Median	20.0	21.8	23.1	23.6	26.1
Cash flow-sales ratio	Mean	7.1%	4.9%	-1.1%	1.2%	1.0%
(per cent)	Median	7.3%	7.2%	7.6%	7.2%	9.2%
Tobin's Q	Mean	0.86	1.05	1.55	1.39	2.33
	Median	0.95	1.00	1.26	1.07	1.87
D (R&D Missing)		0.425	0.409	0.337	0.335	0.330
D (Q missing)		0.261	0.166	0.337	0.335	0.330
D (Negative earning)		0.078	0.103	0.193	0.200	0.173

Source: Hall (1999)

Table 3
U.S. Manufacturing Sector by Industry

Industry	Number of firms	Percent of firms
Food	346	5.83
Tobacco	9	0.15
Textiles	151	2.54
Apparel	192	3.23
Wood	91	1.53
Furniture	96	1.62
Paper & Pulp	135	2.27
Printing	217	3.65
Chemicals	299	5.04
Oil	67	1.13
Rub & Plastics	207	3.49
Leather	45	0.76
SCG	118	1.99
Prim Metals	198	3.33
Fab. Metals	304	5.12
Machinery	522	8.79
ElecMachinery	894	15.06
Auto & Transport Eq.	294	4.95
Sci Instruments	777	13.09
Misc Manufacturing	190	3.20
Drugs	347	5.84
Computer Equip.	437	7.39
Total	5,938	100.00

Table 4
Mergers & Acquisitions by Year

Year	Number of firms	Percent of firms
1977	38	6.10
1978	49	7.87
1979	58	9.31
1980	46	7.38
1981	42	6.74
1982	23	3.69
1983	38	6.10
1984	34	5.46
1985	44	7.06
1986	56	8.99
1987	48	7.70
1988	44	7.06
1989	30	4.82
1990	23	3.69
1991	8	1.28
1992	13	2.09
1993	21	3.37
1994	8	1.28
Total	623	100.00

Table 5
Probability of acquisition for buyers
Number of acquisitions (559)

Variable	Coefficient	Standard error	Z	P> z
Log employment	0.430	0.021	20.31	0.000
R&D – sales ratio	- 3.816	1.900	- 2.00	0.045
D (no R&D)	- 0.165	0.125	- 1.31	0.189
D (R/S>0.5)	0.878	0.992	0.89	0.376
Log of capital-labor ratio	0.033	0.045	0.74	0.458
Log Tobin’s Q	0.313	0.076	4.11	0.000
D (Q missing)	- 2.862	0.581	- 4.92	0.000
D (Q>10)	- 0.295	0.557	- 0.53	0.595
Log (cash flow/sales) c.f.>0	0.188	0.073	2.56	0.010
Log (cash flow/sales) c.f.<0	0.273	0.098	2.78	0.005
D (cash flow negative)	- 0.340	0.203	- 1.68	0.093
Log pseudo-likelihood = -2567.1205			Number of obs. =	44,614
			Wald chi2(20) =	719.32
			Prob > chi2 =	0.0000
			Pseudo R2 =	0.1314

* a full set of two year dummies was included

Table 6
Probability of acquisition for sellers
Number of acquisitions (648)

Variable	Coefficient	Standard error	Z	P> z
Log employment	0.017	0.019	0.93	0.355
R&D – sales ratio	0.956	0.851	1.12	0.261
D (no R&D)	- 0.215	0.094	- 2.29	0.022
D (R/S>0.5)	- 1.801	0.775	- 2.33	0.020
Log of capital-labor ratio	0.022	0.044	0.50	0.614
Log Tobin’s Q	- 0.169	0.059	- 0.28	0.776
D (Q missing)	- 0.954	0.169	- 5.64	0.000
D (Q>10)	- 1.562	1.012	- 1.54	0.123
Log (cash flow/sales) c.f.>0	- 0.089	0.053	- 1.68	0.093
Log (cash flow/sales) c.f.<0	- 0.081	0.054	- 1.52	0.129
D (cash flow negative)	0.932	0.285	3.27	0.001
Log pseudo-likelihood = -3206.7067			Number of obs. =	45,254
			Wald chi2(20) =	214.93
			Prob > chi2 =	0.0000
			Pseudo R2 =	0.0542

* a full set of two year dummies was included

Table 7a

Changes in TFP at Merger Controlling for Propensity to Merge

Estimated propensity to merge	Change in TFP Growth		
	No Acq.	Acq.	KW test
0.00-0.01 (61)	0.66%	-0.10%	1.54 (.214)
0.01-0.03 (135)	0.61%	-0.22%	2.87 (.094)
0.03-0.05 (69)	0.42%	0.62%	1.35 (.246)
0.05-0.07 (58)	0.68%	0.44%	0.84 (.359)
0.07-0.10 (72)	0.83%	1.28%	0.12 (.725)
> 0.10 (89)	1.14%	2.17%	4.00 (.046)
All (484)	0.65%	0.66%	0.22 (.641)

Table 7b

Changes in R&D at Merger Controlling for Propensity to Merge

Estimated propensity to merge	Change in real R&D			Change in R/S		
	No Acq.	Acq.	KW test	No Acq.	Acq.	KW test
0.00-0.01 (38)	15.5%	11.9%	0.10 (.755)	-0.0016	-0.0102	4.87 (.027)
0.01-0.03 (70)	19.1%	8.5%	4.38 (.004)	0.0016	-0.0006	16.77 (.000)
0.03-0.05 (41)	17.5%	6.0%	5.29 (.001)	0.0017	0.0000	4.12 (.004)
0.05-0.07 (40)	19.4%	17.3%	0.38 (.537)	0.0019	0.0027	2.80 (.090)
0.07-0.10 (56)	17.5%	20.6%	0.02 (.880)	0.0019	0.0016	0.96 (.328)
> 0.10 (76)	13.9%	23.0%	4.52 (.003)	0.0013	0.0011	0.51 (.476)
All (321)	16.6%	15.2%	0.25 (.619)	-0.0005	-0.0004	1.29 (.257)

Table 8a
Means of cumulative abnormal returns (CARs) over the 3 days around and over the 30 days preceding the announcement date

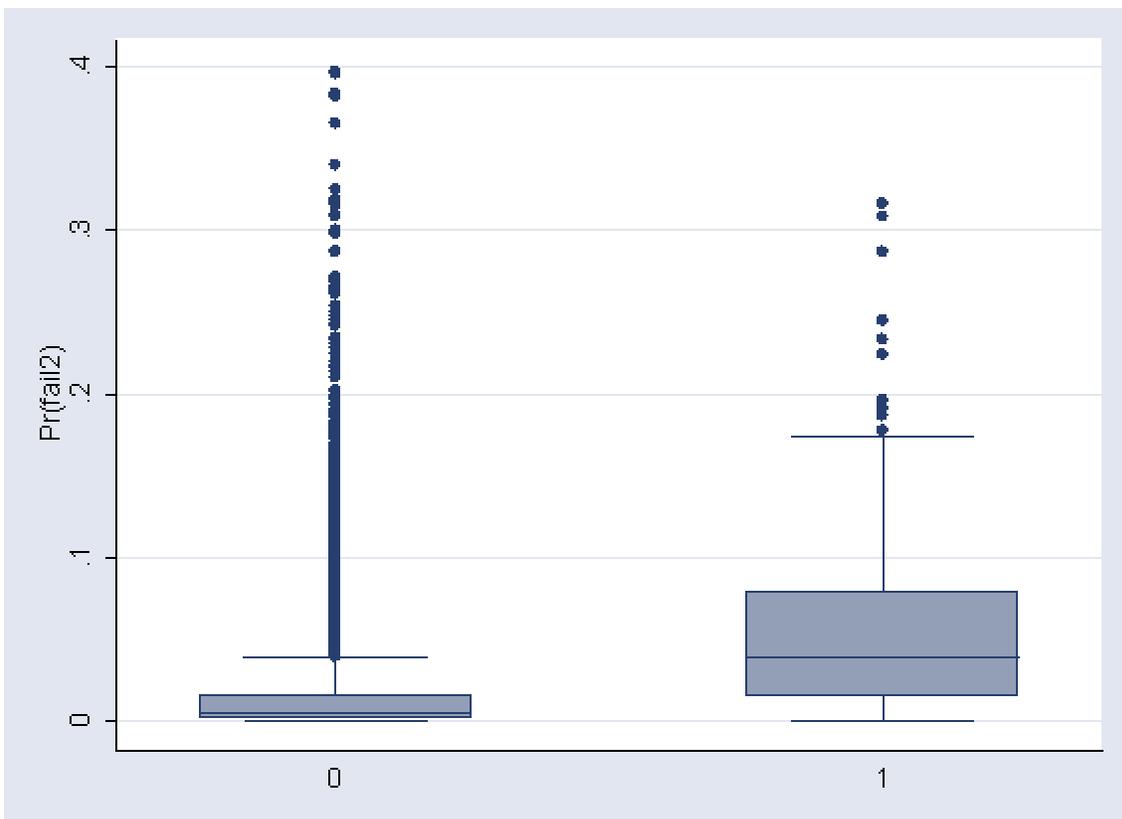
Probabilities (number of observations)	Mean CAR		Mean CAR	
	days: -1, 0, 1	P> z	days: (-30)-(-2)	P> z
0.000 – 0.011 (62)	0.26%	0.3576	0.49%	0.7642
0.011 – 0.017 (130)	- 0.06%	0.8336	0.33%	0.7718
0.017 – 0.026 (70)	0.09%	0.7490	0.97%	0.5686
0.026 – 0.040 (55)	- 0.59%	0.1052	2.84%	0.0050
0.040 – 0.058 (56)	- 0.34%	0.2302	1.95%	0.0446
0.058 – 0.167 (54)	0.00%	0.7794	2.13%	0.0104
All (427)	- 0.08%	0.7794	1.22%	0.0930

Table 8b
Means of cumulative abnormal returns (CARs) over the 3 days around and over the 100 days preceding the announcement date

Probabilities (number of observations)	Mean CAR		Mean CAR	
	days: -1, 0, 1	P> z	days: (-100)-(-2)	P> z
0.000 – 0.011 (62)	0.29%	0.3078	0.18%	0.9282
0.011 – 0.017 (130)	-0.03%	0.9204	1.40%	0.4840
0.017 – 0.026 (70)	0.02%	0.9442	0.11%	0.9522
0.026 – 0.040 (55)	-0.53%	0.1260	3.23%	0.1052
0.040 – 0.058 (56)	-0.29%	0.3030	1.45%	0.4654
0.058 – 0.167 (54)	-0.01%	0.8026	2.31%	0.2460
All (427)	-0.07%	0.8650	1.37%	0.4966

Figure 1

Box Plots for the probabilities that a firm will make an acquisition

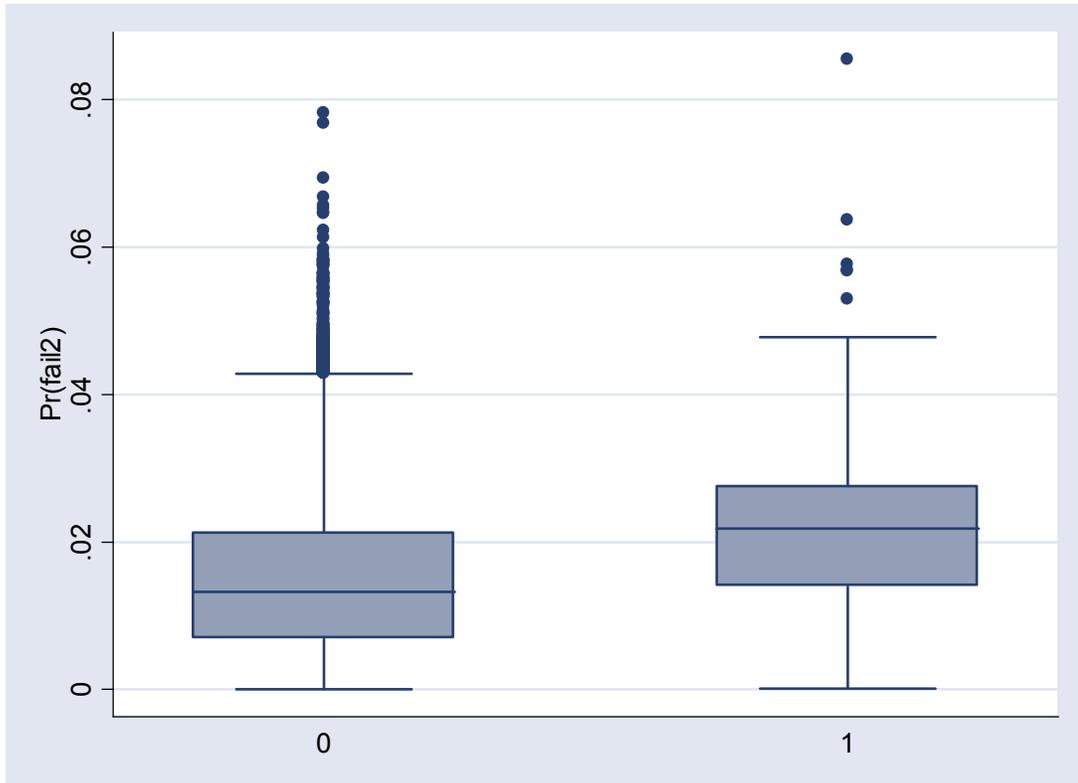


1 = acquiring firms

0 = not acquiring firms

Figure 2

Box Plots for the probabilities that a firm will be acquired



1 = acquired firms

0 = not acquired firms

Chapter 3

Agency costs, synergy gains and bargaining in corporate vs. financial takeovers

Abstract

Private equity takeovers, an increasingly important share of the global M&A market, have been accused to "cheat shareholders out of a fair price" [The Economist, February, 2007]. This paper analyzes this claim by looking at the differences in takeover premia paid by corporate and financial investors (predominantly private equity firms) for 226 quoted companies in the UK between 2004 and 2007. I document that financial investors pay only 70 per cent of the premia paid by corporate investors, 16 per cent and 23 per cent, respectively. The gap is significant and of approximately the same magnitude when I control for several characteristics of the targets and differences in the hypothesized sources of takeover premia: agency costs, bargaining power and synergy gains. Agency costs and bargaining considerations motivated by differences in the ownership structures of targets and bidders are found to be an important source of takeover premia as suggested by the principal-agent literature. Moreover, financial synergies, but not real synergies, account for a significant part of the premia as well. When I split the sample along the type of investors, most of the model's coefficients change signs between the two regressions suggesting important differences in the merger pricing behavior of industrial corporations and private equity firms. The results of the analysis support the argument that private equity firms pay significantly less for their targets and this does not seem to be a result of lower value generated by this type of takeovers. However, it can not be rejected that the observed pattern is not exclusive to private equity firms but holds for private firms in general.

1 Introduction

Private equity firms¹ are becoming an increasingly important player in the global mergers and acquisitions market with a purchasing power that allows them to consider multi-billion dollar transactions. This might have contributed to the resurrection of a decades-old debate over their operations. In the 1980s, private equity groups (or leveraged buyout firms, LBOs) were criticized for taking over companies using mainly debt borrowed against the acquired companies' assets, breaking them up and selling them in pieces at undeservedly high profit margins. Although tighter credit markets and competition from new entrants in the 1990s have forced them to add value to the acquired firms, private equity firms are still accused of "cheating the targets' shareholders out of a fair price", i.e. buying the companies well under value - often with the help of the management - and later selling them at a much higher price. Recent refusals of private equity bidders by some large public companies evidence how the latter are starting to respond to shareholder pressure to get a higher price for their firms. "Shareholders are increasingly asking why they are selling at a price less than what it will be worth in the future" (The Economist, February 10th - 16th 2007).

The premium paid on top of the target's share price that prevails before the acquisition² has been seen as the value assigned to the transaction by stock market investors. The goal of this paper is to investigate and to explain the differences, if any, in the price premia paid for public companies by corporate³ and financial investors⁴.

The question has both practical and theoretical importance. On the practical side, the ma-

¹For the purposes of this study the term "private equity firms" refers to the firms that through investing professionally managed funds gain significant control in private and, increasingly, in public companies. From the different types of private equity investments this paper focuses on leveraged buyout financing that enables publicly traded companies to be taken private.

²In this study takeover, merger and acquisition are used as synonyms.

³Corporate investors are companies with primary activities that are not fund management or financial investments. A vast majority of the public companies - that are usually analyzed by M&A studies - belong to this category.

⁴Note that the group of financial investors, that I consider here, is not restricted to private equity firms (although the latter take up about 70 per cent of the financial investors in the sample). It also includes other financial services firms, such as investment banks and trusts, other financial institutions, as well as individual investors

majority shareholders considering the exit from their companies would be interested in achieving the largest premia possible from the potential bidders. From a theoretical point of view, the question is interesting because there does not seem to be a strong reason for such an unambiguous difference between the premia paid by corporate and financial investors. Considering the three main drivers of takeover premia established in the literature - real and financial synergy gains, reduced agency costs and the respective bargaining power of target and bidder - it will be later argued that smaller synergy gains but larger reduction in agency costs and stronger bargaining power can reasonably be expected from private equity bidders. The overall effect of the expected differences in the three components of the takeover premia is theoretically ambiguous. What's more, contrarily to the accusation, the few empirical studies that report takeover premia in private equity acquisitions (e.g. Kaplan (1989) and Smith (1990)) show that premia paid to selling shareholders in the 1970s and 1980s were 30 to 50 per cent of the pre-buyout market value; much larger than the 20-30 per cent found for corporate takeovers.

In this study I analyze the differences in the abnormal share price returns, i.e. in the takeover premia, of 226 companies that have been the targets of takeovers in the UK during 2004-2007. The sample consists of 144 corporate and 82 financial bidders. Descriptive statistics do not reveal major differences in the characteristics of the targets acquired by these two groups of investors. Yet, I find that corporate investors pay significantly higher takeover premia, at 23 per cent, than financial investors (16 per cent). The 23 per cent takeover premia paid by corporate investors is closer to the lower end of the 20-30 per cent range reported in a field survey by Andrade et al. (2001).

Following the literature, I analyze whether differences in the determinants of takeover premia - agency costs, bargaining power and synergy gains - can explain the premia gap. I find that it is persistent even when controlling for these factors and various characteristics of the target and the deal. When I split the sample and run the regressions separately for financial and corporate bidders most of the model's coefficients change signs between the two regressions suggesting important differences in the merger pricing behavior of industrial corporations and private equity firms. The model's fit is better for corporate investors and the results indicate that, among others, financial synergies, but not real synergies, as well as the change in the ownership structure of the target firm are strong determinants of the takeover premia.

Interestingly, the strong negative coefficient of the ownership change runs against the traditional prediction of the principal-agent literature. This literature suggests that following the merger, the more concentrated shareholder structure of the bidder allows for the reduction of the target's agency costs thereby increasing the value created by the takeover. A plausible explanation for the opposite result, also derived from but previously not emphasized by the principal-agent literature, is that the owners of the bidder company provide stronger incentives for their management to bargain a larger share of the created value for themselves.

Most of the model's coefficients change signs and except for the target's size they are insignificant when the sample is restricted to the financial investors. The smaller sample size - 65 takeovers have complete data necessary to be included in the analysis - may partly account for the insignificance of the variables. Interestingly, the improvement in the ownership structure of the target has a positive coefficient in this regression. Although it is again insignificant, it is not unreasonable to believe that the reduction in agency costs is so much larger in private equity takeovers than in corporate takeovers that its effect offsets that of the bargaining considerations. Moreover, neither real nor financial synergies explain a significant part of the variation in the takeover premia paid by private equity firms. The latter finding is surprising, because due to the well-known debt burden of private equity takeovers, it could be expected that private equity bidders take a strong interest in the financial stringency of their targets.

Overall, it appears that the premia gap can not be explained by differences in the conventionally established determinants of takeover premia. While it is reasonable to expect that measurement error in the data of the private investors is larger, this only affects one variable - i.e. the ownership structure - in the regressions. Therefore, data problems do not seem to be the reason for the model's opposing predictions for the two groups of investors. At the same time, as the robustness checks reveal, it can not be excluded that not only private equity groups but private firms in general pay a significantly lower premia than their public counterparts. My results contribute to the literature on M&A's that has traditionally focused on corporate takeovers (e.g. Bruner, 2002; Andrade et al., 2001; Loughran and Vjih, 1997). There has been little empirical research on the differences in the merger pricing behavior of corporate and financial bidders. As private equity is becoming an increasingly significant player on the global M&A market, a better modeling of their pricing behavior could be an important research

agenda.

The findings are also relevant for the agency cost (principal-agent problem) literature. A vast theoretical and applied work show that management works more efficiently under tighter control, i.e. more concentrated ownership structure reduces agency costs. Mergers affect agency costs through changing the ownership structure of the target firms. The more concentrated the bidder the more value will be created through improving the agency situation of a dispersed target which will be positively reflected in the targets' premia. However, the relative ownership structure of target and bidder may also affect the respective bargaining power of the two firms. A more concentrated bidder provides stronger incentives to its management to bargain a larger share of the created value for themselves which will reduce the premium captured by the target. My findings suggest that for corporate investors the reduction in agency costs is more than offset by the bargaining power effect of the ownership structure. The opposite relationship seems to hold for financial investors.

The rest of the paper is structured as follows. Section 2 summarizes the theoretical and the empirical literature related to this study. Section 3 presents the data and descriptive statistics, and Section 4 explains the results and the robustness tests. The final section concludes.

2 Some theory and related literature

Corporate takeovers create value as documented by a number of studies (e.g. Matsusaka, 1993; Andrade et al., 2001; Jovanovic, 2004).⁵ Targets' share prices on average increase by 20-30 per cent on the announcement date of the mergers. At the same time the acquirer's shareholders hardly break even, therefore the abnormal change in the targets' share prices (takeover premia) represent the value created by the merger. Synergy gains and better management have been identified as the two key sources of the value creation (Slusky and Caves, 1991; Gondhalekar et al. 2004). That is, a merger between two firms can create value by combining the two firms'

⁵As there are very few studies estimating merger premia in private equity transactions the discussion here will start from the vast literature on mergers between public firms (that would be classified as corporate takeovers in this paper). Most of it directly applies to financial transactions with the differences summarized toward the end of the section.

related assets, technologies and financial resources, and by replacing an inefficient management with a more capable one.

Following Slusky and Caves (1991) I describe the takeover premia, PR, with the following expression:

$$PR = (BRES[X_i]/MV) * B(Z_i), \quad (1)$$

where BRES stands for the buyer's reservation value, MV is the market value of the target before the acquisition announcement is made and B is the bargaining power function.

2.1 Synergy gains

The first part of the expression relates the reservation value of the buyer to the market value of the target preceding the acquisition. According to the discussion above, the reservation value of the buyer exceeds the market value of the target by the synergy and agency gains (X_i) that can be expected from the takeover. A number of papers have pointed at the importance of synergistic gains in mergers (Jensen and Ruback, 1983; Shelton, 1988). Bradley et al. (1988) find that the average synergy gain in their sample represents about 7 per cent of the combined firm's equity value, which is on the lower end of the range - 7-14 per cent - reported more recently by Devos et al. (2004). Real synergy gains can be thought of as the increase in the operating revenues or decrease in the operating costs of the combined firm that would not be realized absent the merger. Performance improvement can be expected if the combined firm can utilize the related technologies or common inputs better than the two firms separately. Significant cost savings can be achieved if, for example, the two firms serve a common set of customers or use similar distribution channels. Thus, potential real synergies resulting from the merger can be measured from the closeness of the activities of the acquiring and acquired enterprises. Financial synergies, on the other hand, may arise if there is a significant discrepancy between the two firms' financial stringency. Financial ratios, such as gearing or liquidity, can be used to assess the potential gains from this source.

2.2 Agency costs

It has been argued that further value can be created by replacing the target's existing management with a more efficient one. A merger however does not necessarily involve a change in the management, especially not in case of management buyouts backed by private equity firms. What seems more relevant to my study is that it changes the owners of the target and often, as a result, its ownership structure as well.

A vast literature on the principal-agent problem, or agency cost, is devoted to the analysis of how the ownership structure of a firm affects its performance. In a closely held firm, there is little distinction between ownership and control. On the contrary, in a public firm with dispersed shareholders, the two are separated, and monitoring the management entails significant efforts and costs. In their classic analysis, Jensen and Meckling (1976) argued that managerial equity ownership helps to align the interests of the manager and minority shareholders. Since then, several empirical studies have estimated the relationship between managerial ownership and performance. Morck et al. (1988), for example, estimate a hump-shaped relation, Ang et al. (2002) find a positive relation and Nagar et al. (2002) report a U-shaped relation. While the level of managerial ownership is a natural starting point to assess the agency problem of a firm, the concentration of the ownership of outside shareholders seems to be an equally valuable proxy that captures the same information. Using the predictions of the principal-agent literature it can be argued that through changing the ownership structure of the acquired firm, mergers attenuate or mitigate the agency cost of the target contributing to the overall gain (or loss) created by the merger.

2.3 Bargaining power

The second part of expression (1), $B(Z_i)$, measures the relative bargaining power of the target and the bidder, which determines how the created value is distributed between the two of them. In Slusky and Caves (1991), this depends on the presence of actual and competing bidders that, in the limit, can drive the premia up until the total value created by the merger is captured by the target. Of course, the argument should hold symmetrically, i.e. bidders may have alternative targets which improves their bargaining position. Unfortunately, we rarely

have information on the availability of alternative targets. In contrast, most existing datasets report if the deal featured actual competing bidders. This happens in only three takeovers in my dataset therefore further explanation for the distribution of the gains between the target and the bidder may prove valuable. Slusky and Caves (1991) argue that other relevant factors affecting the bargaining power of the two firms are of little practical importance due to the lack of data to capture their effects. Although they discuss the predictions of the agency literature in relation to the efficiency gains resulting from the merger, they do not suggest that differences in the ownership structures of the merger participants would affect their bargaining positions. Yet, the analogy with the bargaining phase is apparent. A firm with a concentrated ownership structure facing the other side with dispersed owners, that have little control over the operations of their managers, will be better able to secure a significant part of the overall value for themselves. Therefore, the agency situation of a firm does not only affect its performance but may provide stronger bargaining incentives for its management.⁶

Note, that the bargaining power effect of the ownership structures on the targets' takeover premia runs against the agency effect. When a bidder with a tight ownership takes over a target with dispersed owners agency costs of the target will be reduced and premia accruing to the target increases, *ceteris paribus*. At the same time, such a bidder's bargaining power will be stronger therefore it will retain a larger share of the created value to itself. The overall effect of the ownership structure on the takeover premia is therefore ambiguous.

2.4 Expected differences between corporate and financial takeovers

In light of the above discussion what differences shall we expect in the value creation of the two types of takeovers analyzed in this study? First, real synergies - increase in cash flows from the combination of the two firm's assets and technologies - may be less substantial in financial takeovers where the bidder often does not possess assets and technologies related to those of the targets. Although there might be potential synergy gains generated with the private investor's portfolio companies, synergies are usually more widely available for corporate investors than for financial investors (Smit and De Maesenerie, 2005). On the other hand, significant financial

⁶This effect will be called the bargaining incentive or the bargaining power effect of the ownership structure.

synergies can be expected from private equity transactions due to the private equity firms' better access to financial resources and their well-known strategy to lever-up their targets. Mergers by corporate investors, on the contrary, generate considerable real synergy gains as evidenced by a number of studies mentioned earlier. On the issue of agency improvements, private equity firms themselves give us some guidance. An important way in which they claim to add value to the acquired companies is corporate governance. Unlike the owners of public companies, who tend to be too remote and spread to spend time and money monitoring a business, they keep a close interest in the companies they own and their executives benefit from the discipline of incentives and a high debt burden. The better ownership structure that they bring to the acquired company may result in significant performance improvement. What part of it will be captured by the target depends on the bargaining positions of the two firms, itself a function of their respective ownership function. While the three sources of value creation and distribution discussed above may dominate in corporate takeovers, there might be other factors more relevant for private equity transactions. Smit and De Maesenerie (2005), for example, classifies the sources of value creation in public-to-private transactions into two groups: the sources that are common to all private equity firms and hence are bid away during the competitive process, and the ones that are idiosyncratic to the bidder and therefore retained by the private equity groups. They argue that the first group, the one that is reflected in the takeover premia, is composed of the traditional sources of value creation discussed above. It seems justified therefore to focus on these sources and ignore the factors belonging to the idiosyncratic group. Note, however, that this study also ignores the sources of value destruction that have been argued to be important in public-to-private transactions. Pagano et al. (1998) suggest that a company loses all of the followings when going private: access to capital on public markets, liquidity and portfolio diversification, and monitoring provided by the stock market. At the same time, there are a number of advantages: reduction in stock exchange fees and the cost of dissemination of accounting information, increased confidentiality that may be crucial for competitive advantage, etc. Smit and De Maesenerie (2005) argue that "the importance of the disadvantages and costs are significant and are not offset by the benefits for most public firms" contradicting the evidence that private equity firms are now a major player in M&A markets. Indeed, studies that report takeover premia in public-to-private

acquisitions (Kaplan, 1989 and Smith, 1990) show that premia paid to selling shareholders in these transactions in the 1970s and 1980s are 30 to 50 per cent of the pre-buyout market value. These results suggest that the average premium paid by financial investors could be much larger than that found for corporate takeovers (20-30 per cent). However, competition on the market for private equity takeovers has increased considerably since the 1970s and 1980s, therefore it is quite possible that the above figures overestimate the current premia paid by financial investors. Contrarily to previous research, this study analyzes the investments by corporate and financial investors during the same period in the same geographical environment thereby allowing for a direct comparison of the two types of transactions.

3 Data and empirical analysis

3.1 Data

The initial sample consists of all UK domestic mergers and acquisitions in which the target is a public company, are announced - not necessarily completed - between January 2004 and January 2007 and the final stake of the bidder is at least 50 per cent. The source of information was the Dealogic M&A Global data files. The choice of period was severely restricted by the fact that the companies had to be matched with ownership information from other sources. The sample thus selected involves 226 transactions, which conveniently allows for the individual tracking of the companies and only requires ownership information that are not older than 3-4 years. Dealogic collects a vast amount of information about these transactions including descriptive (name, industry classification, nationality of the companies involved), performance related (sales revenue, operating profit, financial ratios) and deal related information (announcement date, bid value, presence of competing bidders, announcement returns, etc.). Summary statistics for the most relevant variables are provided in Appendix 1.

It is interesting to see whether there is any notable difference in the characteristics of the targets selected by the two groups of bidders considered in this study. Therefore, I first classify bidders into groups of corporate and financial investors. Management buy-outs (MBOs) and financial buy-outs (FBOs) belong to the financial investor category. Additionally, any deal where the bidder is a financial investor, financial advisor, private equity or private investor and

the target is an industrial corporation is grouped to the financial investor category. In contrast, if a bidder operates in the financial services sector and acquires another such firm, the deal is considered as a corporate takeover. This is because although the bidder is technically a financial investor, the potential gains from the deal (in particular, synergy gains) are expected to resemble more to those of corporate takeovers. The sample constructed this way contains 144 corporate mergers and 83 financial investments.

Table 1 shows that while the number of private equity takeovers have grown steadily during 2004-2006, corporate investors were more active in 2005 than in the following year. Only three deals that met the above criteria were announced in the first month of 2007, according to Dealogic.

Table 2 shows that the most active industry in M&As is "Information, finance and insurance, and real estate". Many targets of private equity and the vast majority of the firms acquired by corporate corporations operate in this industry.⁷ Other desirable target industries are trading, transportation and manufacturing. As for the bidders, both groups operate predominantly in the information, finance and insurance, and real estate industry. It can be inferred from Table 2 that corporate investors tend to acquire firms that operate in their own industry; this is less the case for financial investors. The observed pattern supports the argument that real synergy gains are more relevant for corporations than for private equity bidders.

All of the targets are public companies that allows for the comparison of takeover premia calculated from their share prices. In contrast, as can be seen in Table 3, bidders include a significant number of private companies. About 43 per cent of corporate investors are private firms, while, as expected, only 10 per cent of financial investors are publicly held.

Previous studies have emphasized the role of competing bidders in driving up the premium paid for target companies. Dealogic reports that in only 3 out of 181 transactions - there is no information about competitors in the remaining 53 cases - were rivals present during the bidding process (Table 4). Interestingly, these 3 cases belong to private equity bidders. The low number of actual competitors suggest that the distribution of gains may be governed by some factors other than competition.

⁷Note that this industry classification - i.e. the first digit of the New American Industry Classification System (NAICS) - is quite broad. Narrower classifications will be used during the regression analysis.

Another variable that researchers have found to be important in explaining the variation in takeover premia is the method of payment. Table 5 shows that 58 per cent of corporate takeovers were settled using only cash, but pure share and combined share and cash financing are also regular. In contrast, 94 per cent of the financial deals are recorded as pure cash financing. As private equity firms are well known for financing their takeovers by mainly debt, this might reflect a problem of observability. Indeed, there are only two cash/debt cases reported in the full sample, and they belong to industrial deals.

Table 6 reports the statistics of the target characteristics and some deal related information. The data reveal a number of surprises. First, in my sample financial investors acquired somewhat larger companies than corporate investors. Although it has been argued that private equity firms now possess a purchasing power comparable to that of public corporations, one would still expect that on average they target smaller firms. Second, the profitability of private equities' targets is better than that of industrial investors. Private equity has been seen by many as an important means for restructuring underperforming businesses. However, this sample suggests that they target fundamentally strong businesses according with recent research summarized by Thornton (2007).

Targets' liquidity ratios are roughly the same in the two types of deals, while financial investors seem to favor lower gearing ratios. The latter is what we expect if private equity firms aim at leveraging up their targets.

One criterion in selecting the sample was that bidders' final stake was to exceed 51 per cent. The reasons for this were ownership consideration - 51 per cent constitutes majority shareholding which entitles the shareholder to change board, hire and fire management, etc. - and the expected abnormal return in share prices. As a result, a decisive majority of the cases involve 100 per cent final stake. Also, the acquired stake is 100 per cent in a very high number of cases, but the sample includes transactions where the acquired stake can be as low as 6.63 per cent. These are outliers from the perspective of this study (no significant abnormal returns can be expected in these cases) and need to be taken care of during the analysis.

Dealogic also collects share prices for the targets on various dates around the acquisition date and computes 1-month, 1-week and 1-day takeover premia from these information. Summary statistics of these variables and the mean tests are reported in the next section.

3.2 Ownership structure

Although Dealogic collects a vast amount of information on mergers and acquisitions, they do not report the ownership structure of the companies engaged in the transaction. The information was collected from the FAME dataset that collects ownership information on UK companies from three different sources: BVD, the annual return and the registry. The information also includes shareholder history going back at least a few years in the cases considered here. However, it appears that a decreasing number of shareholders are reported for more dated transactions, which may affect the analysis. But the distribution of deals over time is very similar for the two types of investors, therefore this bias is not likely to affect the two groups differently.

Following the literature (e.g. Leech and Leahy, 1991) concentration ratios were constructed to capture the ownership structures of targets and bidders. These are the holding of the major shareholder and the sum of holdings of the three major shareholders. The ten-shareholder-concentration figure was dismissed because for a significant number of companies information was only to be found on at most three owners. Table 7 reports the summary statistics and distribution of these variables.

Targets' ownership structures appear to be less concentrated than those of the bidders in both types of transactions. The difference is however much larger in case of financial transactions. This supports the argument that agency improvements in private equity transactions can be expected to be much larger than in corporate takeovers. The ownership change resulting from the transaction, calculated as the difference between the bidder's and the target's holdings shows that targets in this sample became more concentrated after the acquisition. The mean of the change is only 2 per cent for corporate investors but it is much larger, 34 per cent, for financial investors. Still, it appears that private equity firms do not target especially dispersed shareholder structures. In fact, the mean ownership structure of private equity's targets is somewhat more concentrated than that of corporate targets.

Surprisingly, for about 15 targets FAME reports 100 per cent holding for 1 owner immediately preceding the transaction. This is despite the fact the only public, quoted targets were included in the sample. Although this suggests some inconsistencies between the Dealogic and

the FAME datasets, the constructed variables can still be good proxies for the agency situation of these firms.

3.3 Empirical analysis

To test the claim that financial bidders pay lower premia than corporate bidders I compare the targets' average cumulative returns reported by Dealogic. The reported measures can be considered as a proxy for the abnormal returns usually estimated by previous studies. These studies use the so-called market model to estimate the expected return conditioning on the merger not taking place. The market model relates the targets' share price returns during a few months preceding the announcement to the movement of the stock market index. The abnormal returns are then calculated as the difference in the actual returns observed at the announcement dates and the expected return from the share on that day as predicted by the market model. For completeness, I have calculated the abnormal returns from the actual share prices of the targets using this method.⁸ The abnormal returns using the market model are very close to the returns reported by Dealogic, therefore in the followings I work with the Dealogic series. The fact that there is so little difference between the cumulative returns reported by Dealogic and those calculated using the market model is partly due to the substantial fraction of the companies in the dataset that are illiquid. There are only a few changes in the share prices of these companies over the course of several months. Therefore there is no significant relationship between the movements of the share prices and the movements of the market index. The cumulative abnormal returns for these companies is simply a summation of the changes in their share prices during a selected period preceding the announcement date, and this is exactly what Dealogic reports.

Dealogic provides 1-month premia estimates for 204 targets in the sample. Three of these are larger than 100 per cent and are eliminated from the analysis.⁹ The remaining 201 targets earned 21 per cent cumulative return on average during the month preceding the merger an-

⁸The share prices of the targets were collected from Datastream. The results are available from the author on request.

⁹Targets in these three cases gained 194 per cent, 174 per cent and 142 per cent abnormal returns. The first two figures belong to industrial takeovers, while the last to a private equity transaction.

nouncement (Table 8). The corresponding figure for corporate investors is 23 per cent, while that for financial investors is substantially lower at 16 per cent. A t-test for the difference in the means of the two groups reported in the Appendix confirms that the gap is significant at the 5 per cent level. The result on the premia paid by corporate investors is in the range (20-30 per cent) reported by previous studies. The fact that it is closer to the lower bound could suggest that it is becoming more and more difficult to create value through takeovers as firms are constrained to operate more efficiently under increasing pressure from competition, which leaves little scope for efficiency improvement. It could also be the case, that in contrast to previous practice reported by earlier studies corporate investors are now securing part of the overall value to themselves.

Dealogic also reports 1-week and 1-day returns for the takeovers in the sample. Although the magnitude of these returns are smaller, similar conclusions can be drawn regarding the difference between the premia paid by corporate and financial investors. The last set of columns in Table 8 show that targets gained from corporate investors 20 and 16 per cent abnormal return in the one-week and 1-day periods preceding the announcement of the mergers. The corresponding figures are 16 and 12 per cent in case of financial investors. Despite the significant difference in the premia paid by the two groups, the timing of the abnormal share price increases shows similarities. For both groups, 70 per cent of the 1-month premia accrue to the target on the day of the announcement.

The fact that in my sample financial investors paid significantly lower premia than corporate investors does not necessarily mean that financial investors "play an unfair game". The finding could be simply explained if, for example, takeovers executed by private equity firms generated less value than mergers between two large corporations. In a regression analysis I control for the previously discussed determinants of takeover premia and different target characteristics to shed light on this important question.

The inclusion of private companies in a merger analysis has the drawback that measurement error in the information about them or lack of data complicates the analysis. In the following regressions this translates into limited controls for bidder characteristics. The specification

that I finally use is:

$$PR = \beta_1 + \beta_2 * FIT + \beta_3 * T.Liquidity + \beta_4 * Diff.Owner.Str.(B-T) + \beta_5 * Fin.Inv + \beta_6 * T.PMargin + \beta_7 T.Empl. + \beta_8 * T.Illiquid.Stock + \beta_9 * AllCash + \epsilon \quad (2)$$

It is assumed that ϵ is identically, normally distributed. The regression is run using OLS and robust standard errors. The primary coefficient of interest is β_5 which shows the difference between the premia paid by financial investors and corporate investors once different characteristics of the deal and target are controlled for. It is expected that relatedness of the activities of the two firms (Fit) and target's size ($T.Empl$) increase takeover premia due to larger potentials for real synergies. Potential financial synergies may be picked up by the target's liquidity ratio, therefore it's coefficient is also expected to be positive. Similarly, previous studies (Huang and Walkling, 1987; and Slusky and Caves, 1991) observed larger premia when the transactions were financed by cash only ($All.Cash$). Visual inspection of the data suggests that changes in the share prices of illiquid targets ($T.Illiquid.Stock$) are larger than changes in those of their liquid counterparts. On the contrary, more profitable targets ($T.PMargin$), *ceteris paribus*, are expected to reap less premia because there is less scope for improving their profitability.

As explained earlier, no sign can be predicted for the effect of the change in the ownership structure captured by β_4 in equation (2). Note, that this variable was transformed into a dummy because of the observed impreciseness of the ownership data. The dummy takes on the value 1 if the change is larger than 30 per cent and takes on the value (-)1 if the change is larger than (-) 30 per cent¹⁰ in absolute terms. A positive change means that the new structure is more concentrated than the one preceding the takeover. The principal-agent literature suggests that this would reduce agency costs thereby increasing the value created by the merger as well as the target's premium. At the same time, the bidder's management would be more efficient in the bargaining phase thus would secure a larger share of the value created to themselves. This would have the opposite effect on the target's premium.

¹⁰The robustness of the results with the dummy is also tested using the original variable (see Robustness tests).

4 Results and robustness tests

4.1 Full sample

Table 9 reports the results for the regressions using the 1-month premia as the dependent variable. The results indicate that even after controlling for the assumed sources of takeover premia and various characteristics of the target and the deal, the premia gap is significant at the 10 per cent level. The magnitude (7 per cent) is the same as that suggested by the simple mean test. Hence, differences in the previously identified sources of the takeover premia do not seem to be the reason for the significantly lower premia paid by financial investors.

The proxy for larger synergies has a negative coefficient which is not significant. This disappointing finding accords well with that of Slusky and Caves (1991). It could be the case that the Fit variable does not capture the relevant industry classification, which will be checked by the robustness tests. This is plausible because, in line with the expectations, targets' size (the other synergy related variable) has a positive effect on the premia which is significant at the 10 per cent level. An increase of 10 thousand in the target's employees increases the premium by 6 percentage point.

Target's liquidity ratio turns out to be the most significant variable in the regression, suggesting that either financial synergies or the link between liquidity and firm growth (Evans and Jovanovic, 1989) is factored into the takeover premia. The effect is substantial, a 1 percentage point increase in the liquidity ratio increases the premium by 1.4 percentage point.

The coefficient on the targets' profitability has a negative sign, in line with the reasoning that there is less scope to add value to firms that are already profitable, however, the effect is not significant. Similarly, the ownership change has a non-significant, negative effect on the premia, that may suggest that the bargaining power effect offsets the agency improvement effect.

Transactions that are financed by cash only result in 9 percentage points higher target premia (significant at the 10 per cent level). However, the dummy seems to capture the difference between cash and share financing (or a combination of the two), and appears to be unaffected by the debt component of the payment that can be excessively high in private equity takeovers.

As evidenced by the visual inspection of the data, the share price increases of targets with illiquid stocks are 6 percentage points higher (significant at the 10 per cent level) than those of liquid companies. The result is reasonable given that investors typically discount the value of companies the investment in which is difficult to terminate. The findings with the 1-week and the 1-day abnormal returns - shown in Table 10 and Table 11 of Appendix 2 - are qualitatively the same with some differences in the significance of the coefficients.

4.2 Sub-samples

Overall, it appears that the difference in takeover premia paid by corporate and financial investors can not be explained by differences in the sources of takeover premia identified by previous studies. I analyze the two groups separately to shed further light on this finding. The last two columns of Table 9 in Appendix 2 display the results of the regression when the sample is restricted first to corporate investors and then to financial investors.

It can be seen that the model fits the data of corporate investors much better than those of financial investors. In particular, for corporate investors most of the variables are strongly significant. The ownership change becomes a strong explanatory variable with a negative coefficient. This suggests that the bargaining effect of the difference in ownership structure more than offsets its agency effect. All of the variables have the same effect as before splitting the sample albeit with larger magnitudes.

Interestingly, most of the coefficients change signs compared to the previous regression, except for the targets' profit margin, when I restrict the sample to the financial investors. Note, however, that only the targets' size is significant in this regression. A smaller sample size - 65 takeovers have complete data necessary to be included in the analysis - may partly account for the insignificance of the variables. The negative coefficient on the targets' size suggests that private equity firms pay less premia to the shareholders of larger corporations. As private equity firms have only recently entered the arena of multibillion dollar deals this could reflect the increased cost of securing the bid value, the loss of diversification possibility when investing a large part of their fund in only one company and their caution when valuing such large corporations. But then it is surprising that on average they target larger firms than

corporate investors.

Relatedness has a positive coefficient which would be in line with the synergy gain argument. However, even if we ignore that the coefficient is insignificant, the result is dubious, because, intuitively, real synergy gains appear less of a reason for this type of takeovers. At the same time, the liquidity ratio of the target - that is meant to capture financial synergies - does not seem to have any effect on the takeover premia in this regression. Other financial variables, such as gearing ratio, may be better in explaining the variation in takeover premia paid by financial bidders.

Interestingly, the improvement in the ownership structure of the target has a positive coefficient in this regression. Although it is again insignificant, it is not unreasonable to believe that the reduction in agency costs is so much larger in private equity takeovers than in corporate takeovers that its effect offsets that of the bargaining considerations.

Finally, the variables controlling for the bid's financing structure and the liquidity of the target's stock have no effect on the takeover premia. This is unsurprising, though, because an overwhelming majority of these deals are recorded to be financed by only cash and most of the targets of private equity have illiquid stocks.

Apart from the smaller sample size, the worse performance of the model in case of financial investors could also be due to measurement error in the data of financial bidders. However, the measurement error mostly affects the ownership variable because the rest of the information included in the analysis comes from the public targets. Overall, it appears that the premia gap can not be explained by differences in the previously identified sources of takeover premia and in the respective bargaining power of targets and bidders. The instability of the model's coefficients when the regressions are run separately for corporate and financial investors suggests that takeovers by the latter are motivated by different factors.

4.3 Robustness tests

Several variables could play a part in merger pricing that have not been taken into account in the regressions presented above. The primary reason for not including more variables in the basic regressions is that there are many missing values in the dataset, therefore the inclusion

of every new control restricts the sample size. Nevertheless, it is important to use them in the robustness tests.

Several financial variables have been linked to firm growth by different studies. The gearing ratio (debt/assets), for example, has been reported to have a negative effect on growth (Lang et al., 1996, Aivazian et al., 2005). This variable is also of particular interest due to private equity firms better access to debt financing. For both reasons it would be expected that the higher gearing ratio of the target decreases the takeover premium. Table 12 however shows that the gearing ratio has no explanatory variable in any of the three regressions (full sample, corporate takeovers, financial takeovers).

As a further robustness check I analyze whether instead of the corporate vs. financial classification it is the public vs. private classification that matters. The bidders in the group of financial investors are mainly private companies while those in the corporate group include public and private companies as well. I use the subgroup of corporate investors and include a dummy for the public status to infer whether the premia gap exists in the public-private classification. The results in Table 13 suggest that public companies pay about 7 percentage points higher premia than private companies although the coefficient is not significant. The striking similarity between the coefficient of the public-private status and that of the financial investor dummy suggests that it can not be excluded that not only private equity firms but private companies in general pay lower premia than their public counterparts. Differences in the previously identified sources of takeover premia and in the bargaining power of the two parties do not seem to explain this premia gap.

Finally I test whether different construction of some variables change the conclusions of the basic regressions. For example, the disappointing performance of the dummy for real synergies may be the result of the classification of firms into the same industry. The first digit of the industry classification used for the construction of Fit may reflect a category that is too large to identify potential synergy gains between targets and bidders. In Table 14 I show the results of the regressions with dummies that take on the value 1 if the first 2, 3, 4 and 5 digits of the industry classification codes of the two firms are the same. The results suggest that the coarser the industry classification, the less significant the Fit variable, therefore the previously found negative effect with the largest group classification may be misleading. Real synergy gains, as

expected from the merger of two companies with the same activities, do not seem to be priced in the takeover premia.

Another constructed variable that may need some further tests is the ownership change dummy. In order to have a more precise idea of its effect in Table 15 I show the results of the regressions with the original variable, i.e. the largest shareholder's stake of the bidder minus the largest shareholder's stake of the target. Results in the first column show that a 1 percentage point increase in this difference decreases the takeover premia by 0.05 percentage point. The decrease is 3 times larger, i.e. 0.15 percentage point, when we consider only corporate investors. For financial investors, the coefficient indicates a positive effect (0.05 percentage point), which is insignificant.

As previously mentioned, my sample includes all takeovers where the final stake of the acquirer is larger than 50 per cent. This means that the acquired stake can be as low as 7 per cent. In table 16 I show that the conclusions from the basic regressions do not change when I restrict the sample to deals where the acquired stake is at least 50 per cent.

5 Conclusion

This paper has compared takeover premia in traditional corporate and financial (predominantly private equity) takeovers to analyze the claim that private equity firms "cheat shareholders out of a fair price". In 226 mergers and acquisitions during 2004-2007 in the UK corporate investors have paid a significantly larger (on average 23 per cent) takeover premia than financial investors (on average 16 per cent). The 23 per cent premia found for corporate investors is in the range that emerged from field surveys, while the 16 per cent paid by financial investors is much lower than the 30-50 per cent reported by previous studies for the 1970s and 1980s. In parallel with the growth of private equity funds it seems that the value created by private equity takeovers has substantially fallen. The data do not reveal important differences either in the characteristics of the targets of the two types of investors or the competitive environment in which these transactions have been conducted. The results of the regression analysis suggest that there is no significant difference in the primary sources of value creation and distribution in the two types of deals - agency and synergy gains and bargaining power - that could explain

the premia gap. The gap could be explained if the cost of private equity deals were much larger than those of traditional takeovers, however, this would go against the evidence that private equity firms have become a major player on the global M&A market.

The results do not reject the hypothesis that private equity firms may use their target's management to buy the companies under value. However, the empirical model's coefficients are not robust to splitting the data by investor type. The model has a better fit when only traditional corporate mergers are considered and gives opposite predictions for the takeover premia in the two types of investments. This could be the result of the smaller and inherently less precise dataset available for financial investors, or could suggest different merger pricing behavior for the two types of investors. It can not be excluded, though, that the difference is not between corporate and financial takeovers, but lies in the public vs. private classification.

References

Aivazian, V. A., Ge, Y., and J. Qiu (2005): The impact of leverage on firm investment: Canadian evidence, *Journal of Corporate Finance*, 11, 277-291.

Andrade, G., Mitchell, M., and E. Stafford (2001): New evidence and perspectives on mergers, *Journal of Economic Perspectives*, 15, 103-120.

Ang, J. S., Cole, R. A. and J. W. Lin (2000): Agency costs and ownership structure, *Journal of Finance*, 55, 29-52

Bradley, M., Desai, A. and E. H. Kim (1988): Synergistic gains form corporate acquisitions and their division between the stockholders of target and acquiring firms, *Journal of Financial Economics*, 21, 3-40.

Bruner, R. F. (2002): Does M&A pay? A survey of evidence for the decision maker, *Journal of Applied Finance*, 12, 48-68.

Devos, E., Kadapakkam, P-R. and S. Krishnamurthy (2004): Are There Synergy Gains in Mergers?

Evans, D., and B. Jovanovic (1989): An estimated model of entrepreneurial choice under liquidity constraint, *Journal of Political Economy*, 97, 808-827.

Gondhalekar, V. B., Sant, R. R., and S. P. Ferris (2004): The price of corporate acquisition:

determinants of cash takeover premia, *Applied Economics Letters*, 11, 735-739.

Huang, Y.-S. and R. A. Walkling (1987): Target abnormal Returns Associated with Acquisition Announcements: Payment, Acquisition Form, and Managerial Resistance, *Journal of Financial Economics*, 19, 329-349.

Jensen, M. C. and R. S. Ruback (1983): The market for corporate control: The scientific evidence, *Journal of Financial Economics* 11, 5-50.

Jensen, M. C. and W. Meckling (1976): Theory of the firm: managerial behavior, agency costs and ownership structure, *Journal of Financial Economics*, 3, 305-360.

Jovanovic, B. and S. Braguinsky (2004): Bidder Discounts and Target Premia in Takeovers, *American Economic Review*, 94, 46-56.

Kaplan, S. (1989): The effects of management buy-outs on operating performance and value. *Journal of Financial Economics*, 24, 217-254.

Lang, L., Ofek, E. and R. Stulz (1996): Leverage, investment, and firm growth, *Journal of Financial Economics*, 40, 3-30.

Leech, D. and J. Leahy (1991): Ownership Structure, Control Type Classification and the Performance of Large British Companies, *Economic Journal*, 101, 1418-1437.

Loughran, T. and A. M. Vijh (1997): Do long-term shareholders benefit from corporate acquisitions?, *Journal of Finance*, 52, 1765-1790.

Matsusaka, John G. (1993): Takeover motives during the conglomerate merger wave, *Rand Journal of Economics*, 24, 357-379.

Morck, R., Shleifer, A. and R. Vishny (1988): Management ownership and market valuation, *Journal of Financial Economics*, 20, 293-315.

Nagar, V., Petroni, K. and D. Wolfenzon (2002): Ownership Structure and Firm Performance in Closely-Held Corporations, NYU working paper.

Pagano, M., Panetta, F. and L. Zingales (1998): Why do Companies Go Public? An Empirical Analysis, *Journal of Finance*, 53, 27-64.

Shelton, L.M. (1988): Strategic Business Fits and Corporate Acquisition: Empirical Evidence, *Strategic Management Journal*, 9, 279-287.

Smit, H. T. J. and W. De Maeseneire (2005): The Role of Investor Capabilities in Public-to-Private Transactions, Universiteit Gent, Working Paper.

Smith, A. (1990): Corporate ownership structure and performance: the case of management buyouts, *Journal of Financial Economics*, 27, 143-164.

Slusky, A. R. and R. E. Caves (1991): Synergy, Agency, and the Determinants of Premia Paid in Mergers, *The Journal of Industrial Economics*, 39, 277-296.

The Economist, February 10th - 16th, 2007, 81-83.

Thornton, P. (2007): Inside the Dark Box: Shedding Light on Private Equity, The Work Foundation, www.theworkfoundation.com

Appendix I – Summary statistics

Table 1: Distribution of transactions by year of announcement

	Industrial investor	Financial investor
2004	36	25
2005	61	25
2006	46	30
2007	1	2
Total	144	82

Table 2: Distribution of transactions by industry classification

	Industrial investor		Financial investor	
	Target	Bidder	Target	Bidder
1. Agriculture	1	3	-	-
2. Mining, utilities, construction	12	12	6	2
3. Manufacturing	20	10	13	3
4. Wholesale and retail trade, transportation	14	17	16	2
5. Information, finance & insurance, real estate	89	91	29	68
6. Educational services, health care	-	-	2	-
7. Arts, entertainment, accommodation etc.	5	10	16	7
8. Other services	3	1	-	-
9. Public administration				
Total	144	144	82	82

Table 3: Distribution of transactions by bidder's public status

	Industrial investor	Financial investor
Public	99	7
Private	43	68
N/A	2	7
Total	144	82

Table 4: Distribution of transactions by presence of competing bidders

	Industrial investor	Financial investor
No	116	55
Yes	-	3
N/A	28	24
Total	144	82

Table 5: Distribution of transactions by primary payment method

	Industrial investor	Financial investor
Ca/S	3	-
Cash	84	77
Cash/Asld	1	-
Cash/Asld/Cond	1	-
Cash/Debt	2	-
Cash/Notes	1	-
Cash, Share	10	1
Notes	-	1
Share	34	2
Share, Cash	8	1
Total	144	82

Table 6: Summary statistics of selected variables

Group	Industrial investor			Financial investor		
	Observations	Mean	Std. dev.	Observations	Mean	Std. dev.
Target employees	128	1707	5545	77	1954	4265
Target profit	122	-0.40	23.27	73	6.31	18.51
Target liquidity	123	2.06	3.85	69	1.91	4.53
Target gearing	113	145.82	505.52	65	127.38	302.47
Bid value GBP	143	221.48	606.23	78	207.85	320.16
Acquired stake	143	92.84	17.62	81	87.24	22.15
Final stake	143	98.37	7.52	81	98.12	8.64

Table 7: Summary statistics of ownership variables

Group	Industrial investor			Financial investor		
	Observations	Mean	Std. dev.	Observations	Mean	Std. dev.
Target largest owner's share	144	29.51	25.53	82	44.39	35.72
Target 3 largest owners' share	144	42.59	26.44	82	56.96	30.39
Bidder largest owner's share	144	31.13	29.91	82	78.28	31.68
Bidder 3 largest owners' share	144	43.63	29.82	82	89.97	22.48
Ownership change B-T (largest)	144	1.80	33.23	82	33.89	39.77
Ownership change B-T (3 largest)	144	1.03	33.54	82	33.01	35.14

Appendix II – Tests and Regression results

Table 8: T-tests for differences in takeover premia paid by financial vs. industrial investors (outliers, i.e. larger than 100 per cent takeover premia, dropped from the sample)

Group	1-month premium		1-week premium		1-day premium	
	Observations	Mean (std. error)	Observations	Mean (std. error)	Observations	Mean (std. error)
Financial	71	16.08 (2.44)	71	12.90 (2.96)	71	11.35 (2.32)
Industrial	130	23.41 (2.36)	130	20.45 (2.25)	130	16.06 (2.28)
Combined	201	20.82 (1.77)	201	17.78 (1.81)	201	14.40 (1.69)
Difference		-7.34 (3.67)		-7.55 (3.75)		-4.71 (3.53)
H0: difference = 0, Ha: difference < 0						
T	-1.9980		-2.0134			-1.3331
P < t	0.0235		0.0227			0.0920

Regressions with ownership change dummy

Table 9: Regression with 1-month premium*

1-month premium	Full sample	Industrial investors	Financial investors
Fit dummy	-5.34 (4.20)	-10.67 (6.09)	2.29 (5.56)
Target liquidity (%)	1.42 (0.34)	1.51 (0.35)	0.23 (3.25)
Target employees (10.000)	6.11 (3.47)	12.02 (4.06)	-7.58 (2.60)
Target profit margin (%)	-0.16 (0.11)	-0.15 (0.14)	-0.18 (0.18)
Financial investor dummy	-6.91 (4.13)	-	-
Bidder – target shares of largest owner dummy	-4.41 (2.86)	-9.27 (3.07)	2.76 (4.52)
All cash dummy	11.83 (5.33)	15.91 (2.79)	-8.48 (10.97)
Illiquid stock dummy	6.68 (4.18)	10.21 (5.32)	-3.62 (4.92)
Constant	11.53 (6.68)	9.68 (8.35)	29.08 (12.53)
Observation number	161	103	58
R ²	0.1066	0.1957	0.0866

* Dependent variable: 1-month takeover premium. Robust standard errors in parenthesis.

Table 10: 1-week premium*

1-week premium	Full sample	Industrial investors	Financial investors
Fit dummy	-5.56 (4.41)	-8.31 (6.18)	-0.60 (6.75)
Target liquidity (%)	2.00 (0.48)	2.03 (0.50)	1.04 (3.71)
Target employees (10.000)	4.71 (3.29)	8.24 (4.31)	-4.18 (2.68)
Target profit margin (%)	-0.21 (0.10)	-0.18 (0.11)	-0.31 (0.24)
Financial investor dummy	-3.67 (4.41)	-	-
Bidder – target shares of largest owner dummy	-2.90 (2.86)	-4.30 (3.22)	-1.71 (5.70)
All cash dummy	2.91 (5.24)	5.71 (5.70)	-13.78 (11.83)
Illiquid stock dummy	2.30 (4.29)	6.37 (5.39)	-11.44 (6.21)
Constant	15.81 (7.11)	12.73 (8.75)	40.03 (14.08)
Observation number	161	103	58
R ²	0.1082	0.1499	0.1031

* Dependent variable: 1-week takeover premium. Robust standard errors in parenthesis.

Table 11: 1-day premium*

1-day premium	Full sample	Industrial investors	Financial investors
Fit dummy	-3.40 (4.93)	-7.46 (7.84)	3.40 (5.44)
Target liquidity (%)	2.28 (0.54)	2.32 (0.58)	1.49 (3.18)
Target employees (10.000)	2.33 (3.43)	5.75 (4.62)	-4.27 (2.08)
Target profit margin (%)	-0.21 (0.08)	-0.18 (0.09)	-0.32 (0.21)
Financial investor dummy	-1.96 (4.73)	-	-
Bidder – target shares of largest owner dummy	-2.39 (2.41)	-4.83 (2.90)	-1.48 (4.08)
All cash dummy	3.73 (5.45)	7.06 (5.78)	-18.37 (13.52)
Illiquid stock dummy	6.34 (4.23)	7.86 (5.62)	1.35 (3.75)
Constant	7.04 (6.63)	6.43 (8.47)	27.57 (13.91)
Observation number	161	103	58
R ²	0.1177	0.1432	0.2073

* Dependent variable: 1-day takeover premium. Robust standard errors in parenthesis.

Appendix III – Robustness tests

Table 12: Gearing variable*

1-month premium	Full sample	Industrial investors	Financial investors
Fit dummy	-5.81 (4.32)	-12.66 (6.28)	2.66 (5.83)
Target liquidity (%)	1.49 (0.34)	1.49 (0.39)	0.31 (3.49)
Target employees (10.000)	5.82 (3.59)	14.15 (5.97)	-6.96 (2.93)
Target profit margin (%)	-0.21 (0.12)	-0.21 (0.16)	-0.18 (0.19)
Financial investor dummy	-6.94 (4.23)	-	-
Bidder – target shares of largest owner (%)	-3.3 (2.93)	-8.05 (3.37)	2.80 (4.57)
All cash dummy	11.76 (5.44)	16.79 (5.83)	-8.42 (11.57)
Illiquid stock dummy	6.56 (4.45)	9.52 (5.43)	-1.87 (5.25)
Target gearing	0.01 (0.01)	-0.01 (0.02)	0.00 (0.00)
Constant	11.55 (7.06)	11.16 (8.72)	26.53 (13.46)
Observation number	152	95	57
R ²	0.1255	0.2336	0.0920

* Dependent variable: 1-month takeover premium. Robust standard errors in parenthesis.

Table 13: Bidder public-private status*

1-month premium	Industrial investors
Fit dummy	-14.33 (7.08)
Target liquidity (%)	1.52 (0.32)
Target employees (10.000)	11.11 (3.97)
Target profit margin (%)	-0.27 (0.17)
Financial investor dummy	-
Bidder – target shares of largest owner dummy	-11.44 (4.45)
All cash dummy	15.75 (6.64)
Illiquid stock dummy	11.61 (5.50)
Public status dummy	6.55 (6.54)
Constant	8.54 (12.63)
Observation number	104
R ²	0.2286

* Dependent variable: 1-month takeover premium. Robust standard errors in parenthesis.

Table 14: Ownership change variable*

1-month premium	Full sample	Industrial investors	Financial investors
Fit dummy	-5.52 (4.22)	-10.45 (6.18)	2.55 (5.35)
Target liquidity (%)	1.41 (0.35)	1.51 (0.34)	0.18 (3.19)
Target employees (10.000)	5.99 (3.52)	12.36 (3.84)	-7.27 (2.51)
Target profit margin (%)	-0.16 (0.11)	-0.16 (0.13)	-0.19 (0.18)
Financial investor dummy	-7.14 (4.25)	-	-
Bidder – target shares of largest owner (%)	-0.05 (0.04)	-0.15 (0.05)	0.05 (0.06)
All cash dummy	11.60 (5.31)	15.98 (5.64)	-8.38 (10.94)
Illiquid stock dummy	6.65 (4.17)	10.55 (5.28)	-3.57 (4.92)
Constant	11.76 (6.65)	9.12 (8.26)	28.54 (12.51)
Observation number	161	103	58
R ²	0.1043	0.1993	0.0908

* Dependent variable: 1-month takeover premium. Robust standard errors in parenthesis.

Table 15: Fit variable*

1-month premium	Fit – 2 digits	Fit – 3 digits	Fit – 4 digits	Fit – 5 digits
Fit dummy	-1.94 (4.03)	-3.81 (4.19)	-3.69 (4.29)	-5.27 (4.37)
Target liquidity (%)	1.40 (0.35)	1.41 (0.36)	1.42 (0.36)	1.45 (0.36)
Target employees (10.000)	6.07 (3.44)	5.96 (3.32)	5.65 (3.39)	5.68 (3.35)
Target profit margin (%)	-0.16 (0.11)	-0.16 (0.11)	-0.16 (0.11)	-0.16 (0.11)
Financial investor dummy	-5.89 (4.21)	-6.14 (4.25)	-5.94 (4.14)	-6.13 (4.07)
Bidder – target shares of largest owner (%)	-4.70 (2.88)	-4.20 (2.99)	-4.29 (3.01)	-4.03 (2.97)
All cash dummy	11.88 (5.37)	11.30 (5.39)	11.60 (5.35)	11.58 (5.32)
Illiquid stock dummy	6.62 (4.19)	6.61 (4.23)	6.46 (4.25)	6.88 (4.19)
Constant	8.59 (6.55)	9.46 (6.51)	9.13 (6.25)	9.06 (6.01)
Observation number	161	161	161	161
R ²	0.0984	0.1020	0.1014	0.1053

* Dependent variable: 1-month takeover premium. Robust standard errors in parenthesis.

Table 16: Acquired stake*

1-month premium	Full sample	Industrial investors	Financial investors
Fit dummy	-4.82 (4.08)	-9.59 (5.74)	2.55 (5.76)
Target liquidity (%)	1.43 (0.34)	1.53 (0.34)	0.09 (3.30)
Target employees (10.000)	6.58 (3.55)	11.96 (4.05)	-7.70 (2.80)
Target profit margin (%)	-0.15 (0.11)	-0.14 (0.14)	-0.18 (0.18)
Financial investor dummy	-5.99 (3.97)	-	-
Bidder – target shares of largest owner dummy	-4.21 (2.87)	-9.28 (3.07)	2.93 (4.66)
All cash dummy	11.34 (5.27)	15.47 (5.65)	-8.42 (11.12)
Illiquid stock dummy	6.47 (4.13)	9.78 (5.24)	-3.49 (5.23)
Constant	11.06 (6.67)	9.16 (8.27)	28.99 (12.80)
Observation number	157	101	56
R ²	0.1049	0.1947	0.0838

* Dependent variable: 1-month takeover premium. Robust standard errors in parenthesis.

Chapter 4

Rebate or bait? A model of regret and time inconsistency in consumer behavior*

July 14, 2007

Abstract

In this paper we develop a theory of time inconsistency and regret that is motivated by evidence on a "price discrimination" technique widespread in the United States, namely mail-in-rebate promotions. Our model combines partial naivete about future self-control problems and the sunk-cost effect (regret). We assume that agents deviating from their past choices suffer a certain emotional disutility from having brought a bad decision in the past and that this emotional disutility is negatively related to the length of the period between the choice made and the deviation from it. In the context of our application the model explains why in a multi period setting a large number of consumers respond to the rebate offers intending to redeem the rebate and then fail to provide the necessary effort when it comes to collect their money. Moreover, consumer failure to accomplish a task planned in the past (e.g. redeeming the rebate) is more likely when the deadline of completion is longer. This prediction is supported by experimental studies on various forms of procrastination and by field and experimental evidence on mail-in-rebates. We review a number of areas for which the theory may have important implications.

*This Chapter is coauthored with Francesco Drago.

1 Introduction

Consumers repeatedly make mistakes, get frustrated and feel regret over their choices. They fear paying too much and they particularly dislike the thought that they could be better off having made a different choice in the past. Their concern is not at all unfounded as several recent studies demonstrate that consumers could be better off by making wiser choices: i.e. paying for their gym usage on a pay-per-visit basis instead of subscribing for a membership, avoiding expensive add-ons by learning firms' pricing strategies or ignoring irrelevant information when placing a bid at E-bay's auctions (DellaVigna and Malmendier, 2006; Gabaix and Laibson, 2006; Ahlee and Malmendier, 2005). These are just a few examples that belong to an emerging field recently termed as "the economics of paying too much".

A few studies examined how rational firms respond to deviations from standard preferences and biases in decision making. Membership charges in health clubs and subscription fees in the magazine market have found to be particularly profitable devices when consumers have present biased preferences. Our study adds a new item, mail-in-rebate promotions, to the toolkit applied by profit-maximizing firms as a response to such biases. We argue that the evidence on mail-in-rebate promotions can hardly be explained by standard theories and offer a theory of regret and time inconsistency that accounts for the observed consumer behavior.

In this paper mail-in-rebates are defined as "a money-refund offer available to consumers who mail in a proof-of-purchase and other forms to a manufacturer who mails back a portion of the price paid by the consumer". Ample empirical and anecdotal evidence show that a large number of consumers respond to mail-in-rebate promotions fully intending to get back part of the purchase price but fail to exert the necessary effort to collect their money subsequently. "Manufacturers love rebates because redemption rates are close to none...they get people into stores, but when it comes time to collect, few people follow through. And this is just what the manufacturer has in mind" (Greenman, 1999). Indeed, low redemption rates are among the most important factors contributing to the profitability of the rebate programs (Dhar and Hoch, 1996). This phenomenon - i.e. consumers buy the product anticipating the reward but they fail to redeem the rebate afterwards - is termed in the literature as "breakage" or "slippage" (Jolson et al. 1987; Bulkeley, 1998). A recent marketing study involving numerous

companies that frequently work with rebate programs reports a mean breakage rate over 60 percent (PMA, 2005). Similar rates were found by Jolson et al. (1987).

High breakage rates are crucial for the viability of companies that base their entire business strategy on free-after-rebate promotions. Several firms on the Internet offer free-after-rebate opportunities - including, for example, Sony and Microsoft - and according to fulfillment companies even these deals receive only 50 percent redemptions on average. Failure to redeem is especially puzzling in this case, since the prices of the free-after-rebate products are often widely inflated. A wireless phone, for example, that can be found without promotion for 70 dollars may be offered for as much as 700 dollars under a free-after-rebate program.

Even such high failure rates would not be worrisome if they were a result of consumers' uncertainty about their future costs of redemption, an uncertainty that they rationally take into account when they respond to the rebate offer. However, there are serious concerns about the validity of this argument. First, manufacturers would have no reason to choose rebates featuring delayed incentives over a coupon scheme (or over instant rebates) if consumers did not overestimate their probability of redemption. In fact, coupons would be a cheaper way to provide money refunds to consumers who have correct beliefs about their cost of redemption. Second, if consumers had uncertainty about their future redemption costs, manufacturers would need to compensate them for the probability of no redemption (as well as for the fact that the reward is delayed) in the form of higher rebate sizes. Simple calibration demonstrates that in order not to claim such significant rewards, consumers would need to have implausibly high cost realizations when it came to collect their money. (Note the contradiction with the concept behind rebate programs, i.e. to offer lower prices to consumers who have low costs of redemption.) Moreover, the above reasoning cannot explain the existence of the various strategies that manufacturers use to make redemptions harder, such as delayed rebate programs¹ as these would have an equal (detering) effect on the decision to buy the promoted product.

We provide an explanation based on present-biased preferences and naivete about future self-control that rationalizes the empirical evidence connected to rebate promotions. Our focus on present biased preferences is justified by the following reasons. First, rebates feature impor-

¹Under delayed rebate programs the necessary forms must be sent to the manufacturer no earlier than e.g. 150 days and no later than 180 days from the date of purchase.

tant characteristics, in particular they involve immediate costs (the effort to redeem) and future benefits (the face value of the rebate), that have been shown to lead to suboptimal choices in the presence of such preferences. Second, a simple hyperbolic discounting model where agents have a certain degree of naivete about future willpower can replicate well the results that some consumers will fail to redeem even if at the time of purchase they are fully convinced that they will carry through. Third, it has been demonstrated that firms effectively tailor their pricing strategies to profit from this type of consumer bias.

For example, in the health clubs surveyed by DellaVigna and Malmendier (2006) the majority of consumers subscribe for membership even if the comparatively high per-visit charge would be cheaper for most of them given their rare gym attendance. They show that present biased preferences play an important role in driving consumers opt for the ex post more expensive fee structure. When agents have $\beta - \delta$ preferences (Phelps and Pollack, 1967; Laibson, 1997) they give higher weight to the cost of exercising when the time comes to work-out as compared to when they plan their future gym attendance. While fully naive consumers are unaware of this inconsistency the partially naive and sophisticated ones take their future self-control problems into account (Akerlof, 1991; ODonoghue and Rabin, 1999, 2001).

The mechanism is very similar in the case of rebate promotions. Our present biased consumers give lower weight to the cost of redemptions when they purchase their product as compared to when the time comes to exert the effort. Due to their perfect foresight sophisticated consumers will only choose the promoted product if they truly redeem the rebate subsequently. However, naives and partially naives overestimate their future willpower and thus the attractiveness of the promotion. The higher is the difference between their perceived and their actual future willpower the higher the probability is that they fail to collect their money.

Mail-in-rebate promotions are especially interesting as firms have various opportunities to influence consumer behavior. Breakage rates can be manipulated by varying rebate values, the length of the rebate redemption periods, the extent of the effort required to redeem, just to name a few. Variation in the redemption deadlines has been shown to produce some nontrivial behavioral consequences but has not yet been formally investigated (Ho, Lim and Camerer, 2005). Silk's (2005) experimental study has found that a significantly higher proportion of consumers buy the rebated product but a significantly lower proportion succeed in redeeming the

rebates when the redemption deadline is longer. Field data also confirm the lower redemption rates under the less restrictive conditions (PMA, 2005).²

What can account for the deteriorating performance under longer deadlines? Several psychological and economic studies show that our present actions are not independent of our past choices (Thaler, 1980; Thaler and Johnson, 1990; Shefrin, 2000). The widely observed sunk cost effect essentially states that once we invested in something we are reluctant to pull out because of the loss that we will make. In other words, we regret having brought a bad decision in the past, and this regret confers an emotional disutility on the self that is considering deviating from the past decision. This may be especially relevant in the rebate context since consumers partly choose this option to prove that they are smart shoppers. The choice of not redeeming therefore brings not only a monetary but also a self-image loss.

The connection between deadlines and the sunk cost effect is provided by the findings of several researchers, which suggest that the sunk cost effect is not persistent through time (Thaler, 1999; Gourville and Soman, 1998). When the investment is still fresh in our minds, it effectively influences our subsequent choices, but it ceases to have an effect, when past payments fade from our memory. Thus, the initial investment can be thought of as a commitment device, albeit imperfect for agents with severe self-control problems and decreasingly effective for those who engage in endless procrastination.

We incorporate the sunk-cost effect in the simple quasi-hyperbolic model. In line with the empirical evidence, we assume that the sunk cost effect (or the emotional disutility from not redeeming) is decreasing over time (see section 2). In the model a quasi-hyperbolic consumer who responds to a rebate promotion has a finite number of periods to redeem, which is fixed by the deadline she faces. At each period, when she decides whether to complete the task or not, the emotional disutility that would result from no redemption turns out to be an incentive to exert the effort. We first address the problem of how the deadline affects the consumer's behavior given that she responded to the rebate promotion. We show that, for sufficiently long deadlines, a sophisticated consumer always completes the task early independently of the deadline. On the other hand an exponential (time-consistent) consumer completes the task immediately, irrespective of the deadline.

²Evidence for better performances under shorter deadlines is given in the next section.

For naives, the longer deadlines deteriorate their performance. The intuition for this result rests on the fact that a naive might unforeseeably and repeatedly delay to complete the task. At present, she delays because she believes she will accomplish the task with a tolerable delay in the future. However, she underestimates her future tolerable delay, so when the future arrives she tolerates more delay than she expected. In the framework of O'Donoghue and Rabin (2001) the lack of deadlines allows this process to continue infinitely leading to procrastination.³ In our case we first derive an important result without assuming the presence of deadline. A naive may fail to complete the task altogether, but cannot fool herself for ever. When she has delayed the task for long enough, i.e. the sunk cost effect has ceased to provide enough incentives to redeem, she realizes that she will not complete the task; at that time she suffers the emotional disutility. The presence of the deadline has an important role in the course of this process: if at the time of the last opportunity the sunk cost effect is large enough to induce the naive consumer to exert the effort, she will redeem. We show, that even with a mild degree of naivete the probability of redemption is decreasing with the deadline. We emphasize that in an exponential model the introduction of the sunk cost effect would not deliver the results we obtain with our model.

When we allow consumers to have some uncertainty about their future self-control problems, the model provides predictions about the purchase behavior that are in line with the empirical evidence. In particular, when consumers understand that with some probability they overestimate their future will-power, longer deadlines decrease the extent of emotional disutility that consumers eventually suffer when the uncertainty resolves. This in turn leads more consumers to buy the rebated product under the extended deadline condition. Hence, in the presence of uncertainty consumers try to avoid choices that involve more ex-post regret as in previous regret theories (e.g. Loomes and Sugden, 1982).

This paper contributes to two different strands of literature. First, by describing another example of how firms exploit naivete we contribute to the growing literature on behavioral industrial organization and more specifically to its subfield, "economics of paying too much". In so doing we offer an explanation for the overall pattern of rebate redemptions that has not

³In this terminology, procrastination means that the task is delayed forever, i.e. it is never completed. Delay, on the other hand, does not necessarily mean that the task will not be completed at a future point in time.

yet been explored in the literature. However, our analysis is not limited to rebate promotions. In the last section we discuss several applications where our theory has important implications. Our paper also contributes to the nascent literature on procrastination (O’Donoghue and Rabin, 2001). The latter paper predicts that if the monetary costs and benefits associated to a task do not change over time, imposing any kind of deadlines will induce all agents, even the ones with severe self-control problems, to complete the task in question. As such it does not explain why we frequently observe that people fail to accomplish what they plan even in the presence of deadlines. We provide a more realistic explanation for procrastination. In particular, procrastination needs not to be driven by persistent optimistic future beliefs since a partially naive may understand well before the deadline that she will never complete the task. This also means that, in case of partial naivete, the presence of deadlines is not sufficient for successful task completion in our model.

We proceed in Section 2 by reviewing the relevant empirical and experimental literature. In Section 3 and 4 we set up our model. Section 5 discusses several areas for which the theory has important implications and concludes.

2 Evidence

This paper is motivated by two empirical regularities connected to mail-in-rebate promotions that we mentioned in the introduction: high breakage rates and deteriorating performance under longer deadlines. In order to explain these phenomena we concentrate on time inconsistent preferences and naivete about future self-control. Moreover we assume that consumers’ present choices are not independent of their past actions, in other words, they are subject to the so-called sunk cost effect.

A number of empirical papers provide strong direct and indirect evidence for present biased preferences (this literature is surveyed by Frederick et al. 2002). Recently O’Donoghue and Rabin (1999) and (2001) provided theoretical papers to analyze economic behavior of agents that are only in part aware of their future time inconsistency. This strand of literature emphasizes the role of naivete in such phenomena as procrastination and status quo bias. Procrastination

has also been cited as a potential cause of breakage (e.g., Jolson et al. 1987; Soman 1998, Silk 2005).

As we mentioned above, a mild degree of naivete explains why consumers participate in rebate programs and systematically fail to redeem. Naivete in time-inconsistency explains several empirical regularities in many fields such as retirement plans (Madrian and Shea, 2001), health club attendance (DellaVigna and Malmendier, 2006), magazine subscription (Olster and Morton, 2005), credit card use (Ausubel and Shiu, 2004), among others.

Besides investigating the causes of the failure of redemption our study focuses on the question why fewer consumers redeem the rebates when the redemption deadline is longer. The phenomenon that in the longer deadline condition more agents fail to complete a task has already been advocated by experimental studies. Shafir and Tversky (1992) offered students 5 dollars to return a long questionnaire by a given date. The students were randomly assigned to one of three deadline groups - the first group had 5 days, the second group 3 weeks, and the third group had no deadline. The respective rates of returning the questionnaire for the three groups were 60, 42 and 25 percent. Thus, the more time people had to complete the task, the less likely they were to do it. In an experiment Ariely and Wertenbroch (2002) show that people are willing to self-impose costly deadlines in order to overcome procrastination. These self-imposed deadlines turned out to be effective in improving task performance even if not optimal to achieve maximum accomplishment. The experiment of Silk (2005) was specifically designed to replicate consumers' redemption behavior. Six hundred undergraduate students were offered the following choice options: purchase two movie tickets at 13 dollars less a mail-in rebate or purchase two movie tickets for a discounted price of 11 dollars or decline both offers⁴. Participants were randomly assigned to the six experimental conditions which resulted from the two different rebate sizes (6 or 9 dollars) and the three redemption deadlines (1 day, 7 days or 21 days). The results confirmed that the number of purchases was positively correlated with the rebate size and a significantly higher proportion of consumers purchased when the redemption period was longer. However, a significantly lower proportion of consumers applied for the rebate when the redemption period was 21 days than when it was either 7 days or 1

⁴The movie tickets were valid for any movie and show time for a year and a regular price of a ticket at the box office was 12 dollars.

day. Participants reported that the most important reasons for failing to apply for the rebate were that they had procrastinated until the deadline passed, the savings did not seem worth the effort required or they simply forgot to redeem. Field evidence also confirms that longer deadlines exhibit lower redemption rates (PMA, 2005).

In order to capture the effect of deadlines we draw on the extensive psychological literature on the so-called "sunk cost effect" and its decreasing relevance over time. Various researchers have found that costs that have been paid (or "sunk") in the past affect people's behavior in ways that are not consistent with economic theory (Thaler 1980; Thaler and Johnson 1990; Shefrin 2000, Eyster 2002). The following example of Thaler (1980) clarifies the point:

"A family pays 40 dollars 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." Thaler (1980, p. 11).

If they were thinking rationally, at the point of deciding whether to attend the game the family should ignore how much they paid for the tickets and should only take into account the future costs and benefits of attending the game. This economic logic is justified by the argument that because no action (current or future) can avert or reduce a sunk cost, no sunk cost can be attributed to or have any relevance to current or future action.

In a series of experiments, Arkes and Blumer (1985) demonstrated that monetary sunk costs matter to individuals when making decisions. In hypothetical situations and in situations involving actual money, subjects in these experiments made decisions not on the basis of expected marginal costs and benefits, but on the basis of expenses already incurred.

A few empirical studies provide evidence of the sunk cost effect in stock markets by observing trade volume decreases when prices fall. The decline in volume with lower prices is attributed in part to the reluctance of asset owners to sell at a loss or at a price less than one previously prevailing (e.g. see Thaler, 1980 and Shefrin, 2000). Investors treat the price paid for their investment, even though it is sunk, as a cost they would like to recover (Chip, Heath and Lang, 1999).

Although sunk costs influence subsequent decisions, they truly sink in our memories after

a sufficiently long time. The gradual reduction in the effect of past expenditures is termed "payment depreciation" by Gourville and Soman (1998). They obtained usage data from the members of a health club that charges the fees to its members twice a year. Attendance rates were highest in the month in which the fees were paid and then declined over the next five months, only to jump again when the next bill came out. Their conclusion is that the members felt driven to justify their fee by using the health club, a feeling that decreases over time.

The decreasing sunk-cost effect is consistent with the findings of DellaVigna and Malmendier (2006) according to which monthly payment members were much more likely to continue using a health club after their first year. These members were still exercising at the end of the year, when the yearly members had stopped attending. By paying each month, they were reminding themselves to attend often, making the club more useful to them. Hence they were more willing to sign up again at the end of the year.

Similar findings are reported by Arkes and Blumer (1985). In their experiment people who were ready to buy season tickets to a campus theater group were randomly placed into three groups: one group paid full price, one group got a small (13 percent) discount, and one group received a large (47 percent) discount. The experimenters then monitored how often the participants attended plays during the season. In the first half of the season, those who paid full price attended significantly more plays than those who received discounts, but in the second half of the season there was no difference among the groups.

The fact that the family in Thaler's example decide to attend the basketball game despite the snowstorm can be interpreted that they are not willing to suffer the regret that they would feel for having made a bad decision in the past. As mentioned in the introduction, our paper is also related to the literature on regret theory (Bell, 1982; Fishburn, 1982; Loomes and Sugden, 1982 and 1983). In these papers, regret is a psychological reaction to a choice of a prospect whose performance is worse than the performance of a forgone prospect. A recent study that incorporates this aspect of consumer choice is Eyster (2002). In his model people dislike the thought of being better off having made different choices in the past. Instead of simply regretting their mistakes, they may bring suboptimal decisions in the present in order to make their past choices look better. That is, they rationalize their past choices through their current choices. They exhibit a preference for consistency, which may give rise to the

sunk cost effect. In Eyster's formulation a choice situation is required to involve meaningful uncertainty about the future states of the world in order for the preference of consistency to arise. We show that when agents discount the future hyperbolically and are (partially) naive about their future self-control, the possibility of time inconsistency gives rise to preference for consistency even without uncertainty about the future.

3 Preliminaries of the Model

In this section we set up a three period model of consumer behavior and we allow consumers to have time inconsistent preferences and overconfidence about their future self-control problems. At time 0 the consumer decides whether to buy a product that includes a rebate promotion. We distinguish consumers based on their willingness to pay for the product in question. In particular, consumers who have a high willingness to pay would buy the product even without the rebate promotion, their valuation of the product, v , is higher than or equal to its shelf price, p . Those consumers who have low willingness to pay do not find it worth to pay the original price for the product, for them $p > v$.

We formulate the decision problem of whether to reply to the rebate offer by comparing the payoff consequences of the consumers' choices to their next best alternative. For the consumers who have high willingness to pay the next best alternative is to buy the product without a rebate promotion. Given that the purchase price is the same with and without the rebate promotion, their immediate payoff from deciding for the promoted product is zero. On the other hand, consumers who would not buy the product without the rebate promotion incur a cost $p - v$ at time 0.

We denote the $t = 0$ payoff consequence of replying to the rebate promotion by Δ , which, according to the above discussion equals $|\min(v - p, 0)|$. To simplify the discussion here we assume that a consumer who purchased the product with the rebate offer is required to send back the proof-of-purchase and the other forms to the manufacturer at a given future date, $t = 1$. The analysis of the consumers' behavior when they have multiple opportunities to complete this task is deferred to the next section.

Upon replying to the rebate offer the consumer decides whether to redeem the rebate or

not. If she decides to redeem she incurs cost c at $t = 1$ and receives a portion of the purchase price, $B > \Delta$, at $t = 2$.⁵ In case she decides not to redeem, she attains payoff 0 at time 1 and at time 2. Similarly, the payoffs of her next best alternative (irrespective of whether it is buying the product without a rebate promotion or not buying at all) is zero in both periods.

Suppose the individual has $\beta - \delta$ time preferences (Phelps and Pollack, 1967 and Laibson, 1997). In this case the present value of a flow of future utilities $(u_t)_{t \geq 0}$ is

$$u_0 + \beta \sum_{t=1}^{\infty} \delta^t u_t.$$

For $\beta < 1$ we have quasi-hyperbolic discounting which implies time-inconsistency: the discount factor between the current period and the next one is $\beta\delta$, while the discount factor between any two periods in the future is δ . Assume that $\beta < 1$ and that the consumer at each period has a belief about future self-control $\hat{\beta} \geq \beta$. Following O'Donoghue and Rabin (2001) we term as sophisticated the consumer who has beliefs $\hat{\beta} = \beta$, naive the consumer for whom $\hat{\beta} = 1$, and partially naive the one with $\beta < \hat{\beta} < 1$. Given these preferences a consumer will buy the promoted product if and only if at time 0 the following two conditions are satisfied:

(Buying conditions)

$$-\Delta + \beta\delta(-c + \delta B) \geq 0$$

$$-c + \hat{\beta}\delta B \geq 0$$

The first condition requires that at $t = 0$ the net present value of purchasing the product conditional on redeeming the rebate in period 2 be non-negative. This formulation implies that a consumer whose valuation of the product is lower than its prevailing price would not buy the rebated product if she did not believe to redeem later. The second condition shows that whether the consumer believes to redeem at $t = 1$ depends on her perceived future self control problem. A consumer who bought the promoted product will fail to redeem if at time 1 the following condition does not hold:

⁵Throughout the analysis we assume that consumers receive the face value of the rebate, B , a certain period (e.g. 6 weeks) after they have mailed the various forms to the manufacturer. Our results carry over to the case when manufacturers first collect all the rebate claims and send the checks to the consumers at a certain date after the redemption deadline.

(Redemption condition)

$$-c + \beta\delta B \geq 0$$

It is clear from the above conditions that sophisticated consumers will never fail to redeem due to their perfect foresight ($\hat{\beta} = \beta$). Therefore they only buy the promoted product if they redeem subsequently. On the other hand, naifs and partially naifs may purchase the product fully intending to redeem the rebate, but may fail to act at $t = 1$ according to their $t = 0$ wishes. The higher is the discrepancy between the actual and perceived future self control problems, in other words the more confident is the consumer about her future willpower, the more probable it is that she will fail to complete the task in the future. This simple benchmark is useful to derive further results.

4 The Model

We assume that the consumers are subject to the so called sunk-cost effect and we explore the behavioral consequences of this assumption in the context of our application. We first introduce the sunk-cost effect in the framework above, then we extend the analysis to the case when consumers have several opportunities to complete the task. Consider a consumer who has bought the good with the rebate promotion $t = 0$, and has a valuation of the product that is lower than the shelf-price ($\Delta > 0$). If she does not redeem the rebate at $t = 1$ she suffers an emotional disutility from having brought a bad decision in the past. We assume the emotional disutility to be proportional to the net cost paid at time 0, Δ .⁶ The following expressions state the payoffs associated to the two choices, i.e. to accomplish or not the task.

$$-c + \beta\delta B \quad \text{in case of redemption}$$

$$-\gamma\Delta \quad \text{in case of no redemption,}$$

⁶Although this is a very natural assumption, we emphasize that our results are not sensitive to the assumed relation between Δ and the emotional disutility.

where $\Delta = |\min(v - p, 0)|$ as explained in the previous section and γ is a discount factor at which emotional disutility decays over time. We discussed the behavioral foundations for the assumption that the emotional disutility decreases over time in section 2. We assume that the consumer perfectly predicts how the emotional disutility decreases irrespective of the degree of naivete.⁷ This formulation has the implication that the consumer accomplishes the task if and only if:

(Redemption condition with sunk-cost effect)

$$-c + \beta\delta B + \gamma\Delta \geq 0.$$

The above condition reflects the incentive provided by the sunk cost effect. For a consumer whose valuation of the product is lower than its shelf price ($\Delta > 0$) the thought that she could have been better off by not responding to the rebate offer effectively increases the utility from redemption. Hence, redemption generates two sources of utility: a standard physical utility, and another, emotional one, which results from being consistent with past decisions. The emotional utility can also be interpreted as ego-utility which is the counterpart of the self-image loss in case of failure.

In order to account for the sunk cost effect, the optimal conditions for a consumer to reply to the rebate promotion need to be modified. The consumer finds the rebate promotion attractive if and only if:

(Buying conditions with sunk-cost effect)

$$-\Delta + \beta\delta[-c + \delta B + \gamma\Delta] \geq 0$$

$$-c + \hat{\beta}\delta B + \gamma\Delta \geq 0$$

These modified buying conditions reflect that the consumer anticipates that in case of redemption she will experience the emotional utility from being consistent with past decisions (or, from

⁷As it will be clear, assuming naivete about how the emotional disutility affects future utility just strengthens our results.

avoiding regret resulting from past mistakes). This utility is experienced at $t = 1$ and hence it is discounted by $\beta\delta$. To clarify the point, we set up a simple example that we will carry on throughout the analysis. Suppose that a partially naive consumer with self-control problems $\beta = 0.65$ and beliefs $\hat{\beta} = 0.7$ considers buying a DVD-player that includes 17 dollars rebate offer. Assume also that she values the product 5 dollars less than its original price and it costs her 11 dollars to complete and mail the required forms. Moreover, suppose that $\delta = \gamma = 0.90$. In this case, $c = 11$, $\Delta = 5$, and $B = 17$. The two buying conditions are satisfied.⁸ In this example, the partially naive consumer will reply to the rebate promotion at $t = 0$ and will redeem the rebate at $t = 1$ because also the redemption condition is satisfied.⁹ A sophisticated with $\hat{\beta} = 0.65 = \beta = 0.65$ will behave similarly as the second buying condition is satisfied also for $\beta = 0.65$.¹⁰ Now we turn to the case when consumers have several opportunities to redeem preceding the deadline T .

4.1 Redemption decisions with deadlines

In the followings we utilize the framework provided by O'Donoghue and Rabin (2001) to analyze choice and procrastination. In each period t the consumer decides whether to do the task in t . We denote with a_t the action taken at period t ; if $a_t = 1$ she accomplishes the task and if $a_t = 0$ she does not. We denote with the strategy $\alpha(t, \hat{\beta}) = (\hat{a}_{t+1}^t, \hat{a}_{t+2}^t, \dots, \hat{a}_T^t)$ the consumer's period- t belief about when she will redeem the rebate, given that she does not accomplish the task in period t , i.e. given that $a_t = 0$. For example, $\hat{a}_{t+2}^t = 1$ means that at period t the consumer believes to redeem two periods later, i.e. at $t + 2$. Although the consumer can complete the task at most once, a strategy specifies an action for all periods. The strategy $\alpha(t, \hat{\beta}) = (0, 0, \dots, 0)$ plays an important role in the analysis, saying that the consumer believes that she will not redeem by the deadline T . We denote such a strategy with $\alpha(t, \hat{\beta}) = \emptyset$. Let $V_t[\alpha(t, \hat{\beta}), \beta, \delta]$ be the period- t preferences over current actions conditional on following strategy $\alpha(t, \hat{\beta})$:

⁸The first buying condition is $-5 + 0.65 \cdot 0.90[-11 + 0.90 \cdot 17 + 0.90 \cdot 5] = 0.148$, while the second condition is $-11 + 0.90 \cdot 0.70 \cdot 17 + 0.90 \cdot 5 = 4.21$.

⁹In particular, $-11 + 0.90 \cdot 0.65 \cdot 17 + 0.90 \cdot 5 = 3.445$.

¹⁰For the sophisticated, the second buying condition is equivalent to her redemption condition, which is the same redemption condition of the partially naive.

$$\begin{aligned}
V_t[\alpha(t, \hat{\beta}), \beta, \delta] &= -c + \beta\delta B + \gamma^t \Delta \quad \text{if } a_t = 1 \\
V_t[\alpha(t, \hat{\beta}), \beta, \delta] &= \beta\delta^{\tau-t}(-c + \delta B + \gamma^{t+\tau} \Delta) \quad \text{if } a_t = 0, \tau \equiv \min \{d > 0 | \hat{a}_{t+d}^t = 1\} \text{ exists} \\
V_t[\alpha(t, \hat{\beta}), \beta, \delta] &= -\gamma^t \Delta \quad \text{if } a_t = 0, \quad \text{and } \alpha(t, \hat{\beta}) = \emptyset
\end{aligned}$$

The consumer at each period t chooses an action that maximizes her current preferences $V^t[\alpha(t, \hat{\beta}), \beta, \delta]$ given her current beliefs about strategy $\alpha(t, \hat{\beta})$. In the first case, the consumer redeems at time t , incurs the immediate cost and experiences the delayed benefit and the emotional utility. In the second case, the consumer does not redeem at time t but thinks to redeem τ periods later so that both the cost and the benefits are discounted by β . In the third case, the consumer does not redeem at time t and understands that she will not redeem by the deadline T . The consequence is that she suffers the emotional disutility from having made a bad decision in the past.

We require that the individual strategy $\alpha(t, \hat{\beta})$ be optimal given the individual beliefs $\hat{\beta}$: $\alpha(t, \hat{\beta}) = \arg \max V_t^{t'}[\alpha(t, \hat{\beta}), \hat{\beta}, \delta]$ for all t' , where $V_t^{t'}[\alpha(t, \hat{\beta}), \hat{\beta}, \delta]$ are the period- t preferences the consumer perceives will have in the future $t+t'$. That is, the action taken at each period is optimal given the perceived future behavior and incentives that are determined by the perception $\hat{\beta}$. The optimal strategy $\alpha(t, \hat{\beta})$ is essentially equivalent to the perfect perfection strategy of O'Donoghue and Rabin (2001).¹¹ We characterize the optimal strategy and we prove some of the results derived in this section in the Appendix.

In order to clarify how the optimal beliefs $\alpha(t, \hat{\beta})$ are formed, let $\tilde{V}_t[\tau, \beta, \delta]$ be the values for the period- t intertemporal utility from completing the task in period $\tau \geq t$. For $\tau = t$, $\tilde{V}_t[\tau, \beta, \delta] = -c + \beta\delta B + \gamma^t \Delta$, while for $\tau > t$, $\tilde{V}_t[\tau, \beta, \delta] = \beta\delta^{\tau-t}[-c + \delta B + \gamma^{t+\tau} \Delta]$. In our example if we set $t = 1$ we have:

¹¹In this case, however, we do not need to impose the restriction of external consistency according to which a person's beliefs must be consistent across periods, i.e. a person's belief of what she will do in period τ must be the same in all $t < \tau$. In this paper we deal with a finite period model so that the external consistency always holds.

$$\begin{aligned}
\tilde{V}_t[\tau = 1, \beta, \delta] &= -11 + (.65)(.90)17 + (.90)5 &= 3.445 \\
\tilde{V}_t[\tau = 2, \beta, \delta] &= (0.65)(.90)[-11 + (.90)17 + (.90)^2 5] &= 4.8847 \\
\tilde{V}_t[\tau = 3, \beta, \delta] &= (.65)(.90)^2[-11 + (.90)17 + (.90)^3 5] &= 4.1830 \\
\tilde{V}_t[\tau = 4, \beta, \delta] &= (.65)(.90)^3[-11 + (.90)17 + (.90)^4 5] &= 3.592 \\
\tilde{V}_t[\tau = 5, \beta, \delta] &= (.65)(.90)^4[-11 + (.90)17 + (.90)^5 5] &= 3.0929
\end{aligned}$$

As it is clear from this example, due to the presence of β , delaying the task to period 2 is preferable to doing the task now (at $t = 1$). Whether the consumer actually delays now depends on how she perceives future behavior. In particular, she delays now if, based on her perceived future behavior, she believes that she will accomplish the task in the future with no more than a maximum tolerable delay. The maximum tolerable delay that the consumer tolerates in the above example is 3 periods because after the 4th period the value of redemption is lower than the value if she completes the task today (at $t = 1$).¹² Hence, the consumer delays at $t = 1$ if she believes to redeem in no more than 4 periods. For a consumer with self-control β and beliefs $\hat{\beta}$ we denote the maximum tolerable delay at time t with $d_t(\beta|\hat{\beta})$:

$$d_t(\beta|\hat{\beta}) \equiv \max \left\{ d \in \{0, 1, 2, \dots, T\} \mid -c + \beta\delta B + \gamma^t \Delta < \beta\delta^d (-c + \delta B + \gamma^{d+t} \Delta) \right\}$$

The definition of the maximum tolerable delay incorporates two effects that work in opposite directions. The first is created by the presence of β and leads to prefer a high delay (because it makes the cost less salient), the other is created by the the long-run discount factor δ and renders the delay less pleasant (because it decreases the perceived utility of B in the future). After a critical date the computation of the maximum tolerable delays at different points in time may deliver $d_{t+1}(\beta|\hat{\beta}) \geq d_{t+2}(\beta|\hat{\beta}) \geq d_{t+3}(\beta|\hat{\beta}) \geq \dots \geq d_T(\beta|\hat{\beta})$ (see the Appendix). Some of the inequalities will hold with certainty if at any point in time the maximum tolerable delay is greater than 0. This is because at the deadline (at the last opportunity of redemption) the consumer will surely not tolerate any further delay, i.e. $d_T(\beta|\hat{\beta}) = 0$. Another factor

¹²The maximum tolerable delay above is 3 if $T > 3$, otherwise the maximum tolerable delay is set by the deadline itself.

contributing to the decreasing maximum tolerable delay is provided by the decreasing payoff from redemption, a result of the sunk-cost effect. It is important to note, however, that the changes in the maximum tolerable delays do not affect the overall analysis. The predictions of the model hold for any kind of patterns that the maximum tolerable delays may exhibit.

Now, in order to understand whether a consumer with a maximum tolerable delay greater than 0 is willing to delay the redemption, we need to know what she thinks about her future behavior. This in turn is determined by her beliefs about her maximum tolerable delay in the future. At time t we denote the perceived maximum tolerable delay in the future $t + \tau$, with $d_t^\tau(\hat{\beta}|\hat{\beta})$:

$$d_t^\tau(\hat{\beta}|\hat{\beta}) \equiv \max \left\{ d \in \{0, 1, \dots, T\} \mid -c + \hat{\beta}\delta B + \gamma^{t+\tau}\Delta < \hat{\beta}\delta^d(-c + \delta B + \gamma^{d+t+\tau}\Delta) \right\}$$

Given her perceived future maximum tolerable delay the consumer sets her redemption strategy by assigning an action (0 or 1) to each future period. Before describing the consumer's redemption strategy we need to introduce the following notation:

$$n_{\hat{\beta}} \equiv \min \left\{ n \in [1, 2, \dots, T] \mid -c + \hat{\beta}\delta B + \gamma^{n_{\hat{\beta}}+1}\Delta \leq 0 \right\}$$

We term $n_{\hat{\beta}}$ the perceived payoff-worth deadline since the consumer foresees that at any $t > n_{\hat{\beta}}$ the net future payoff from doing the task is negative. Note that $n_{\hat{\beta}}$ depends on $\hat{\beta}$, so that the true payoff-worth deadline may differ from the perceived one, if $\hat{\beta}$ is not equal to β . We term the true payoff-worth deadline n_β and calculate using β in place of $\hat{\beta}$ in the above definition.

Now consider the sophisticated with $\beta = \hat{\beta}$ and $n_{\hat{\beta}} = n_\beta$. Sophistication also implies that $d_t^\tau(\hat{\beta}|\hat{\beta}) = d_{t+\tau}(\beta|\hat{\beta})$ which means that her perceived future maximum delay is equal to the actual future maximum tolerable delay. We first analyze the case when $n_{\hat{\beta}} = n_\beta < T$. In the example above for the sophisticated with $\beta = \hat{\beta} = 0.65$, $n_\beta = 14$ as at period 15th, $-c + \hat{\beta}\delta B + \gamma^{15}\Delta$ is negative.¹³ The optimal strategy of the sophisticated is found by solving

¹³The relevant calculations are $-c + \hat{\beta}\delta B + \gamma^{14}\Delta = -11 + (.65)(.90)17 + (.90)^{14}5 = 0.08881$ and at period 15th $-c + \hat{\beta}\delta B + \gamma^{15}\Delta = -11 + (.65)(.90)17 + (.90)^{15}5 = -0.02554$

the problem by backward induction. Since in period t she redeems only if she predicts that she will delay more than $d_t^\tau(\hat{\beta}|\hat{\beta})$ periods, her strategy will feature completing the task in each $d_t^\tau(\hat{\beta}|\hat{\beta}) + 1$ period. In the above example, assume for simplicity that the actual maximum tolerable delay is constant for all the periods.¹⁴ Applying this reasoning, the sophisticated plans to complete the task in each 4th period calculating backwards from the perceived payoff-worth deadline $n_{\hat{\beta}} = 14$. Accordingly, her strategy is to redeem in periods 2nd, 6th, 10th, 14th. Given this strategy, the sophisticated passes the opportunity to redeem in period 1 since she believes to redeem in period 2 and she can tolerate a 3 period delay ($d_t(\beta|\hat{\beta}) = 3$). In period 2, she predicts that if she delays now she will redeem in period 6. As this would lead to a delay of 4 periods, that she can not tolerate, she redeems in period 2.

Note that this outcome crucially depends on the fact that $d_t^\tau(\hat{\beta}|\hat{\beta}) = d_{t+\tau}(\beta|\hat{\beta})$. The analysis is very similar in case $T < n_{\hat{\beta}} = n_{\beta}$ with the only difference that the calculation goes from T backwards. Now we state the first result for sophisticated and exponential consumers.

Proposition 1 (*Behavior of exponential and sophisticated consumers*) *An exponential consumer completes the task immediately independently of the deadline as she has a maximum tolerable delay equal to zero. A sophisticated may have a positive maximum tolerable delay which defines her redemption strategy $\alpha(t, \hat{\beta})$. Given that for a sophisticated $d_t^\tau(\hat{\beta}|\hat{\beta}) = d_{t+\tau}(\beta|\hat{\beta})$, it never happens that she fails to redeem by the deadline T . For $T > n_{\hat{\beta}} = n_{\beta}$, the optimal strategy of a sophisticated is independent of the redemption deadline, T . For both $T < n_{\hat{\beta}} = n_{\beta}$ and $T > n_{\hat{\beta}} = n_{\beta}$ the sophisticated will redeem at the earliest redemption date set by her strategy.*

As we anticipated in the introduction a purely exponential model with sunk-cost effect would not deliver redemption failure, which is confirmed by Proposition 1. On the contrary, as exponentials have a maximum tolerable delay equal to zero not only do not they procrastinate, but they do not even allow for any delay. The presence of hyperbolic discounting implies a behavior of the sophisticated consumer that is observationally non-equivalent to that of an exponential consumer. At the same time, the presence of the sunk-cost effect renders

¹⁴In reality, as we mentioned before, the maximum tolerable delay may vary over time. Considering this would only complicate the exposition and would not affect our results qualitatively. The proof is deferred to the Appendix.

exponentials and sophisticated more similar with respect to the timing of redemption. The presence of the emotional disutility generates a taste for consistency that makes the maximum tolerable delay lower and thus closer to that of the exponentials.

Now consider the case of the partially naive for whom $\hat{\beta} > \beta$. Note that for a partially naive the perceived payoff-worth deadline, $n_{\hat{\beta}}$, may not be the same as the true one n_{β} . This is because she overestimates her future gains from redemption. Under general conditions $n_{\hat{\beta}} > n_{\beta}$, and this will be our focus in the subsequent analysis.

We again start with binding payoff-worth deadline, i.e. $n_{\hat{\beta}} < n_{\beta} < T$. Consider the previous example with a partially naive having $\hat{\beta} > \beta$ and $\hat{\beta} = 0.7 > \beta = 0.65$. Her actual tolerable delay, $d_t(\beta|\hat{\beta})$, is the same as that of the sophisticated before, i.e. 3. However, her perceived future maximum delay, $d_t^{\tau}(\hat{\beta}|\hat{\beta})$, can be lower than her actual future maximum delay, $d_{t+\tau}(\beta|\hat{\beta})$, due to the fact that $\hat{\beta} > \beta$. In order to see this, we calculate her perceived future ($t + 1$) value of redemption from the perspective of period t , denoted with $\tilde{V}_t^{t+1}[\tau, \hat{\beta}, \delta]$, as follows. $\tilde{V}_t^{t+1}[\tau, \hat{\beta}, \delta] = -c + \hat{\beta}\delta B + \gamma^{t+1}\Delta$ for $\tau = t + 1$ and $\tilde{V}_t^{t+1}[\tau, \hat{\beta}, \delta] = \hat{\beta}\delta^{\tau-t+1}[-c + \delta B + \gamma^{t+\tau}\Delta]$ for $\tau > t + 1$. For $t = 1$:

$$\begin{aligned} \tilde{V}_t^{t+1}[\tau = 2, \hat{\beta}, \delta] &= -11 + (.7)(.90)17 + (.90)^25 &= 3.76 \\ \tilde{V}_t^{t+1}[\tau = 3, \hat{\beta}, \delta] &= (.7)(.90)[-11 + (.90)17 + (.90)^35] &= 5.0053 \\ \tilde{V}_t^{t+1}[\tau = 4, \hat{\beta}, \delta] &= (.7)(.90)^2[-11 + (.90)17 + (.90)^45] &= 4.2981 \\ \tilde{V}_t^{t+1}[\tau = 5, \hat{\beta}, \delta] &= (.7)(.90)^3[-11 + (.90)17 + (.90)^55] &= 3.7009 \end{aligned}$$

It is easy to see that the perceived future tolerable delay is $d_t^{\tau}(\hat{\beta}|\hat{\beta}) = 2$ that is lower than the true one, $d_{t+\tau}(\beta|\hat{\beta}) = 3$. Moreover, the perceived deadline for the partially naive is $n_{\hat{\beta}} = 27$ as at period 28 her perceived future value from redemption becomes negative.¹⁵ Her strategy will feature completing the task in each $d_t^{\tau}(\beta|\hat{\beta}) + 1$ period. Accordingly, the strategy $\alpha(t, \hat{\beta})$ of the consumer is to redeem in periods 3rd, 6th, ..., 27th (assuming again that the delay does not change over time). Hence, the partially naive will not redeem in the first two periods because she believes to complete the task in period 3. However, in period 3 she again fails

¹⁵The relevant calculations in this case are $-c + \hat{\beta}\delta B + \gamma^{27}\Delta = -11 + (.7)(.90)17 + (.90)^{27}5 = 0.000075$ at period 27 and $-c + \hat{\beta}\delta B + \gamma^{28}\Delta = -11 + (.7)(.90)17 + (.90)^{28}5 = -0.028$ at period 28th.

to redeem because, based on her strategy, she believes to do the task in period 6 and she can actually tolerate (and prefers) 3 periods of delay, ($d_t(\beta|\hat{\beta}) = 3$). The same reasoning applies to each redemption date set by her strategy except for the last one. At the perceived deadline ($n_{\hat{\beta}} = 27th$) she understands that she will not redeem in the future. However, she fails to complete the task even at the last opportunity because the net payoff from redemption has already become negative ($-c + \beta\delta B + \gamma^{27}\Delta$ is negative and equal to -0.2882). This is because the actual deadline until which redeeming is profitable (the payoff-worth deadline, n_{β}) is the same as that of the sophisticated, i.e. 14. The failure of the partially naive is thus the result of two forces: unforeseeable and repeated delays and overestimation of the deadline until which there are enough incentives to complete the task. Finally, at the perceived payoff worth-deadline 27, $\alpha(t, \hat{\beta}) = \circlearrowleft$ and she experiences the emotional disutility $\gamma^{27}\Delta = 0.05814 \cdot 5 = 0.2907$. The analysis is essentially the same when $n_{\beta} < T < n_{\hat{\beta}}$, with the difference that the calculations go from the redemption deadline, T , backwards. The partially naive will delay the task until T , at which point it is not profitable to complete the task. Now it is possible to state the second result.

Proposition 2 (*Behavior of partially naive with binding payoff-worth deadline*) *Assume $T > n_{\hat{\beta}} > n_{\beta}$. A partially naive for whom $d_t^{\tau}(\hat{\beta}|\hat{\beta}) + 1 \leq d_{t+\tau}(\beta|\hat{\beta})$ delays to complete the task until $n_{\hat{\beta}}$. At that period there are no incentives to complete the task, she does not redeem and she suffers the emotional disutility $\gamma^{n_{\hat{\beta}}}\Delta$. The extent of the emotional disutility that a partially naive eventually suffers is decreasing with the degree of naivete, $\hat{\beta} - \beta$. For $n_{\beta} < T < n_{\hat{\beta}}$, the partially naive for whom $d_t^{\tau}(\hat{\beta}|\hat{\beta}) + 1 \leq d_t(\beta|\hat{\beta})$ delays until T and at that period she does not complete the task. In this case, the extent of the emotional disutility she suffers is independent of $\hat{\beta} - \beta$ and decreasing with T . The behavior of a partially naive for whom $d_t^{\tau}(\hat{\beta}|\hat{\beta}) + 1 = d_t(\beta|\hat{\beta}) + 1$ is the same as that of a sophisticated.*

This result differs from that of O'Donoghue and Rabin (2001) in two important aspects. There, without deadlines, partially naives may delay a task infinitely without ever realizing that they will never complete it. While our theoretical prediction delivers the result that the agent may fail to complete the task, here this behavior needs not be enforced by persistent optimistic beliefs in future willpower. Moreover, in the framework of O'Donoghue and Rabin (2001) the

introduction of a deadline would lead all agents to accomplish the task, since even partially naives, for whom $d(\hat{\beta}|\hat{\beta}) + 1 < d(\beta|\hat{\beta})$, will succeed to do it at the last opportunity. As such, that framework can not explain why in reality we frequently observe that people fail to accomplish what they plan even in the presence of deadlines. In our case, the sunk cost effect creates a payoff-worth deadline until which it is profitable to complete the task. In principle, we should also get the result that at the payoff-worth deadline even partially naives succeed in providing the necessary effort. However, due to the fact that their $\hat{\beta} > \beta$, partially naives overestimate the length of the payoff-worth deadline because they overweight their future benefit from redemption. This will lead them to delay the task until a misperceived payoff-worth deadline, at which point they do not find it profitable to complete the task. Interestingly, a partially naive may give up to complete the task even before the redemption deadline, T , expires.

Now we consider the case where $T < n_\beta < n_{\hat{\beta}}$, i.e. the redemption deadline is binding. In the example above, assume that $T = 12$ (maintaining that $n_\beta = 14$ and $n_{\hat{\beta}} = 27$). The strategy of the partially naive who has a $d_t(\hat{\beta}|\hat{\beta}) = 2$ is to accomplish the task in periods 3rd, 6th, ..., 9th, 12th. Using the logic above, since $d_t^\tau(\hat{\beta}|\hat{\beta}) + 1 \leq d_{t+\tau}(\beta|\hat{\beta})$, she will delay until period 12. At that period she will complete the task because $-c + \beta\delta B + \gamma^{12}\Delta > 0$.

Proposition 3 (*Behavior of partially naives with binding redemption deadline*) *Assume $T < n_\beta < n_{\hat{\beta}}$. A partially naive for whom $d_t^\tau(\hat{\beta}|\hat{\beta}) + 1 \leq d_{t+\tau}(\beta|\hat{\beta})$ delays to complete the task until T . At that period she does redeem. A partially naive for whom $d_t^\tau(\hat{\beta}|\hat{\beta}) + 1 = d_{t+\tau}(\beta|\hat{\beta}) + 1$ redeems at the earliest redemption date of her strategy.*

An implication of the last two propositions is that for a partially naive with a certain degree of naivete $\hat{\beta} - \beta$, the probability of redemption is decreasing with the deadline. This theoretical prediction matches the empirical evidence discussed in the introduction. The main results concerning delay and failure in redemption are driven by naivete in future self-control problems. In the example we have discussed even a small degree of naivete, i.e. $\hat{\beta} - \beta = 0.7 - 0.65$, led the consumer to postpone the task until the last possible period.

4.2 Buying decision with deadlines

In this sub-section we study the buying conditions given that the consumer can accomplish the task in any of the periods preceding the deadline T . At period 0 the consumer computes her maximum tolerable delay that she believes to have in the future, $d_t^r(\hat{\beta}|\hat{\beta})$. Accordingly, her strategy is to complete the task in each $d_t^r(\hat{\beta}|\hat{\beta}) + 1$ periods and she expects to complete the task at the earliest redemption date of her strategy. Define $\mu^* = n_{\hat{\beta}} - z^*(d_t^r(\hat{\beta}|\hat{\beta}) + 1)$, where $z^* \equiv \max \left\{ z \in [1, 2, \dots, T - 1] | n_{\hat{\beta}} - z d_t^r(\hat{\beta}|\hat{\beta}) \geq 1 \right\}$, for $T < n_{\hat{\beta}}$.¹⁶ Then, the buying conditions are:

(Buying conditions with sunk-cost effect)

$$\begin{aligned} -\Delta + \beta \delta^{\mu^*} [-c + \delta B + \gamma^{\mu^*} \Delta] &\geq 0 \\ -c + \hat{\beta} \delta B + \gamma^{\mu^*} \Delta &\geq 0 \end{aligned}$$

As we documented in section 2, with longer deadline more consumers reply to the rebate promotion. As such, however, these conditions do not provide any insight about the relationship between the extent of the deadline and the incentive to buy. At period 0, both the sunk-cost effect and the extent of the deadline can have an effect on the decision to buy if the consumer anticipates how her emotional disutility varies with the deadline in case of no redemption. As we showed in the previous sub-section, in case of failure the emotional disutility is decreasing with the length of the redemption deadline (if the deadline is binding) and decreasing with $\hat{\beta}$ (in case of binding payoff-worth deadline). Therefore if the consumer takes this into account, longer deadlines should generate more incentives to buy. In order to formalize this intuition we need to introduce uncertainty at period 0 about future willpower. Note that uncertainty about the "state of nature" in the future is a key and necessary assumption of regret theories for the analysis of anticipated regret on current actions.

In particular, at period 0 consider a consumer with self-control problems β and beliefs $\hat{\beta} > \beta$ and suppose that her perceived maximum tolerable delay is such that the two buying conditions above are satisfied. However, she believes with probability q that in replying to the rebate promotion she is making a mistake because her future willpower is such that she will

¹⁶In the definition of μ^* , if $T > n_{\hat{\beta}}$ it is necessary to substitute T for $n_{\hat{\beta}}$.

delay until the deadline T and at that period she will fail to redeem. On the other hand, with probability $1 - q$ she believes that she will find enough incentives to redeem at the earliest redemption date of her strategy. Hence, the first buying condition in this case is the following:

(First buying condition with sunk-cost effect with uncertainty and $T < n_{\hat{\beta}}$)

$$-\Delta + q \left\{ \beta \delta^{\mu^*} [-c + \delta B + \gamma^{\mu^*} \Delta] \right\} + (1 - q) \left\{ -\gamma^T \Delta \right\} \geq 0$$

It is possible to show that this condition is more easily satisfied for longer T , so that, for a given degree of naivete, there will be a positive relationship between the extent of deadline and the probability to reply to the rebate promotion. As before, in the case of binding payoff-worth deadline, it is necessary to substitute in the conditions $n_{\hat{\beta}}$ for T in the definition of μ^* and the first buying condition is:

(First buying condition with sunk-cost effect with uncertainty and $T > n_{\hat{\beta}}$)

$$-\Delta + q \left\{ \beta \delta^{\mu^*} [-c + \delta B + \gamma^{\mu^*} \Delta] \right\} + (1 - q) \left\{ -\gamma^{n_{\hat{\beta}}} \Delta \right\} \geq 0$$

Now we can state the last results of the model.

Proposition 4 (Buying decision with uncertainty) *Assume that the redemption deadline is binding, i.e. $T < n_{\beta} < n_{\hat{\beta}}$. A longer redemption deadline increases the consumer's incentives to reply to the rebate promotion.*

This proposition matches the empirical finding on rebates documented in section 2: more consumers are willing to respond to a rebate promotion when the deadline is longer. In line with regret theory where in the presence of uncertainty an agent chooses an action to minimize her ex-post regret, our consumer prefers to respond to a rebate promotion associated to a longer deadline because this may cause lower regret in the future.

5 Discussion

The above model shows how rational firms can strategically design the features of investment goods (the initial investment, the effort cost, the benefit and the deadlines) to achieve various

goals. In particular by offering shorter deadlines the firm can maximize the number of candidates who complete the given task in question. By offering longer deadlines, the firm can achieve higher participation and lower success rates in a program. In this section we review some areas where this device can prove to be important.

Similarly to mail-in-rebates, lottery games feature the characteristics of the above described investment goods. Going to the lottery office, paying the price of the lottery ticket and filling it out can be considered as the initial investment. After the announcement of the winner numbers the agent can collect the amount gained within a certain period. It is natural to assume that the cost of collecting the reward and the size of the reward does not change over time. While previous models would predict that all participants succeed in claiming their prize, if not earlier than at the deadline, our own experience confirms that failure happens regularly. Our model not only implies that failure happens, but predicts that failure rates are higher under milder conditions, e.g. longer deadlines.

A different area for which our model has important predictions involves the various courses offered by educational institutions. In order to achieve a specific degree or a license in some field subjects are often required to complete a course and to sit for a final exam. In many cases the agent has several opportunities to complete the exam, e.g. the practical exam to get a driving license. In case of failing to pass the exam the substantial initial investment (the cost of the course and the effort to follow the classes) is lost. Considering that self control problems may not imply that a person is less able to drive a car safely, based on the predictions of our model it could be argued that the deadline in this case should not be set by private firms. Indeed, in practice we observe that the process of getting a driving license is under state regulation.

Our model also predicts that the same person will succeed in collecting a relatively low reward if the deadline is short and at the same time may forgo a much higher gain if the deadline is long enough. The implications of our theory are not limited to situations that involve deadlines. We can also explain a very general phenomenon, i.e. a person may procrastinate a task forever, with allowing her to realize at a certain point that she will never succeed in completing it. We believe that this is a more realistic description of why we fail to pick up a book that we ordered and already paid for, or why we never collect our gloves that we forgot in the movie theater (even if we swore to get them back the moment we realized our loss).

Appendix

In this Appendix we characterize the optimal strategy of the consumer. It is important to note that in this paper, using the terminology of O'Donoghue and Rabin (2001), we analyzed only tasks that are $\hat{\beta}$ -worthwhile. In the context of our application this implies that there exists at least a period $t < T$ for which $-c + \hat{\beta}\delta B + \gamma^t \Delta > 0$. In the followings we will consider the case of redemption deadlines that are not binding, i.e. $T > n_{\hat{\beta}}$ as this is the one where the consumer may fail to redeem. The results concerning the optimal strategy, when it is necessary to substitute T for $n_{\hat{\beta}}$, generalize to the case of binding redemption deadlines. Recall that $\alpha(t, \hat{\beta}) = (\hat{a}_{t+1}^t, \hat{a}_{t+2}^t, \dots, \hat{a}_T^t)$ is the consumer's period- t belief about when she will accomplish the task and that $V_t[\alpha(t, \hat{\beta}), \beta, \delta]$ are the period- t preferences over current actions conditional on following strategy $\alpha(t, \hat{\beta})$. We first start with the following Lemma.

Lemma 1 *For any pair (B, T) , for all t either $\alpha(t, \beta) = \emptyset$ or there exists a unique cyclical strategy $\alpha(t, \hat{\beta})$. The strategy is to redeem every $d_t^{\tau}(\hat{\beta}|\hat{\beta}) + 1$ periods.*

Proof of Lemma 1. Given the definition of $\alpha(t, \hat{\beta})$, the proof of the first part of Lemma 1, i.e. the strategy is either $\alpha(t, \beta) = \emptyset$ or $\alpha(t, \beta) \neq \emptyset$ is trivial. The uniqueness of the strategy when this is different from the null one is based on the argument that there exists a unique equilibrium in the presence of a finite horizon.

The other part of the proof follows O'Donoghue and Rabin (2001). Given the definition of $d_t^{\tau}(\hat{\beta}|\hat{\beta})$, for any $d' \in \{t + \tau + 1, t + \tau + 2, \dots, t + \tau + d_t^{\tau}(\hat{\beta}|\hat{\beta})\}$, if $\hat{a}_{t^*}^t = 1$ and $\hat{a}_{t^*-d}^t = 0$ for all $d \in \{1, 2, \dots, d' - 1\}$, then the action that maximizes the preference at $t^* - d'$ from the perspective of period t is $a = 0$. For $d' = d_t^{\tau}(\hat{\beta}|\hat{\beta}) + 1$, if $\hat{a}_{t^*}^t = 1$ and $\hat{a}_{t^*-d}^t = 0$ for all $d \in \{1, 2, \dots, d' - 1\}$, then the action that maximizes the preference at $t^* - d'$ from the perspective of period t is $a = 1$. It follows that the optimal strategy $\alpha(t, \hat{\beta})$ features completing the task every $d_t^{\tau}(\hat{\beta}|\hat{\beta}) + 1$ periods and not completing the task otherwise.

QED

If the perceived maximum tolerable delay resulting from the expressions given in the text, $d_t^{\tau}(\hat{\beta}|\hat{\beta})$, and the maximum tolerable delay, $d_t(\beta|\hat{\beta})$, are such that these are greater than $n_{\hat{\beta}}$, then

$d_t^\tau(\hat{\beta}|\hat{\beta})$ and $d_t(\beta|\hat{\beta})$ are equal to $n_{\hat{\beta}} - (t + \tau)$ and to $n_{\hat{\beta}} - t$, respectively. This is the case whenever the maximum tolerable delay, $d_t(\beta|\hat{\beta})$, is such that $-c + \hat{\beta}\delta B + \gamma^{t+d_t(\beta|\hat{\beta})} < 0$ and the perceived maximum tolerable delay in the future $t + \tau$, $d_t^\tau(\hat{\beta}|\hat{\beta})$, is such that $-c + \hat{\beta}\delta B + \gamma^{t+\tau+d_t^\tau(\hat{\beta}|\hat{\beta})} < 0$. It follows that there might exist a critical period after which the maximum tolerable delays are decreasing.

Define $\tilde{d}_t^\tau(\hat{\beta}|\hat{\beta})$ and $\tilde{d}_t(\beta|\hat{\beta})$ as the counterpart of $d_t^\tau(\hat{\beta}|\hat{\beta})$ and $d_t(\beta|\hat{\beta})$ when the conditions $-c + \hat{\beta}\delta B + \gamma^{t+d_t(\beta|\hat{\beta})}\Delta \geq 0$ and $-c + \hat{\beta}\delta B + \gamma^{t+\tau+d_t^\tau(\hat{\beta}|\hat{\beta})}\Delta \geq 0$ are not satisfied. We can term $\tilde{d}_t^\tau(\hat{\beta}|\hat{\beta})$ and $\tilde{d}_t(\beta|\hat{\beta})$ as the unconstrained maximum tolerable delays.

Lemma 2 *The last date of the redemption strategy, $\alpha(t, \hat{\beta})$, is equal to the perceived payoff worth deadline, $n_{\hat{\beta}}$. The last but one date of the redemption strategy is $n_{\hat{\beta}} - (d_t^{\tau^*}(\hat{\beta}|\hat{\beta}) + 1)$ where $t + \tau^*$ is the first period backward from $n_{\hat{\beta}}$ such that $\tilde{d}_t^{\tau^*}(\hat{\beta}|\hat{\beta}) = d_t^{\tau^*}(\hat{\beta}|\hat{\beta})$.*

Proof of Lemma 2. The last date of the redemption strategy is the deadline because at that date the agent cannot tolerate any further delay. The last but one date of the redemption strategy must be $n_{\hat{\beta}} - (\tilde{d}_t^{\tau^*}(\hat{\beta}|\hat{\beta}) + 1)$ with τ^* as defined in the lemma. Any other period after $t + \tau^*$ but $n_{\hat{\beta}}$, indeed, can not be a date of the redemption strategy because in that period $d_t^\tau(\hat{\beta}|\hat{\beta}) < \tilde{d}_t^\tau(\hat{\beta}|\hat{\beta})$ and $d_t(\beta|\hat{\beta}) < \tilde{d}_t(\beta|\hat{\beta})$ as $d_t^\tau(\hat{\beta}|\hat{\beta})$ and $d_t(\beta|\hat{\beta})$ are equal to $n_{\hat{\beta}} - (t + \tau)$ and to $n_{\hat{\beta}} - t$, respectively. Hence, for any period after $t + \tau^*$ the agent can always tolerate a delay until $n_{\hat{\beta}}$. The first date in which she cannot tolerate a delay is $n_{\hat{\beta}} - (d_t^{\tau^*}(\hat{\beta}|\hat{\beta}) + 1)$. The result follows.

QED

Lemma 3 *If for all periods of the redemption strategy that precede $n_{\hat{\beta}}$, $d_{t+\tau}(\beta|\hat{\beta}) + 1 \geq d_t^\tau(\hat{\beta}|\hat{\beta})$, then the consumer fails to redeem.*

Proof of Lemma 3. If $d_{t+\tau}(\beta|\hat{\beta}) + 1 \geq d_t^\tau(\hat{\beta}|\hat{\beta})$ then the consumer is a partially naive that misperceives also the perceived payoff worth deadline, i.e. $n_{\hat{\beta}} > n_{\beta}$. At any date of the redemption strategy after n_{β} , the consumer fails to redeem because she does not find enough incentives, i.e. $-c + \beta\delta B + \gamma^{n_{\hat{\beta}}+1}\Delta < 0$. Hence, it is necessary to show that the consumers procrastinate at least until n_{β} . Assume $d_{t+\tau}(\beta|\hat{\beta}) + 1 \geq d_t^\tau(\hat{\beta}|\hat{\beta})$. Since the consumer prefers to

delay today, $V_t[\alpha(t, \hat{\beta}), \beta, \delta] = \beta\delta^{\tau+d-t}(-c + \gamma^{d+t+\tau}\Delta)$ for some $d \in [t+\tau+1, t+\tau+2, \dots, \tau + t + d_t^-(\hat{\beta}|\hat{\beta}) + 1]$ and, since the current action is not to redeem, this expression must be equal or greater than $\beta\delta^{d_t^-(\hat{\beta}|\hat{\beta})+1}(-c + \delta B + \gamma^{d_t^-(\hat{\beta}|\hat{\beta})+1+t+\tau}\Delta)$. Note that the consumer plans to redeem at $t + \tau$ only if $-c + \beta\delta B + \gamma^{t+\tau}\Delta \geq \beta\delta^{d_t^-(\hat{\beta}|\hat{\beta})+1}(-c + \delta B + \gamma^{d_t^-(\hat{\beta}|\hat{\beta})+1+t+\tau}\Delta)$ but actually redeems at $t + \tau$ only if $-c + \beta\delta B + \gamma^{t+\tau}\Delta \geq \beta\delta^{d_{t+\tau}(\beta|\hat{\beta})+1}(-c + \delta B + \gamma^{d_{t+\tau}(\beta|\hat{\beta})+1+t+\tau}\Delta)$. Since the definition of $d_{t+\tau}(\beta|\hat{\beta})$ implies that $-c + \beta\delta B + \gamma^{t+\tau}\Delta < \beta\delta^d(-c + \delta B + \gamma^{d+t+\tau}\Delta)$, for any date of the redemption strategy τ and for any $d_{t+\tau}(\beta|\hat{\beta})$, if $d_{t+\tau}(\beta|\hat{\beta}) + 1 \geq d_t^-(\hat{\beta}|\hat{\beta})$ then the optimal action is not to redeem at any date of the redemption strategy $\alpha(t, \hat{\beta})$. The result follows.

QED

References

Ahlee, H. and Malmendier, U. (2005). Do Consumers Know Their Willingness to Pay? Evidence from eBay Auctions. Mimeo, University of Stanford.

Akerlof, G. A. Procrastination and Obedience. (1991) *American Economic Review* 81(2), 1-19.

Ausubel, L.M. and Shui H. (2004). Time Inconsistency in the Credit Card Market. Mimeo, University of Maryland.

Ariely, D. and Wertenbroch, K. (2002). Procrastination, Deadlines and Performance: Self-Control By Precommitment, *Psychological Sciences*, 13, 219-224.

Arkes, H. R. and Blumer, C. (1985). The psychology of sunk cost, *Organizational Behavior and Human Decision Processes*, 35(1), 124-140.

Bell, D. E. (1982). Regret in Decision Making Under Uncertainty, *Operations Research*, 30(5), 961-981.

Bulkeley, W. M. (1998). Rebates' Secret Appeal to Manufacturers: Few Consumers Actually Redeem Them, *The Wall Street Journal*, February 10, B1.

Chip, H., Heath, S. and Lang, M. (1999). Psychological Factors and Stock Option Exercise, *Quarterly Journal of Economics*, 64 (2), 601-28.

DellaVigna, S. and Malmendier, U. (2006). Paying Not to Go to the Gym, *American Economic Review*, forthcoming.

Dhar, S. K. and Hoch, S. J. (1996). Price Discrimination Using In-Store Merchandising, *Journal of Marketing*, 60, 17-30.

Eyster, E. (2002). Rationalizing the Past: A Taste for Consistency, Mimeo, Nuffield College, Oxford.

Fishburn, P. C. (1982). Nontransitive Measurable Utility, *Journal of Mathematical Psychology*, 26, 31-67.

Gabaix, Xavier and Laibson, David (2006) "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets", *Quarterly Journal of Economics*, 121 (2).

Greenman, C. (1999). The Trouble with Rebates, *The New York Times*, September 16, G1.

- Gourville, J. T. and Soman, D. (2002). Pricing and the Psychology of Consumption, *Harvard Business Review*, Vol. 80, No. 9.
- Gourville, J. T. and Soman, D. (1998). Payment depreciation: the effects of temporally separating payments from consumption, *Journal of Consumer Research*.
- Ho, T., Lim, N and Camerer, C. (2005) "Modeling the Psychology of Consumer and Firm Behavior with Behavioral Economics", *Journal of Marketing Research*, forthcoming.
- Jolson, M., Weiner, J. and Rosecky, R. (1987). Correlates of Rebate Proneness, *Journal of Advertising Research*, 27, 33-43.
- Laibson, D. I. (1997). Golden Eggs and Hyperbolic Discounting, *Quarterly Journal of Economics*, 112, 443-477.
- Loomes, G. and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92, 805-824.
- Loomes, G. and Sugden, R. (1983). A rationale for preference reversal. *American Economic Review*, 73, 428-432.
- Madrian, B. and Shea, D.F. (2001). The Power of Suggestion: Inertia in 401(k) Participation and Saving Behavior, *Quarterly Journal of Economics*, 116, 1149-1525.
- O'Donoghue, T. D. and Rabin, M. (1999a). Doing It Now or Later, *American Economic Review*, 89(1), 103-133.
- O'Donoghue, T. D. and Rabin, M. (2001). Choice and Procrastination, *Quarterly Journal of Economics*, 116, 121-159.
- Olster, S. and Morton, F. (2006). Behavioral Biases Meet the Market: The Case of Magazine Subscription Prices, *Berkeley Electronic Journals Economic Analysis and Policy*, Advances: Vol. 5, No. 1
- Phelps, E. S., and Pollak, R. (1968). On Second-Best National Saving and Game-Equilibrium Growth, *Review of Economic Studies*, 35, 185 - 199.
- PMA, (2005). Annual Report on Rebate Promotion, Promotion Marketing Association.
- Shafir, E. and Tversky, A. (1992). Choice under Conflict: the Dynamics of Deferred Decisions, *Psychological Science*, 3, 358-361.
- Shefrin, H. (2000). Recent Development in Behavioral Finance, *Journal of Private Portfolio Management*, 107-17.

Silk, T. G. (2005). Why Do We Buy but Fail to Redeem? The Influence of Pre-Purchase, Post-Purchase, and Repeat-Purchase Factors On Consumers' Subjective and Objective Probability of Redeeming Mail-in Rebates, Mimeo, University of Florida.

Soman, D. (1998). The Illusion of Delayed Incentives: Evaluating Future Effort-Money Transactions, *Journal of Marketing Research*, 35, 427-437.

Thaler, R. H. (1980). Toward a positive theory of consumer choice, *Journal of Economic Behavior and Organization*, 1 (1980), 39-60.

Thaler, R. H. (1985). Mental accounting and consumer choice. *Marketing Science* 4, 199-214.

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, 12, 183-206.

Thaler, R. H. and Johnson, E. J. (1990). Gambling with the House Money and Trying to Break Even: the Effects of Prior Outcomes on Risky Choice, *Management Science*, 36(6), 643-660.