

Psychology and Economics: Evidence from the Field*

Stefano DellaVigna
UC Berkeley and NBER
sdellavi@berkeley.edu

Forthcoming, *Journal of Economic Literature*

Abstract

The research in Psychology and Economics (a.k.a. Behavioral Economics) suggests that individuals deviate from the standard model in three respects: (i) non-standard preferences; (ii) non-standard beliefs; and (iii) non-standard decision-making. In this paper, I survey the empirical evidence from the field on these three classes of deviations. The evidence covers a number of applications, from consumption to finance, from crime to voting, from charitable giving to labor supply. In the class of non-standard preferences, I discuss time preferences (self-control problems), risk preferences (reference dependence), and social preferences. On non-standard beliefs, I present evidence on overconfidence, on the law of small numbers, and on projection bias. Regarding non-standard decision-making, I cover framing, limited attention, menu effects, persuasion and social pressure, and emotions. I also present evidence on how rational actors—firms, employers, CEOs, investors, and politicians—respond to the non-standard behavior described in the survey. Finally, I briefly discuss under what conditions experience and market interactions limit the impact of the non-standard features.

*I would like to thank Roger Gordon (the editor), three exceptionally careful referees, Dan Acland, Malcolm Baker, Brad Barber, Nicholas Barberis, Dan Benjamin, Saurabh Bhargava, Colin Camerer, David Card, Raj Chetty, James Choi, Sanjit Dhami, Constanca Esteves, Ernst Fehr, Shane Frederick, Drew Fudenberg, David Hirshleifer, Eric Johnson, Lawrence F. Katz, Georg Kirchsteiger, Jeffrey Kling, Howard Kunreuther, David Laibson, George Loewenstein, Erzo F.P. Luttmer, Rosario Macera, Ulrike Malmendier, MichelAndre Marechal, John Morgan, Ted O'Donoghue, Ignacio Palacios-Huerta, Joshua Palmer, Vikram Pathania, Matthew Rabin, Ricardo Reis, Uri Simonsohn, Rani Spiegler, Bjarne Steffen, Justin Sydnor, Richard Thaler, Jeremy Tobacman, Michael Urbancic, Ebonya Washington, Kathryn Zeiler, and Jonathan Zinman for useful comments and suggestions. Thomas Barrios, Charles Lin, and Anitha Sivasankaran provided excellent research assistance. I also want to thank the students of my Psychology and Economics graduate class who over the years helped shape the ideas in this paper. Finally, I would like to express all my gratitude to David Laibson and Matthew Rabin for their exceptional generosity in sharing their insights with the next generations of behavioral economists. This paper would not have been possible without them.

1 Introduction

The core theory used in economics builds on a simple but powerful model of behavior. Individuals make choices so as to maximize a utility function, using the information available, and processing this information appropriately. Individuals' preferences are assumed to be time-consistent, affected only by own payoffs, and independent of the framing of the decision.

Laboratory experiments in both the psychology and the economics literature raise serious questions about these assumptions, though. In the laboratory, individuals are time-inconsistent (Thaler, 1981), show a concern for the welfare of others (Charness and Rabin 2002, Fehr and Gächter, 2000), and exhibit an attitude toward risk that depends on framing and reference points (Kahneman and Tversky, 1979). They violate rational expectations, for example by overestimating their own skills (Camerer and Lovallo, 1999) and overprojecting from the current state (Read and van Leeuwen, 1998). They use heuristics to solve complex problems (Gabaix, Laibson, Moloche, and Weinberg, 2006) and are affected by transient emotions in their decisions (Loewenstein and Lerner, 2003).

Unclear from these experiments, though, is how much these deviations from the standard theory in the laboratory affect economic decisions in the field. In markets people hone their behavioral rules to match the incentives they face and sort into favorable economic settings (Levitt and List, 2007). This is likely to limit the impact of deviations from the standard model in markets. However, other forces are likely to increase the impact. Firms often have incentives to accentuate the deviations of consumers to profit from them (DellaVigna and Malmendier, 2004). In addition, important economic decisions such as the choice of retirement savings or a house purchase are taken seldom, with limited scope for feedback and sorting.

The objective of this paper is to summarize a growing list of recent papers that document aspects of behavior in market settings that also deviate from the forecasts of the standard theory. This research area is known as Psychology and Economics (or Behavioral Economics). The evidence suggests deviations from the standard theory in each step of the decision-making process: 1) non-standard preferences, 2) incorrect beliefs, and 3) systematic biases in decision-making. For each of these three steps, I present an example of the laboratory evidence, introduce a simple model if available, and summarize the strength and weaknesses of the field evidence. Since the focus of the paper is on the field evidence, I do not survey the laboratory evidence or the theoretical literature.

To fix ideas, consider the following stylized version of the standard model, modified from Rabin (2002a). Individual i at time $t = 0$ maximizes expected utility subject to a probability

distribution $p(s)$ of the states of the world $s \in S$:

$$\max_{x_i^t \in X_i} \sum_{t=0}^{\infty} \delta^t \sum_{s_t \in S_t} p(s_t) U(x_i^t | s_t). \quad (1)$$

The utility function $U(x|s)$ is defined over the payoff x_i^t of player i and future utility is discounted with a (time-consistent) discount factor δ .

The first class of deviations from the standard model in (1) is non-standard preferences, discussed in Section 2. I focus on three dimensions: time preferences, risk preferences, and social preferences. With respect to time preferences, the findings on self-control problems, for example in retirement savings, challenge the assumption that the discount factor, δ , is time-consistent. With respect to risk preferences, the evidence, for example in insurance decisions, suggests that the utility function $U(x_i|s)$ depends on a reference point r : the utility function becomes $U(x_i|r, s)$. With respect to social preferences, the evidence, for example on charitable giving, suggests that the utility function depends also on the payoff of other people x_{-i} : the utility is $U(x_i, x_{-i}|s)$. The research on non-standard preferences constitutes the bulk of the empirical research in Psychology and Economics.

The second class of deviations from the standard model in (1) is non-standard beliefs $\tilde{p}(s) \neq p(s)$, reviewed in Section 3. Systematic overconfidence about own ability can help explain managerial behavior of CEOs. Non-Bayesian forecasting rationalizes ‘gambler’s fallacy’ behavior in lotteries and overinference from past stock returns. The overprojection of current tastes on future tastes can explain aspects of the purchase of seasonal items.

The third class of deviations from the standard model is non-standard decision-making, discussed in Section 4. For given utility $U(x|s)$ and beliefs $p(s)$, individuals resort to heuristics (Tversky and Kahneman, 1974) instead of solving the complex maximization problem (1). They are affected by the framing of a decision problem, for example in investment decisions. They simplify a complex decision by being inattentive to less salient features of a problem, from asset allocation to purchase decisions. They use sub-optimal heuristics when choosing from a menu of options X_i , such as for savings plans or loan terms. They are also subject to social pressure and persuasion, for example in their workplace performance and in voting decisions. Finally, they are affected by emotions, as in the case of investment decisions.

While I organize the deviations in three separate classes, the three types of deviations are often related. For example, persuasion leads to different decisions through the change in beliefs that it induces.

The first part of the paper provides evidence that these deviations affect the behavior of individual decision-makers, such as consumers and small investors. But are these deviations large enough to matter for our theories of how markets and institutions work? In Section 5, I provide evidence on how rational actors respond to these behavioral anomalies. In particular, I discuss the response of firms, employers, managers, investors, and politicians. These agents

appear to have changed their own behavior in ways that would be puzzling given the standard theory but that are consistent with utility-maximizing responses to the documented behavioral anomalies. This provides indirect support for the lab and field evidence of the anomalies.

Given this evidence, I expect that the documented deviations from the standard model will be increasingly incorporated in economic models. Indeed, features such as time inconsistency and reference dependence have become common assumptions. In the concluding Section, I present remarks on why these deviations matter also in the field; I also discuss directions for future research in Psychology and Economics.

This overview differs from other surveys of Psychology and Economics (Rabin, 1998; Rabin, 2002a; Mullainathan and Thaler, 2001; Camerer, 2005) because it focuses on empirical research using non-laboratory data. A number of caveats are in order. First, this paper does not provide an overview by field of application, since it is organized instead by psychological principles; the interested reader can consult as a starting point the book chapters in Diamond and Vartiainen (2007). Second, the emphasis of the paper is on (relatively) detailed summaries of a small number of papers for each deviation. As such, the survey provides a selective coverage of the field evidence, though it strives to cover all the important deviations.¹ Finally, this overview undersamples studies in Marketing and provides a partial coverage of the research in Behavioral Finance, probably the most developed application of Psychology and Economics, for which a comprehensive survey of the empirical findings is available (Barberis and Thaler, 2004).

2 Non-standard Preferences

2.1 Self-Control Problems

The standard model (1) assumes a discount factor δ between any two time periods that is independent of when the utility is evaluated. This assumption implies time consistency, that is, the decision maker has the same preferences about future plans at different points in time.²

Laboratory Experiments. Experiments on intertemporal choice, summarized in Loewenstein and Prelec (1992) and Frederick, Loewenstein, and O'Donoghue (2002), have cast doubt on this assumption. This evidence suggests that discounting is steeper in the immediate future than in the further future. For example, the median subject in Thaler (1981) is indifferent between \$15 now and \$20 in one month (for an annual discount rate of 345 percent) and between

¹This overview does not discuss deviations from the standard model that are widely documented in experiments but not in the field, such as will-power exhaustion and the availability heuristics.

²Strictly speaking, the standard model merely assumes time consistency, not a constant discount factor δ . Still, most of the evidence in this Section—the adoption of costly commitments or behavior that differs from the plans—directly violates time consistency and hence also this more general version of the standard model.

\$15 now and \$100 in ten years (for an annual discount rate of 19 percent).³ The preference for immediate gratification captured in these studies appears to have identifiable neural underpinnings. Intertemporal decisions involving payoffs in the present activate different neural systems than decisions involving only payoffs in future periods (McClure et al., 2004).

Intertemporal preferences with these features capture *self-control problems*. When evaluating outcomes in the distant future, individuals are patient and make plans to exercise, stop smoking, and look for a better job. As the future gets near, the discounting gets steep, and the individuals engage in binge eating, light another (last) cigarette, and stay put on their job. Preferences with these features therefore induce time inconsistency.

Model. Laibson (1997) and O’Donoghue and Rabin (1999a) formalized these preferences using (β, δ) preferences⁴, building on Strotz (1956), Phelps and Pollak (1968), and Akerlof (1991). Labelling as u_t the per-period utility, the overall utility at time t , U_t , is

$$U_t = u_t + \beta\delta u_{t+1} + \beta\delta^2 u_{t+2} + \beta\delta^3 u_{t+3} + \dots$$

The only difference from the standard model (with δ as the discount factor) is the parameter $\beta \leq 1$, capturing the self-control problems. For $\beta < 1$, the discounting between the present and the future is higher than between any future time periods, capturing the main finding of the experiments. For $\beta = 1$, this reduces to the standard model.

A second key element in this model is the modelling of expectations about future time preferences. O’Donoghue and Rabin (2001) allow the agent to be partially naive (that is, overconfident) about the future self-control problems. A partially naive (β, δ) agent expects in the future period $t + s$ to have the utility function

$$\hat{U}_{t+s} = u_{t+s} + \hat{\beta}\delta u_{t+s+1} + \hat{\beta}\delta^2 u_{t+s+2} + \hat{\beta}\delta^3 u_{t+s+3} + \dots$$

with $\hat{\beta} \geq \beta$. The agent may be sophisticated about the self-control problem ($\hat{\beta} = \beta$), fully naive ($\hat{\beta} = 1$), or somewhere in between. This model, therefore, combines self-control problems with a form of overconfidence, naiveté about future self-control.

Other models have been proposed to capture self-control problems, including axiomatic models that emphasize preferences over choice sets (Gul and Pesendorfer, 2001) and models of the conflict between two systems, a planner and a doer (Shefrin and Thaler, 1981 and

³The laboratory experiments on time preferences face at least three issues: (i) most experiments are over hypothetical choices, including Thaler (1981); (ii) in the experiments with real payments, issues of credibility regarding the future payments can induce seeming present bias; (iii) the discounting should apply to consumption units, rather than to money (in theory, over monetary outcomes, only the interest rate should matter). While none of the experiments fully addresses all three issues, the consistency of the evidence suggests that the phenomenon is genuine.

⁴These preferences are also labelled quasi-hyperbolic preferences, to distinguish them from (pure) hyperbolic preferences, and present-biased preferences.

Fudenberg and Levine, 2006, among others). For lack of space, and since most applied work has referred to the (β, δ) model, we refer only to this latter model in what follows.

As an example of how the (β, δ) model operates, consider a good with immediate payoff (relative to a comparison activity) b_1 at $t = 1$ and delayed payoff b_2 at $t = 2$. An investment good, like exercising or searching for a job, has the features $b_1 < 0$ and $b_2 > 0$: the good requires effort at present and delivers happiness tomorrow. Conversely, a leisure good, like consumption of tempting food or watching TV, has the features $b_1 > 0$ and $b_2 < 0$: it provides an immediate reward, at a future cost.

How often does the agent *want* to consume, from an ex ante perspective? If the agent could set consumption one period in advance, at $t = 0$, she would consume if $\beta\delta b_1 + \beta\delta^2 b_2 \geq 0$, or

$$b_1 + \delta b_2 \geq 0. \tag{2}$$

(Notice that β cancels out, since all payoffs are in the future)

How much does the agent *actually* consume at $t = 1$? The agent consumes if

$$b_1 + \beta\delta b_2 \geq 0. \tag{3}$$

Compared to the desired consumption, therefore, a (β, δ) agent consumes too little investment good ($b_2 > 0$) and too much leisure good ($b_2 < 0$). This is the self-control problem in action. In response, a sophisticated agent looks for commitment devices to increase the consumption of investment goods and to reduce the consumption of leisure goods.

Finally, how much does the agent *expect* to consume? The agent expects to consume in the future if

$$b_1 + \hat{\beta}\delta b_2 \geq 0, \tag{4}$$

with $\hat{\beta} \geq \beta$. Compared to the actual consumption in (3), the agent overestimates the consumption of the investment good ($b_2 > 0$) and underestimates the consumption of the leisure good ($b_2 < 0$). Naiveté therefore leads to mispredictions of future usage.

I now present evidence on the consumption of investment goods (exercise and homeworks) and leisure goods (credit card take-up and life-cycle savings) that can be interpreted in light of this simple model.

Exercise. DellaVigna and Malmendier (2006) use data from three US health clubs offering a choice between a monthly contract X_M with lump-sum fee L of approximately \$80 per month and no payment per visit, and a pay-per-visit contract X_p with fee p of \$10. Denote by $E(x_M) |_{X_M}$ the expected number of monthly visits under the monthly contract X_M . Under the standard model, individuals choosing the monthly contract must believe that $pE(x_M) |_{X_M} \geq L$, or $L/E(x_M) |_{X_M} \leq p$: the price per expected attendances under the monthly contract should be lower than the fee under payment-per-usage. Otherwise, the individual should have chosen the pay-per-usage treatment. DellaVigna and Malmendier (2006), however, find that health

club users that choose the monthly contract X_M attend only 4.8 times per month. These users pay \$17 per visit even though they could pay \$10 per visit, a puzzle for the standard model. A model with partially naive (β, δ) members suggests two explanations for this finding. The users may be purchasing a commitment device to exercise more: the monthly membership reduces the marginal cost of a visit from \$10 to \$0, and helps to align actual attendance in (3) with desired attendance in (2). Alternatively, these agents may be overestimating their future health club attendance, as in (4). Direct survey evidence on expectation of attendance and evidence on contract renewal are most consistent with the latter interpretation.⁵

Homework and Deadlines. Ariely and Wertenbroch (2002) present evidence on homework completion and deadlines. The subjects are 51 professionals enrolled in a section of a semester-long executive education class at Sloan (MIT), with three homeworks as a requirement. At the beginning of the semester, they set binding deadlines (with a cost of lower grades for delay) for each of the homeworks. According to the standard model, they should set deadlines for the last day of the semester: there is no benefit to setting early deadlines, since they do not receive feedback on the homeworks, and there is a cost of lower flexibility. (A maximization without constraints is always preferable to one with constraints.) According to a model of self-control, instead, the deadlines provide a useful commitment device. Since homework completion is an investment good ($b_2 > 0$), individuals spend less time on it than they wish to ex ante (compare equations (2) and (3)). A deadline forces the future self to spend more time on the assignment. The results support the self-control model: 68 percent of the deadlines are set for weeks prior to the last week, indicating a demand for commitment.⁶

This result leaves open two issues. First, do the self-set deadlines improve performance relative to a setting with no deadlines? Second, is the deadline setting optimal? If the individuals are partially naive about the self-control, they will under-estimate the demand for commitment (equation (4)). In a second (laboratory) experiment, Ariely and Wertenbroch (2002) address both issues. Sixty students complete three proofreading assignments within 21 days. The control group can turn in each assignment at any time within the 21 days, a first treatment group can choose three deadlines (as in the class-room setting described above), and a second treatment group faces equal-spaced deadlines. The first result is that self-set deadlines indeed improve performance: the first treatment group does significantly better than the control group, detecting 50 percent more errors (on average, 105 versus 70) and earning substantially more as a result (on average, \$13 versus \$5). The second result is that the deadline setting is not optimal: the group with equal-spaced deadlines does significantly better than the other groups, on average detecting 130 errors and earning \$20. This provides evidence of

⁵In Section 5, I discuss how the contracts offered by health club companies are consistent with the assumption of naive (β, δ) consumers (DellaVigna and Malmendier, 2004).

⁶Ariely and Wertenbroch (2002) also compare the performance in this section to the performance in another section with equal-spaced deadlines, with results similar to the ones described below. However, the students are not randomly assigned to the two sections.

partial naiveté about the self-control problems.

Credit Card Take-up. Ausubel (1999) provides evidence on credit card usage using a large-scale field experiment run by a credit card company. The company mailed randomized credit card offers, varying both the pre-teaser and the post-teaser interest rates. For example, compared to an offer of 6.9% interest rate for six months and 16% thereafter (the control group), the treatment group ‘Pre’ received a lower pre-teaser rate (4.9% followed by 16%); the treatment group ‘Post’, instead, received a lower post-teaser rate (6.9% followed by 14%). For each offer, Ausubel (1999) observes the response rate and 21 months of history of borrowing for the individuals that take the card. Across these offers, the average balance borrowed in the first 6 months is about \$2,000, while the average balance in the subsequent 15 months is about \$1,000.⁷ Given these borrowing rates, the standard theory predicts that the increase in response rate for treatment ‘Pre’ (relative to the control group) should be smaller than for treatment ‘Post’: neglecting compounded interest, $6/12 * 2% * \$2,000$ is smaller than $15/12 * 2% * \$1000$ (the comparison would only be more favorable for the ‘Post’ treatment if we could observe the balances past 21 months). Instead, the increase in take-up rate for the ‘Pre’ treatment (386 people out of 100,000) is 2.5 times larger than the increase for the ‘Post’ treatment (154 people out of 100,000). Individuals over-respond to the pre-teaser interest rate. Ausubel’s interpretation of this result is that individuals (naively) believe that they will not borrow much on a credit card, past the teaser period. These findings are consistent with underestimation of future consumption for leisure goods, as in (4).

Life-Cycle Savings. The (β, δ) model of self-control can also help explain puzzling features of life-cycle accumulation and credit-card borrowing data. Building on Laibson (1997) and Angeletos et al. (2001), Laibson, Repetto and Tobacman (forthcoming) estimate a fully-specified model of life-cycle accumulation with liquid and illiquid saving. They show that the (β, δ) model can reconcile two facts: high credit card borrowing (11.7 percent of annual income) and substantial illiquid wealth accumulation (216 percent of annual income for the median consumer of age 50-59).⁸ Standard models have a hard time explaining both facts, since credit card borrowing implies high impatience, which is at odds with substantial wealth accumulation. The model with self-control problems predicts high spending on liquid assets, but also a high demand for illiquid assets, which work as commitment devices.

Ashraf, Karlan, and Yin (2005) document directly the demand for illiquid savings as a commitment device, and its effect. They offer an account with a commitment device to 842 randomly determined households in the Philippines with a pre-existent bank account. Access to funds in these accounts is constrained to reaching a self-specified savings goal or a self-

⁷Of course, the differences in interest rates will affect the borrowing directly, through incentive and selection effects. However, these differences are small enough in the data that we can, to a first approximation, neglect them in these calculations.

⁸The figures (from Laibson et al., 2006) refer to high-school graduates.

specified time period. A control group of 466 households from the same sample is offered a verbal encouragement to save but with no commitment. The results reveal a sizeable demand for commitment, and an impact of commitment on savings. In the treatment group, 202 of 842 households take up the commitment savings product. In the 842 treatment households, savings in the bank after six months are significantly more likely to increase, compared to the 466 control households that received a pure encouragement: the share of households with increased savings is 33.3 percent in the treatment and 27.7 in the control.⁹ The comparison includes individuals in the treatment group that do not take up the commitment savings product; the treatment-on-the-treated estimate is larger by a factor of 842/202. Benartzi and Thaler (2004), described in Section 5 below, provide evidence of substantial demand for commitment devices in retirement savings in the US.

Skiba and Tobacman (2008) examine the role of self-control in the demand for payday loans, one of the fastest-growing sources of credit in the US, with ten million households borrowing in 2002. These loans provide (typically) two weeks of liquidity for annualized (compounded) interest rates of over 7000%. Using an administrative data set from a payday lender, Skiba and Tobacman (2008) fit a life-cycle consumption model to the observed borrowing behavior. While an exponential time discounting model with a high discount rate can rationalize the observed borrowing at a high interest rate, it has difficulty explaining the relatively low default rate. Defaulting borrowers have on average already repaid or serviced five payday loans, making interest payments of 90 percent of their original loan’s principal. Impatient but time-consistent borrowers would take advantage of default sooner, or borrow less. The observed facts are most consistent with a model of partially naive (β, δ) consumers. These consumers borrow today expecting to borrow less in the future and procrastinate defaulting (which is assumed to have immediate monetary or stigma costs).

Default Effects in 401(k)s. The impact of the status quo (default) in retirement savings is the final set of findings bearing on self-control problems, as I discuss below.¹⁰ Madrian and Shea (2001) consider the effect on the contribution rates in 401(k)s of a change in default. Before the change, the default is non-participation in retirement savings; after the change, the default is participation at a 3% rate in a money market fund. In both cases, employees can override the default with a phone call or by filing a form; also, in both cases, contributions receive a 50 percent match up to 6% of compensation. Madrian and Shea (2001) find that the change in default has a very large impact: one year after joining the company, the participation rate in 401(k)s is 86% for the treatment group and 49% for the control group.

Choi et al. (2004) show that these findings generalize to six companies in different industries with remarkably similar effect sizes. This finding is not limited to retirement choices in the

⁹These figures refer to the total bank balance across all accounts for a household, that is, they are not due to switches of savings from an ordinary account to the account with commitment device.

¹⁰Samuelson and Zeckhauser (1988) is an early paper documenting default effects.

U.S.. Cronqvist and Thaler (2004) examine the choice of retirement funds in Sweden after the privatization of social security in the year 2000. They find that 43.3 percent of new participants choose the default plan, despite the fact that the government encouraged individual choice, and despite the availability of 456 plans. Three years later, after the end of the advertisement campaign encouraging individual choice, the proportion choosing the default plan increased to 91.6 percent. Overall, the finding of large default effects is one of the most robust results in the applied economics literature of the last ten years.¹¹

What explains the large default effect for retirement savings? Transaction costs alone are unlikely to explain default effects. Employees can change their retirement decisions at any time using the phone or a written form. Such small transaction costs are dwarfed by the tax advantages of 401(k) investments, particularly in light of the 50 percent match (up to 6% of compensation) in place at the Madrian and Shea (2001) company. At a mean compensation of about \$40,000, the match provides a yearly benefit of \$1,200, assuming a discount rate equal to the interest rate. It is hard to imagine transaction costs of this size.

O'Donoghue and Rabin (1999b and 2001) show that self-control problems, combined with naiveté, can explain the observed default effect even for small transaction costs. Consider a naive (β, δ) agent that has to decide when to undertake a decision with immediate disutility from transaction costs $b_1 < 0$ and delayed benefit $b_2 > 0$, such as enrolling in retirement savings. This agent would rather postpone this activity, given the self-control problems, as in equation (3). Moreover, this agent is (incorrectly) convinced that if she does not do the activity today, she'll do it tomorrow, as in (4). This agent postpones the activity day-after-day, ending up never doing it. O'Donoghue and Rabin (2001) show that, in the presence of naiveté, even a small degree of self-control problems can generate (infinite) procrastination. These papers distinguish between 'procrastination' and 'delay': procrastination is a delay that ex ante the agent does not anticipate. O'Donoghue and Rabin (1999b) presents calibrations in a deterministic set-up, which DellaVigna and Malmendier (2006) extend to the case of stochastic transaction costs. Both papers show that naive (β, δ) agents accumulate substantial delays in a costly activity (respectively, signing up for a 401(k) and cancelling a health club membership).

Both self-control problems *and* naiveté are required to explain the observed default effects. Unlike naive agents, sophisticated (β, δ) agents do not exhibit large default effects for reasonable parameter values (O'Donoghue and Rabin, 2001). While sophisticated agents would like to postpone activities with immediate costs, they realize that doing an activity now is better than postponing it for a long time.¹²

¹¹Default effects matter in other decisions, such as contractual choice in health-clubs (DellaVigna and Malmendier, 2006), organ donation (Johnson and Goldstein, 2003; Abadie and Gay, 2006), and car insurance plan choice (Johnson et al, 1993).

¹²There are a number of alternative interpretations of the observed default effects, such as inattention and limited memory about 401(k) investment. However, most explanations are unlikely to match the magnitudes of the delay. For example, consider a sophisticated agent with limited memory. Being rational, she is aware that,

If procrastination of a financial transaction is indeed responsible for the default effects in Madrian and Shea (2001) and in Choi et al. (2004), we should expect that, if individuals were forced to make an active choice at enrollment, they would display their true preferences for savings. In this case, they bear the transaction cost whether they invest or not, and hence investing does not have an immediate cost, i.e., $b_1 = 0$. In this situation, the short-run self does not desire to postpone the choice. Choi et al. (2005) analyze a company that required its employees to choose the retirement savings at enrollment. Under this Active Decision plan, 80% of workers enrolled in a 401(k) within one year of joining the company. Later, this company switched to a no-investment default, and the one-year enrollment rate declined to 50%. Requiring workers to choose, therefore, produces an enrollment rate that is only slightly lower than under the automatic enrollment in Madrian and Shea (2001).¹³

Welfare. These studies have welfare and policy implications. They suggest that savings rates for retirement in the US may be low due to a combination of procrastination and defaults set to no savings. The (β, δ) model implies that the individuals are likely to be happier with defaults set to higher savings rates. A change in policy with defaults set to automatic enrollment is an example of cautious paternalism (Camerer et al., 2003), in that it helps substantially individuals with self-control problems and inflicts little or no harm on individuals without self-control problems. These individuals can switch to a different savings rate for a low transaction cost. In Section 5, we present the results of a plan with automatic enrollment and other features designed to increase savings (Benartzi and Thaler, 2004). An alternative design could be based on the requirement to make an active choice, as in Choi et al. (2005). Social Security is a commitment device to save, albeit one that consumers cannot opt out of, and that thus can hurt consumers with no self-control problems.

Summary. A model of self-control problems with partial naiveté can rationalize a number of findings that are puzzling to the standard exponential model: (i) excessive preference for membership contracts in health clubs; (ii) positive effect of deadlines on homework grades and preference for deadlines; (iii) near-neglect of post-teaser interest rates in credit-card take-up; (iv) liquid debt and illiquid saving in life-cycle accumulation; (v) demand for illiquid savings as commitment devices; (vi) default effects in retirement savings and in other settings.

The partially-naive (β, δ) model, therefore, does a good job of explaining qualitative patterns across a variety of settings involving self-control. A frontier of this research agenda is to establish whether one model can fit these different facts not just qualitatively, but also quantitatively. Quantitative calibrations, for example, are crucial to understand default effects.

if she does not sign up at employment, she will not sign up until the next time she remembers. If the anticipated delay is long enough, the agent prefers to sign up immediately, hence exhibiting very limited default effects. A model of limited memory with added naiveté explains the findings in a similar way to the naive (β, δ) model.

¹³The effect of the Active Decision may also be due to a deadline effect for naive (β, δ) employees, who know that the next occasion to enroll will not be until several months later.

Pursuing this agenda further, a few papers have estimated values for the time preference parameters. Laibson, Repetto, and Tobacman (forthcoming) estimate annual time preference parameters ($\beta = .70, \delta = .96$) on life-cycle accumulation data. Paserman (forthcoming), building on DellaVigna and Paserman (2005), uses job search data to estimate ($\beta = .40, \delta = .99$) for low-wage workers and ($\beta = .89, \delta = .99$) for high-wage workers. Both papers assume sophistication. Skiba and Tobacman (2008) allow for partial naiveté and estimate ($\beta = .53, \hat{\beta} = .90, \delta = .45$) for the sample of payday loan borrowers.¹⁴

2.2 Reference Dependence

The simplest version of the standard model as in (1) assumes that individuals maximize a global utility function over lifetime consumption $U(x|s)$.

Laboratory Experiments. A set of experiments on attitude toward risk call into question the assumption of a global utility function. An example (using hypothetical questions) from Kahneman and Tversky (1979) illustrates the point. A group of 70 subjects is asked to consider the situation: “In addition to whatever you own, you have been given 1,000. You are now asked to choose between A: (1,000, .50), and B: (500),” where (1,000, .50) indicates a lottery that assigns .50 probability of 1,000 and .50 probability of 0. A different group of 68 subjects is asked to consider: “In addition to whatever you own, you have been given 2,000. You are now asked to choose between C: (-1,000, .50), and D: (-500).” The allocations A and C are identical, and so are B and D. However, in the first group only 16 percent of the subjects choose A, in contrast with 69 percent of subjects choosing C in the second group. Clearly, framing matters (see also Section 4.1).

Choices in lotteries with real payoffs display similar violation of the standard theory. In Fehr and Goette (2007), 27 out of 42 subjects prefer 0 Swiss Franks for sure to the lottery (-5, $p = .5$; 8, $p = .5$). Under the standard model, this implies an unreasonably high level of risk aversion (Rabin, 2000). A subject that made this choice for all wealth levels would also reject the lottery (-31, $p = .5$; $\infty, p = .5$), which offers an infinite payout with probability .5.

Model. Kahneman and Tversky (1979), in the second most cited article in economics since 1970 (Kim, Morse, and Zingales, 2006), propose a reference-dependent model of preferences that, unlike the standard model, can fit most of the experimental evidence on lottery choice. According to prospect theory, subjects evaluate a lottery $(y, p; z, 1 - p)$ as follows: $\pi(p)v(y - r) + \pi(1 - p)v(z - r)$. Prospect theory is characterized by: (i) *Reference Dependence*. The value function v is defined over differences from a reference point r , instead of over the overall wealth; (ii) *Loss Aversion*. The value function $v(x)$ has a kink at the reference

¹⁴In Paserman (2006) (respectively, Skiba and Tobacman, 2008), the model is estimated at the weekly (bi-weekly) level, so the β parameter refers to the one-week (two-week) discounting. The δ parameter is the annualized equivalent.

point and is steeper for losses ($x < 0$) than for gains ($x > 0$); (iii) *Diminishing Sensitivity*. The value function v is concave over gains and convex over losses, reflecting diminishing sensitivity to outcomes further from the reference point; (iv) *Probability weighting*. The decision-maker transforms the probabilities with a probability-weighting function $\pi(p)$ that overweights small probabilities and underweights large probabilities.

The four features of prospect theory are designed to capture the evidence on risk-taking, including risk-aversion over gains, risk-seeking over losses, and contemporaneous preference for insurance and gambling. It can also capture framing effects as in the example above. Lottery A is evaluated as $\pi(.5)v(1,000)$ and hence, given the concavity of $v(x)$ for positive x and given $\pi(.5) \approx .5$, is inferior to lottery B, valued $v(500)$. Conversely, lottery C is evaluated as $\pi(.5)v(-1,000)$ and, given the convexity of $v(x)$ for negative x , is preferred to lottery D.

The large majority of the follow-up literature, however, adopts a simplified version of prospect theory incorporating only features (i) and (ii). The subjects maximize $\sum_i p_i v(x_i|r)$, where $v(x|r)$ is defined as

$$v(x|r) = \begin{cases} x - r & \text{if } x \geq r; \\ \lambda(x - r) & \text{if } x < r, \end{cases} \quad (5)$$

where $\lambda > 1$ denotes the loss aversion parameter. Prospect theory, even in the simplified version of expression (5), can explain the aversion to small risk exhibited experimentally, as in the example above from Fehr and Goette (2007). A prospect-theoretic subject evaluates the lottery $(-5,.5; 8,.5)$ as $.5\lambda * (-5) + .5 * 8 = 4 - 2.5\lambda$. This subject prefers the status-quo for $\lambda > 8/5$. (The experimental evidence from Tversky and Kahneman (1992) suggests $\lambda \approx 2.25$).

A fifth feature of reference-dependent preferences is *Narrow Framing* (Barberis, Huang, and Thaler, 2006; Rabin and Weizsäcker, forthcoming). According to the standard economic model, a decision-maker offered a gamble integrates the risk induced by the gamble with the other sources of uncertainty she faces. For example, a bike messenger in Fehr and Goette (2007) offered the lottery $(-5,p = .5; 8,p = .5)$ should aggregate these risky earnings with the (highly volatile) earnings from the job, fluctuations in the value of the assets, etc., to determine how the lottery affects consumption utility. If she did so, even if she evaluated the lottery according to prospect theory, she would be very likely to accept the gamble (counterfactually). The background risk implies that the bike messenger is unlikely to be near the kink at the reference point; as such, she is affected only to a limited extent by the loss aversion. Barberis, Huang, and Thaler (2006) formalizes this argument showing that even non-expected utility theories that display first-order risk aversion (including prospect theory) do not accommodate the observed risk-taking behavior in the laboratory, unless one assumes narrow framing.

A decision-maker with Narrow Framing, instead, considers each risk in isolation and evaluates a lottery as if it were the only determinant of consumption. Indeed, this assumption is implicit in evaluating the lottery in Fehr and Goette (2007) as $.5\lambda * (-5) + .5 * 8$ (as we did above). Importantly, this assumption is routinely used in experimental papers (without ex-

PLICIT reference to narrow framing) to recover consumer preferences from observed behavior, in that all income from outside the experiment is ignored, e.g., Holt and Laury (2002). I discuss further the role of framing in Section 4.1.

I assume narrow framing in the following applications to economic phenomena, and discuss the application of reference dependence also to settings that do not involve risk (such as the endowment effect and labor supply).

Endowment Effect. A finding consistent with prospect theory and inconsistent with the standard model is the so-called endowment effect, an asymmetry in willingness to pay (WTP) and willingness to accept (WTA). In a laboratory experiment, Kahneman, Knetsch, and Thaler (1990) randomly allocate mugs to one group of subjects. They use an incentive-compatible procedure to elicit the WTA for subjects that received the mug, and the WTP for subjects that were not allocated the mug. According to the standard theory, the two valuations should on average be the same. The median WTA of \$5.75, however, is twice as large as the median WTP of \$2.25. Since theoretically wealth effects could explain this discrepancy, in a different experiment Kahneman, Knetsch and Thaler introduce choosers, alongside buyers and sellers. Choosers, who are not endowed with a mug, choose between a mug and a sum of money; the experimenters elicit the price that induces indifference. Their choice is formally identical to the choice of the sellers (except for the fact that the choosers are not endowed with the mug); hence, according to the standard theory, the sum of money that makes them indifferent should correspond to the WTA of sellers. Instead, in this experiment the median WTA for sellers is \$7.12, while the price for choosers is \$3.12 (and the WTP for buyers is \$2.87). The asymmetry between WTA and WTP, replicated in a number of studies, has important implications for economics such as low volume of trades in markets and inconsistencies in the elicitation of contingent valuations in environmental decisions.

The endowment effect is predicted by a reference-dependent utility function with loss-aversion $\lambda > 1$. We assume that the subjects do not exhibit loss aversion with respect to money.¹⁵ Assume that the utility of the subjects is $u(1)$ if they received a mug, and $u(0)$ otherwise, with $u(1) > u(0)$. Consider subjects with a piece-wise linear utility function (5), where the reference point r depends on whether the subjects were assigned a mug. Subjects with the mug have reference point $r = 1$ and assign utility $u(1) - u(1) = 0$ to keeping the mug and utility $\lambda[u(0) - u(1)] + p_{WTA}$ to selling the mug for the sum p_{WTA} . Subjects without the mug have reference point $r = 0$ and assign value $u(1) - u(0) - p_{WTP}$ to getting the mug at price p_{WTP} and utility $u(0) - u(0) = 0$ to keeping the status-quo. The prices that make both groups of subjects indifferent between having and not having the mug are

$$p_{WTA} = \lambda[u(1) - u(0)] \quad \text{and} \quad p_{WTP} = u(1) - u(0),$$

¹⁵Subjects in Kahneman, Knetsch, and Thaler (1990) do not exhibit the endowment effect with respect to tokens of monetary value. In the presence of loss aversion λ also with respect to money, the endowment effect is magnified: $p_{WTA} = \lambda p_C = \lambda^2 p_{WTP}$.

hence $p_{WTA} = \lambda p_{WTP}$. A loss-aversion parameter $\lambda = 5.75/2.25$ fits the evidence in Kahneman et al. (1990). Notice that choosers choose a mug if $u(1) - u(0) \geq p_C$, and hence $p_C = p_{WTP}$ with referent-dependent preferences, approximately as observed.

Plott and Zeiler (2005) criticize this set of experiments on the ground that the endowment effect may be due to procedural features, such as limited experience and lack of anonymity. They elicit the WTP and WTA for a mug after granting anonymity and providing practice. In addition in 2 of 3 sessions they provide extensive training consisting of 14 rounds of elicitation of WTA and WTP for lotteries. In contrast to Kahneman et al. (1990), they find no evidence of the endowment effect for mugs, with a mean WTA of \$5.56 and a mean WTP of \$6.62. Plott and Zeiler interpret this result as suggesting that, once one allows for experience and anonymity, the endowment effect disappears. However, an alternative interpretation is that their procedure does not generate an endowment effect because they did not vary the ‘endowment’ sufficiently. Namely, both WTA and WTP groups are given a mug at the beginning of the experiment (though only the WTA group is told that they own it).¹⁶ In addition, while Plott and Zeiler (2005) do not obtain any endowment effect for mugs, they find a significant endowment effect in the lottery rounds, as Isoni, Loomes, and Sugden (2008) point out. In Isoni, Loomes, and Sugden (2008)’s replication of the Plott and Zeiler (2005) procedure, the mean WTA is larger than the mean WTP by a factor of between 1.02 and 2.19, a substantial endowment effect. Isoni et al. (2008) also find some evidence (albeit insignificant) of an endowment effect for mugs (mean WTA over mean WTP of 1.19). Additional research will be needed to explain the discrepancies between these findings. In any case, the Kahneman et al. (1990) results are relevant to a number of economic decisions where anonymity is not perfect and experience is limited, such as buying or selling a house.

A different set of papers in this literature examines the endowment effect in the field, namely a sports card fair in List (2003 and 2004). By selection, the participants in a sports card fair have at least some experience with cards, but some subjects are more experienced than others.¹⁷ The design is as follows. List (2003) randomly assigns sports memorabilia A or B as compensation for filling out a questionnaire. After the questionnaire is filled out, the participants are asked whether they would like to switch their assigned memorabilia for the other one. Since the objects are chosen to be of comparable value, the standard model predicts trade about 50 percent of the time. Instead, subjects with below-average trading experience switch only 6.8 percent of the time. This provides evidence of the endowment effect in a natural economic setting where the subjects are familiar with the objects being traded and

¹⁶Consistent with this hypothesis, Knetsch and Wong (2007) replicate Plott and Zeiler (2005)’s results if they follow the same endowment procedure, but obtain a strong endowment effect if they place objects in front of the ‘owners’ only. (Although the latter treatment changes the wording too, making it difficult to disentangle the effects.)

¹⁷In List (2003), experienced subjects are those with self-reported frequency of card trading of at least 6 times a month, the mean in the sample.

with trading itself. This suggests that familiarity with the object and procedures is unlikely to be responsible for the differences between the results of Kahneman et al. (1990) and Plott and Zeiler (2005).

List (2003) then examines the behavior of subjects with high trading experience. These subjects are not only familiar with the objects but also trade cards at least 6 times a month. Unlike the less experienced traders, these subjects switch 46.7 percent of the time, displaying no endowment effect. The difference between the two groups is not due to the fact that inexperienced traders are approximately indifferent between the two memorabilia, and hence willing to stick to the status quo. In another treatment eliciting WTA and WTP, the WTA is substantially larger than the WTP for inexperienced subjects (18.53 versus 3.32), but not for experienced subjects (8.15 versus 6.27). Next, List (2003) attempts to test whether the difference between the two groups is due to self-selection of subjects without the endowment effect among the frequent traders, or is a causal effect of trading experience. In a follow-up study performed months later, the endowment effect decreases in the trading experience accumulated in the intervening months, supporting the latter interpretation. Finally, and most surprisingly, List (2004) shows that the more experienced card traders also display substantially less endowment effect with respect to other goods, such as chocolates and mugs.

Overall, the evidence in List (2003 and 2004) suggests that the endowment effect is a feature of trading behavior that market experience can temper.¹⁸ This evidence leaves open (at least) two interpretations. One interpretation is that experience with the market leads individuals to become aware of their loss aversion, and counteract it: experience mitigates loss aversion. A very different interpretation is that experience does not affect loss aversion, but it impacts the reference-point formation. Assume that experienced traders expect to trade the object that they are assigned with probability .5, independent of which group they are assigned to. As in Köszegi and Rabin (2006), subjects have a stochastic reference point, $r = 1$ with probability .5 and $r = 0$ otherwise. For individuals assigned the good, the (expected) value of keeping the good is $.5 * [u(1) - u(0)] + .5 [u(1) - u(1)] = .5 [u(1) - u(0)]$; the (expected) value of selling the good $.5 * [u(0) - u(0) + p_{WTA}] + .5 [\lambda (u(0) - u(1)) + p_{WTA}] = .5 [\lambda (u(0) - u(1))] + p_{WTA}$. This implies $p_{WTA} = .5 (1 + \lambda) [u(1) - u(0)]$. It is easy to show with similar calculation that

$$p_{WTP} = .5 (1 + \lambda) [u(1) - u(0)] = p_{WTA}.$$

If experienced subjects have rational expectations about their reference point (Köszegi and Rabin, 2006), they exhibit no endowment effect, even if they are loss-averse. The follow-up literature should consider carefully the determination of the reference point.

Housing Market. For home-owners who are deciding the sale price of a house, the initial purchase price is likely to be a salient reference point. Loss aversion induces the home-owners that would sell at a loss to ask for a higher sale price. To illustrate this point, consider

¹⁸In the Conclusion, I discuss the role of experience further.

a model in which a higher sale price P has two effects: (i) it increases the utility of sale $U(P)$, conditional on a sale; (ii) it lowers the probability $p(P)$ of a sale ($p'(P) < 0$). The home-owner maximizes $\max_P p(P)U(P) + (1 - p(P))\bar{U}$, where \bar{U} is the reservation utility. The optimal price P^* equates the marginal gain of increasing the price, $p(P)U'(P)$, to the marginal cost, $-p'(P)(U(P) - \bar{U})$. Consider a piece-wise linear utility as in (5). For $P \geq P_0$, the owner equates $p(P)$ and $-p'(P)(P - P_0 - \bar{U})$ (notice that \bar{U} could be negative since we have normalized $U(P)$ to equal 0 for $P = P_0$). For $P < P_0$, instead, the owner compares $p(P)\lambda$ and $-p'(P)(\lambda(P - P_0) - \bar{U})$. For a standard risk-neutral agent ($\lambda = 1$), the two conditions coincide, leading to a solution P_{RN}^* . For a loss-averse agent ($\lambda > 1$), however, for a price P below the reference price P_0 , the marginal benefit of a higher price increases discontinuously from $p(P)$ to $p(P)\lambda$, while the marginal cost decreases. Both effects imply that, if in the risk-neutral case the solution P_{RN}^* was to sell at a price lower than P_0 , the loss-averse owner sells at a higher price $P_{LA}^* > P_{RN}^*$. Loss aversion, hence, leads to higher sale prices for units that would sell at a loss because the agent feels a higher marginal disutility of money.

Genesove and Mayer (2001) provide evidence on this phenomenon using data on sales of Boston condominiums in the years 1990-1997. The identification is provided by a housing market boom (1983-1987) followed by a slump (1989-1992). This pattern induces substantial variation in the purchase price P_0 even for otherwise comparable units, depending on the year of purchase. Hence, some of the sellers expect to make a loss relative to the original purchase price when selling at the predicted price \hat{P} , while other sellers expect a gain. The authors then test whether the listing price (relative to the predicted price) is higher for units for which the predicted price \hat{P} falls below the original price P_0 , controlling for the characteristics of the unit. In doing so, the authors face the complication that the predicted sale price \hat{P} is noisily estimated. Importantly, unobservable unit quality would bias the estimates toward finding the result, since unobservably good units are more likely to appear to be selling at a loss (since P_0 will be high for these units) and are more likely to sell above the predicted price. Even after accounting for this bias, the authors find that the listing price for units predicted to sell at a loss is significantly above the predicted value. The magnitude is large: a 1 percent predicted loss translates into a .25 percent higher listing price. This result is larger for units owned by individuals than for units owned by investor, suggesting a mitigating impact of experience as in List (2003 and 2004). The higher listing price translates into both a longer time on the market and a higher final transaction price (though the latter effect is, as expected, smaller relative to the effect on the listing price). These findings are, therefore, consistent with the reference-dependent model with loss aversion outlined above. The authors do not test an additional prediction of the model above, bunching of the listing price at the initial purchase price P_0 .

Finance. Two of the most important applications of reference-dependent preferences are to the field of finance.¹⁹ The first application is to the equity premium puzzle: equity returns

¹⁹Barberis and Thaler (2003) present a more comprehensive survey of these applications.

outperformed bond returns by on average 3.9 percentage points during the period 1871-1993 (Campbell and Cochrane, 1999), a premium too large to be reconciled with the standard model, except for extremely high risk aversion (Mehra and Prescott, 1985). Benartzi and Thaler (1995) use a calibration²⁰ to show that this premium is instead consistent with what loss-averse investors require to invest in stocks, provided that they evaluate their portfolio performance annually. At horizons as short as a year, the likelihood that stocks underperform relative to bonds requires a substantial compensation in terms of returns, given loss aversion. At a longer horizon, the likelihood of underperformance decreases, and the implied equity premium decreases too; the converse is true for a shorter horizon. In a paper that carefully formalizes the idea of Benartzi and Thaler (1995), Barberis, Huang, and Santos (2001) show that reference-dependent preferences can indeed match the observed equity premium. This paper uses the simplified prospect-theory model with piece-wise linear function as in (5), relying on reference dependence and loss aversion for the predictions.

The second application is to the so-called disposition effect, which denotes the tendency to sell ‘winners’ and hold on to ‘losers’. Odean (1998) documents this phenomenon using individual trading data from a discount brokerage house during the period 1987-1993. Defining gains and losses relative to the purchase price of a share, Odean computes the share of realized gains $PGR = (\text{Realized Gains})/(\text{Realized Gains} + \text{Paper Gains})$ to equal .148. The share of realized losses $PLR = (\text{Realized Losses})/(\text{Realized Losses} + \text{Paper Losses})$ equals .098. Odean (1998) shows that the large difference between the propensity to realize gains (PGR) and the propensity to realize losses (PLR) is not due to portfolio rebalancing, to ex-post higher returns for ‘losers’ (if anything, ‘winners’ outperform ‘losers’), or to transaction costs. The disposition effect is puzzling for the standard theory, since capital gain taxation would lead to expect that investors liquidate ‘losers’ sooner. This puzzle is a robust finding, replicated more recently by Ivkovich, Poterba, and Weisbenner (2005), who show that the effect is present in both taxable and tax-deferred accounts (though larger in tax-deferred accounts).

Prospect theory is a possible explanation for this phenomenon. The concavity over gains induces less risk-taking for ‘winner’ stocks, and hence more sales of ‘winners’. The convexity over losses induces more risk-taking for ‘loser’ stocks, and hence more purchases of ‘losers’. Barberis and Xiong (forthcoming), however, point out that this argument does not take into account the kink at the reference point. This kink induces high local risk-aversion that generates a pressure to sell both ‘losers’ and ‘winners’ around the reference point. Reference-dependent investors take this into account and enter the stock market only if expected returns are sufficiently high. For high expected returns, however, ‘winners’ are likely to be further from the reference point than ‘losers’, generating more pressure to sell (due to the closeness to the kink) for ‘winners’, contrary to the disposition effect. Indeed, Barberis and Xiong (forthcoming) simulate a calibrated model of reference-dependent preferences that takes these effects into

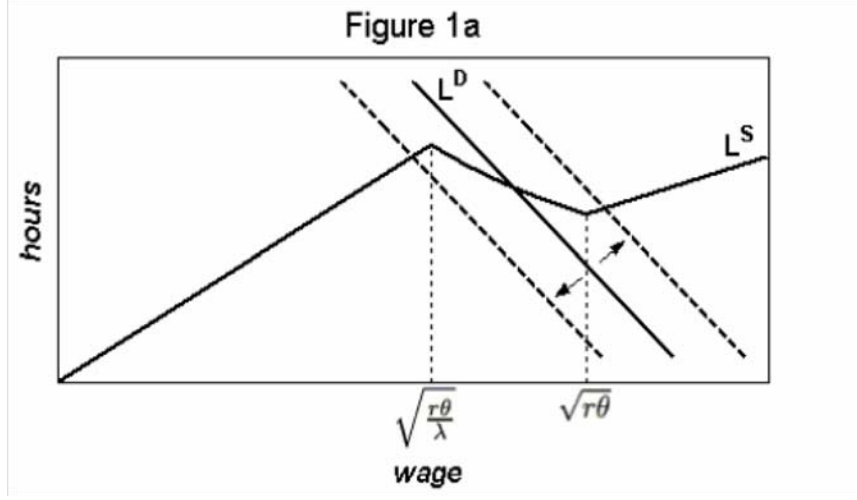
²⁰The calibration uses the loss-aversion parameter estimated from the experiments.

account. They obtain the disposition effect only for certain ranges of the parameters and obtain the opposite pattern for most of the parameters. Their benchmark model assumes that investors, when evaluating the holdings, make no distinctions between realized gains/losses and ‘paper’ gains/losses. Investors, however, may treat the two utility carriers asymmetrically and derive utility (or disutility) only from realized gains and losses. Investors may even go as far as distancing themselves from the paper losses. Karlsson, Loewenstein, and Seppi (2007) show that investors are substantially less likely to look up their holding on the Internet when the stock market is doing poorly. If investors only evaluate losses when they realize them, as Barberis and Xiong (forthcoming) show with an extension of their model, reference-dependent preferences mostly produce the disposition effect patterns.

Labor Supply. As a further application, we consider the response of labor supply to wage fluctuations. This response, in general, reflects a complex combination of income and substitution effects (Card, 1994). Here, we consider a simple case in which income effects can, to a first approximation, be neglected: jobs in which workers decide the labor supply daily, and in which the realization of the daily wage is idiosyncratic. Taxi drivers, for example, decide every day whether to drive for the whole shift or end earlier; the effective wage varies from day-to-day as the result of demand shifters such as weather and conventions. For these occupations, the income effect from (uncorrelated) changes in the daily wage is negligible, and we can neglect it by assuming a quasi-linear model. Each day, workers maximize the utility function $u(Y) - \theta h^2/2$, where the daily earning Y equals hw , h is the number of hours worked, w is the daily wage, and $\theta h^2/2$ is the (convex) cost of effort.

Following the simplified prospect theory formulation in (5), we assume that the utility function $u(Y)$ equals $(Y - r)$ for $Y \geq r$, and $\lambda(Y - r)$ otherwise, where r is a target daily earning. Reference-dependent workers ($\lambda > 1$) are loss-averse with respect to missing the daily target earning. For $\lambda = 1$, this model reduces to the standard model with risk-neutral workers.

In the standard model ($\lambda = 1$), workers maximize $wh - \theta h^2/2$, yielding an upward-sloping labor supply curve $h^* = w/\theta$. As the wage increases, so do the hours supplied, in accordance to the substitution effect between leisure and consumption. A reference-dependent worker ($\lambda > 1$), instead, exhibits a non-monotonic labor supply function (Figure 1a). For a low wage ($w < \sqrt{r\theta/\lambda}$), the worker has not yet achieved the target earnings, and an increase in wage leads to an increase in hours worked ($h^* = \lambda w/\theta$), as in the standard model. For a high wage ($w > \sqrt{r\theta}$), the worker earns more than the target, and the labor supply is similarly upward-sloping, albeit flatter ($h^* = w/\theta$). For intermediate levels of the wage ($\sqrt{r\theta/\lambda} < w < \sqrt{r\theta}$), instead, the worker is content to earn exactly the daily target r . Any additional dollar earned makes it easier to reach the target and leads to reductions in the number of hours worked ($h^* = r/w$); this generates a locally downward-sloping labor supply function.



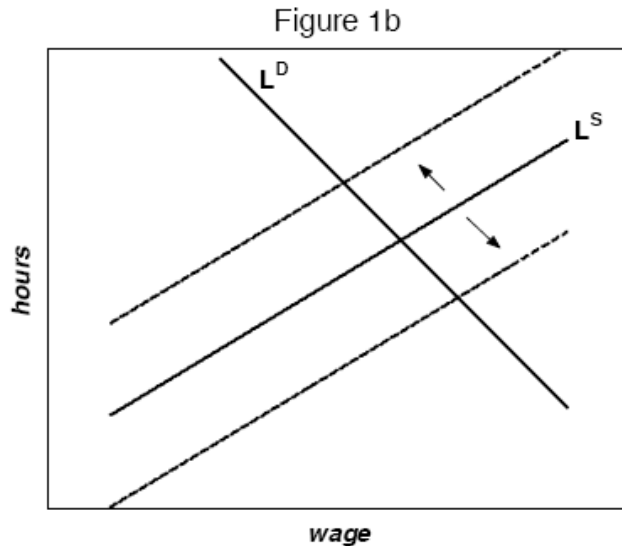
Camerer, Babcock, Loewenstein, and Thaler (1997) use three data sets of hours worked and daily earnings for New York cab drivers to test whether the labor supply function is upward-sloping, as the standard theory above implies, or downward-sloping. Denote by $Y_{i,t}$ and $h_{i,t}$ the daily earnings and the hours worked on day t by driver i . Camerer et al. (1997) estimate the OLS labor-supply equation

$$\log(h_{i,t}) = \alpha + \beta \log(Y_{i,t}/h_{i,t}) + \Gamma X_{i,t} + \varepsilon_{i,t}. \quad (6)$$

Increases in the daily wage, computed as $Y_{i,t}/h_{i,t}$, lead to decreases in the number of hours worked $h_{i,t}$ with estimated elasticities in the three data sets of $\hat{\beta} = -.186$ (s.e. .129), $-.618$ (s.e. .051) and $-.355$ (s.e. .051). The authors conclude that the data reject the standard model (which predicts a positive elasticity), and support a reference-dependent model with daily earnings as the reference point. As Figure 1a shows, though, the labor supply function is not necessarily downward-sloping for target earners, and it is almost certainly not log-linear, unlike in specification (6). Nevertheless, the finding of a negative elasticity is consistent with reference-dependent preferences for shifts in labor demand corresponding to a wage in the interval $\sqrt{\theta r/\lambda} < w < \sqrt{\theta r}$.

Specification (6) is open to two main criticisms. First, a negative elasticity $\hat{\beta}$ is expected if the daily fluctuations in wages for cab drivers are due to shifters of labor supply (like rain that make driving less pleasant), rather than shifters of labor demand. As Figure 1b illustrates, if labor supply shifts across days, the resulting equilibrium points plot out a downward-sloping curve even if the labor supply function is upward-sloping. Camerer et al. (1997) include controls for plausible labor supply shifters such as weekday and rain; they also use interviews of cab drivers to argue that the factors affecting the wage are unlikely to change the marginal cost of driving; however, in the absence of an instrument for labor supply, this objection is a concern. Second, specification (6) suffers from division bias, which biases downward the

estimate of β . Since the daily wage is computed as the ratio of daily earnings and hours worked, and since hours worked is the left-hand-side variable in (6), any measurement error in $h_{i,t}$ induces a mechanical downward bias in $\hat{\beta}$. Camerer et al. (1997) address this objection by instrumenting the daily wage of worker i by the summary statistics of the daily wage of the other workers on the same shift. The estimates of β are still negative, though noisier.



Farber (2005) uses a different data set of 584 trip sheets for 21 New York cab drivers and estimates a hazard model that does not suffer from division bias. For any trip t within a day, Farber (2005) estimates the probability of stopping as a function of the number of hours worked $h_{i,t}$ and the daily cumulative earnings to that point, $Y_{i,t}$:

$$Stop_{i,t} = \Phi(\alpha + \beta_Y Y_{i,t} + \beta_h h_{i,t} + \Gamma X_{i,t}),$$

where Φ is the c.d.f. of a standardized normal distribution. The standard theory predicts that β_Y should be zero (since earnings are not highly correlated within a day), while reference dependence predicts that β_Y should be positive. Farber (2005) finds that β_Y is positive ($\hat{\beta}_Y = .015$), but not significantly so. While the author cannot reject the standard model, the point estimates are not negligible: a ten percent increase in $Y_{i,t}$ (about \$15) is predicted to increase the probability of stopping by $15 * .015 = .225$ percentage points, a 1.6 percent increase relative to the average of 14 percentage points. This corresponds to an elasticity between earnings and stopping of .16. These findings do not contradict prospect theory, since Farber (2005) does not test the hypothesis that cab drivers have reference-dependent preferences (Failing to reject the null is different from rejecting the alternative hypothesis of prospect theory, especially in light of the positive point estimates). In a more recent paper, Farber (forthcoming) addresses this issue and tests, using the same data set, a simple model of labor supply which explicitly allows

for reference-dependent preferences with a stochastic reference point. The findings provide weak evidence of reference dependence: the estimated model implies a loss-aversion coefficient λ significantly larger than zero. At the same time, however, the estimated variation across days in the reference daily earning is large enough that reference dependence loses predictive power. It is an open question whether more precise data on the income targets, perhaps obtained via a survey or induced experimentally, would help to identify reference-dependence effects when combined with the labor supply decision.

Given the lack of an instrument for daily wage fluctuations, the evidence on the labor supply of taxi drivers is unlikely to settle the debate on reference dependence and labor supply. Fehr and Goette (2007) provide new evidence using a field experiment on the labor supply of bike messengers. Like taxi drivers, bike messengers choose how long to work within a shift. Fehr and Goette (2007) randomly assign 44 messengers into two groups. Each group receives a 25 percent higher commission for the deliveries for just one month in two different months. This design solves both problems discussed above, since the increase in wage is exogenous, and the wage and the actual deliveries are exactly measured.

Bike messengers in the treatment group respond in two ways to the exogenous (and anticipated) temporary increase in wage: (i) they work 30 percent more shifts; (ii) within each shift, they do 6 percent fewer deliveries. The first finding is consistent with both the standard model and the reference-dependent model with daily targets. (When deciding on which day to work, reference-dependent workers will sign up for shifts on days in which it is easier to reach the daily target.) The second finding is consistent with target earning, and not with the standard model, which predicts an increase in the number of hours worked within each shift. This finding, however, is also consistent with an extension of the standard model in which workers in the treatment group get more tired, and hence do fewer deliveries, because they work more shifts. Fehr and Goette (2007) provide additional evidence in support of reference-dependence combining the field evidence with a laboratory experiment on risk-taking. The bike messengers that display loss aversion in the lab—i.e., they reject a $(-5, .5; 8, .5)$ lottery—exhibit a more negative response (though not significantly so) in their deliveries to the wage increase. This correlation is not predicted by the fatigue explanation, but is predicted by the reference dependence model.

Insurance. A puzzling feature of insurance behavior is the pervasiveness of small-scale insurance. Insurance policies on, for example, the telephone wiring are commonplace despite the fact that, in case of an accident, the losses amount to at most \$50 (Cicchetti and Dubin, 1994). This is a puzzle for expected utility, which implies local risk-neutrality and hence no demand for small-scale insurance (except in the unrealistic case of fair pricing). Sydnor (2006) provides evidence of excess small-scale insurance for the \$36 billion home insurance industry. Since mortgage companies require home insurance, the consumer choice is limited to the level of deductible in a standard menu: \$250 vs. \$500 vs. \$1000. Using a random

sample of 50,000 members of a major insurance company in one year, Sydnor documents that 83% of customers and 61% of new customers choose deductibles lower than \$1000. The modal homeowner chooses a \$500 deductible, thereby paying on average \$100 of additional premium relative to a \$1000 deductible. However, the claim rate is under 5%, which implies that the value of a low deductible is about \$25 in expectation. The standard homeowner, therefore, is sacrificing $\$100 - \$25 = \$75$ in expectations to insure against, at worst, a $\$500 - \$100 = \$400$ risk.

This indicates a strong preference for insuring against small risks that is a puzzle for the standard theory, unless one assumes three-digit coefficients of relative risk aversion. This deviation from the standard model involves substantial stakes. If, instead of choosing a low deductible, homeowners selected the \$1000 deductible from age 30 to age 65 and invested the money in a money market fund, their wealth at retirement would be \$6,000 higher. Sydnor (2006) shows that a calibrated version of prospect-theory can match the findings by the over-weighting of the small probability of an accident and the loss aversion with respect to future losses²¹. The two components of prospect theory each account for about half of the observed discrepancy between the predicted and the observed willingness to pay for low deductibles. Social pressure by the salesmen (who are paid a percentage of the premium as commission) may also contribute to the prevalence of low-deductible contracts.

Employment. Mas (2006) estimates the impact of reference points for the New Jersey police. In the 9 percent of cases in which the police and the municipality do not reach an agreement, the contract is determined by final offer arbitration. The police and the municipality submit their offers to the arbitrator, who has to choose one of the two offers. In theory (Mas, 2006), if the disputing parties are equally risk-averse, the winner in arbitration is determined by a coin toss.²² Mas (2006) exploits this prediction of quasi-random assignment to present evidence on how police pay affects performance for 383 arbitration cases from 1978 to 1995. In the cases in which the offer of the employer is chosen, the share of crimes solved by the police (the clearance rate) decreases by 12 percent compared to the cases in which the police offer is chosen. The author also documents a smaller increase in crime. Lower pay therefore induces the police to devote less effort to fighting crime, a finding consistent with a number of interpretations, including efficiency wages and reciprocity.

Mas (2006) then provides evidence that reference points mediate this effect, which depends more on expected wages than on actual wages. Mas generates a predicted award based on a set of observables as a proxy for the reference point, and computes how the clearance rate responds to differences between the award and the predicted award. The response is significantly higher

²¹Loss aversion could in principle go the other way, since individuals that are loss-averse to paying a high premium may as well prefer the high deductible. Experimental evidence, however, suggests that consumers will adjust their reference point on the premium side, since they are expecting to pay the premium for sure, but cannot adjust the reference point on the future uncertain loss.

²²In reality, the arbitrator rules for the municipality in 34.4 percent of cases, suggesting that the unions are more risk-averse than the employers. The key result on reference dependence is independent of this assumption.

for cases in which the police loses—and hence is on the loss side—than for cases in which the police wins—and hence is on the gain side. This asymmetry is consistent with reference-dependent preferences with loss aversion. Assume for example that the utility function of the police is $[V + v(w|r)]e - \theta e^2/2$, where $v(w|r)$ captures the impact of the wage w on the desirability of effort; assume also that $v(w|r)$ is reference-dependent as in (5). This complementarity between police pay w and effort e in the utility function can be interpreted as reference-dependent reciprocity. The first-order condition, then, implies $e^*(w) = [V + v(w|r)]/\theta$. Given loss aversion in $v(w|r)$, this induces a stronger response of the police on the loss side (w below r) than on the gain side (w above r).

Summary. Reference-dependent preferences help explain: (i) excessive aversion to small risks in the laboratory; (ii) endowment effect for inexperienced traders; (iii) the reluctance to sell houses at a loss; (iv) equity premium puzzle in asset returns; (v) (possibly) the tendency to sell ‘winners’ rather than ‘losers’ in financial markets; (vi) (some evidence of) target earnings in labor supply decisions; (vii) the tendency to insure against small risks; (viii) effort in the employment relationship. I have discussed cases in which the evidence is more controversial (labor supply and endowment effect) and cases in which it is unclear whether reference-dependence is an explanation for the phenomenon (disposition effect). I have also discussed how the original model in Kahneman and Tversky (1979) (and the calibrated version in Tversky and Kahneman, 1992) is rarely applied in its entirety, often appealing just to reference dependence and loss-aversion.

A key issue in this literature is the determination of the reference point r . Often, different assumptions about the reference point are plausible, which makes the application of the theory difficult. Köszegi and Rabin (2006) have proposed a solution. They suggest that the reference point be modeled as the (stochastic) rational-expectations equilibrium of the transaction. In any given situation, this model makes a prediction for the reference point, without the need for additional parameters (though there can often be multiple equilibria, and hence multiple possible reference points). This theory also provides a plausible explanation for some of the puzzles in this literature. For example, as we discussed above, it predicts the absence of endowment effect among experienced traders (List, 2003), even if these traders are loss-averse. Experienced traders expect to trade any item they receive, and hence their reference point is unaffected by the initial allocation of objects.

2.3 Social Preferences

The standard model, in its starkest form as in (1), assumes purely self-interested consumers, that is, utility $U(x_i|s)$ depends only on own payoff x_i .

Laboratory Experiments. A large number of laboratory experiments calls into question the assumption of pure self-interest. I present here the results of two classical experiments,

which I relate to the field evidence below. (i) *Dictator game*. In this experiment (Forsythe et al., 1994) a subject (the dictator) has an endowment of \$10 and chooses how much of the \$10 to transfer to an anonymous partner. While the standard theory of self-interested consumers predicts that the dictator would keep the whole endowment, Forsythe et al. (1994) find that sixty percent of subjects transfers a positive amount. (ii) *Gift Exchange game*. This experiment (Fehr, Kirchsteiger, and Riedl, 1993) is designed to mirror a labor market. It tests efficiency wages models according to which the workers reciprocate a generous wage by working harder (Akerlof, 1982). The first subject (the firm) decides a wage $w \in \{0, 5, 10, \dots\}$. After observing w , the second subject (the worker) responds by choosing an effort level $e \in [.1, 1]$. The firm payoff is $(126 - w)e$ and the worker payoff is $w - 26 - c(e)$, with $c(e)$ increasing and slightly convex. The standard theory predicts that the worker, no matter what the firm chooses, exerts the minimal effort and that, in response, the firm offers the lowest wage that satisfies the participation constraint for the workers ($w = 30$). Fehr et al. (1993) instead find that the workers respond to a higher wage w by providing a higher effort e . The firms, anticipating this, offer a wage above the market-clearing one (the average w is 72). These results have been widely replicated and have given rise to a rich literature on social preferences in the laboratory, summarized in Charness and Rabin (2002) and Fehr and Gächter (2000).

Model. Several models have been proposed to rationalize the behavior in these experiments; we introduce a simplified version of the social preference model in Charness and Rabin (2002), which builds on the formulation of Fehr and Schmidt (1999).²³ In a two-player experiment, the utility of subject 1 is defined as a function of the own payoff (x_1) and other-player's payoff (x_2):

$$U_1(x_1, x_2) \equiv \begin{cases} \rho x_2 + (1 - \rho)x_1 & \text{when } x_1 \geq x_2; \\ \sigma x_2 + (1 - \sigma)x_1 & \text{when } x_1 < x_2. \end{cases} \quad (7)$$

The standard model is a special case for $\rho = \sigma = 0$. The case of baseline altruism is $\rho > 0$ and $\sigma > 0$, that is, player 1 cares positively about player 2, whether 1 is ahead or not. In addition, Charness-Rabin (2002) assume $\rho > \sigma$, that is, player 1 cares more about player 2 when 1 is ahead. Fehr and Schmidt (1999) propose an equivalent representation of preferences²⁴ and assume $0 < \rho < 1$, like Charness-Rabin (2002), but also $\sigma < -\rho < 0$. When player 1 is behind, therefore, she prefers to lower the payoff of player 2 (since she is inequality-averse). These two models can explain giving in a Dictator Game with a \$10 endowment. The utility of giving

²³In these models, players care about the inequality of outcomes, but not about the intentions of the players (though the general model in Charness and Rabin, 2002 allows for the role of intentions). Another class of models (including Rabin, 1993 and Dufwenberg and Kirchsteiger, 2004), based on psychological games, instead assumes that subjects care about the intentions that lead to specific outcomes. A common concept is reciprocity—subjects are nice to subjects that are helpful to them, but not to subjects that take advantage of them. These models also explain the laboratory findings.

²⁴Fehr-Schmidt preferences take the form: $U_1(\pi_1, \pi_2) = \pi_1 - \alpha \min(\pi_2 - \pi_1, 0) - \beta \min(\pi_1 - \pi_2, 0)$; they are equivalent to the preferences in (7) for $\beta = \rho$ and $\alpha = -\sigma$.

\$5 is higher than the utility of giving \$0 if $5 \geq \max((1 - \rho)10, \sigma 10)$, that is, if $\rho \geq .5 \geq \sigma$ (altruism is high enough, but not so high that a player would transfer all the surplus to the opponent.) Fehr and Schmidt (1999) show that model (7) can also rationalize the average behavior in the Gift Exchange game for high enough ρ : altruistic workers provide effort to lower the inequality with the firm; the firm, anticipating this, raises w .

Charitable Giving. The size of charitable giving is suggestive of social preferences in the field. In the US, in 2002, 240.9 billion dollars were donated to charities, representing an approximate 2 percent share of GDP (Andreoni, 2006). Donations of time in the form of volunteer work were also substantial: 44 percent of respondents to a survey reported giving time to a charitable organization in the prior year, with volunteers averaging about 15 hours per month (Andreoni, 2006). Altogether, a substantial share of GDP reflects a concern for others, a finding qualitatively consistent with the experimental findings. However, while social preferences are a leading interpretation for giving, charitable donations may also be motivated by other factors, such as desire for status and social pressure by the fund-raisers.

Even if we take it for granted that giving is an expression of social preferences, it is difficult to use models such as (7) to explain quantitatively the patterns of giving in the field for three reasons. (i) These models are designed to capture the interaction of two players, or at most a small number of players. Charitable giving instead involves a large number of potential recipients, from local schools in Oakland, CA to NGOs in Africa. (ii) The utility representation (7) implicitly assumes that x_1 and x_2 include only the experimental payoffs from, say, the dictator game. In the field, it is difficult to determine to what extent x_1 and x_2 should include, for example, the disposable income. (iii) In one-to-one fund-raising situations (hence side-stepping issue (i)), models such as (7) over-predict giving. Suppose, for example, that $x_1 = \$1,000$ is the disposable income of person 1 and $x_2 = \$0$ is the disposable income of person 2, for example, a homeless person. For $\rho \geq .5 \geq \sigma$, the model predicts that person 1 should transfer $(\$1000 - \$0) / 2 = \$500$, a level of giving much higher than 2 percent of GDP. One has to make ad-hoc assumptions on x_1 to reproduce the observed level of giving. For these reasons, while models of social preferences are very useful to understand behavior in the laboratory, they are less directly applicable to the field, compared to models of self-control and of reference-dependence. Andreoni (2006) overviews models that better predict patterns of giving, such as models of warm glow.

There are, however, field settings which resemble more closely the laboratory set-up. When a fund-raiser contacts a person directly, the situation resembles a dictator game, except for the lack of anonymity. Field experiments in fund-raising, including List and Lucking-Reilly (2002) and Landry, Lange, List, Price, and Rupp (2006), estimate the effect on giving of variables such as the seed money (the funds raised early on), the use of a lottery, and the identity of the solicitor. Charitable giving is increasing in the seed money (List and Lucking-Reilly, 2002) presumably because of signaling of quality of the charity, and in the attractiveness

of female solicitors for door-to-door fund-raising, especially for male respondents (Landry et al., 2006). The latter result implies that giving in door-to-door fund-raising is not purely the result of altruism, suggesting a more instrumental view of giving. Overall, however, these field experiments do not yet answer the key question of what motivates most giving, genuine social preferences or more instrumental reasons, such as social pressure. Questions such as this are likely to be the focus of future research.

Workplace Relations. Workplace relations between employees and employer can be upset at the time of contract renewal, and workers may respond by sabotaging production. Krueger and Mas (2004) examine the impact of a three-year period of labor unrest at a unionized Bridgestone-Firestone plant on the quality of the tires produced at the plant. The workers went on strike in July 1994 and were replaced by replacement workers. The union workers were gradually reintegrated in the plant in May 1995 after the union, running out of funds, accepted the demands of the company. An agreement was not reached until December 1996. Krueger and Mas (2004) finds that the tires produced in this plant in the 1994-1996 years were ten times more likely to be defective. The increase in defects does not appear due to lower quality of the replacement workers. The number of defects is higher in the months preceding the strike (early 1994) and in the period in which the union workers and the replacement workers work side-by-side (end of 1995 and 1996). While a bargaining interpretation cannot be ruled out, this provides some evidence that negative reciprocity in response to what workers perceive as unfair treatment can have a large impact on worker productivity.

Bandiera, Barankay, and Rasul (2005) test for the impact of social preferences in the workplace among employees. They use personnel data from a fruit farm in the UK and measure changes in the productivity as a function of changes in the compensation scheme. In the first eight weeks of the 2002 picking season, the fruit-pickers were compensated on a relative performance scheme in which the per-fruit piece rate is decreasing in the average productivity. In this system, workers that care about others have an incentive to keep the productivity low, given that effort is costly. In the next eight weeks, the compensation scheme switched to a flat piece rate per fruit. The switch was announced on the day the change took place. Bandiera et al. (2005) find that, after the change to piece rate, the productivity of each worker increases by 51.5 percent; the estimate holds after controlling for worker fixed effects and is higher for workers with a larger network of friends. The result is not due to a change in incentives: the flat piece rate is on average lower than the relative-pay piece rate, which would contribute to lowering, rather than increasing, productivity after the switch. These results can be evidence for altruism; they can, however, also be evidence of collusion in a repeated game, especially since in the field each worker can monitor the productivity of the other workers. To test for these explanations, the authors examine the effect of the change in compensation for growers of a different fruit where the height of the plant makes monitoring among workers difficult. For this other fruit, the authors find no impact on productivity of the switch to piece rate. This

implies that the findings are not due to altruism, but rather to collusion or to a different form of social preferences, reciprocity. According to this latter interpretation, the lack of observability of the behavior of others inhibits not only collusion, but also reciprocal behavior.

Gift Exchange in the Field. The Bandiera et al. (2005) paper underscores the importance of controlling for repeated game effects in tests of social preferences. We now consider field experiments that test for Gift Exchange controlling for these effects. Field experiments (like laboratory experiments) give the researcher more control over the design of an economic situation. Falk (forthcoming) examines the importance of gifts in fund-raising. The context is the mailing of 9,846 solicitation letters in Switzerland to raise money for schools in Bangladesh. One third of the recipients receives a postcard designed by the students of the school, another third receives four such postcards, and the remaining third receives no postcards. The three mailings are otherwise identical, except for the mention of the postcard as a gift in the two treatment conditions. The donations are increasing in the size of the gifts. Compared to the 12.2 percent frequency of donation in the control group, the frequency is 14.4 percent in the small gift and 20.6 percent in the large gift treatment. Conditional on a donation, the average amount donated is slightly smaller in the large-gift treatment, but this effect is small relative to the effect on the frequency of donors. The large treatment effects do not appear to affect the donations at next year's solicitation letter, when no gift is sent. A gift, therefore, appears to trigger substantial positive reciprocity, as in the laboratory version of the Gift Exchange.

Gneezy and List (2006) test the gift exchange with two field experiments in workplace settings. In the first experiment, they hire 19 workers for a six-hour data entry task at a wage of \$12 per hour; in the second experiment, they hire 23 workers to do door-to-door fund-raising for one weekend at a wage of \$10 per hour. In both cases, they divide the workers into a control and a treatment group. The control group is paid as promised, while the treatment group is told after recruitment that the pay for the task was increased to \$20 per hour. The authors test whether the treatment group exerts more effort than the control group, as predicted by the gift exchange hypothesis, or the same effort, as predicted by the standard model. The findings are two-fold. At first, the treatment group exerts substantially more effort, consistent with gift exchange: treated workers log 20 percent more books in the first hour and raise 80 percent more money in the morning hours. The difference however is short-lived: the performances of control and treatment group are indistinguishable after two hours of data entry and after three hours of fund-raising. In these two applications, the increase in wage does not pay for itself. These experiments suggest that the gift exchange may have an emotional component which dissipates over time.

Kube, Maréchal, and Puppe (2008a) use a similar design for a six-hour library work in Germany, with an additional negative gift exchange treatment. This group of subjects, upon showing up, is notified that the pay is 10 Euro per hour, compared to the promised pay of 'presumably' 15 Euro per hour. No one decides to quit. This group logs 25 percent fewer books

compared to the control group, a difference that, unlike in the Gneezy and List (2006) paper, does not decline over time. The group in the positive gift exchange treatment (paid 20 Euro) logs 5 percent more books, an increase which also does not dissipate over time. As in the laboratory findings, negative reciprocity is stronger than positive reciprocity.

List (2006) presents evidence that not everyone reciprocates a generous transfer. Attendees of a sports card fair participate in a field experiment involving buying a card from a dealer. One group is instructed to offer \$20 for a card of good quality (PSA grade 9), while another group is instructed to offer \$65 for a card of top quality (PSA grade 10). The quality of the card can be verified by an expert but is not apparent on inspection. Dealers that are ‘non-local’ (and hence are not concerned with reputation) offer cards of the same average quality to the two groups, displaying no gift-exchange behavior.²⁵ These dealers, however, display gift-exchange-type behavior in laboratory experiments designed to mirror the Fehr et al. (1993) experiment. These findings raise interesting questions on when gift-exchange behavior does and does not arise. One explanation of the findings is that bargaining in a market setting is construed as a situation where norms of gift exchange do not apply, possibly because a transfer of \$60 is not considered a ‘gift’. Hence, the dealers do not display such norms when selling cards, but they do instead when participating in an experiment where the presence of a gift is clearer. More broadly, this suggests that we need to understand the economic settings in which gift-exchange norms apply (such as charitable giving and, to some extent, employment relationships) and the ones where they do not apply (such as market bargaining).

Kube, Maréchal, and Puppe (2008b) provide evidence on the importance of such norms. Within a field experiment along the lines of Gneezy and List (2006) and Kube et al. (2008a), workers are hired to catalogue books for three hours. Relative to a control group of 17 workers hired for the announced hourly wage of 12 Euros, two treatment groups receive an unexpected gift: the first group of 16 students receives a 7 Euro (20 percent) wage increase, while the second group of 15 students receives a thermos bottle worth 7 Euros. This design is motivated by evidence on gift perceptions: subjects in an online survey presented with the experimental design perceive the employer as kinder when the gift is a thermos compared to money. The worker effort is consistent with gift exchange. Compared to the control group, productivity is 30 percent higher in the thermos group but only 6 percent higher in the money group. Interestingly, in the thermos group the relative increase in productivity is larger than the increase in labor costs, suggesting that, unlike in Gneezy and List (2006), gifts can pay for themselves (although, since the market value of the task to the library is not clear, one cannot say for sure). The gift exchange response does not simply depend on the perceived economic value of the gift: in a separate experiment, 172 subjects offered the choice of a 7 Euro payment or the thermos overwhelmingly prefer the monetary payment. Future research will need to

²⁵Dealers that are ‘local’, that is, that attend the fair frequently, offer higher-quality card to the \$65 group, presumably because of reputation-building.

provide more evidence on the psychology of gift-giving, as well as models of it.

Summary. Social preferences help explain: (i) giving to charities; (ii) the response of striking workers to wage cuts; (iii) the response of giving to gifts in fund-raisers; (iv) the response of effort to unanticipated changes in pay, at least in the short-run; (v) the response of effort to non-monetary gifts. However, the research on social preferences displays more imbalance between laboratory and field, compared to the research on self-control and on reference dependence. The models of social preferences which match the laboratory findings are not easily applicable to the field, overpredicting, for example, the amount of giving. It will be important to see more papers linking the findings in the laboratory, which allows the most control on the design, to the evidence in the field; the recent literature on Gift Exchange is a good example. A separate issue is the difficulty of distinguishing in the field social preferences from repeated game strategies (as in Bandiera et al., 2005) and other alternative explanations. For example, social pressure (Section 4.4) can explain regularities in giving, such as the higher effectiveness of high-pressure fund-raising methods (such as phone calls) relative to low-pressure ones (such as mailings). Creative field experiments such as those in this Section can be designed to distinguish different explanations.

3 Non-standard Beliefs

The standard model in (1) assumes that consumers are on average correct about the distribution of the states $p(s_t)$. Experiments suggest instead that consumers have systematically incorrect beliefs $\tilde{p}(s_t)$ in at least three ways: (i) *Overconfidence*. Consumers over-estimate their performance in tasks requiring ability, including the precision of their information; (ii) *Law of Small Numbers*. Consumers expect small samples to exhibit large-sample statistical properties; (iii) *Projection Bias*. Consumers project their current preferences onto future periods.

3.1 Overconfidence

Surveys and laboratory experiments present evidence of overconfidence about ability. In Svenson (1981), 93 percent of subjects rated their driving skill as above the median, compared to the other subjects in the experiment.²⁶ Most individuals underestimate the probability of negative events such as hospitalization (Weinstein, 1980) and the time needed to finish a project (Buehler, Griffin, and Ross 1994). In Camerer and Lovallo (1999), subjects play multiple rounds of an entry game in which only the top c out of n entrants make positive profits. In the luck treatment the top c subjects are determined by luck, while in the skill treatment the top c subjects are determined by ability in solving a puzzle. More subjects enter in the skill

²⁶This finding admits alternative interpretations, such as that each individual may define driving ability in a self-serving way. These interpretations, however, are addressed in the follow-up literature.

treatment than in the luck treatment, indicating that subjects overestimate their (relative) ability to solve puzzles.²⁷

The first example of overconfidence in the field is the naiveté about future self-control by consumers in the choices of health club contracts, credit cards, and 401(k) plans, documented in Section 2.1 (DellaVigna and Malmendier, 2006; Ausubel, 1999; Madrian and Shea, 2001). Naiveté is an example of overconfidence since self-control is a desirable ability.

In a second example, Malmendier and Tate (2005, forthcoming) provide evidence on overconfidence by CEOs about their ability to manage a company. They assume that CEOs are likely to overestimate their ability to pick successful projects and to run companies. As such, these top managers are likely to invest in too many projects, and to over-pay for mergers. To test these hypotheses, Malmendier and Tate identify a proxy for overconfidence, and examine the correlation of this proxy with corporate behavior. In particular, they identify as overconfident CEOs who hold on to their stock options until expiration, despite the fact that most CEOs are heavily under-diversified. They interpret the lack of exercise as overestimation of future performance of their company. In Malmendier and Tate (forthcoming) they find that these CEOs are 55 percent more likely to undertake a merger, and particularly so if they can finance the deal with internal funds. (Overconfident CEOs are averse to seeking external financing, since they deem it overpriced.) Of course, the correlation between option exercise and corporate behavior could be due to alternative reasons, such as insider information of the CEO. However, Malmendier and Tate show that insider information does not appear to be the explanation, since the CEOs that delay exercising stock options do not on average gain money by doing so. Managerial overconfidence provides one explanation for the underperformance of companies undertaking mergers. Malmendier and Tate (2005) use the same proxies to show that overconfidence explains in part the excess sensitivity of corporate investment to the availability of cash flows, a long-standing puzzle in corporate finance.

Overconfidence about own-company performance likely extends to rank-and-file employees. Cowgill, Wolfers, and Zitzewitz (2008) study the prediction markets that Google set up for its own employees (with real payoffs). While securities not related to Google are correctly priced on average, the securities with implications for Google display substantial overconfidence: in two-outcome markets, the share that pays one dollar if the favorable outcome for Google

²⁷A more recent literature including, among others, Krueger (1999), suggests that, while overconfidence is typical for easy tasks such as driving, underconfidence can arise for hard tasks such as playing the piano, a dichotomy known as the ‘hard-easy effect’. As Krueger (1999) suggests, the subjects, when comparing their skills to the skills of others, do not appreciate that others similarly find these tasks respectively easy and hard. This results in overconfidence for easy tasks and underconfidence for hard task. A difficulty in applying the ‘hard-easy effect’ to economics is the practical definition of a hard task. While running a company is arguably one of the hardest tasks one can imagine, it likely seems easy to a CEO that spends most of its time doing it. Sorting implies that economic agents would mostly face tasks that they deem easy, making the underconfidence result less relevant.

occurs trades at 45.6 cents, while the average payoff is only 19.9 cents. While this evidence is specific to Google, survey evidence suggests that this phenomenon is more general. Indeed, overconfidence of employees about own-company performance is a leading explanation for the provision of stock options to rank-and-file employees (Oyer and Schaefer, 2005; Bergman and Jenter, 2007). Stock options have become a common form of compensation: the (Black and Scholes) value of options granted yearly to employees in public companies was over \$400 (about one percent of compensation) in 1999 (Oyer and Schaefer, 2005). Incentive effects are unlikely to explain the issuance, given that the contribution of each individual employee to firm value is very limited. Instead, overconfidence about own-company performance can make stock options an attractive compensation format for employers.²⁸ This form of overconfidence is particularly plausible since the workers that are overconfident about a particular company are more likely to sort into it.

A third example of overconfidence is the tendency to overestimate the precision of own information, which is also a skill. For example, Alpert and Raiffa (1982) ask a group of 100 MBA students to provide answers for 10 numeric queries such as ‘*the number of foreign automobiles imported into the U.S. in 1967 in thousands*’ and ‘*the total egg production in millions in the U.S. in 1965*’. The students are also asked for 98 percent confidence intervals. If the students estimated correctly the precision of their information, their confidence intervals should contain the correct answer in approximately 980 of the 1000 responses. Instead, the intervals contain the correct answer in only 574 of the 1000 cases! The elicitation of 75 percent confidence intervals provides similar evidence of overconfidence.

Odean (1999) provides field evidence consistent with overconfidence using data from a discount broker on all the trades of 10,000 individual investors for the years 1987-1993. If the investors overestimate the precision of their information about individual companies, they will trade too much. Indeed, the investors trade on average 1.3 times per year, with a commission cost for buying or for selling a security of over 2 percent per transaction. In addition to these substantial transaction costs, the individual investors pay a return cost to trading, since the stocks sold over-perform the purchases by about 3 percent over the next year. For individual investors, therefore, overconfidence has a substantial impact on returns. Interestingly, there is a gender differential in over-trading that is consistent with the psychology findings, which suggests that men are more overconfident than women about financial decisions: men trade 45 percent more than women, and hence pay a larger returns cost. (Barber and Odean, 2001).

Overconfidence about the precision of private information also helps explain other anomalies in financial markets, such as short-term positive correlation of returns (momentum) and

²⁸Bergman and Jenter (2007) point out that employees can also purchase shares on the open market, and hence do not need to rely on the company providing them. They examine the conditions under which the company will still offer options to overconfident employees, and provide evidence that option compensation is used most intensively when employees are likely to be overconfident.

long-term negative correlation (long-term reversal) (Daniel, Hirshleifer, and Subrahmanyam, 1998). To explain these phenomena, overconfidence needs to be coupled with self-attribution bias, which is the tendency to discount information that is inconsistent with one’s priors. Overconfidence induces individuals to trade excessively in response to private information; in the long-run, the public information prevails and the valuation returns to fundamentals, inducing a long-term reversal. The self-attribution bias is responsible for momentum: in the short-term, as investors receive additional private information, they interpret as more informative the information that conforms to their beliefs, and hence become even more overconfident. In Sections 3.2 and 4.2 we discuss how the law of small numbers and limited attention provide alternative explanations for these same financial markets anomalies.

Summary. Overconfidence helps explain: (i) patterns in health club contract choice, credit card take-up, and default effects, presented in Section 2.1 (overconfidence about self-control); (ii) value-destroying mergers and investment-cash-flow sensitivity (overconfidence about managerial ability); (iii) stock option compensation packages for rank-and-file employees; (iv) excess trading, momentum, and long-term reversal (overconfidence about precision of information). These applications are settings in which overconfidence is particularly likely according to the laboratory evidence: overconfidence is more common when feedback is noisy (i.e., for stock returns) and the decision-maker has an illusion of control (i.e., for managers).

3.2 Law of Small Numbers

Overconfidence is only one form of non-Bayesian beliefs detected in experiments. Tversky and Kahneman (1974) describe a number of deviations from rational updating, including the overweighting of information that is available and representative. I focus on two phenomena—‘gambler’s fallacy’ and overinference—that are examples of reliance on the availability and representativeness heuristics. To illustrate these phenomena, I use Rabin (2002b)’s model of the law of small numbers. Rabin (2002b) assumes that subjects, observing a sequence of signals drawn from an i.i.d. process, believe (incorrectly) that the signals are drawn from an urn of size $N < \infty$ without replacement. If the distribution of the signals is known, this induces a ‘gambler’s fallacy’ belief: after a draw of a signal, subjects expect the next draw to be a different signal (since the draw is without replacement). For example, suppose that the return to a mutual fund is drawn from an urn with 10 balls, 5 Up and 5 Down, with replacement. After two draws of Up, a rational investor expects the probability of another Up to be .5. However, a believer in the law of small number computes such probability as $3/8 < .5$, since two balls ‘Up’ have already been drawn. This is an example of the representativeness heuristics, in that the sequence ‘Up, Up, Down’ is judged as more representative than the sequence ‘Up, Up, Up’.

Terrell (1994) provides field evidence of the ‘gambler’s fallacy’ in New Jersey’s pick-three-numbers game. The lottery is a pari-mutuel betting system: the fewer individuals bet on a

number, the higher is the expected payout. Terrell (1994) finds that the payout for a number that won one or two weeks before is 33 percent higher than for an average number. Belief in ‘gambler’s fallacy’ leads lottery players to bet less on numbers that won recently, at the cost of a lower expected payoff. This pattern is found in a number of other betting markets, including the Maryland daily-numbers lottery (Clotfelter and Cook, 1993). It is likely to apply also to other situations in which the probabilities are known, but subjects misconstrue the i.i.d. nature of the draws. An example is the forecast of the gender of a third child following two boys (or two girls).

The model in Rabin (2002b) delivers a second testable prediction. In the case of uncertain distribution of signals, the subjects overinfer from a sequence of signals of one type that the next signal will be of the *same* type. While this overinference appears to be the opposite of the ‘gambler’s fallacy’, it is a complementary phenomenon. Consider a mutual fund with a manager of uncertain ability. The return is drawn with replacement from an urn with 10 balls. With probability .5 the fund is well managed (7 balls Up and 3 Down) and with probability .5 the fund is poorly managed (3 Up and 7 Down). After observing the sequence ‘Up, Up, Up’, a rational investor computes the probability that the mutual fund is well-managed as $P(Well|UUU) = .5P(UUU|Well) / [.5P(UUU|Well) + .5P(UUU|Poor)] = .7^3 / (.7^3 + .3^3) \approx .927$. A Law-of-Small-Number investor also applies Bayes Rule but has the wrong model for $P(UUU|Well)$ and $P(UUU|Poor)$. Hence, her forecasted probability for $P(Well|UUU)$ equals $(7/10 * 6/9 * 5/8) / [(7/10 * 6/9 * 5/8) + (3/10 * 2/9 * 1/8)] \approx .972$. Hence, this investor over-infers about the ability of the mutual-fund manager after three good performances. Assume now that the Law-of-Small-Number investor believes that the urn is replenished after 3 periods. When forecasting the performance in the next period, the rational investor expects an Up performance with probability $.927 * .7 + (1 - .927) * .3 \approx .671$, while the Law-of-Small-Number investor expects Up with probability $.972 * .7 + (1 - .972) * .3 \approx .689$, which is higher.

Benartzi (2001) provides field evidence of overinference (also called extrapolation): the degree to which employees invest in employer stock depends strongly on the past performance of the stock. In companies in the bottom quintile of performance in the past ten years, 10.4 percent of employee savings are allocated to employer stock, compared to 39.7 percent for companies in the top quintile. This difference does not reflect information about future returns. Companies with a higher fraction of employees investing in employer stock underperform over the next year relative to companies with a lower fraction.

Barber, Odean, and Zhou (forthcoming) use data on individual trades to show that individual US investors purchase stocks with high past returns, also consistent with overinference. The average stock that individual investors purchase outperformed the stock market in the previous three years by over 60 percent. (Interestingly, the average stock sold also outperformed the stock market, though by a smaller amount, consistent with either a belief in ‘gambler’s fallacy’ or with the disposition effect (Section 2.2).)

Overinference in stock holdings can also induce predictability in asset returns. To the extent that investors overinfer from past performance, stocks with high past returns should get overpriced, and ultimately underperform. DeBondt and Thaler (1985) compare stocks that performed particularly well in the past 3 years (‘winners’) to stocks that did poorly in the past 3 years (‘losers’). The ‘winners’ underperform the ‘losers’ by 25 percentage points over the next three years, again consistent with overinference.

Barberis, Shleifer, and Vishny (1998) apply an alternative model of the law of small number to financial markets. While the draws are i.i.d., investors believe that the draws come from either a ‘mean-reverting’ regime or a ‘trending’ regime; in addition, the investors believe that the first regime is more likely *ex ante*. If investors observe a sequence of identical signals, in the short-run they expect a mean-reverting regime (the gambler’s fallacy); hence, the returns under-react to information, inducing short-term positive correlation (momentum). However, after a longer sequence, the individuals overinfer, as in Rabin (2002b), and expect a ‘trending’ regime; this induces a long-term negative correlation of returns. Hence, the law of small numbers can explain two key features of observed returns, short-term positive correlation and long-term negative correlation (see also the discussion in Section 3.1).

3.3 Projection Bias

A third way in which individuals have systematically incorrect beliefs is that they expect their future preferences to be too close to the present ones; for example, they project current hunger levels on the future. Read and van Leeuwen (1998) asked office workers to choose a healthy snack or an unhealthy snack to be delivered a week later (in the late afternoon). Workers were asked either when they were plausibly hungry (in the late afternoon) or when satiated (after lunch). In the first group, 78 percent chose an unhealthy snack, compared to 42 percent in the second group.

Similarly, individuals under-appreciate the extent to which they adapt to future circumstances. Gilbert, Pinel, Wilson, Blumberg, and Wheatly (1999) ask subjects to forecast their happiness in correspondence of an event, and compare these responses to the responses after the event has occurred. Thirty-three current assistant professors at the University of Texas (1998) forecast that getting tenure would significantly improve their happiness (5.9 versus 3.4 on a 1-7 scale). However, the difference in rated happiness between 47 assistant professors that were awarded tenure by the same university and 20 that were denied tenure is smaller and not significant (5.2 versus 4.7). Similar results apply for happiness forecasts as a function of the election of a Democratic or Republican president, compared to the realized *ex-post* differences. While these are just survey responses (below we provide evidence of impact on behavior), they suggest a consistent pattern of projection of the current preferences.

Loewenstein, O’Donoghue and Rabin (2003) propose a simple model of projection bias.

Assume that utility u is a function of consumption c and of a state variable s , that is, $u = u(c, s)$. The current state is s' and the future state is s . Then, when predicting the future utility $\hat{u}(c, s)$, an individual with projection bias expects utility

$$\hat{u}(c, s) = (1 - \alpha) u(c, s) + \alpha u(c, s') \quad (8)$$

rather than $u(c, s)$. The parameter $\alpha \in [0, 1]$ captures the extent of projection bias, with $\alpha = 0$ denoting the standard case and $\alpha = 1$ the case of full projection bias. This model can capture the mis-prediction of future hunger, as well as the under-appreciation of adaptation.

Conlin, O'Donoghue, and Vogelsang (forthcoming) present evidence of projection bias using a data set of 2 million orders of cold-weather apparel items. They consider the effect of weather at the time of purchase on the probability that an item is returned, conditional on purchase. According to the standard model, colder weather at the time of purchase should not affect the probability of a return, or may affect it negatively (since colder weather at the time of purchase is correlated with colder weather over the subsequent days). Projection bias, instead, makes the opposite prediction. On colder days, individuals overestimate the use that they will make of a cold-weather item, and hence are *ex post* more likely to return the item.²⁹ This prediction holds whether the projection bias regards future utility, as in (8) ('I expect to like cold-weather items very much'), or future weather ('I expect the coming winter to be very cold').

Conlin et al. (forthcoming) find that a reduction in the order-date temperature of 30°F—corresponding to a decrease, for example, from 40°F to 10°F—increases the average return rate of a cold-weather item by 3.96 percent, consistent with projection bias. A simple structural model of projection bias as in (8) implies estimates for $\hat{\alpha} \approx 0.5$, implying that consumers predict future tastes roughly half-way between present tastes and actual future tastes.

4 Non-standard Decision-Making

Even given utility $U(x|s)$ and belief $p(s)$, individuals make non-standard decisions. We analyze: (i) the impact of framing of a decision; (ii) the underweighting (or overweighting) of information because of limited attention; (iii) sub-optimal heuristics used for choices out of menu sets; (iv) social pressure—explicit pressure by others—and persuasion—excess impact of the beliefs of others; (v) emotions.

4.1 Framing

A key tenet of psychology is that the context and the framing of a situation matter. Two equivalent decision problems that are framed differently may elicit different responses. Tversky

²⁹A possible confound is that on colder days more 'marginal' individuals (which are more likely to return) order cold-weather clothing. The standard model modified for this form of heterogeneity makes the same prediction as the projection bias model.

and Kahneman (1981) present a classical example, which we reproduce in the version of Rabin and Weizsäcker (forthcoming). A group of subjects is asked to consider a pair of ‘concurrent decisions. [...] **Decision 1.** Choose between: A. a sure gain of $\text{€}2.40$ and B. a 25% chance to gain $\text{€}10.00$ and a 75% chance to gain $\text{€}0.00$. **Decision 2.** Choose between: C. a sure loss of $\text{€}7.50$ and D. a 75% chance to lose $\text{€}10.00$ and a 25% chance to lose $\text{€}0.00$.’ Of 53 participants playing this lottery for money, 49 percent chooses A over B and 68 percent chooses D over C. Overall, 28 percent of the subjects chooses the combination of A and D. This combined lottery, which amounts to a 75% chance to lose $\text{€}7.60$ and a 25% chance to gain $\text{€}2.40$, however, is dominated by the combined lottery of B and C, which reduces to a 75% chance to lose $\text{€}7.50$ and a 25% chance to gain $\text{€}2.50$. Over a quarter of the subjects, therefore, chooses a dominated lottery when the choice is presented with a narrow framing, that is, with each lottery presented individually. A separate group of 45 subjects is presented the same choice in the broad framing, that is, they are shown the distribution of outcomes induced by the four options. In this group, not surprisingly, none of the subjects chooses the A and D combination. Clearly, the framing of choices matters.

We can understand this first example of framing effects in light of a reference-dependent utility function (Section 2.2) with narrow framing. Individuals evaluate each of the two lotteries *separately*, comparing the outcomes relative to a reference point. The individuals are approximately risk-neutral over gains, inducing the 49 percent choosing A over B, and risk-seeking over losses, hence the 68 percent choosing D over C.³⁰ Importantly, the individuals accept the framing induced by the experimenter and do not aggregate the lotteries, that is, they frame narrowly. This example illustrates a general feature of human decisions: judgments are comparative, and changes in the framing can affect a decision if they change the nature of the comparison, even if they do not affect the underlying economic trade-offs. While there is no field evidence directly corresponding to this framing manipulation, in Section 2.2 I discussed several applications to the field of reference-dependent preferences.

A second example illustrates how the presentation format can affect preferences (in this case about financial options) even aside from its impact on reference points. Benartzi and Thaler (2002) survey 157 UCLA employees that participate in a 403(b) plan and ask them to rate three plans (labelled plans A, B, and C): their own portfolio, the average portfolio, and the median portfolio. For each portfolio, they present the 5th, 50th, and 95th percentile of the projected retirement income from the portfolio (obtained using the Financial Engines retirement calculator). Given revealed preferences, one would expect individuals on average to prefer their own plan to the other plans. However, the own portfolio rating (3.07) is about the same as the average portfolio rating (3.05) and substantially lower than the median portfolio rating (3.86). Indeed, 62 percent of employees gave a higher rating to the median portfolio

³⁰Other versions of this experiment (typically for hypothetical stakes) indicate a higher percentage of subjects choosing A over B, consistent with risk-aversion over gains, as in prospect theory.

than to their own portfolio. Re-framing the decision in terms of ultimate outcomes, therefore, appears to affect preference substantially. However, an alternative interpretation is that these employees never considered the median portfolio in their retirement savings decision, and would have chosen it had it been offered. To address this explanation, Benartzi and Thaler (2002) survey 351 participants in a different retirement plan who were explicitly offered a customized portfolio and actively opted out of it. These employees rate their own portfolio, the average portfolio, and the customized portfolio, similarly re-framed in terms of ultimate income. A majority (61 percent) of the employees prefers the customized portfolio (which they previously turned down) to their own portfolio. The choice of retirement savings, hence, depends substantially on the format of the choices presented. This framing effect presumably reflects the fact that consumers put too little weight on factors that determine ultimate returns, such as fees, or that they do not appreciate the riskiness of their investments.

A third example is the case in which the framing focuses the attention on different aspects of the options. Duflo, Gale, Liebman, Orszag, and Saez (2006) estimate the impact of a match on IRA participation for low- and middle-income households using a field experiment. A random sub-sample of H&R Block customers are offered either no match, a 20 percent, or a 50 percent match on the first \$1,000 contributed to an IRA. The take-up rate increases from 2.9 percent in the control group to 7.7 percent in the 20 percent match group and to 14.0 percent in the 50 percent match group. The authors then compare this substantial increase to the response to a comparable match induced by tax credits in the Saver's Tax Credit program. The effective match rate for IRA contributions decreases from 100 percent to 25 percent at the \$30,000 household income threshold. Duflo et al. (2006) compare the IRA participation for households slightly below the threshold (\$27,500-\$30,000) to households slight above the threshold (\$30,000-\$32,500). To control for other difference between these two income groups, they estimate the difference-in-difference relative to households in the same income groups that are however ineligible for the program. The difference in match rate lowers contributions by only 1.3 percentage points, a much smaller impact relative to the effects in the H&R Block field experiment. While there are a number of differences between the programs, a prominent factor is likely to be the simplicity of the H&R Block match that garnered more attention to the match. The next Section presents further evidence about the impact of limited attention on economics decisions. This example illustrates the importance of considering behavioral factors such as framing in the design of public policy programs.

4.2 Limited Attention

In the starkest form of the standard model, individuals make decisions using all the available information. Since Simon (1955), economists have attempted to relax this strong assumption and have proposed models in which individuals simplify complex decisions, for example by

processing only a subset of information³¹. In economic experiments, the simplifying heuristics include thinking only one step ahead in dynamic problems (Gabaix et al., 2006).

The laboratory studies in psychology indicate that attention is a limited resource. In studies of dichotic listening (Broadbent, 1958), for example, subjects hear different messages in the right ear and in the left ear, and are instructed to attend to one of the messages. When asked about the other message, they remember very little of it. Moreover, in treatments in which they have to rehearse a sentence or a sequence of numbers while listening, their capacity to attend to a message is substantially lower.

I present here a simple model of attention as a scarce resource and derive testable implications. Consider a good whose value V (inclusive of price) is determined by the sum of two components, a visible component v and an opaque component o : $V = v + o$. Due to inattention, the consumer perceives the value to be $\hat{V} = v + (1 - \theta)o$, where θ denotes the degree of inattention, with $\theta = 0$ as the standard case of full attention. The interpretation of θ is that each individual sees the opaque information o , but then processes it only partially, to the degree θ .³² The inattention parameter θ is itself a function of the salience $s \in [0, 1]$ of o and of the number of competing stimuli N : $\theta = \theta(s, N)$. Based on the psychology evidence, I assume that the inattention θ is decreasing in the salience s and increasing in the competing stimuli N : $\theta'_s < 0$ and $\theta'_N > 0$. Inattention is zero for a fully salient signal: $\theta(1, N) = 0$. The consumer's demand is $D[\hat{V}]$, with $D'[x] > 0$ for all x .

This model suggests, broadly speaking, three strategies to identify the inattention parameter θ , which the papers described below undertake. The first is to compute how the valuation \hat{V} responds to a change in o ; the derivative $\partial\hat{V}/\partial o = (1 - \theta)$ can be compared to $\partial\hat{V}/\partial v = 1$ to test for limited attention. Hossain and Morgan (2006) and Chetty, Looney, and Kroft (forthcoming) in the section on alcohol taxes follow this avenue. The second is to examine the response of consumer valuation to an increase in the salience s , $\partial\hat{V}/\partial s = -\theta'_s o$, and test whether it differs from zero. This is the strategy of Chetty et al. (forthcoming) in their field experiment. The third strategy is to vary the number of competing stimuli N , $\partial\hat{V}/\partial N = -\theta'_N o$, and test whether this has an effect. This is the strategy of DellaVigna and Pollet (forthcoming) and Hirshleifer, Lim, and Teoh (2007). All three of these strategies identify a piece of opaque information o with regards to which the decision-makers are not fully attentive.

This research is subject to two caveats. The first caveat is that measuring the salience of information involves a subjective judgment, similar to the judgment involved in setting the reference point in prospect theory. While in most settings (such as the ones in this Section)

³¹Conlisk (1996) provides an early survey of this literature. We discuss the model of inattention by Gabaix and Laibson (2006) in Section 5.

³²An alternative model (Chetty, Looney, and Kroft, 2007) posits that θ is the probability that an individual perceives the opaque signal, rather than the degree to which each individual incorporates the signal. This alternative model leads to similar results but a more cumbersome solution for settings like an auction.

it is rather clear which features are visible and which are opaque, the psychology experiments do not provide a general criterion. The second caveat is that we do not address whether the inattention is rational or not. In general, models of limited attention can be rephrased as rational model with information costs in which less salient information has higher costs of acquisition. In most of the examples below, however, the opaque information is publicly available at a zero or small cost (for example, the information on earnings announcements), making a rational interpretation of the findings less plausible.

Inattention to Shipping Costs. In eBay auctions, the price of an item is more vivid than the shipping cost, because the shipping cost is not listed in the item title and also because historically most purchases have not involved shipping. Define v as the value of the object and o as the negative of the shipping cost: $o = -c$. Since eBay is (essentially) a second price auction, the inattentive bidders bid their value net of the (perceived) shipping cost: $b^* = v - (1 - \theta)c$. The revenue raised by the seller is $b^* + c = v + \theta c$. A \$1 increase in the shipping cost c , therefore, increases revenue by θ dollars. In the case of full attention ($\theta = 0$), increases in the shipping cost have no effect on revenue. Hossain and Morgan (2006) examine these predictions with a field experiment. In the treatment c_{LO} , they auction CDs with a \$4 reserve price and no shipping cost, while in treatment c_{HI} they auction CDs with a \$.01 reserve price and a \$3.99 shipping cost. The change in reserve price guarantees that the two auctions are equivalent for a fully attentive bidder. The average revenue raised in treatment c_{HI} is \$1.79 higher (\$10.16 vs. \$8.37) than in treatment c_{LO} , and is higher for 9 out of 10 CDs³³. These estimates imply substantial inattention: $\hat{\theta} = 1.79/3.99 = .45$. A second set of auctions with higher shipping costs ($c_{LO} = \$2$ and $c_{HI} = \$6$), leads to smaller increase of revenue in the high-shipping cost condition (\$12.87 vs. \$12.15), corresponding to an inattention parameter $\hat{\theta} = 0.72/4 = .18$.

Inattention to Non-Transparent Taxes. Chetty et al. (forthcoming) study whether consumers are inattentive to taxes that are not transparently factored in the price of a good, like indirect state taxes. They use data on the demand for items in a grocery store. Assume that demand D is a function of the visible part of the value v , including the price p , and of the less visible part o , in this case the state tax $-tp$: $D = D[v - (1 - \theta)tp]$. The change in log-demand $\Delta \log D$ from making the tax fully salient ($s = 1$ and hence $\theta = 0$) is (linearizing the demand) $\log D[v - tp] - \log D[v - (1 - \theta)tp] = -\theta tp * D'[v - (1 - \theta)tp] / D[v - (1 - \theta)tp] = \theta t * \eta_{D,p}$, where $\eta_{D,p}$ is the price elasticity of demand. (Since demand D is a function of value minus price, $\eta_{D,p} = -pD'[v - (1 - \theta)tp] / D[v - (1 - \theta)tp]$) Notice that the response is zero for fully attentive consumers ($\theta = 0$). This implies $\theta = \Delta \log D / (t * \eta_{D,p})$. Chetty et al. (forthcoming) manipulate the salience of taxes with a field experiment. In a three-week period, the price tags of certain items make salient the after-tax price, in addition to indicating the pre-tax price. Compared to previous-week sales for the same item, and compared to items for which tax was not made salient, the average quantity sold decreases (significantly) by 2.20 units

³³We exclude CDs that do not sell from this computation; the difference would be \$2.60 if they were included.

relative to a baseline level of 25, an 8.8 percent decline. Since the price elasticity $\eta_{D,p}$ in this sample is estimated to be -1.59 and the tax is 7.375 percent, we can compute $\hat{\theta} = (-.088)/(-1.59 * .07375) \approx .75$. In a separate estimation strategy, Chetty et al. (forthcoming) identify the impact on beer consumption of changes across States and over time in the excise and sales taxes. Since the excise tax is included in the price, while the sales tax is added at the register, inattentive consumers should be more responsive to changes in the excise tax than to changes in the sales tax. Indeed, the first elasticity is substantially larger, leading to an estimate of the inattention parameter of $\hat{\theta} = .94$. Consumer inattention to non-transparent taxes is substantial.

Inattention to Complex Information in Rankings. In other settings, the familiarity of information depends on the simplicity of the data format. Pope (2007) studies the response of consumers to rankings of hospitals and colleges by the *US News and World Report*. Each year, the company constructs a continuous quality score from 0 to 100 largely based on reputation scores, and then creates rankings based on this score. Both the scores and the rankings are published in the yearly report. While the continuous score contains all the information, the rankings are presumably easier to process (“No. 5 Hospital” vs. “Hospital with 89/100 score”). Pope shows that, holding constant the quality score, hospital discharges respond significantly to differences in ranks among hospitals; similarly, college applications respond to differences in ranks among colleges. Pope (2007) also provides a calibration of the inattention or thinking costs necessary to justify this result.

Inattention to Financial News. Limited attention among investors induces under-reaction to newly-released information and hence can explain anomalies such as momentum (Hong and Stein, 1999). Huberman and Regev (2001) examines the case of the company EntreMed, an interesting example of under-reaction to information. On November 28, 1997, Nature prominently features an article reporting positive results on a cure for a type of cancer for a drug patented by EntreMed. On the same day, the New York Times reports an article on the same topic on page A28. Unsurprisingly, the stock price of EntreMed increases by 28 percent. What is surprising is what happens next. On May 4, 1998, the New York Times publishes on the front page an article on EntreMed that is very similar to the article that it had already published in November. Despite the fact that the article contains no new hard information, it leads to a 330 percent one-day return for EntreMed, and to a 7.5 percent one-day return for all bio-tech companies, moving billions in market capitalization. The stock price of EntreMed does not revert to the previous level over the whole next year.

While this is just a case study, it stresses the importance of studying systematically the response to new information. One important setting is the release of quarterly earnings news, and the consequent response of asset prices. To simplify, assume that v is the known information about cash-flows of the company, and that o is the new information contained in the earnings announcement. On the day before the announcement, the company price is $P = v$.

On the day of the announcement, the updated company value is $v + o$. However, since the investors are inattentive, the asset price P responds only partially to the new information: $P = v + (1 - \theta)o$. Over time, as the information makes its way to the inattentive investors (for example through additional articles as in the *EntreMed* case), the price incorporates the full value $v + o$. This implies that the short-run stock return r_{SR} equals $r_{SR} = (1 - \theta)o/v$; the long-run stock return r_{LR} , instead, equals $r_{LR} = o/v$. In this example, a measure of investor attention is $(\partial r_{SR}/\partial o)/(\partial r_{LR}/\partial o) = (1 - \theta)$. (The division by $(\partial r_{LR}/\partial o)$ is a re-normalization which makes the measure scale-invariant) The higher is the inattention, the smaller is the immediate response and the larger is the predictability of stock returns in the days following the announcement, a phenomenon known as post-earnings announcement drift (Bernard and Thomas, 1989). Inattention leads to delayed absorption of information.

While this setting is highly stylized, similar results obtain after allowing for uncertainty and arbitrage, as long as arbitrage is limited by risk aversion and short investor horizons (for example, DellaVigna and Pollet, forthcoming). DellaVigna and Pollet (forthcoming) estimate the empirical counterpart of $(\partial r_{SR}/\partial o)/(\partial r_{LR}/\partial o)$ using the response of returns r to the earnings surprise o . They measure returns in the 2 days surrounding an announcement (r_{SR}) and over the 75 trading days from an announcement (r_{LR}). The immediate response captures 54 percent of the overall response, implying substantial inattention: $\hat{\theta} \approx .46$. If the delayed response is due to attention deficits, the delay should be even stronger when a higher share of investors are distracted (higher θ). DellaVigna and Pollet (forthcoming) use the weekend as a proxy of investor distraction. For announcements made on Friday, indeed, the share of immediate response $(\partial r_{SR}/\partial o)/(\partial r_{LR}/\partial o)$ is 41 percent, implying $\hat{\theta} \approx .59$, consistent with higher inattention before the weekend. This provides an explanation for the observed release of worse earnings on Friday: companies maximizing short-term value release worse news on low-attention days.

In a similar context, Hirshleifer, Lim, and Teoh (2007) analyze the impact of informational overload (high N in the framework above). They find that the incorporation of earnings news into stock prices is 20 percent slower on days in which more announcements take place. Increasing the amount of competing information accentuates the effect of limited attention.

Another related study is Cohen and Frazzini (forthcoming), which analyzes how investors respond to indirect, and hence less salient news (low s in the framework above). They consider companies linked in the supplier-customer chain. When a customer company announces substantial earnings news, the news affects also the supplier, but this indirect effect is less likely to attract attention. Indeed, Cohen and Frazzini (forthcoming) show that suppliers of companies which experience declining stock returns have lower stock returns 1 to 3 months later. They measure the speed of the response of returns to news about the customer company using $(\partial r_{SR}/\partial o)/(\partial r_{LR}/\partial o)$, where r_{SR} is the one-month return and r_{LR} is a seven-month return. They find that for the customer company, 93 percent of the overall response occurs in

the initial month; for the supplier company, instead, only 60 percent of the overall response occurs in the first month, suggesting substantial inattention to indirect links.

A final dimension of salience s is the temporal distance. Holding constant the informativeness, information that is further into the future (or past) is less likely to be salient. In general, it is difficult to control for informativeness, since information that is further away is usually less relevant or less precisely estimated. DellaVigna and Pollet (2007) address this issue by considering future demand shifts due to demographics. Unlike other demand determinants, cohort size shifts are highly predictable even ten years into the future. For example, if a large cohort is born in 2006, school bus companies in 2012 are going to experience a forecastable increase in demand and, if the market is not perfectly competitive, in profits. If investors are perfectly attentive, this increase will be incorporated into returns already in 2006 and stock returns from 2006 to 2012 will not be predictable using demographic information. However, if investors neglect information beyond 5 years into the future, the stock prices will increase only in 2007, and stock returns from 2006 to 2012 will be predictable using public information on demographics. Using data for 48 industries from 1939 to 2003, DellaVigna and Pollet show that the growth rate in demand due to demographics 5 to 10 years ahead forecasts stock returns in an industry positively. These results are consistent with inattention to information further than approximately 5 years into the future.

Summary. Limited attention helps explain the (partial) neglect of: (i) shipping costs in eBay auctions; (ii) non-transparent taxes; (iii) complex information in rankings; (iv) earnings news, specially before weekends and on days with more competing news; (v) news about linked companies; (vi) demand shifts in the distant future. As an example of application to another field, a literature on inattention in macroeconomics developed from the models of sticky information of Mankiw and Reis (2002) and of rational inattention of Sims (2003).

4.3 Menu Effects

In this Section, I consider choices out of a (typically large) menu set, such as for investment options or politicians on a ballot. The evidence in psychology suggests that individuals use (at least) five sub-optimal heuristics to simplify these decisions: (i) *Excess Diversification* (or *1/n Heuristic*); (ii) *Preference for the Familiar*; (iii) *Preference for the Salient*; (iv) *Choice Avoidance*; (v) *Confusion in Implementing the Choices*.

Excess Diversification. Individuals facing a complex choice may simplify it by diversifying excessively across the options. An example in psychology is Simonson (1990). In a first treatment (simultaneous condition), students in a class chose snacks to be consumed over the next three class meetings, one per meeting. In a second treatment (sequential condition), the subjects chose the snack sequentially on each of the three class meetings. In the simultaneous condition, the subjects display excess diversification: 64 percent of subjects chose three

different snacks, while in the sequential condition only 9 percent of subjects made this choice.

Benartzi and Thaler (2001) study whether excess diversification applies to 401(k) investments. As a special case, they study the case of equal diversification across the n available options, the $1/n$ heuristics. They use aggregate data on the 1996 plan assets for 162 companies which offer an average of 6.8 plan options. Lacking individual-level data, they study an aggregate implication of the $1/n$ heuristic. If individuals divide their investments approximately equally across options, their exposure to equity will be increasing in the availability of equity options in the 401(k) plan. Across plans, Benartzi and Thaler estimate the relationship

$$\%Invested\ In\ Equity = \hat{\alpha} + .36(.04) * \%Equity\ Options + \hat{B}X \quad (9)$$

(s.e. in parentheses), where the control variable X is the availability of employer stock in the portfolio. In companies with an equity share that is 10 percentage points higher, the employees invest 3.6 percent more in equity plans. This finding is consistent with a weak form of the $1/n$ heuristic (If the employees followed the $1/n$ heuristics strictly, the coefficient should be 1 rather than .36). A confound is that the equity content of a plan may be designed to cater to the preferences of the employees, resulting in reverse causation.

Huberman and Jiang (2006) investigate the investor diversification using a data set on the individual choice of employees in 647 401(k) plans managed by Vanguard. They estimate specification (9) at the individual level with a large set of individual-level and plan-level controls X . They obtain the relationship $\%Invested\ In\ Equity = \hat{\alpha} + .29(.11) * \%Equity\ Options + \hat{B}X$ for funds with less than 10 options and $\%Invested\ In\ Equity = \alpha + .06(.09) * \%Equity\ Options + \hat{B}X$ for funds with more than 10 options. The relationship predicted by the $1/n$ heuristic, therefore, is present when the number of funds is small (as in the Benartzi and Thaler sample), but not when the number of funds is large. Huberman and Jiang provide additional evidence suggesting that the predictive power of the $1/n$ heuristic is low. In particular, the number of funds chosen by employees hardly responds at all to the number of investment options offered in the plan. (This test differs from the one above as it is not conditional on equity vs. non-equity choices.) There is some evidence of a *conditional 1/n heuristic*: conditional on the allocations chosen, individuals allocate their savings approximately equally. 37 percent of employees follow this behavior among employees investing in 4 funds, 26 percent among employees investing in 5 funds, and 53 percent among employees investing in 10 funds; the behavior is instead not common for non-round numbers. Overall, some employees use a version of the $1/n$ heuristic when the number of investment options is small; when the number is large, other heuristics, which I discuss next, are at play.

Preference for the Familiar. A different heuristic to simplify complex decisions is the choice of a familiar option. This tendency is wide-spread among individual investors. Investors in the USA, Japan, and the UK allocate 94%, 98%, and 82% of their equity investment, respectively, to domestic equities (French and Poterba, 1991). While the preference for own-

country equity may be due to costs of investments in foreign assets, the same pattern appears for within-country investment. Huberman (2001) documents the geographical distribution of the shareholders of the Regional Bell companies. The fraction invested in the own-state Regional Bell is 82 percent higher than the fraction invested in the next Regional Bell company. The preference for the familiar occurs despite substantial costs of under-diversification.

A particularly egregious case is the preference for own-company stock. On average, employees invest 20-30 percent of their discretionary funds in employer stocks (Benartzi, 2001), despite the fact that the employees' human capital is already invested in their company. This choice does not reflect private information about future performance. Companies where a higher proportion of employees invest in employer stock have lower subsequent one-year returns, compared to companies with a lower proportion of employee investment.

The preference for familiar options is consistent with ambiguity aversion. As in the classical Ellsberg (1961) paradox, investors that are ambiguity-averse may prefer an investment with known distribution of returns to an investment with unknown distribution, even if the average returns are the same for the two investments, and despite the benefits of diversification.

Preference for the Salient. Barber and Odean (forthcoming) show that individual investors simplify complex portfolio decisions also by choosing a salient option. Using individual trading data, they show that individual investors are net buyers of companies with unusually high, or low, performance in the previous day, of companies with high trading volume, and of companies in the news. The effects are large: for companies in the highest or lowest decile of the previous day's returns, the Buy-Sell imbalance ($\text{Buy-Sell}/(\text{Buy}+\text{Sell})$) for individual investors is 20 percentage points higher than for companies in the fifth decile. These results suggest that individuals solve the informational overload problem of which stocks to buy by picking companies that stand out. The same problem does not present itself for stock sales, since most investors own only a small number of stocks at any given time. This asymmetric pattern for stock purchases and sales could, however, also be due to short-sale constraint: individual investors would like to sell stocks that are 'salient', but cannot do so due to the short-sale constraints. Consistent with the inattention explanation, Barber and Odean (forthcoming) show that the Buy-Sell imbalance patterns are similar for stock that the individual investors already own, and that hence they could easily sell.

The preference for the salient takes different forms in different contexts. In the choice of candidates on a ballot, the first politician on the list stands out. A long-standing literature in political science, going back to Bain and Hecock (1957), examines the effect of the order of candidates on a ballot. Among the most convincing studies, Ho and Imai (2008) provides evidence that the order matters even when it is randomly determined. They exploit the natural experiment induced by the California voting system which, since 1975, explicitly randomizes the ballot order of candidates across Assembly Districts. They show that in the 1998 and 2000 general elections a minor party candidate experiences on average a 10 percent increase

in votes when first on the list. The effect is instead very small for candidates of the major parties, suggesting that irrelevant information is used as a tie-breaker for cases in which the decision-maker has less information. In primary elections, in which candidates are on average less known, the effect is stronger: the impact of being first in the list is on average a 20 percent increase, roughly 1.6 percent of the party vote.

Choice Avoidance. The fourth heuristic used to deal with difficult decisions is perhaps the most surprising: avoiding the choice altogether, possibly in favor of the default action. In a field experiment, Iyengar and Lepper (2000) compare the behavior of consumers in an upscale grocery store, where at some times consumers were offered the opportunity to taste 6 jams (the simple-choice treatment), while at other times the tasting included 24 jams (the difficult-choice treatment). They find that in the 24-jam condition, more consumers stop to sample jams (145 versus 104 customers), but substantially *fewer* buy jams (4 versus 31 customers). This finding is surprising in light of the fact that in the standard model more choice can only lead to increased purchases. The explanation for this counter-intuitive phenomenon is that when the choice is difficult, for example because the choice set is large, individuals find the decision stressful and look for ways to avoid the choice.

Bertrand et al. (2006) examine the impact of a small or large menu set in the context of a field experiment on the mailing of 50,000 loan offers in South Africa. The authors randomize, among other things, the format of the table illustrating the use of the loan. The small-table format lists only one loan size as an example, while the big-table format presents four different loan sizes. The finding is consistent with the choice avoidance results. The take-up in the small-table format is .6 percentage points larger compared to a baseline of 8 percentage points, an effect size equivalent to a reduction of the (monthly) interest rate by 2.3 percentage points.

Choi, Laibson, and Madrian (2006) also provide evidence from a field experiment that a smaller number of investment options increases participation in 401(k) plan. Participation increases by 10 percentage points when non-participating employees receive a card that allows them, if mailed back, to enroll in a default plan (3 percent contribution in a balanced fund).³⁴

Confusion. A final category of behavior, confusion, differs from the previous heuristics in that it does not reflect a preference, whether to avoid difficult choices or for salient options, but simply an error in the implementation of the preferences. As such, it differs from most behavioral phenomena which reflect a directional bias. A first setting is the choice of a political candidate among those in a ballot. Shue and Luttmer (forthcoming) consider California voters in the 2003 recall elections and exploit the random variation in the placement of candidates on

³⁴The increase may be due to a reminder effect of the card. However, in other settings, reminders, and more generally financial education, do not have such large effects. For example, Choi, Laibson, and Madrian (2005) sent a survey including 5 questions on the benefits of employer match to 345 employees that were not taking advantage of the match. A control group of 344 employees received the same survey except for the 5 specific questions. The treatment had no significant effect on the savings rate.

the ballot, similarly to Ho and Imai (2008). They find that the vote share of minor candidates is significantly higher for candidates whose name on the ballot is adjacent to the name of a major candidate. While this phenomenon could be due to a spill-over in attention, confusion is a more likely explanation: the effect of horizontal adjacency (a name to the right or to the left of the major candidate) is almost entirely due to adjacency on the confusing side. For example, in the sequence *Bubble, Candidate A, Bubble, Schwarzenegger, Bubble, Candidate B*, it is Candidate B that benefits from the presence of a major candidate, since some voters mistake its bubble for the bubble of Schwarzenegger. Candidate A does not benefit, nor do candidates located at a diagonal adjacency. Further, the spill-over of votes is larger for more confusing voting methods (such as punch-cards) and for precincts with a larger share of lower-education demographics, that are more likely to make errors when faced with a large number of options. This method allows for a measure of confusion. Across different voting methods, about 1 in 300 voters meaning to vote for a major candidate instead vote for a minor candidate. The phenomenon hence is small but not irrelevant. Importantly, it can have an aggregate effect, since confusion is likely to have a different prevalence among the voters of different major candidates.

Interestingly, Rashes (2001) identifies a similar phenomenon in the trades of two companies, MCI and MCIC. The ticker for the MCI communication company is MCIC, while MCI is the ticker for a little-known closed-end mutual fund, Massmutual Corporate Investors. Some investors attempting to trade shares of the larger communication company confuse tickers and trade the MCI company instead, resulting in a .56 correlation between the two trading volumes. This occurs despite the difference in fundamentals: the MCIC company, for example, has only a .03 correlation in volume with the communication company AT&T. The mis-trading causes a smaller, but still significant correlation of stock returns. Arbitrage moderates the impact of confusion on stock returns, but does not fully eliminate it. Using the correlation in volume and the average volume of trade for the two companies, one can compute the incidence of confusion among MCIC investors: about 1 in 2000 trades are placed in error, a confusion rate smaller than, but of the same order of magnitude, the confusion rate displayed by California voters.³⁵

Summary. When choosing from a large menu of options, decision-makers: (i) (to same extent) diversify excessively across the options; (ii) choose familiar options, such as own-country or own-company stock; (iii) choose salient options in investment choice or at the ballot; (iv) avoid the choice and do not invest (or do not purchase); (v) display some confusion in implementing their choices.

³⁵Assume that the volume (number of trades) of MCIC $V_{MCIC,t}$ equals a constant α plus the shares traded due to confusion, $s * V_{MCI,t}$, where s is the share of investors of MCI that incorrectly trade MCIC: $V_{MCI,t} = \alpha + sV_{MCIC,t}$. Given a simple correlation coefficient between the daily volumes $V_{MCI,t}$ and $V_{MCIC,t}$ of .56, we can infer $\hat{s} \hat{s} = .56 * s.d. (V_{MCI,t}) / s.d. (V_{MCIC,t}) \approx .56 * 10^{-3}$.

4.4 Persuasion and Social Pressure

Persuasion. In the standard model, individuals take into account the incentives of the information provider. The neglect of these incentives can lead to excess impact of the beliefs of the information provider, which I label *persuasion*. An example from a laboratory experiment is Cain, Loewenstein, and Moore (2005). The subjects are paid for the precision of the estimates of the number of coins in a jar. Since they see the jar only from a distance, they have to rely on the advice of a second group of subjects, the advisors, that inspect the jar from close. The two experimental treatments vary the incentives for the advisors. In a first treatment, the advisors are paid for how closely the subjects guess the number of coins; in a second treatment, the advisors are paid for how high the subjects' guess is. Despite the fact that the incentives are common-knowledge, the estimate of the subjects is 28 percent higher in the second treatment. The subjects do not discount enough for the conflict of incentives of the advisors.

In a financial setting, Malmendier and Shanthikumar (2007) analyze how small and large investors respond to recommendation by analysts. Analyst forecasts are notoriously biased upward—94.5 percent of recommendations are Hold, Buy, or Strong Buy—, and affiliated analysts are even more biased. Malmendier and Shanthikumar (2007) show that large investors take into account this bias, and discount the information: for example, they respond to a Hold recommendation by selling the shares of a company, and they discount heavily positive recommendations by affiliated analysts. Small investors, instead, are subject to persuasion. They follow the recommendations literally—for example holding a stock in response to a Hold recommendation—and do not discount for the additional distortions due to analyst affiliation.

In a political setting, DellaVigna and Kaplan (2007) tests whether the information provided by a news source convinces on average its audience. They exploit the geographical variation in the introduction in the cable programming of the Fox News Channel, a more conservative channel relative to the pre-existing news sources (CNN and the networks). They show that Fox News availability in the town cable programming in 2000 is largely idiosyncratic, conditional on a set of controls. Using the voting data for 9,256 towns, they find that the vote share for Republicans in 2000 is half-a-percentage point higher in the towns offering Fox News. They estimate that Fox News convinced 5 to 30 percent of the audience that was not already Republican, depending on the audience measure. The effect is of about the same size for the Presidential candidates and for the US Senate candidates, which Fox News does not cover. This indicates that the effect of Fox News extends beyond the candidates covered to the general political beliefs of the voters. The impact of Fox News can be a temporary effect for Bayesian voters that are learning about the bias of Fox News, or a persuasion effect for non-rational voters that do not take sufficiently into account the political orientation of Fox News.

Social Pressure. A separate reason for excess impact of the beliefs of others is the pressure to conform, or social pressure (Akerlof, 1991). Two classical laboratory experiments

illustrating the power of social pressure are Asch (1951) and Milgram (1963). In one of the Asch (1951) experiments, the subjects are shown two large white cards with lines drawn on them: the first card has three lines of substantially differing length on them, while the second card has only one line. The subjects are asked which of the lines in the second card is closest in length to the line in the first card. In a control treatment, the subjects perform the task in isolation and achieve 98 percent accuracy. In the high-social-pressure treatment, the subjects choose the line of comparable length after 4 to 8 subjects (who, unbeknownst to them, are confederates) unanimously choose the wrong answer. On average, over a third of subjects give the wrong answer to avoid disagreeing with the unanimous judgment of the other participants. While this result could be interpreted as social learning, the learning is unlikely to be about the length of the line, but possibly about the rules of the experiment. It should also be pointed out that the subjects were not paid for accuracy.

In the Milgram (1963) experiment, a group of subjects is told that their task is to monitor the learning of another subject (a confederate), and to inflict electric shocks on this subject when he makes an error. Encouraged by the experimenter, 62 percent of the subjects escalate the electric shocks up to a level of 450 Volts, despite hearing the subject scream in pain. This proneness to obedience comes as a surprise to the subjects themselves. When a different group of 40 subjects is provided with a description of the experiment and asked to predict how far subjects would go in inflicting shocks, no one predicts that 450 Volts would be reached.

In the field, social pressure is hard to distinguish from rational diffusion of information. In some studies, however, the social pressure motive is evident. Garicano, Palacios-Huerta, and Prendergast (2005) measure the length of extra-time that referees assign at the end of a game of soccer; in the extra-time the teams can score goals. They find that referees on average give twice as much extra time (4 minutes versus 2 minutes) when the extra time is bound to advantage the local team (1 goal behind) than when it is bound to hurt it (1 goal ahead). The effect is larger when stakes are higher (toward the end of the season) and when the social pressure is larger (larger attendance at the game). Referees respond significantly to pressure by the local public, despite official rules on what determines the length of extra-time.

Some of the *peer effect* literature also points to the importance of social pressure. Falk and Ichino (2006) measure the effect of peer pressure on task performance. High-school students in Switzerland were recruited to perform a one-time job for a flat payment; they were instructed to stuff letters into envelopes for 4 hours. The control group of 8 students did the task individually, while the treatment group of 16 students worked in pairs (but each student was instructed to stuff the envelopes individually). Students in the treatment group stuffed significantly more envelopes (221 vs. 190), and coordinated the effort within group: the within-pair standard-deviation of output is significantly less than the (simulated) between-pairs standard deviation.

While the results of Falk and Ichino (2006) could also be due to social learning, Mas and Moretti (forthcoming) presenting direct evidence of social pressure. They find that high-

productivity cashiers in a supermarket chain increase the productivity of co-workers that are present in the same shift. The effect is not due to exchange of information, such as on a price tag. The positive peer effect occurs only when the more productive co-worker is behind and therefore can observe the other worker's productivity. The effect is quite large: a one percent increase in the average permanent productivity of the workers behind increases the productivity of the peer by .23 percent; the effect is even larger for co-workers that are working at a closer distance. There is no effect of a highly-productive co-worker in front.

4.5 Emotions

Some of the previous phenomena, such as self-control problems, social preferences in giving, and projection bias in food purchase are likely mediated (at least partially) by emotional states, respectively temptation, empathy, and hunger. In Section 2.3, for example, we discussed how the transient effect of a 'gift' in the Gneezy-List (2006) field experiment points to the role of emotions in gift-exchange behavior. A large literature in psychology suggests that emotions play an important role in decision-making, and that different emotions operate very differently (Loewenstein and Lerner, 2003). In this Section, I consider two examples of emotions, mood and arousal, for which field evidence is available.

In psychology studies, even minor mood manipulations have a substantial impact on behavior and emotions. For example, on sunnier days, subjects tip more at restaurants (Rind, 1996) and express higher levels of overall happiness (Schwarz and Clore, 1983). In the field, mood fluctuations induced by the weather affect stock returns, despite the fact that daily weather fluctuations are unlikely to affect fundamentals. Days with higher cloud cover in New York are associated with lower aggregate US stock returns (Saunders, 1993). Hirshleifer and Shumway (2003) extend this analysis to 26 countries between 1982 and 1997 using the weather of the city where the stock market is located. They find a negative relationship between cloud cover (detrended from seasonal averages) and aggregate stock returns in 18 of the 26 cities. Days with completely covered skies have daily stock returns .09 percent lower than days with sunny skies, five percent of a standard deviation. After controlling for cloud cover, other weather variables such as rain and snow are unrelated to returns. If mood is the channel for these effects, other mood-altering events should have similar effects. Indeed, international soccer matches impact the daily stock returns for the losing country (Edmans, Garcia, and Norli, 2007). Compared to a day with no match, a loss lowers daily returns (significantly) by .21 percent. (Surprisingly, a win has essentially no effect). More important matches, such as World Cup elimination games, have larger effects. The effect does not appear to depend on whether the loss was expected or not. International matches in other sports have a consistent, though smaller, effect.

The effect of these mood-altering events on returns is likely due to (i) an impact on risk aversion or perception of volatility, or (ii) a projection of the mood to economic fundamen-

tals. The evidence above does not allow to distinguish these two effects. Mood induced by atmospheric factors can also induce subtler changes in behavior. Simonsohn (2007) examines the role of weather on the day of campus visit to a prestigious university. Students visiting on days with more cloud cover are significantly *more* likely to enroll. Simonsohn suggests that higher cloud cover induces the students to focus more on academic attributes versus social attributes of the school, a hypothesis supported by laboratory experiments.

A second set of laboratory experiments suggests that emotional arousal has an important short-run effect on decisions. In one experiment, subjects that are sexually aroused as part of the treatment report a substantially higher willingness to engage in behavior that may lead to date rape (Ariely and Loewenstein, 2005). In other experiments, subjects exposed to violent video clips are more likely to display more aggressive behavior, such as aggressive play during a hockey game, compared to a control group watching non-violent clips (Josephson, 1987). The short-run impact on violence is not due to imitation, since the violent movie did not involve sport scenes. Hence, this effect is likely due to arousal induced by exposure to movie violence.

Dahl and DellaVigna (2007) provide field evidence on this effect and estimate the short-run impact of exposure to media violence on violent crime. They exploit the time-series variation in movie violence at the box office and compare days in which the blockbuster movies are violent to days in which the blockbuster movies are non-violent. They find that on days in which exposure to media violence is higher, violent crime is *lower*. In particular, in the night following the exposure (midnight to 6AM), for every million people exposed to violent movies, violent crime is 2 percent lower. The results hold after controlling flexibly for seasonality, weather, and are robust to placebo specifications.

What explains this result, apparently in contrast with the laboratory findings? The key factor is the difference in treatments. In the laboratory experiments, treatment and control groups are required to watch either a violent or non-violent movie. Hence, the experiments capture the direct impact of movie violence, holding everything else constant. In the natural experiment in the field, instead, consumers optimally choose between violent movies and their other favorite activity. Hence, the estimated impact in the field captures the net effect of exposure to the movies, *compared* to the impact of this alternative activity. The subpopulation attending violent movies would, on average, be doing an even more dangerous activity, such as drinking at a bar, if they were not attending the theater. The two different sets of results, therefore, can be reconciled. Exposure to movie violence can lower violent behavior relative to the foregone alternative activity (the field findings), even if it increases violent behavior relative to exposure to non-violent movies (the laboratory findings). Indeed, one can use the field estimates to infer the direct effect of violent movies (as in the laboratory) if one can control for selection. After accounting for selection along observable dimensions (age and gender), Dahl and DellaVigna (2007) provide some evidence that indeed the direct effect of violent movies is to induce more violent crime, like in the laboratory experiments.

This study suggests that one ought to be careful about directly comparing the results of laboratory and field studies (Levitt and List, 2007). Results that appear at first contradictory can be reconciled in light of differences in design. In this particular case, both the laboratory and the field estimates provide useful evaluations of policy-relevant parameters. The field results provide evidence on the short-run impact of a policy that restricts the amount of media violence available in the theaters. The laboratory evidence provides evidence on the ban of violent material embedded in non-violent programming.

5 Market Response

In the previous Sections, I have documented how consumers deviate from the standard model in their choices of credit cards, clothing items, eBay bidding strategies, giving, health clubs, housing prices, insurance contracts, loans, and lotteries. I have discussed how workers make non-standard effort, labor supply, and retirement savings decisions. I have provided evidence of disposition effect, inattention, and overtrading among investors. Finally, I documented evidence of salience effects and confusion among voters.

This evidence is just the first step towards a better understanding of markets where agents display non-standard preferences and beliefs. This evidence raises a natural question: how do markets and institutions respond to these non-standard features? An important test for Psychology and Economics is whether it helps to understand markets and institutions, in addition to explaining individual behavior.

This Section discusses how rational actors respond to the non-standard features of other agents. Profit-maximizing firms respond to the non-standard features of consumer behavior in their contract design and pricing (Behavioral Industrial Organization). Employers tailor their employment contracts to the non-standard behavior of the employees (Behavioral Labor Economics). In response to the non-standard behavior of investors, rational investors alter their trading strategies, and firm managers alter the capital structure (Behavioral Finance and Behavioral Corporate Finance). Politicians change their behavior to respond to voter biases (Behavioral Political Economy). Finally, policy-makers can use the findings in Psychology and Economics to inform the design of institutions and of policy (Behavioral Institutional Design).

Before I proceed, I discuss an important caveat. If consumers have non-standard features, why should one expect firms, employers, financial operators, and politicians to not have them? Experience is a key difference. Unlike individual consumers, firms can specialize, hire consultants, and obtain feedback from large data sets and capital markets. Firms are also subject to competition. Compared to consumers, therefore, firms are less likely to be affected by biases (except for principle-agent problems), and we expect them to be close to profit maximization. In addition, even when, despite the reasons above, firms still have non-standard features, they still have incentives to respond to the non-standard features of consumers. Similar arguments

apply for employers, institutional investors, top managers, and politicians.

Behavioral Industrial Organization. The interaction between consumers with biases and rational, profit-maximizing firms is the central theme of the growing literature in behavioral industrial organization, surveyed in Ellison (2006). While this literature is mostly theoretical, the papers surveyed here also make predictions about observed pricing.

DellaVigna and Malmendier (2004) consider the profit-maximizing pricing with $(\beta, \hat{\beta}, \delta)$ consumers with self-control problems. A (monopolistic) firm sells a product which, as in Section 2.1, has immediate payoff b_1 and delayed payoff b_2 . The set-up covers investment goods such as exercise ($b_1 < 0$ and $b_2 > 0$) and leisure goods such as gambling ($b_1 > 0$ and $b_2 < 0$). The immediate payoff b_1 is stochastic with c.d.f. F . The firm produces the good at marginal cost c and sells it using a two-part tariff, with a lump-sum fee L and a unitary price p . DellaVigna and Malmendier (2004) show that the profit-maximizing price p^* satisfies

$$p^* - c = - \left(1 - \hat{\beta}\right) \delta b_2 \frac{f(\hat{\beta}\delta b_2 - p^*)}{f(\beta\delta b_2 - p^*)} - \frac{F(\hat{\beta}\delta b_2 - p^*) - F(\beta\delta b_2 - p^*)}{f(\beta\delta b_2 - p^*)}. \quad (10)$$

For standard agents ($\beta = \hat{\beta} = 1$), the two terms on the right-hand side of (10) are zero: the firm prices at marginal cost, $p^* = c$, to align the incentives of the consumers. For sophisticated agents with self-control problems ($\beta = \hat{\beta} < 1$), only the first term in (10) is non-zero: the firm prices investment goods below marginal cost ($p^* < c$) and leisure goods above marginal cost ($p^* > c$) to provide a commitment device—the pricing increases the consumption of investment goods and lowers the consumption of leisure goods. The deviation from marginal cost pricing, $-(1 - \hat{\beta}) \delta b_2$, is exactly the difference in how much the current self and the future selves value the delayed payoff b_2 ; hence, the firm offers a perfect commitment device. For fully naive agents with self-control problems ($\beta < \hat{\beta} = 1$), only the second term in (10) is non-zero: the firm again prices investment goods below marginal cost and leisure goods above marginal cost again, but for a different reason—it takes advantage of consumer overestimation (underestimation) of the consumption of investment (leisure) goods. The deviation from marginal cost pricing is indeed a function of the mis-estimation of consumption $F(\hat{\beta}\delta b_2 - p^*) - F(\beta\delta b_2 - p^*)$. These results generalize to the case of perfect competition, since competition only alters the equilibrium fee L^* . This theory rationalizes the presence of contracts with no payment per visit in health clubs ($b_2 > 0$), the presence of high interest rates but no annual fees for credit cards ($b_2 < 0$), and cheap room rates and buffets for gamblers in Las Vegas ($b_2 < 0$).

Eliaz and Spiegel (2006) generalize this analysis to allow for heterogeneity in naiveté and a more general form of time-inconsistency of preferences. They show that firms offer two types of contracts: perfect commitment devices that cater to time-inconsistent agents that are sufficiently sophisticated, and contracts that take advantage of the consumers that are sufficiently naive. Interestingly, the fully sophisticated agents do not exert any informational externality on the naïve types. Thus, the provision of the perfect commitment device does not

reduce the gains that the monopolist can extract from naïve types.

Gabaix and Laibson (2006) analyze the pricing with boundedly-rational consumers that do not pay attention to hidden features of products, that they call add-ons. In equilibrium, firms charge above-marginal cost prices for the add-ons. As in DellaVigna and Malmendier (2004), the firms respond to the misprediction of future purchases. This model provides an explanation for high (hidden) fees on bank accounts and credit cards. Gabaix and Laibson (2006) also discuss how markets do not generally provide incentives for de-biasing naive consumers.

Heidhues and Köszegi (2005) study the pricing of a monopolist when consumers have reference-dependent preferences and the reference point is the rational expectations equilibrium (Köszegi and Rabin, 2006). Consumers are loss-averse with respect to both lower quality and higher price, relative to the reference point. The main predictions are sticky prices (despite no menu costs) and sales, two common features of pricing. In equilibrium, even if costs are stochastic, firms adjust prices seldom in response to cost shifts because consumers suffer more from price increases than they benefit from price cuts. In addition, firms offer random sales because the expectation of sales increases the likelihood of purchases at high prices.

These papers point to a dichotomy in the welfare effects of the market response. If the agents have non-standard preferences, such as self-control problems or loss aversion, but have rational expectations, the firms provide welfare-maximizing contracts. The contracts offer first-best commitment devices against the self-control problem (DellaVigna and Malmendier, 2004; Eliaz and Spiegler, 2006) or lower the probability of falling below the reference point (Heidhues and Köszegi, 2005). If, instead, the agents have non-rational expectations, such as about the self-control or about the inattention, the profit-maximizing contract is likely to magnify the bias. Firms take advantage of the wrong expectations in the consumption of the tempting good (DellaVigna and Malmendier, 2004; Eliaz and Spiegler, 2006) or of the add-on (Gabaix and Laibson, 2006).

Behavioral Labor Economics. Contracting within a firm is also consistent with this framework. Kahneman, Knetsch, and Thaler (1986) present suggestive evidence using a survey that workers display loss aversion with respect to nominal wage losses, but not with respect to real wage losses. For example, 62 percent of respondents find unfair a wage cut of 7% in the presence of no inflation, but only 22 percent of respondents find unfair a 5% increase in salaries in presence of 12% inflation. Bewley (1999) documents similar patterns in a series of interviews. In response to a dislike for nominal wage cuts, a profit-maximizing employer should set wages such that nominal wage cuts would be rare. Card and Hyslop (1997) provide evidence on this prediction using CPS data. They consider the distribution of year-to-year changes in the nominal log wage, $\log w_t - \log w_{t-1}$. In the presence of aversion to nominal wage losses, we expect a discontinuity in the distribution at $\log w_t - \log w_{t-1} = 0$. Rather than introducing small cuts in the nominal wages that may lower morale and productivity, the employer keeps wages constant ($\log w_t - \log w_{t-1} = 0$), compensating possibly by firing more workers. Card

and Hyslop indeed show that a substantial fraction of the distribution of $\log w_t - \log w_{t-1}$ is missing for negative values, despite the presence of measurement error in the wage that tends to attenuate this finding. This is an example of a market response to a bias which is likely to maximize utility for the biased agents. The observed distribution of wages is such that the employees suffer only rarely the cost of nominal wage cuts.

Behavioral Finance. In asset markets, arbitrage in principle is likely to limit the importance of behavioral biases such as inattention and overconfidence for price formation. If an irrational agent believes that a (fair) coin will land on tails sixty percent of the time, arbitrage by well-informed agents will keep the odds of tails around fifty percent. In actual financial markets, however, several factors limit the impact of arbitrage. DeLong et al. (1991) considers the case of a mis-pricing that is stochastic, persistent, and correlated by so-called noise traders. If arbitrageurs are risk-averse and have a limited investment horizon, the noise traders affect the equilibrium price, despite arbitrage. If noise-traders are, for example, bullish about dot-coms, they will bid the price of dot-com shares higher. The arbitrageurs do not know whether the mis-pricing will get even worse in the next period, and given their short horizons (they have to liquidate the shares next period) they cannot short the shares aggressively enough. DeLong et al. (1991) also shows that the noise traders are not driven out of the market; under some conditions, in fact, they outperform the rational traders (since they take more risk).

The recent research in behavioral finance builds on the noise-trade models to capture the limits of arbitrage, and hence the relevance of non-standard behavior for asset prices. At the same time, this literature moved beyond these models in making explicit the source of ‘noise trading’. In Sections 3.1 and 4.2, for example, we discussed models of overconfidence and limited attention, which make specific predictions about the non-standard behavior and hence the effect on returns. The evidence on this class of models is summarized in Shleifer (2000) and Barberis and Thaler (2003).

Behavioral Corporate Finance. In corporate finance, the standard theory assumes that managers maximize company value subject to agency problems, given the demands of rational investors and creditors. A recent theory, known as *market timing*, expands this framework and assumes that investors may have an irrationally high or low valuation of the company. The CEO rationally responds to the mis-valuation through the equity issuance and merger decisions. CEOs provide additional shares to investors and undertake mergers when the shares are most likely to be over-priced, lowering the welfare of the biased investors. Market timing can explain the systematic underperformance of initial public offerings (IPOs) in the 3-5 years following the IPO (Loughran and Ritter, 1995). According to this interpretation, managers of private companies go public when the shares of their companies are over-priced, hence the underperformance of IPOs. Baker, Ruback, and Wurgler (2006) reviews the evidence supporting this theory. This theory complements the standard theory that issuance decisions respond to investment opportunities.

Behavioral Political Economy. Another setting in which we expect an asymmetry between rational and biased agents is politics. While politicians are experienced agents facing high-stake incentives and significant competition, voters make infrequent low-stake decisions—whether to vote and for whom. Therefore, we expect political settings to be well-described by the interaction of rational politicians and voters with non-standard preferences, such as imperfect memory and limited attention.

Eisensee and Stromberg (2007) provides an example of politicians responding to a bias of voters, inattention. They consider the decision by US ambassadors to release US aid in the days following a natural disaster in the country. Ambassadors presumably are more likely to release aid if they, or the government, get credit for their generosity. To capture this phenomenon, Eisensee and Stromberg exploit variation in voter inattention due to the presence of major news items in the US television evening news, or due to major sporting event like the Olympics. They find that the probability of USAID relief is 15 percent lower for disasters occurring on days with a 2 standard deviation higher intensity of news in the US media. Similarly, the probability of relief is 30 percent lower in the period of the Olympics. On days in which the American public is less likely to notice the US generosity, generous acts are less likely to take place. This is consistent with politician response to limited attention of voters.

Behavioral Institutional Design. While firms, investors, managers, and politicians may respond to biases by exploiting them, the response to biases need not be predatory. Societal rules and institutions can be designed to counter-act the effect of consumer biases and improve the welfare of consumers. Benartzi and Thaler (2004)’s Save More Tomorrow (SMarT) plan is an example of one such institutional design for 401(k) savings. In a SMarT plan, the contribution rate is set to increase at each future wage increase up to a capped level. While savings increases are the default, employees can opt out of the plan at any time. This plan is an attractive commitment device to individuals with self-control problems, since the default applies to future savings rates, rather than current ones. In addition, the plan is designed with an eye to individuals that are averse to nominal wage cuts (see above), since the increases in contribution rates occur at the time of pay increases.

Benartzi and Thaler (2004) provide evidence on three implementations of this plan. In the earliest implementation, the plan is offered to 207 employees that accept to meet with a financial consultant, but do *not* accept to increase the savings rate immediately, as recommended by the consultant. Of these 207 individuals, 162 individuals accept the SMarT plan, indicating a wide-spread demand for commitment. In this subset of 162 individuals, the contribution rate increases from 3.5 percent to 13.6 percent in just four years. This increase includes the 32 individuals who opted out of the plan by the fourth year. The early results from the other two implementations of the SMarT plan indicate that the take-up of the plan is lower if it is offered as an option via mail, as opposed to with an in-person meeting. The effects conditional on take-up are, however, similarly large.

A simple change in defaults, hence, can go a long way toward addressing under-saving. The large impact of the SMarT plan, which comes at very limited administrative cost, is particularly noticeable when compared to the limited impact of other policies to increase retirement savings such as financial education. Also, while this plan is designed to benefit individuals with self-control problems, it does not hurt individuals with time-consistent preferences, since these individuals can switch at any time. The success of this plan, as well as the research on default effect in 401(k) savings, have motivated the bill on *Automatic Savings and Pension Protection Act* that Congress enacted in 2006. This law gives incentives to companies to adopt 401(k) plans with automatic enrollment and automatic increases in savings.

While the evidence in Psychology and Economics can have important policy implications, such as in this case, other considerations limit the policy reach of this evidence. First, unlike in the Benartzi and Thaler (2004) case, welfare-enhancing policies can be impractical—for example, it is more difficult to use defaults to help people exercise more. Second, political economy considerations suggest caution in the implementation of policies (Glaeser, 2006). Nevertheless, behavioral phenomena should be taken into account alongside standard phenomena in the policy design. Future research will tell whether the 2006 Savings Act will remain an isolated application of Psychology and Economics to policy or will be the first of several.

6 Conclusion

In this survey, I summarized the field evidence on three classes of deviations from the standard model: non-standard preferences, non-standard beliefs, and non-standard decision-making. I discussed how rational agents in the market respond to these non-standard features. As this survey documented, deviations from the standard model are not confined to laboratory decisions. Most phenomena that are important in laboratory experiments also affect decisions in a variety of economic settings. Hence, I expect that economists will increasingly take behavioral phenomena into account in their analysis.

Why don't market forces eliminate non-standard behavior? While a full discussion of this objection is beyond the scope of this article, I address two related arguments, one on experience and another on aggregation. A first argument is that experience reduces non-standard behavior. Indeed, experience appears to mitigate the endowment effect (List, 2003 and 2004). Palacios-Huerta and Volji (2007) provide concordant evidence on the effect of experience on the ability to perform backward induction. They consider the centipede game. Chess players, who have to routinely perform backward induction-type reasoning, come close in their play to the predictions of backward induction, in sharp contrast to college students.

However, it would be wrong to conclude, based on this evidence, that behavioral phenomena should not matter in the field. I list four reasons. (i) In a number of economic decisions, feedback is infrequent (such as in house purchases) or noisy (such as in financial investments),

and hence most individuals are inexperienced. (ii) Experience can exacerbate a bias if individuals are not Bayesian learners. Haigh and List (2004) use a simple investment game and show that professional investors display significantly more myopic loss aversion (see Section 2.2) than students. A possible explanation is that the short-term incentives in the workplace teach these investors to frame problems narrowly, contrary to the prediction of the standard theory. (iii) In principle, debiasing by experienced agents can be a substitute for direct experience. However, as Gabaix and Laibson (2006) show, experienced agents such as firms typically have little or no incentive to debias individuals. (iv) Finally, not all non-standard features should be mitigated by experience. Experience should not affect social preferences any more than it should affect preferences for the characteristics of cars.

A second argument is that, even if experience or debiasing do not eliminate the biases, the biases will not affect aggregate market outcomes. The argument is made forcefully in financial markets: given arbitrage, the rational investors set prices. However, as we discussed, the limits to arbitrage (DeLong et al., 1991) imply that individuals with non-standard features will in general affect stock prices. In addition, in most settings, there is no plausible incentive to eliminate a bias and hence the effect of non-standard behavior aggregates linearly. If a share of the population procrastinates saving for retirement, the aggregate savings rate will reflect proportionally the under-saving by this group. This is true unless a different institutional design is put in place, such as the SMarT plan (Benartzi and Thaler, 2004). (Notice that this plan was designed by academics, not by firms—though firms ultimately adopted it).

Finally, the papers on behavioral IO indicate that the non-standard features, instead of having no impact, can in fact have a disproportionate impact on market outcomes. Lee and Malmendier (2007) provide an example regarding overbidding in eBay auctions. They define a case of overbidding when the final auction price is higher than a posted price for the same good available on eBay itself. They focus on an item for which the posted price is stable and essentially always available, and hence should be an upper bound for the bids in a rational model. The authors show that 42 percent of auctions end at a price above the posted price, a conclusion robust to the inclusion of shipping costs, to differences in item quality and in seller reputation. The key aggregation point is that this behavior is generated by many fewer than 42 percent of overbidders. In fact, only 17 percent of bidders ever overbid. The auction design is such that the overbidders are disproportionately likely to determine the final price.

To conclude, a natural question is what empirical research in Psychology and Economics will look like in the future. Methodologically, I expect future research to continue using mostly the methods encountered in this overview, field experiments (such as List, 2003 and Falk, forthcoming), natural experiments (such as Madrian and Shea, 2001 and DellaVigna and Kaplan, 2007), and inference from menu choice (such as DellaVigna and Malmendier, 2006 and Sydnor, 2006), in addition, of course, to the laboratory experiments. These methodologies will increasingly provide, in addition to reduced-form estimates, also structural estimates of the

parameters (such as in Laibson et al., 2006 and Conlin et al., forthcoming). These estimates will allow us to address a number of open questions, such as whether models of (β, δ) preferences can explain choice in different decisions for fixed parameters β , $\hat{\beta}$, and δ . This estimation would benefit from the availability of data sets with multiple decisions by the same individual. While individuals are likely to differ in their preferences and beliefs, we expect the same individual to behave consistently if the existing models capture the behavior accurately.³⁶ New, more parsimonious models of the phenomena presented in this survey will likely also emerge, as Fudenberg (2006) predicts.

As for the topics, future research is likely to reduce the imbalance across fields in economics and across topics in psychology. While the research in behavioral finance and consumption-savings has been very active, relatively few studies, instead, have tackled mortgage markets, development, and political decisions, fields ripe for exploration. The applications to public policy are also likely to become more common, along the lines of the 2006 *Automatic Savings and Pension Protection Act* inspired by the research on default effects in 401(k) savings. I also expect a continued interest in the market interaction between standard and non-standard agents, as in Section 5. In addition, while Psychology and Economics has focused mostly so far on inexperienced agents, future research is likely to document also the judgmental biases, such as overconfidence, of experienced agents such as managers and politicians.

Future research is also likely to explore psychological phenomena that have been largely neglected. For example, emotions, automatic processing, and implicit discrimination are likely to matter for economic decisions such as divorce, judicial sentencing, and policing.

To conclude, in the ten years since Rabin (1998)'s *Journal of Economic Literature* survey, the field of Psychology and Economics has accomplished much, in the mind of this writer. More parsimonious models and a boom of evidence from the field have contributed to integrate the laboratory findings in Psychology and Economics more into mainstream economics. Still, much remains to be done. Our knowledge of behavioral deviations is partial, with disparities by field, limited use of certain methodologies, and few applications of important psychological phenomena. Hopefully, ten years from now, we will be able to assess quantitatively which psychological factors matter in which decisions.

³⁶In a laboratory experiment, Fisman, Kariv, and Markovits (2007) use repeated decisions on giving to another subject to identify types of subjects with different social preferences. Their results suggest substantial heterogeneity.

References

- [1] Abadie, Alberto, and Sebastien Gay. 2006. “The Impact of Presumed Consent Legislation on Cadaveric Organ Donation: A Cross-Country Study” *Journal of Health Economics*, 25(4): 599–620.
- [2] Akerlof, George A. 1982. “Labor Contracts as Partial Gift Exchange.” *Quarterly Journal of Economics*, 97(4): 543–569.
- [3] Akerlof, George A. 1991. ”Procrastination and Obedience”, *American Economic Review*, 81(2): 1-19.
- [4] Alpert, M., and H. Raiffa. 1982. “A Progress Report on the Training of Probability Assessors.” In *Judgment Under Uncertainty: Heuristics and Biases*, ed. D. Kahneman, P. Slovic and A. Tversky, 294–305. Cambridge: Cambridge University Press.
- [5] Andreoni, James. 2006. “Philanthropy.” In *Handbook of Giving, Altruism, and Reciprocity*, ed. Serge-Christophe Kolm and Jean Mercier Ythier, 1201–1265. Amsterdam: Elsevier/North-Holland.
- [6] Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. “The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation.” *Journal of Economic Perspectives*, 15(3): 47–68.
- [7] Ariely, Dan and George Loewenstein. 2005. “The Heat of the Moment: The Effect of Sexual Arousal on Sexual Decision Making”, *Journal of Behavioral Decision Making*, 18(1), 1–12.
- [8] Ariely, Dan and Klaus Wertenbroch. 2002. “Procrastination, Deadlines, and Performance: Self-control by Precommitment.” *Psychological Science*, 13(3): 219-224.
- [9] Asch, Solomon E. 1951. “Effects of Group Pressure Upon the Modification and Distortion of Judgment.” In H. Guetzkow (Ed.), *Groups, leadership, and men*. Pittsburgh: Carnegie Press.
- [10] Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. “Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines.” *Quarterly Journal of Economics*, 121(2): .
- [11] Ausubel, Lawrence. 1999. “Adverse Selection in the Credit Card Market.” Working Paper, University of Maryland.
- [12] Bain, Henry and Donald S. Hecock. 1957. *Ballot Position and Voter’s Choice*. Detroit: Wayne State University Press.

- [13] Baker, Malcolm, Richard Ruback, and Jeffrey Wurgler. 2006. "Behavioral Corporate Finance: A Survey." In *The Handbook of Corporate Finance: Empirical Corporate Finance*, ed. Espen Eckbo. New York: Elsevier/North Holland.
- [14] Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2005. "Social Preference and The Response To Incentives: Evidence From Personnel Data", *Quarterly Journal of Economics*, 120(3): 917-962.
- [15] Barber, Brad and Terrance Odean. 2001. "Boys will be Boys: Gender, Overconfidence, and Common Stock Investment," *Quarterly Journal of Economics*, 116(1), 261-292.
- [16] Barber, Brad and Terrance Odean. Forthcoming. "All that Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors." *Review of Financial Studies*.
- [17] Barber, Brad, Terrance Odean, and Ning Zhou. Forthcoming. "Systematic Noise," *Journal of Financial Markets*.
- [18] Barberis, Nicholas, Ming Huang, and Tano Santos. 2001. "Prospect Theory and Asset Prices", *Quarterly Journal of Economics*, 116(1): 1-53.
- [19] Barberis, Nicholas, Ming Huang, and Richard Thaler. 2006. "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing" *American Economic Review*, 96(4): 1060-1090.
- [20] Barberis, Nicholas, Andrei Shleifer and Robert Vishny. 1998. "A model of investor sentiment", *Journal of Financial Economics*, 49: 307-345.
- [21] Barberis, Nicholas and Richard Thaler. 2003. "A Survey of Behavioral Finance." in *Handbook of the Economics of Finance*.
- [22] Barberis, Nicholas and Wei Xiong. Forthcoming. "What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation." *Journal of Finance*.
- [23] Benartzi, Shlomo. 2001. "Excessive Extrapolation and the Allocation of 401(k) Accounts to Company Stock" *Journal of Finance*, 56(5):1747-1764.
- [24] Benartzi, Shlomo and Richard H. Thaler. 1995. "Myopic Loss Aversion and the Equity Premium Puzzle." *Quarterly Journal of Economics* 100, 73-92.
- [25] Benartzi, Shlomo and Richard H. Thaler. 2001. "Naive Diversification Strategies in Defined Contribution Saving Plans." *American Economic Review*, 91(1): 79-98.
- [26] Benartzi, Shlomo and Richard H. Thaler. 2002. "How Much is Investor Autonomy Worth?" *Journal of Finance*, 57(4): 1593-1616.

- [27] Benartzi, Shlomo and Richard H. Thaler. 2004. "Save More Tomorrow (TM): Using Behavioral Economics to Increase Employee Saving." *Journal of Political Economy*, 112(S1): S164-S187.
- [28] Bergman, Nittai and Dirk Jenter. 2007. "Employee Sentiment and Stock Option Compensation." *Journal of Financial Economics*, 84(3), 667-712.
- [29] Bernard, Victor L. and Jacob K. Thomas. 1989. "Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?," *Journal of Accounting Research*, 27: 1-36.
- [30] Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman. 2006. "What's Psychology Worth? A Field Experiment in the Consumer Credit Market." Working paper.
- [31] Bewley, Truman F. 1999. *Why Wages Don't Fall During a Recession*, Harvard University Press.
- [32] Broadbent, Donald E. 1958. *Perception and communication*, Pergamon Press, New York, NY.
- [33] Buehler, R., Dale Griffin, and M. Ross. 1994. "Exploring the 'Planning Fallacy': Why People Underestimate their Task Completion Times." *Journal of Personality and Social Psychology*, 67: 366-381.
- [34] Cain, Daylain, George Loewenstein, and Don Moore. 2005. "The Dirt on Coming Clean: Perverse Effects of Disclosing Conflicts of Interest," *Journal of Legal Studies*, 34, 1-25.
- [35] Camerer, Colin. 2005. "Behavioral Economics" *World Congress of the Econometric Society*, London, 18-24.
- [36] Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cab Drivers: One Day At a Time", *Quarterly Journal of Economics*, 112(2), 407-441.
- [37] Colin Camerer, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue, and Matthew Rabin. 2003. "Regulation For Conservatives: Behavioral Economics And The Case For 'Asymmetric Paternalism'" *University of Pennsylvania Law Review*, 151, 1211-1254.
- [38] Camerer, Colin, and Lovallo. 1999. "Overconfidence and Excess Entry: An Experimental Approach," *American Economic Review*, 89(1): 306-318.
- [39] Campbell, John Y., and John Cochrane. 1999, "By Force of Habit: A Consumption-based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy*. 107, 205-251.

- [40] Card, David. 1994 “Intertemporal Labor Supply: An Assessment,” in C. Sims, ed., *Advances in Econometrics: Proceedings of the Econometric Society Sixth World Congress* (New York: Cambridge University Press).
- [41] Card, David, and Dean Hyslop. 1997. “Does inflation grease the wheels of the labor market?” in Christina D. Romer and David H. Romer (eds.), *Reducing Inflation: Motivation and Strategy*. NBER, *Studies in Business Cycles*, 30. Chicago: University of Chicago Press.
- [42] Charness, Gary, and Rabin, Matthew. 2002. “Understanding Social Preferences with Simple Tests.” *Quarterly Journal of Economics* 117(3), 817-869.
- [43] Chetty, Raj, Adam Looney, and Kory Kroft. 2007. Forthcoming. “Salience and Taxation: Theory and Evidence”, *American Economic Review*.
- [44] Choi, James J., David Laibson, and Brigitte C. Madrian. 2005. “\$100 Bills on the Sidewalk: Suboptimal Saving in 401(k) Plans” Working Paper.
- [45] Choi, James J., David Laibson, and Brigitte C. Madrian. 2006. “Reducing The Complexity Costs of 401(k) Participation Through Quick Enrollment” NBER Working Paper No. 11979.
- [46] Choi, James J., David Laibson, Brigitte C. Madrian and Andrew Metrick. 2004. “For Better or For Worse: Default Effects and 401(k) Savings Behavior,” in David Wise editor *Perspectives in the Economics of Aging*. Chicago, IL: University of Chicago Press, 81-121.
- [47] Choi, James J., David Laibson, Brigitte Madrian, and Andrew Metrick. 2005. “Optimal Defaults and Active Decisions,” NBER Working Paper.
- [48] Cicchetti, Charles J and Jeffrey A. Dubin. 1994. “A Microeconomic Analysis of Risk Aversion and the Decision to Self-Insure.” *The Journal of Political Economy*, 102(1): 169-186.
- [49] Clotfelter, Charles T. and Philip J. Cook. 1993. “The ‘Gambler’s Fallacy’ in Lottery Play.” *Management Science*, 39(12), 1521-1525.
- [50] Cohen, Lauren, and Andrea Frazzini. Forthcoming. “Economic Links and Predictable Returns”, *Journal of Finance*.
- [51] Conlin, Michael, Ted O’Donoghue, and Timothy J. Vogelsang. Forthcoming. “Projection Bias in Catalog Orders.” *American Economic Review*.
- [52] Conlisk, John. 1996. “Why Bounded Rationality?”, *Journal of Economic Literature*, 34(2): 669–700.

- [53] Cowgill, Bo, Justin Wolfers, and Eric Zitzewitz. 2008. “Using Prediction Markets To Track Information Flows: Evidence from Google” Working Paper.
- [54] Cronqvist, Henrik, and Richard H. Thaler. 2004 “Design Choices in Privatized Social-Security Systems: Learning from the Swedish Experience.” *American Economic Review Papers and Proceedings*, 94(2): 424–428.
- [55] Dahl, Gordon and Stefano DellaVigna. 2007. “Does Movie Violence Increase Violent Crime?”, Working Paper.
- [56] Daniel, Kent D., David Hirshleifer, and Avanidhar Subrahmanyam. 1998. “Investor psychology and security market under- and over-reactions”, *Journal of Finance*, 53: 1839-1886.
- [57] De Bondt, Werner F. M. and Richard Thaler. 1985. “Does The Stock Market Overreact?” *Journal of Finance*, 40(3): 793-805.
- [58] DellaVigna, Stefano and Ethan Kaplan. 2007. “The Fox News Effect: Media Bias and Voting.” *Quarterly Journal of Economics*, 122(3): 1187-1234.
- [59] DellaVigna, Stefano and Ulrike Malmendier. 2004. “Contract Design and Self-Control: Theory and Evidence.” *Quarterly Journal of Economics*, 119: 353-402.
- [60] DellaVigna, Stefano and Ulrike Malmendier. 2006. “Paying Not to Go to the Gym.” *American Economic Review*, 96(3): 694-719.
- [61] DellaVigna, Stefano and M. Daniele Paserman. 2005. “Job Search and Impatience.” *Journal of Labor Economics*, 23(3): 437-466.
- [62] DellaVigna, Stefano and Joshua Pollet. 2007. “Demographics and Industry Returns” *American Economic Review*, 97: 1167-1702.
- [63] DellaVigna, Stefano and Joshua Pollet. Forthcoming. “Investor Inattention and Friday Earnings Announcements” *Journal of Finance*.
- [64] DeLong, Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. 1990. “Noise Trader Risk in Financial Markets.” *Journal of Political Economy* 98(4): 703-738.
- [65] Diamond, Peter and Hannu Vartiainen. 2007. *Behavioral Economics and Its Applications*, Princeton University Press.
- [66] Duflo, Esther, William Gale, Jeffrey Liebman, Peter Orszag, and Emmanuel Saez. 2006. “Saving Incentives For Low- And Middle-Income Families: Evidence From A Field Experiment With H&R Block,” *Quarterly Journal of Economics*, 121(4), 1311-1346.

- [67] Dufwenberg, Martin and Kirchsteiger, Georg. 2004. "A Theory of Sequential Reciprocity", *Games and Economic Behavior*, 47, 268-98.
- [68] Edmans, Alex, Diego Garcia, and Oyvind Norli. 2007. "Sports Sentiment and Stock Returns" *Journal of Finance*, 62(4).
- [69] Eisensee, Thomas and David Strömberg. 2007. "News Floods, News Droughts, and U.S. Disaster Relief" *Quarterly Journal of Economics*, 122(2).
- [70] Eliaz, Kfir, and Ran Spiegler. 2006. "Contracting with Diversely Naive Agents" *Review of Economic Studies*, 73(3): 689–714.
- [71] Ellison, Glenn. 2006. "Bounded Rationality in Industrial Organization" in Blundell, Newey and Persson (eds.), *Advances in Economics and Econometrics: Theory and Applications*, Ninth World Congress, Cambridge University Press.
- [72] Ellsberg, Daniel. 1961. "Risk, Ambiguity, and the Savage Axioms", *Quarterly Journal of Economics*, 75: 643-69.
- [73] Falk, Armin. Forthcoming. "Gift Exchange in the Field." *Econometrica*.
- [74] Falk, Armin, and Andrea Ichino. 2006. "Clean Evidence of Peer Effects," *Journal of Labor Economics*, 24(1): 39-57.
- [75] Farber, Henry S. 2005. "Is Tomorrow Another Day? The Labor Supply of New York City Cab. Drivers," *Journal of Political Economy*, 113(1): 46-82.
- [76] Farber, Henry S. Forthcoming. "Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers." *American Economic Review*.
- [77] Fehr, Ernst and Simon Gächter, 2000. "Fairness and Retaliation - The Economics of Reciprocity", *Journal of Economic Perspectives*, 14, 159-181.
- [78] Fehr, Ernst and Lorenz Goette. 2007. "Do Workers Work More if Wages are High? Evidence from a Randomized Field Experiment." *American Economic Review*, 97(1).
- [79] Fehr, Ernst, Georg Kirchsteiger and A. Riedl. 1993. "Does Fairness prevent Market Clearing? An Experimental Investigation", *Quarterly Journal of Economics*, 108(2): 437-460.
- [80] Fehr, Ernst, and Schmidt, Klaus. 1999. "A Theory of Fairness, Competition and Cooperation," *Quarterly Journal of Economics* 114(3): 817-868.
- [81] Fisman, Raymond, Shachar Kariv, and Daniel Markovits. 2007. "Individual Preferences for Giving", *American Economic Review*, 97(5), 1858-1876.

- [82] Forsythe, Robert, Joel L. Horowitz, N.E. Savin, and Martin Sefton. 1994. "Fairness in Simple Bargaining Experiments" *Games and Economic Behavior*, 6(3): 347–369.
- [83] Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40(2): 351–401.
- [84] French, Kenneth R. and Poterba, James M. 1991. "Investor Diversification and International Equity Markets" *American Economic Review*, 81: 222–226.
- [85] Fudenberg, Drew. 2006. "Advancing Beyond *Advances in Behavioral Economics*," *Journal of Economic Literature*, 44: 694–711.
- [86] Fudenberg, Drew and David Levine. 2006. "A Dual Self Model of Impulse Control." *American Economic Review*, 96: 1449–1476.
- [87] Gabaix, Xavier, David Laibson, Guillermo Moloche and Stephen Weinberg. 2006. "Information Acquisition: Experimental Analysis of a Boundedly Rational Model", *American Economic Review*, 96(4): 1043–1068.
- [88] Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics*, 121(2): 505–540.
- [89] Garicano, Luis, Ignacio Palacios-Huerta, and Canice Prendergast. 2005. "Favoritism Under Social Pressure," *The Review of Economics and Statistics*, 87(2): 208–216.
- [90] Genesove, David and Christopher Mayer. 2001. "Loss Aversion and Seller Behavior: Evidence from the Housing Market," *Quarterly Journal of Economics*, 116(4), 1233–1260.
- [91] Gilbert, Daniel T., Elizabeth C. Pinel, Timothy D. Wilson, Stephen J. Blumberg, and Thalia P. Wheatley. 1998. "Immune neglect: A source of durability bias in affective forecasting." *Journal of Personality and Social Psychology*, 75(3): 617–638.
- [92] Glaeser, Edward. 2006. "Paternalism and Psychology" *The University of Chicago Law Review*, 73(1): 133–156.
- [93] Gneezy, Uri, and John List. 2006. "Putting Behavioral Economics To Work: Testing For Gift Exchange In Labor Markets Using Field Experiments", *Econometrica*, 74(5), 1365–1384.
- [94] Gul, Faruk and Wolfgang Pesendorfer. 2001. "Temptation and Self-Control," *Econometrica*, 69(6): 1403–1435.

- [95] Haigh, Michael, and John A. List. 2004. "Do Professional Traders Exhibit Myopic Loss Aversion? An Experimental Analysis," *Journal of Finance*, 60(1), 523-534.
- [96] Heidhues, Paul, and Botond Köszegi. 2005. "The Impact of Consumer Loss Aversion on Pricing." Working Paper.
- [97] Hirshleifer, David A., Sonya S. Lim, and Siew Hong Teoh. 2007. "Driven to Distraction: Extraneous Events and Underreaction to Earnings News" Working Paper.
- [98] Hirshleifer, David A. and Tyler Shumway. 2003. "Good Day Sunshine: Stock Returns and the Weather", *Journal of Finance*, 58(3), 1009-1032.
- [99] Ho, Daniel E., and Kosuke Imai. 2008. "Estimating Causal Effects of Ballot Order from a Randomized Natural Experiment: California Alphabet Lottery, 1978-2002." *Public Opinion Quarterly*, 72(2).
- [100] Holt, Charles A., and Susan K. Laury. 2002. "Risk aversion and incentive effects." *American Economic Review*, 92(5), 1644-1655.
- [101] Hong, Harrison and Jeremy C. Stein. 1999. "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets", *Journal of Finance*, 54(6): 2143-2184.
- [102] Hossain, Tanjim and John Morgan. 2006. "...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay," *Advances in Economic Analysis & Policy*, 6(2): 1429-1429.
- [103] Huberman, Gur. 2001. "Familiarity Breeds Investment," *Review of Financial Studies*, 14(3): 659-680.
- [104] Huberman, Gur, and Tomer Regev. 2001. "Contagious Speculation and a Cure for Cancer: A Nonevent that Made Stock Prices Soar." *Journal of Finance*, 56(1): 387-396.
- [105] Huberman, Gur, and Wei Jiang. 2006. "Offering vs. Choice in 401(k) Plans: Equity Exposure and Number of Funds." *Journal of Finance*, 61(2): 763-801.
- [106] Ivkovich, Zoran, James Poterba and Scott Weisbenner. 2005 "Tax-Motivated Trading by Individual Investors." *American Economic Review*, 95(5): 1605-1630.
- [107] Iyengar, Sheena and Mark Lepper. 2000. "When Choice is Demotivating: Can One Desire Too Much of a Good Thing?" *Journal of Personality and Social Psychology*, 79(6): 995-1006.
- [108] Johnson, Eric J. and Daniel Goldstein. 2003. "Do Defaults Save Lives?" *Science*, 302(5649), 1338-1339.

- [109] Johnson, Eric J., John Hershey, Jacqueline Msezaros, and Howard Kuhnreuther. 1993. "Framing, Probability Distortions, and Insurance Decisions." *Journal of Risk and Uncertainty*, 7(1): 35–51.
- [110] Josephson, Wendy L. 1987. "Television Violence and Children's Aggression: Testing the Priming, Social Script, and Disinhibition Predictions." *Journal of Personality and Social Psychology*, 53, 882–890.
- [111] Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1986. "Fairness and the assumptions of economics." *Journal of Business*, 59(4): S285-S300.
- [112] Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1990. "Experimental Tests of the Endowment Effect and the Coase Theorem." *Journal of Political Economy*, 98(6): 1325-1348.
- [113] Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk" *Econometrica*, 47(2): 263-292.
- [114] Kim, E. Han, Adair Morsem and Luigi Zingales. 2006. "What Has Mattered To Economics Since 1970", *Journal of Economic Perspectives*, 20(4): 189-202.
- [115] Köszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences" *Quarterly Journal of Economics*, 121(4): 1133-1166.
- [116] Krueger, Alan and Alexandre Mas. 2004. "Strikes, Scabs and Tire Separations: Labor Strife and the Production of Defective Bridgestone/Firestone Tires," *Journal of Political Economy*, 112(2): 253-89.
- [117] Kruger, Justin. 1999. "Lake Wobegone Be Gone! The 'Below-Average Effect' and The Egocentric Nature of Comparative Ability Judgments" *Journal of Personality and Social Psychology*, 77(2): 221-232.
- [118] Kube, Sebastian, Michel André Maréchal, and Clemens Puppe. 2008a. "Do Wage Cuts Damage Work Morale? Evidence From a Natural Field Experiment" University of Zurich, IEW Working Paper.
- [119] Kube, Sebastian, Michel André Maréchal, and Clemens Puppe. 2008b. "The Currency of Reciprocity – Gift-Exchange at the Workplace" Working Paper.
- [120] Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 112(2): 443-477.
- [121] Laibson, David, Andrea Repetto, and Jeremy Tobacman. Forthcoming. "Estimating Discount Functions with Consumption Choices over the Lifecycle," *American Economic Review*.

- [122] Landry, Craig E., Andreas Lange, John A. List, Michael K. Price, and Nicholas G. Rupp. 2006. "Toward An Understanding Of The Economics Of Charity: Evidence From A Field Experiment," *Quarterly Journal of Economics*, 121(2): 747-782.
- [123] Lee, Han and Ulrike Malmendier. 2007. "The Bidder's Curse", Working Paper.
- [124] Levitt, Steven and John A. List. 2007. "What Do Laboratory Experiments Measuring Social Preferences Tell Us About The Real World?" *Journal of Economic Perspectives*, 21(2).
- [125] List, John A. and David Lucking-Reiley. 2002. "The Effects of Seed Money and Refunds on Charitable Giving: Experimental Evidence from a University Capital Campaign." *Journal of Political Economy*, 110(1): 215-233.
- [126] List, John A. 2003. "Does Market Experience Eliminate Market Anomalies?" *Quarterly Journal of Economics*, 118(1): 41-71.
- [127] List, John A. 2004. "Neoclassical Theory Versus Prospect Theory: Evidence from the Marketplace," *Econometrica*, 72(2): 615-625.
- [128] List, John A. 2006. "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions", *Journal of Political Economy*, 114(1): 1-37.
- [129] Loewenstein, George and Jennifer Lerner. 2003. "The Role of Affect in Decision Making," in *Handbook of Affective Science*, eds. Richard J. Dawson, Klaus R. Scherer and H. Hill Goldsmith, Oxford: Oxford University Press, 619-42.
- [130] Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias In Predicting Future Utility," *Quarterly Journal of Economics*, 118(4): 1209-1248.
- [131] Loewenstein, George and Drazen Prelec. 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Quarterly Journal of Economics*, 107(2): 573-597.
- [132] Loughran, Tim and Jay R. Ritter. 1995. "The New Issues Puzzle." *Journal of Finance*, 50(1): 23-51.
- [133] Madrian, Brigitte C. and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics*, 116(4): 1149-1187.
- [134] Malmendier, Ulrike, and Devin Shanthikumar. 2007. "Are Small Investors Naive about Incentives?", *Journal of Financial Economics*, 85(2), 457-489.

- [135] Malmendier, Ulrike, and Geoffrey Tate. 2005. "CEO Overconfidence and Corporate Investment." *Journal of Finance*, 60 (6): 2661-2700.
- [136] Malmendier, Ulrike, and Geoffrey Tate. Forthcoming. "Who Makes Acquisitions? CEO Overconfidence and the Market's Reaction," *Journal of Financial Economics*.
- [137] Mankiw, N. Gregory and Ricardo Reis. 2002. "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve", *Quarterly Journal of Economics*, 117(4): 1295-1328.
- [138] Mas, Alexander. 2006. "Pay, Reference Points, And Police Performance", *Quarterly Journal of Economics*, 121(3): 783-821.
- [139] Mas, Alexander and Moretti, Enrico. Forthcoming. "Peers At Work", *American Economic Review*.
- [140] McClure, Samuel M., David Laibson, George Loewenstein and Jonathan D. Cohen. 2004. "Separate Neural Systems Value Immediate and Delayed Monetary Rewards." *Science*, 306.
- [141] Mehra, R., and E. Prescott. 1985, "The equity premium: a puzzle", *Journal of Monetary Economics*, 15: 145-161.
- [142] Milgram, Stanley. 1963. "Behavioral study of obedience." *Journal of Abnormal and Social Psychology*, 67: 371-378.
- [143] Mullainathan, Sendhil and Richard H. Thaler. 2001. "Behavioral Economics" in *International Encyclopedia of Social Sciences*, Pergamon Press, 1094-1100.
- [144] Odean, Terrance. 1998. "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance*, 53(5): 1775-1798.
- [145] Odean, Terrance. 1999. "Do investors trade too much?" *American Economic Review*, 89: 1279-1298.
- [146] O'Donoghue, Ted and Matthew Rabin. 1999a. "Doing It Now or Later," *American Economic Review*, 89(1): 103-124.
- [147] O'Donoghue, Ted and Matthew Rabin. 1999b. "Procrastination in Preparing for Retirement", in *Behavioral Dimensions of Retirement Economics*, Henry Aaron, editor, The Brookings Institution.
- [148] O'Donoghue, Ted and Matthew Rabin. 2001. "Choice and Procrastination," *Quarterly Journal of Economics*, 116(1), 121-160.

- [149] Oyer, Paul and Scott Schaefer. 2005. “Why do Some Firms Give Stock Options to All Employees?: An Empirical Examination of Alternative Theories.” *Journal of Financial Economics*, 76(), 99-133.
- [150] Palacios-Huerta, Ignacio and Oscar Volji. 2007. “Field Centipedes”, Working paper.
- [151] Paserman, M. Daniele. Forthcoming. “Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation”, *Economic Journal*.
- [152] Phelps, E. and R. A. Pollak. 1968 “On Second-best National Saving and Game-equilibrium Growth”, *Review of Economic Studies*, 35: 185–199.
- [153] Plott, Charlie R., and Kathryn Zeiler. 2005. “The Willingness to Pay/Willingness to Accept Gap, the ‘Endowment Effect’ Subject Misconceptions and Experimental Procedures for Eliciting Valuations”, *American Economic Review*, 95: 530-545.
- [154] Pope, Devin. 2007. “Reacting to Rankings: Evidence from ‘America’s Best Hospitals and Colleges’”, Working Paper.
- [155] Rabin, Matthew. 1993. “Incorporating Fairness into Game Theory and Economics”, *American Economic Review*, 83, 1281-1302.
- [156] Rabin, Matthew. 1998. “Psychology and Economics.” *Journal of Economic Literature*, 36(1): 11-46.
- [157] Rabin, Matthew. 2000. “Risk Aversion and Expected-Utility Theory: A Calibration Theorem,” *Econometrica* 68(5), 1281-1292.
- [158] Rabin, Matthew. 2002a. “A Perspective on Psychology and Economics” *European Economic Review*, 46(4-5): 657-685.
- [159] Rabin, Matthew. 2002b. “Inference by Believers in the Law of Small Numbers.” *Quarterly Journal of Economics*, 117(3): 775-816.
- [160] Rabin, Matthew and Georg Weizsäcker. Forthcoming. “Narrow Bracketing and Dominated Choices” *American Economic Review*.
- [161] Rind, B. 1996. “Effect of beliefs about weather conditions on tipping”. *Journal of Applied Social Psychology*, 26(2), 137–147.
- [162] Michael S. Rashes. 2001. “Massively Confused Investors Making Conspicuously Ignorant Choices (MCI–MCIC).” *Journal of Finance*, 56(5): 1911-1927.

- [163] Read, Daniel and Barbara van Leeuwen. 1998. "Predicting hunger. The Effects of Appetite and Delay on Choice." *Organizational Behavior and Human Decision Processes*, 76, 189-205.
- [164] Samuelson, William and Richard Zeckhauser. 1988. "Status-Quo Bias". *Journal of Risk and Uncertainty*.
- [165] Saunders, E. M. J., 1993. "Stock Prices and Wall Street Weather", *American Economic Review*, 83, 1337-1345.
- [166] Schwarz, Norbert, & Clore, G. L. 1983. "Mood, Misattribution, and Judgments of Well-being: Informative And Directive Functions of Affective States". *Journal of Personality and Social Psychology*, 45(3), 513-523.
- [167] Shefrin, Hersh M. and Richard H. Thaler. 1981. "An Economic Theory of Self-Control." *Journal of Political Economy*, 89(2), 392-406.
- [168] Shleifer, Andrei. 2000. *Inefficient Markets: An Introduction to Behavioral Finance*. Oxford University Press.
- [169] Shue, Kelly and Erzo Luttmer. Forthcoming. "Who Misvotes? The Effect of Differential Cognition Costs on Election Outcomes," *American Economic Journal: Economic Policy*.
- [170] Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics*, 69(1), 99-118.
- [171] Simonsohn, Uri. 2007. "Weather to Go To College", Working Paper.
- [172] Simonson, Itamar. 1990. "The Effect of Purchase Quantity and Timing on Variety-Seeking Behavior," *Journal of Marketing Research*, 27(2), 150-162.
- [173] Sims, Christopher. 2003. "Implications of Rational Inattention", *Journal of Monetary Economics*, 50(3), 665-690.
- [174] Skiba, Paige and Jeremy Tobacman. 2008. "Payday Loans, Uncertainty, and Discounting: Explaining Patterns of Borrowing, Repayment, and Default" Working paper.
- [175] Strotz, R. H. 1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies*, 23: 165-180.
- [176] Svenson, Ola. 1981. "Are we all less risky and more skillful than our fellow drivers?" *Acta Psychologica* 47: 143-148.
- [177] Sydnor, Justin. 2006. "Sweating the Small Stuff: Risk Aversion in Home Insurance." Working paper.

- [178] Terrell, Dek. 1994. "A Test of the Gambler's Fallacy – Evidence from Pari-Mutuel Games," *Journal of Risk and Uncertainty*, 8(3): 309-317.
- [179] Thaler, Richard H. 1981. "Some Empirical Evidence on Dynamic Inconsistency," *Economics Letters*, 8: 201-7.
- [180] Tversky, Amos and Daniel Kahneman. 1974. "Judgement under uncertainty: Heuristics and biases," *Science*, 125: 1124–1131.
- [181] Tversky, Amos and Daniel Kahneman. 1981. "The Framing of Decisions and The Psychology of Choice," *Science*, 211(4481), 453-458.
- [182] Tversky, Amos, and Daniel Kahneman. 1992. "Advances in prospect theory: cumulative representation of uncertainty", *Journal of Risk and Uncertainty*, 5: 297-323.
- [183] Weinstein, Neil D. 1980. "Unrealistic Optimism about Future Life Events." *Journal of Personality and Social Psychology*, 39: 806-820.