5

Information and Self-Control

ISABELLE BROCAS AND JUAN D. CARRILLO

1. INTRODUCTION

Nobody doubts that the behavior of agents often violates the most basic premises of rationality imposed in economics. Examples of such violations are provided in some of the essays in this volume. They suggest that looking for a purely rational explanation is sometimes either hopeless or absurd. On the other hand, labeling as irrational any conduct for which the explanation is not immediate, seems also a quite unsatisfactory strategy. In this essay, we show that if we keep the traditional ingredients of economic analysis and enrich these models in a few (and fairly natural) directions, then it is possible to rationalize some puzzling patterns of behavior.

There is considerable evidence that most smokers overestimate the risk of lung cancer, most entrepreneurs underestimate the risk of failure of their business activities, and most people overestimate the risk of transmission of the HIV virus during unprotected sexual intercourse. Should we deduce from these systematic biases that individuals are intrinsically optimistic in some situations (entrepreneurship) and intrinsically pessimistic in others (risks associated with pleasurable activities)? Can we conclude that there is a lack of rationality in the agents' capacity to process information? The answer is 'not necessarily'.

Impulse buying has also received a great deal of attention, especially in the management literature (see e.g., Bell, 1976). With the development of credit opportunities for the purchase of consumption goods, habits have been radically modified. Impulse buying has become so recurrent that some individuals are forced to freeze (in the literal sense) their own credit cards, in order to limit their spending. Is this behavior driven entirely by emotions and urges? Does it mean that agents are irrational and unable to anticipate the long run damaging consequences of this behavior? Again, the answer is 'not necessarily'.

The object of the exercise we propose is not to demonstrate the power of formal model-building or to reaffirm our confidence in the orthodox economic methodology. On the contrary, if our hypotheses are correct, then analyzing problems

We thank Andrew Caplin, Mathias Dewatripont and Itzhak Gilboa for their helpful comments. Isabelle Brocas gratefully acknowledges financial support from the European Commission (TMR network FMRX-CT98-0203).

under this light has three main advantages. First, it offers unambiguous predictions concerning the class of situations in which a population of rational agents may look like intrinsic optimists, intrinsic pessimists, and impulsive. These behaviors, widely documented in psychology, have for a long time resisted economic analyses. In addition, one can use comparative static analysis to explore the impact of external factors on the magnitude of these apparent biases in judgments and behaviors. Second, the new ingredients of these enriched models may prove useful in explaining other apparently unrelated phenomena. Some of them are briefly discussed below, but clearly there is room for further research along these lines. Third, and maybe most important, once the underlying motivations for these original behaviors are understood, it is possible to conduct welfare analyses that will result in policy recommendations. In fact, identifying general patterns of behavior in order to isolate inefficiencies is probably the most standard methodology employed by economists. This, together with the belief that the ultimate goal is to offer solutions in order to avoid suboptimal behaviors, are the two main differences with their fellow psychologists.

The plan of the paper is as follows. In Section 2, we discuss general issues about incentives to acquire information and their aggregate effects on the behavior of agents. In Section 3, we rationalize the willingness of agents to keep optimistic or pessimistic beliefs and we study its welfare implications. In Section 4, we analyze why agents may act impulsively. Some final remarks are presented in Section 5.

2. THE ROLE OF INFORMATION IN DECISION-MAKING

2.1. Unbiased Individual Beliefs and Biased Aggregate Behavior

The literature on self-perception suggests that most individuals hold excessively optimistic beliefs about their own qualities. One of the most popular among these studies argues that 90 percent of individuals consider themselves to be better than the average driver (see e.g., Köszegi, 2000, for a brief overview of this literature). Can we conclude on the basis of this observation that individuals are not rational processors of information? Of course we cannot. In fact, assuming that individuals with imperfect knowledge incorporate information in a rational way (or, formally stated, that they update information according to the Bayes rule), only implies that the beliefs in the population cannot be biased *on average*. Therefore, it does not preclude, for example, a situation in which the vast majority of individuals (say, 90 percent) has a slightly positive bias in their perceived ability, and the small minority (say, 10 percent) has a significant negative bias. The point of this observation is not to claim that Bayesian updating is a perfect predictor for the information processing behavior of agents. It only suggests that the Bayesian rationality imposes fewer constraints than one might think at first sight. Technically, under the Bayesian theory only the first order moment of the belief cannot be manipulated through the acquisition or avoidance of information. By contrast, higher order moments (and in particular the skewness of the distribution of beliefs) can be manipulated by rational agents.

This simple observation may have very powerful economic implications. Suppose that individuals who drive more carefully have a lower belief about their own ability. As long as this negative relation between perceived ability and careful driving is linear, then the average attention paid by drivers is given by the *fixed* average belief about ability in the population. It will therefore be non-manipulable, that is, independent of the information acquired by individuals. On the contrary, suppose that there is a threshold in the perceived ability below which high attention is paid and above which low attention is paid. As already noted, the information possessed by each individual affects the skewness of the distribution of beliefs in the population. Therefore, in this second situation, the average attention paid by individuals will crucially depend on the beliefs held by each person. In particular, by strategically manipulating the information conveyed to a population of *rational* individuals, it may be possible to choose between inducing an average attention that is arbitrarily close to the high level and one that is close to the low level.

2.2. Self-Manipulation of Beliefs

It is clear, from the previous example, what the gains are for an individual (self-interested or social planner) by restricting the access to information to other agents and therefore manipulating their beliefs. A more intriguing issue is why any person would refrain from obtaining some information relevant for his own decision-making. From now on, we will call the situation in which a rational agent deliberately avoids free information that would reduce his uncertainty before acting as the 'self-manipulation of beliefs'.[1] There are at least three reasons for such behavior in the literature.

First, the simplest way to make self-manipulation desirable is to explicitly assume that individuals derive utility (directly or indirectly) from their beliefs. This issue has been explored in a number of papers. Caplin and Leahy (1999, 2002, Chapter 4, this volume) introduce anticipatory feelings in a standard expected utility framework, by allowing utility to depend on beliefs about the future states of the world. Information, then, becomes more desirable in some situations than in others, as it affects not only optimal decision-making but also the evolution of beliefs and therefore the anticipatory utility itself.[2] Köszegi (2000) also incorporates beliefs in the utility function by assuming that agents derive utility from having positive views about themselves. In his setting, individuals are likely to avoid information when they are satisfied with their current beliefs (self-image protection motive) and to seek evidence when they are dissatisfied (self-image enhancement motive). Last, Rabin (1995) considers the case of an individual without moral concerns, but with a self-restraining mechanism

[1] Obviously, one reason for avoiding information is that it may be costly. However, the question here is why would an individual be willing to pay not to get information.

[2] See also Gilboa and Schmeidler (1988) and Geanakoplos, Pearce and Stacchetti (1989).

that forbids him to undertake actions whenever he believes that they exert a negative externality on others. In his model, the beliefs indirectly affect utility and therefore refusing information can be optimal. If negative beliefs are avoided, the self-restraining mechanism does not become operational, and the agent can engage in immoral (but optimal from his perspective) actions.

Second, from the economics of information literature, it is by now well-known that obtaining (public or private) information can be hurtful in multi-person situations. For example, Hirshleifer (1971) points out that the revelation of public information may destroy the possibility of a mutually beneficial insurance agreement. Crémer (1995) shows that it can be optimal for an employer to commit not to acquire private information about the circumstances that affect the performance of his subordinate: the ex-post costs of coarser information can be outweighed by the ex-ante benefits due to the subordinate's higher incentives to put more effort into his job. Note that in this model, the uncertainty is about a characteristic of the environment. However, if we postulate that individuals have an imperfect knowledge about some of their own personal traits, then it follows that the interacting agents may also benefit from avoiding private and public information about their own attributes.

Third, avoiding free information has also proved to be optimal in situations in which an individual has conflicting preferences. Intuitively, analyzing an individual with an intrapersonal conflict of preferences is not very different from analyzing a multi-person situation in which the preference of each person conflicts with those of his peers. It is, therefore, quite natural that individuals with internally conflicting goals may decide to manipulate their own beliefs, whenever this option is available. Intrapersonal disputes can be of a different nature. According to Thaler and Shefrin (1981), the agent is, at every period, composed of a planner who is interested in the agent's long-run utility and a doer who is only concerned with short-run payoffs. Since actions are taken by the doer, the planner restricts the set of alternatives in order to mitigate the doer's willingness to satiate immediate gratification at the expense of long-run welfare. Loewenstein (1996) argues that agents are often out of control due to visceral factors such as emotions, moods, and other somatic influences. These unanticipated states also result in discrepancies between the optimal and the realized behavior. In Bodner and Prelec (2002, Chapter 6, this volume), an agent infers his preferences from his own decisions. In other words, by means of his actions, the decision-making body signals to the mind what his true concerns are. Last, there is a growing literature starting with Strotz (1956), which argues that individuals discount short-term events relatively more heavily than long-term events. Formally, according to this theory, the discounting of future payoffs would be best approximated by hyperboles and not by the traditional exponential functions.[3] This type of discounting generates an intrapersonal conflict of preferences, because the optimal

[3] See Ainslie (1992) or Loewenstein and Prelec (1992) for a comprehensive theoretical and empirical discussion of hyperbolic discounting.

dynamic plan of actions of an individual at some date may no longer be optimal when reconsidered some time later. This intrapersonal conflict, in turn, creates severe welfare losses for the individual. As first pointed out by Carrillo and Mariotti (2000), a strategic self-manipulation of beliefs can alleviate this utility loss. In Section 3 we will review the main findings of this paper and the subsequent literature, relate them to the motivating examples presented in Section 1, and explain the implications for economics and psychology. The attention will be centered on the collection/avoidance of information of an individual with hyperbolic discounting. There are two main reasons for restricting the analysis to this case. First, the types of situations we have in mind are closely related to the individual's tendency to excessively satiate immediate gratification at the expense of future welfare, and excessively delay unpleasurable activities that are profitable in the long-run. Second, hyperbolic discounting has received wide support from the empirical and experimental literature (both in economics and in psychology) and, at the same time, it is a well-defined departure from the standard postulates of rationality.[4] However, there are definitely other interesting insights to be developed about the reasons and consequences of self-manipulation of beliefs in the various other settings described above.

3. STRATEGIC IGNORANCE AND SELF-CONTROL

3.1. The Value of Ignorance

We first start by considering a population of time-consistent utility maximizing individuals, who are choosing whether to invest in a risky project (the decision is binary). The investment requires some immediate cost and yields a delayed benefit that is proportional to the agent's talent. If agents know their own capacity perfectly, only those whose talent is above a certain threshold will choose to invest. If agents do not know their capacity but can learn it at no cost, then they will acquire this information and, just like before, act according to their revealed talent.

The problem is less trivial for an agent with time-inconsistent preferences. Given that,under hyperbolic discounting, immediate payoffs are overweighed relative to more distant ones, the agents' willingness to invest now depends on the date of reference. More precisely, the talent threshold above which an agent finds it profitable to invest in the next period (and get the benefits in two periods) is below the threshold above which he finds it profitable to invest in the current period (and get the benefits next period). In other words, for an intermediate level of talent, the agent would like to invest in a future date, but when the time of exerting effort arrives, the cost is overweighed and the agent reconsiders his decision. This is the standard procrastination problem first analyzed by

[4] It is partly for this reason that hyperbolic discounting has become, among economists, one of the least controversial departures. Naturally, this is not to say that all researchers accept its validity. See for instance Mulligan (1996) and Rubinstein (2000).

Akerlof (1991) and further developed by O'Donoghue and Rabin (1999) and others. Naturally, rational agents perfectly anticipate this change in their preferences. However, in the absence of exogenous commitment devices, nothing can be done to impose the previously optimal choice.

The situation becomes most interesting when we consider time-inconsistent preferences, imperfect self-knowledge, and the possibility of learning. As shown by Carrillo and Mariotti (2000), the combination of these three ingredients together with the binary nature of the investment decision can help in explaining some of the apparently irrational behaviors mentioned in Section 1. Consider an individual with an imperfect knowledge about his talent, who can acquire, at no cost, all the information about his ability on the date before making his investment decision. Recall that this agent faces a very specific conflict of preferences: some investments considered profitable the date before exerting effort are not undertaken when this date arrives. Because of this conflict, the agent may be interested in refusing information about his own talent. For instance, suppose that the expected talent is sufficiently high so that, in case of not getting information in the current period, the agent (who then maximizes expected utility) strictly prefers to invest in the next date. By learning his true talent, the agent may realize that his true ability takes an intermediate value, in which case he would like to invest in the next period, but when this date arrives he prefers not to do it anymore. To avoid such inefficient procrastination, the agent can rationally decide not to acquire the information. In other words, time-inconsistent preferences generate an endogenous cost of learning equal to the cost of inefficient procrastination. Naturally, ignorance has also some costs: if the agent's true talent is indeed sufficiently low, then by ignoring it and investing the agent is adopting a truly suboptimal behavior. Overall, under time-inconsistent preferences ignorance can act as a commitment device against the unavoidable change of preferences. The optimal learning decision will trade-off the benefits of avoiding inefficient procrastination and the usual costs of ignorance for decision-making.

Note that a necessary condition for the agent to remain ignorant is that, in case of not acquiring information, he must unambiguously prefer to invest in the next period. This is the key result of our analysis. Again, the reason is that the goal of ignorance is to avoid inefficient procrastination and induce future self-investment. Obviously, ignorance can then be only an effective mechanism if the agent is willing to invest given his current beliefs.

Now, consider an entire population of individuals, each of them with an imperfect knowledge about their ability. According to our previous analysis, two different behaviors can be observed depending on the shape of each individual's distribution of beliefs. Some agents will learn their ability and invest only if they are sufficiently talented. Others will decide to remain ignorant and invest for sure. Overall, after making their learning decision, the skewness in the posterior distribution of beliefs in the population will be such that the vast majority of agents will consider themselves talented enough to invest. Interestingly, this biased

aggregate behavior of the population, driven by a desire to avoid procrastination, occurs even though each agent is rational and has an unbiased belief about his own talent.[5]

3.2. Predicting Aggregate Behavior

We have seen that time-inconsistent preferences together with imperfect self-knowledge and the possibility of learning, predicts an aggregate tendency in the population to keep an optimistic attitude towards investment. Interestingly, the same mechanism predicts that the population will exhibit an aggregate tendency to keep excessively negative attitudes towards pleasurable but long-run damaging activities.

The crucial difference between a decision to invest in a risky business and a decision to consume something good but with negative side-effects is that, in the latter the benefits of the activity come earlier than the costs while in the former it is the opposite. This may sound unimportant, however it has a crucial impact on the incentives of agents to gather information. In order to appreciate the difference, consider the decision of a time-inconsistent individual to smoke. Due to his intrapersonal conflict, rather than a consideration of the health damages provoked by cigarettes, the optimal consumption pattern may be to smoke in the current period and then stop forever. Naturally, and again without commitment devices, the agent will, in this case, end up smoking at every date; that is, the conflict of preferences induces a future excessive consumption from the current perspective. Now, suppose the agent has only an imperfect knowledge about the effects of tobacco on his own health. The only reason not to acquire information is to avoid an inefficient future consumption. Hence, a necessary condition for the agent's willingness not to acquire information is that the current beliefs must be sufficiently negative. In particular, given his current knowledge, the agent must prefer not to smoke not only in the future, but even in the present. As usual, these benefits of ignorance have to be compared to the costs of making choices with restricted information. Overall, in the smoking case, the population will consist of a fraction of agents who are perfectly informed about the side effects of smoking and consume accordingly, and a fraction of agents who retain their pessimistic beliefs and abstain.[6]

As a general conclusion, the vast majority of the population will believe that their health is too fragile to resist the effects of tobacco, and their talent high

[5] As discussed in Section 2.1, if the optimal level of investment were positive for all agents but proportional to their talent, then there would be no aggregate biased behavior. Yet, imposing a binary choice is not necessary either: an aggregate bias would occur for any non-linear relation between the talent parameter and the level of investment.

[6] The analysis predicts that only non-smokers may refrain from learning the true dangers of tobacco. This is because, in our example, individuals do not decide the amount of cigarettes they consume. If we relax the binary decision assumption, then some smokers may also decide to remain uninformed.

enough to succeed in risky businesses. Hence, time-inconsistency together with the endogenous decision to acquire information can explain both a feeling of optimism (in entrepreneurship) and a feeling of pessimism (for pleasurable but damaging activities), without any assumption of irrationality. The key, to understand which situation applies, is to determine whether the costs come, on an average, earlier or later than the benefits in the activity considered. These testable predictions can be pushed further: as the gap between the costs and the benefits of the activity increases, the strategic ignorance becomes a more useful self-disciplining device, and therefore it is more likely to observe an aggregate bias in the behavior of the population. Obviously, we do not believe that this approach can explain all the biases in judgment. For example, agents have a tendency to keep positive views about issues related to their own intelligence or leadership capacity. These are cases in which there is no clear delay between the costs and benefits of the activities involved, and therefore our analysis cannot be applied. What this theory offers, instead, is one possible explanation to a specific class of problems.

3.3. Welfare Implications and Social Interactions

In the previous analysis, each individual was studied in isolation. However, in many situations, the learning choices and the actions undertaken by each person affect the rest of the population. Therefore, if we want to understand the aggregate economic effects of strategic ignorance and its welfare implications, it is essential to take into account these endogenous interactions.

Brocas and Carrillo (1999) provide a first step towards this generalization. They consider the same investment decision as in the example above, except that the investment now requires some capital and individuals are cash constrained. This slight enrichment of the economic setting is sufficient to ensure that, the behavior resulting from the agents' intrapersonal conflict affects the welfare of all the individuals in the population. In other words, the agents are linked to each other by the inconsistency of their preferences.

The argument of the paper is as follows. If individuals are cash constrained, they need to resort to banks in order to finance their investment. In the standard (time-consistent) situation, agents learn their own ability to succeed, and those who are sufficiently talented apply for a loan. The banks are competitive and do not know the ability of the applicants, but they infer from their willingness to invest that it must be above a certain threshold. The interest rate of the money is then endogenously determined and inversely proportional to the expected talent of applicants (high talent implies a low risk of failure, therefore a low risk for banks of not being repaid, which in turn makes a low interest rate possible). By contrast, as already noted, the time-inconsistent agents have a tendency to procrastinate. To avoid this inefficient behavior they remain ignorant and retain their initial beliefs, as long as these are sufficiently positive. Naturally, the banks correctly anticipate that the agents are now excessively confident about their

chances of success, so they increase the interest rate in order to counter this optimism. Overall, the endogenous decision of each agent to refuse information and apply for a loan is individually optimal, but it exerts a negative externality on the rest of the population by affecting the interest rates. The paper shows that, in some situations, the ex-ante welfare of all the individuals in the population is increased if the government forces them to learn their talent before deciding whether to invest. This measure eliminates the prevailing overconfidence of potential investors, and therefore induces banks to lower the interest rates. Naturally, the intervention is so efficient, that the costs induced by learning are small relative to the gains of a low interest.

In our view, this investment story is a good example of an analysis that can be of interest for both psychologists and economists. First, it suggests that it is possible to expand our understanding of human behavior and, in particular, of some apparently irrational conducts such as the willingness to keep optimistic prospects. Second, it indicates that understanding the reasons for such human behavior opens the door to the design of welfare improving policy measures.

3.4. Available Information and Unavoidable Information

The literature on self-manipulation of beliefs provides an interesting implication about the value of information. There exist two different practical measures to encourage individuals to acquire information. The first one is to increase the sources of information at their disposal, that is to make more information available. The second one is to give incentives to incorporate the existing information, that is to make the available information unavoidable. When the public interest is congruent with the private interest of the agents, then both options are socially desirable and it is unclear which one has the greatest merits. However, in the case of a conflict between private and public interests, the outcomes can be radically different: the first alternative may be negative for social welfare and the second one positive.

Rabin's (1995) paper was the first to discuss these opposing effects. An agent who acts subject to a moral rule will try to avoid the beliefs where this moral constraint becomes binding. Having a better access to information then, increases his capacity to manipulate his own knowledge and, in particular, to avoid negative beliefs about the morality of his preferred action. Therefore, an agent with an unlimited access to information is more likely to exhibit immoral conduct than an agent with restricted access. On the contrary, forcing an agent to acquire information reduces his self-manipulation capacity, which is likely to improve his moral conduct.

These same differences between available and unavoidable information operate in the case of time-inconsistent individuals. Allowing better access to information, increases the capacity of potential investors to manipulate their beliefs and retain their optimistic prospects. On the contrary, forcing them to acquire some news destroys the possibility of using strategic ignorance in their

own private interest. The main difference with Rabin's model is that this second alternative may not only be socially better, but even privately better as long as it is simultaneously imposed on all individuals in the society. As a policy implication, it means that governments should not facilitate the access to information to potential investors, but, rather, encourage the use of information (e.g., by requiring an accurate market study with every loan application).

3.5. Other Applications

The basic self-control and information gathering model developed above, has been extended in a number of directions. Carrillo (1998) analyzes the decision of an individual to undertake activities that are pleasant in the short-run, but with a long-run negative effect on welfare. Examples include not only the ingestion of addictive substances, such as tobacco (as in the example discussed in Section 3.2), alcohol, or marijuana, but also the engagement in non-addictive activities like gambling, eating fat-food, or having extramarital relations. The basic postulate is that the pleasure derived from these activities is individual-specific and therefore it can only be evaluated after repeated exposure. In other words, there is learning through consumption. With these premises, the paper studies why abstinence and excessive exposure are both more frequently observed than moderate behavior, and reaches the following conclusions. As long as the agent is engaged in the activity, he learns about its net payoff. For example, by repeatedly smoking the agent updates the instantaneous pleasure and the negative effects on health. A time-inconsistent individual may then fall into a state of beliefs in which he wishes to have a high current consumption and a moderate future consumption (one pack of cigarettes today and a couple of cigarettes from tomorrow onwards). However, given the impossibility of commiting to a future behavior, he ends up consistently overconsuming. When beliefs are such that the expected intertemporal payoff under continual excesses (one pack a day) is smaller than the under-sustained abstention, the agent prefers not to consume at all. Abstention is only a second-best solution. By construction, moderate consumption at every period would be preferable, but this is not feasible due to hyperbolic discounting. Moreover, the strategy of abstaining is the agent's way of not learning. Therefore, it is his only possible commitment strategy to avoid sinful temptations in the future. To sum up, this theory explains not only convergence to but also persistence of abstention as the agent's optimal self-commitment device. It is important to realize that the 'complete abstention' strategy is not imposed, on the grounds that it acts better than any other personal rule as a focal point (as sometimes suggested in the literature). Rather, it endogenously becomes a basin of attraction due to its learning properties.

Benabou and Tirole (2000*a*) introduce imperfect memory and the possibility of increasing or decreasing the probability of remembering past events in the framework of Carrillo and Mariotti (2000). The paper shows that multiple equilibria might concur in the intrapersonal game, when memory management is costless.

The idea is as follows. An individual does not want to remember information only if the news is likely to induce a future behavior suboptimal from the current perspective (i.e., if there is a high probability that the news will be 'bad'). However, the individual cannot fool himself, so this mechanism of strategically not remembering past signals is perfectly anticipated. Therefore, as the degree of censoring is increased, a lack of information is more likely to be interpreted in the future as 'bad news', which itself increases the incentives to censor information in a first place. This self-fulfilling mechanism immediately results in the coexistence of equilibria, with no repression, full repression, and partial repression of information.

The issue of willpower has also been studied in a similar context by Benabou and Tirole (2000*b*). The paper assumes that agents have an imperfect knowledge about their own capacity to resist impulses and short-run temptations. According to their current information, individuals have to decide whether to engage in a 'willpower-dependent' activity or not and, if they do, whether to resist or succumb to temptation. One of the main results is that in order to preserve their self-reputation, agents may end up adopting a compulsive, self-defeating behavior.[7]

4. IMPULSIVE BEHAVIOR AND SELF-CONTROL

We now analyze another puzzling pattern of behavior, namely the agent's willingness to engage in impulsive consumption. In Section 2.2, we argued that some authors (e.g., Loewenstein, 1996) have identified emotions and drives as the causes for some irrational, impulsive conducts. Here, by contrast, we provide a rational view for this behavior in a specific class of situations. We determine the type of circumstances in which impulsive actions are likely to occur and the possible remedies.

4.1. Exotic Holidays: An Example of Impulsive Behavior

Vacation time has arrived and I, a cash constrained and hyperbolic agent, have to decide where to go. This year and the next one are my two last opportunities to spend an entire month of holidays. I am, therefore, quite tempted to become indebted and travel to a distant exotic island, even though I do not know exactly how nice this trip will be. After considering the delayed costs and the current (thus overweighed) expected benefits, it turns out that the net present value (NPV) of this alternative is slightly negative. What should I do? There are three different variants of this situation, depending on the amount of information that I will receive about how enjoyable these vacations can be, if this year I refrain from making the trip. First, suppose that I will not get any extra news. Then, given its negative NPV, I go neither this year nor the next one. Second, suppose that a friend with very similar tastes to mine is traveling to that same place in half a year.

[7] We refer to Benabou and Tirole (2002, Chapter 8, this volume) for an exhaustive analysis of memory management and willpower.

My optimal decision is to wait for his news, learn with accuracy how pleasurable the trip is, and then decide whether to go next year. Last and most interesting, consider the case in which I will get some meager and noisy information. If the news is slightly negative, it is optimal not to go at all. By contrast, if the news is slightly positive and I do not go this year, then I will choose to go the next year. Naturally, this decision has expected benefits from next year's perspective (this is why I follow that option). However, from the current perspective, the benefits of a future holiday are not overweighed, and therefore the overall NPV of this future decision may be strongly negative. Anticipating such an inefficient behavior and in the absence of other commitment devices, my optimal decision may be to go this year, even if I know that the costs offset the expected benefits. Overall, the argument in favor of enjoying the current holidays is something like 'since I will go next year if I am given the smallest reason to go, I might as well go now'.

4.2. The General Case

The previous example suggests that impulsive behavior may not be the result of an irrational and unanticipated reaction to an urge or any other external cue. Instead, a fully rational individual with a self-control problem may engage in apparently suboptimal consumption decisions, only as a commitment against future actions that are even more inefficient.

This idea was first suggested by O'Donoghue and Rabin (1999) for consumption decisions in deterministic environments. Then, Brocas and Carrillo (1998) studied the role played by information in explaining this impulsive behavior. Their theory rests on three basic ingredients. First, as usual, a self-control problem. Second, the possibility of undertaking an irreversible (or partially irreversible) consumption decision that yields a short-run net benefit and a long-run net cost. Third, imperfect information about the net payoff of consuming and a per-period exogenous revelation of information, as long as the agent has not made his irreversible choice. Credit purchase is probably the clearest example of the situations we have in mind. However, the setting is sufficiently general to be adapted to other problems such as the protection of the environment and the financing of public projects by official bodies.

In the absence of an intrapersonal conflict, the situation is characterized by a fundamental tradeoff: delaying consumption is costly because future payoffs are discounted, but it allows the acquisition of valuable information for future decision-making.[8] However, as we know, information is not always valuable if

[8] There is an extensive literature on this subject. The basic result is the existence of a cutoff value at each period, above which the agent consumes and below which he waits until the next date. The difference between the payoff at this cutoff and the expected payoff of never consuming is called the option value of waiting (OVW), and it is always non-negative. Moreover, when the consumption horizon is finite, then the OVW strictly decreases as the number of periods in which it can be exerted decreases and it becomes equal to zero in the last period. We refer to Dixit and Pindyck (1994) and the references therein for a review.

agents face a self-control problem. Therefore, the nature of the tradeoff is modified under time-inconsistency. The paper shows that the amount of information transmitted between periods is the key feature that will determine whether the agent will behave in an impulsive manner or not. The main idea is as follows. Under hyperbolic discounting, current benefits are overweighed relative to future costs. Due to the irreversibility of consumption, the optimal current decision of an individual depends on his future behavior in the case of postponing consumption (which determines his outside option). However, when the agent evaluates this outside option, he does not have to determine the effect on the current welfare of future choices optimal from the current perspective, but rather the effect on the current welfare of future choices optimal from the future perspective. Naturally, this future behavior will be a function of the signals received by the agent, which suggests that the amount of information transmitted between two periods must be a crucial variable. The most interesting situation arises when the flow of information is 'small', that is, when there are little variations from period to period in the expected profitability of consumption. Suppose that the agent's current beliefs are such that consuming has a current, slightly negative NPV. If he refrains from consuming and the information revealed is 'moderately positive', he will plunge into a state of beliefs where future consumption is desirable from the future perspective but highly undesirable from the current one. In order to avoid this situation, the individual prefers to rush and consume in the present period also with a negative (but at least close to zero) NPV. Summing up, consumption takes place not because of its intrinsic value but only as a commitment device against future choices.

The inefficiency of this two-period analysis is magnified in a multi-period situation. Indeed, a future consumption with a negative NPV from a future perspective is even more damaging to welfare from the current viewpoint. Hence, the agent is all the more willing to consume with the negative payoffs that he anticipates a future inefficient choice. Last, it is important to realize that the value of information is not monotonic in the amount of information transmitted. If the inter-period flow of news is sufficiently important, then the intrapersonal conflict is not a major problem: an agent who waits for information will learn, with a high probability, either that consumption is worthless or desirable both from a past and a current viewpoint. In other words, information is on an average valuable because it is very unlikely that it directs the agent to the state of beliefs where present and future preferences are in conflict. On the other extreme, if no information ever flows in, the agent consumes either in the first period or never at all, but again there is no inefficiency.

4.3. Welfare Implications and Social Interactions

The main lesson of the previous analysis is that self-control problems may induce agents to adopt inefficient consumption decisions. The reason is the existence of a range of payoffs, such that consumption is profitable from the current perspective

but highly detrimental from a past perspective. Naturally, this undesirable behavior is more likely to occur the greater the range of payoffs in which the intrapersonal conflict occurs and also, for a given range, the greater the chances of believing that payoffs are in this region. It is then possible to analyze comparative statics about the kind of situations in which agents will act impulsively. First and trivially, impulsive conduct will occur more often the stronger the agent's intrapersonal conflict and the bigger the delay between the costs and the benefits of the consumption goods. Second, the chances of eventually believing that payoffs are in the inefficient region are highest when the information transmitted between periods is small, and when the consumption decision can be delayed for a long period of time. Last, in decisions involving high stakes (i.e., both high current benefits and high delayed costs) the salience of current payoffs are more likely to trigger the self-defeating behavior.

Once the reasons for the impulsive behavior have been identified, it is possible to offer solutions in order to avoid or at least mitigate the inefficiencies. The most radical measure is to stay away from information. This is quite difficult to achieve: some news is often unavoidable and, once noticed, the information processing cannot be stopped. Moreover, information is sometimes extremely useful. Otherwise, a typical solution to self-control problems is to commit with third parties on future decisions (e.g., to freeze the credit card). One possible type of commitment is to make promises that are costly to break (e.g., due to reputation concerns as in Carrillo and Dewatripont (2000)). However, such commitments need to be renegotiation-proof and they also entail some implicit costs.

5. CONCLUDING REMARKS

The purpose of this essay has been to show that by enlarging the standard economic setting, it is possible to account for a number of puzzling patterns of behavior. The main new ingredients incorporated in the analysis are quite natural: hyperbolic discounting of future payoffs, imperfect knowledge about some of the agent's own attributes, and various possibilities of improving this self-knowledge. The combination of these factors is sufficient to explain some apparent biases in judgments and behaviors, like the tendency to stay optimistic or pessimistic prospects in some activities, to stick to restrictive personal rules of behavior, and to act impulsively. Naturally, some other behaviors could also be analyzed under this light.

Before concluding, we would like to insist on two points. First, as discussed all along in the paper, the explanations offered to these pervasive human behaviors are valid only for a specific class of situations. That is, our arguments make sense only if some precise conditions are met (activities with a natural delay between costs and benefits, imperfect information, etc). If individuals exhibit a similar behavior in settings where these conditions are not satisfied, then we need another (may be complementary) explanation. Second, rationalizing human behavior is not a goal in itself. Our objective is to improve the understanding of human

conduct in order to offer solutions for inefficient actions. Building a model as standard as possible with a minimum set of assumptions and which captures the main motivations of the individual is, in our view, an important step to reach this goal. Naturally, this is not always feasible, so one should not be opposed to incorporating some departures from rationality whenever it is appropriate.

REFERENCES

Ainslie, G. (1992). *Picoeconomics*. Cambridge: Cambridge University Press.

Akerlof, G.A. (1991). Procrastination and obedience. *American Economic Review*, *81*, 1–19.

Bell, D. (1976). *The coming of post-industrial society: A venture in social forecasting*. New York: Basic Books.

Benabou, R. and Tirole, J. (2000*a*). *Self-confidence: Intrapersonal strategies*, mimeo, Princeton.

——— and——— (2000*b*). *Personal rules*, mimeo, Princeton.

——— and ——— (2001). Self-knowledge and self-regulation and economic approach, this volume.

Bodner R. and Prelec, D. (2002). In I. Brocas and J.D. Carrillo (Eds.), *The psychology of economic decisions*. Oxford: Oxford University Press.

Brocas, I. and Carrillo, J.D. (1998). *A theory of haste with applications to impulse buying and destruction of the environment*, CEPR D.P. 2027, London.

——— and Carrillo, J.D. (1999). *Entrepreneurial optimism and excessive investment*, mimeo, ECARES, Brussels.

Caplin, A. and Leahy, J. (1999). Psychological expected utility theory and anticipatory feelings. *Quarterly Journal of Economics*, in press.

——— and ——— (2002). *Behavioral policy*. In I. Brocas and J.D. Carrillo (Eds.), *The psychology of economic decisions*. Oxford: Oxford University Press.

Carrillo, J.D. (1998). *Self-control, moderate consumption, and craving*, CEPR D.P. 2017, London.

——— and Mariotti, T. (2000). Strategic ignorance as a self-disciplining device. *Review of Economic Studies*, *67*, 529–44.

——— and Dewatripont, M. (2000). *Promises, promises . . .*, mimeo, ECARES, Brussels.

Crémer, J. (1995). Arm's length relationships. *Quarterly Journal of Economics*, *110*, 275–95.

Dixit, A.K. and Pindyck, R.S. (1994). *Investment Under Uncertainty*. New Jersey: Princeton University Press.

Geanakoplos, J., Pearce, D., and Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior*, *1*, 60–79.

Gilboa, I. and Schmeidler, D. (1988). Information dependent games: Can common sense be common knowledge? *Economics Letters*, *27*(3), 215–21.

Hirshleifer, J. (1971). The private and social value of information and the reward from incentive activity. *American Economic Review*, *61*, 561–74.

Köszegi, B. (2000). *Self-image and information acquisition*, mimeo, MIT.

Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes*, *65*, 272–92.

——— and Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *Quarterly Journal of Economics*, *107*, 573–98.

Mulligan, C. (1996). *A logical economist's argument against hyperbolic discounting*, mimeo, University of Chicago.

O'Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review 89*, 103–24.

Rabin, M. (1995). *Moral preferences, moral constraints, and self-serving biases*, mimeo, UC Berkeley.

Rubinstein, A. (2000). *Is it Economics and Psychology?*: *The case of hyperbolic discounting*, mimeo, Tel Aviv and Princeton.

Strotz, R.H. (1956). Myopia and inconsistency in dynamic utility maximisation. *Review of Economic Studies*, *23*, 166–80.

Thaler, R.H. and Shefrin, H.M. (1981). An economic theory of self-control. *Journal of Political Economy*, *89*, 392–406.