Comments on "Behavioral" Decision Theory*

Ran Spiegler†

January 2011

1 Introduction

The developments in decision theory summarized in the pair of excellent surveys extend the boundaries of orthodox rational choice theory, incorporating types of behavior formerly considered less-than-fully rational and therefore outside the scope of the theory. More and more psychological phenomena are nowadays modelled as being consistent with utility maximization over *some* domain of objects of choice, and this means that they are treated as "rational" at some level. The method - which Lipman and Pesendorfer dub as "Krepsian", after David Kreps's pioneering works in the late 1970s - is to start with an intuition (possibly backed by systematic observations) suggesting how choice behavior in general can be related to a certain "unconventional" psychological "force", and to look for a specific domain of choice objects such that the force could be effectively elicited from revealed (complete and transitive) preferences over this domain.

For example, if the decision theorist is interested in behavioral manifestations of anxiety, then studying preferences over the domain of temporal lotteries can be illuminating because the decision maker's preference for early resolution of uncertainty may reveal the anxiety that this uncertainty generates in him (Kreps and Porteus (1978)). Similarly, the domain of choice sets ("menus") is appropriate for identifying the force of "desire for flexibility" (Kreps (1979)). Finally, observe that the focus on Anscombe-Aumann acts (as opposed to Savage acts) in the ambiguity literature is entirely consistent with this methodology. Anscombe-Aumann acts are appropriate

*This is a discussion of the Decision Theory surveys in the Econometric Society 2010 World Congress. I thank Ayala Arad, Eddie Dekel, Kfir Eliaz, Tzachi Gilboa, Barton Lipman, Wolfgang Pesendorfer and Ariel Rubinstein for extremely helpful discussions.

†Tel Aviv University and University College London. E-mail: r.spiegler@ucl.ac.uk. URL: http://www.tau.ac.il/~rani.

choice objects, not because of the technical simplifications they allow, but because the distinction they make between "roulette" and "horse" lotteries enables us to elicit the decision maker's attitudes to risk in distinction to his attitudes to uncertainty.

Thus, "finding the right domain" is crucial for the decision-theoretic exercise in this tradition. Even when we have a strong intuition that observed choices in a particular domain are heavily influenced by a particular psychological force, the domain may be ill-suited for the "Krepsian" exercise, for two, diametrically opposed reasons. First, observed choices in the natural domain may be inconsistent with maximization of a utility function over that domain. Second, they may be consistent with an entirely standard utility function, thus failing to elicit the distinct psychological force we are interested.

For instance, the model of self-control preferences due to Gul and Pesendorfer (2001) has the property that observed choices from menus rationally trade-off the "commitment" utility function u and the "temptation" utility function v, such that the decision maker chooses from menus as if he consistently maximizes the utility function $u + v$. Therefore, while his behavior in this domain reflects temptations, it does so in a way that cannot be distinguished from standard rational choice. In contrast, choices from menus in the extended model of temptation-driven preferences due to Dekel, Lipman and Rustichini (2009) typically violate the weak axiom of revealed preferences. Thus, in both cases the basic domain of choice alternatives is ill-suited for the decision-theoretic exercise of mapping observed choices into a utility maximization model. It is only when we consider choices in a different domain - specifically, choices *between* menus - that we can elicit temptation as a behaviorally distinct phenomenon without abandoning utility maximization as a modeling device.

My objective in the present comments is to draw attention to some limitations of this methodology of searching for a domain of choice objects such that revealed preferences over this domain would satisfy completeness and transitivity, and elicit a psychological phenomenon of interest. I will show that these desiderata sometimes clash, and that the way decision theorists tend to respond to this clash can cause important aspects of the phenomenon to go out of focus.

Before I begin, I wish to make two remarks about the scope of these comments. First, I refrain from discussing the specific classes of models that receive the closest attention in the two surveys. Rather, I address other models that follow the same methodology. Second, I adopt a descriptive interpretation of decision models, as opposed to a normative one. The expansion of decision theory away from a narrow definition of rationality has made it harder for decision theorists to claim both types of

interpretation at the same time, and as a result theorists have often faced the need to choose one over the other. The two surveys reflect this dilemma: Gilboa and Marinacci adopt an explicitly normative interpretation of decision models, whereas Lipman and Pesendorfer implicitly take the descriptive approach. The present discussion shares Lipman and Pesendorfer's "behavioral" orientation.

2 Preferences over Menus

The model of self-control preferences over menus due to Gul and Pesendorfer (2001) was a highly successful application of the "Krepsian" methodology. Success breeds emulation. And so in the last decade, decision theorists have been attracted to the basic framework of preferences over menus, and employed it to illuminate a variety of psychological phenomena (belief distortions, contemplation costs, ex-post regret, internalization of a social pressure to act altruistically).

This proliferation of applications highlighted an important feature of this approach. When a decision theorist analyzes preferences over menus instead of (or in conjunction with) choices from menus, he shifts his analytic focus, at least in part, *from the choice behavior itself to its ex-ante anticipation.* For example, Sarver (2008) does not ask directly how choice from a menu reflects regret, but how the anticipation of ex-post regret governs choices between menus. Similarly, Dillenberger and Sadowski (2010) do not ask directly how altruistic and selfish motives determine an agent's allocation decision, but how his anticipation of a future pressure to act altruistically shapes his preferences over sets of feasible allocations.

I argue that this shift of focus often creates a number of interpretational difficulties. In what follows, I illustrate these difficulties with a mix of phenomena that have been studied in the literature and some that have not (yet), in which case my discussion is inevitably speculative. I also ignore the possibility that switching to some other "Krepsian" domain would solve the problems that I highlight. These aspects of the presentation serve expositional purposes; I hope that they will not cause the reader to dismiss the examples as anecdotal, as I believe that they are characteristic of the methodology.

2.1 Infinite-Regress Problems

According to a common interpretation of preference-over-menus models, preferences over observed choices between menus reflect the decision maker's anticipation of a cer-

tain behavioral "force" that affects choices from menus, but the preferences over menus themselves are not affected by this force. When this interpretation is not convincing - i.e., when choices between menus are plagued by the same effects as choices between ultimate choice objects - the impulse is to shift our gaze up another level, to choices between menus of menus, and so forth. But this leads to an obvious infinite regress problem, because the difficulty is not likely to disappear in this process of domain enrichment.

The most well-known example of such an infinite-regress problem is computational complexity. For instance, revealed preferences over choice objects can be incomplete because computational complexity constrains the decision maker's ability to compare and rank all pairs of choice alternatives. Menus of choice alternatives can serve as a natural alternative domain. For example, a decision maker who anticipates the complexity of ranking alternatives issue may prefer smaller menus. The infinite-regress critique of this approach is that the problem of ranking menus is even more complex than the task of ranking the ultimate choice objects. Therefore, assuming complete and transitive preferences over menus requires special pleading (for instance, that one hyper-rational agent determines the choice set that another, boundedly rational agent will subsequently face).

Is there a similar infinite-regress problem in the two-period model of self-control preferences (Gul and Pesendorfer (2001))? The common interpretation of this model is that self-control costs are incurred only at the second-period stage of choosing from menus, but are not present during the first stage of choosing between menus. Unlike the case of computational complexity, here it makes perfect sense to assume that people are tempted by choice alternatives only when they are close and tangible (Noor (2009) makes a similar argument). When I am offered a new credit card, I feel much less tempted by the prospect of future consumption than when I stroll in a shopping mall armed with this credit card and observe a fancy, unaffordable gadget. This dual attitude to temptations lies at the heart of the phenomenon.

Ex-post regret

Sarver (2008) studies preferences over menus that reflect the anticipation of ex-post regret. The key axiom he imposes implies that when the decision commits to a singleton menu in the first stage, he does not anticipate any ex-post regret arising from this decision. For example, consider the following pair of scenarios. In the first scenario, I go for dinner at a restaurant that serves both fish and steak, and I choose fish over steak. In the second scenario, I make a reservation for a future dinner at a restaurant that serves only fish, choosing it over another restaurant that serves only meat. Sarver's

model assumes that my decision to eat fish generates a sense of ex-post regret in the first scenario, but not in the second.

I find this duality harder to defend than in the case of self-control. One tends to regret things for which one feels *responsible*. If the decision maker can trace a consequence to his own past decision, it should not matter much how far back that decision goes. It should not matter for the sense of ex-post regret whether the decision maker committed to a singleton menu $\{a\}$ in period 1, or chose a larger set A in period 1 and proceeded to select $a \in A$ in period 2. As I eat my fish, why should my sense of ex-post regret ("I could have eaten steak") differ in the above pair of scenarios?[1]

Sarver acknowledges this problem (see Section 4.3 in his paper), but argues that we can separate regret arising from decisions made at different stages, and that he only axiomatizes the preferences that reflect regret associated with choices *from* menus. This is an interesting argument, but I am not convinced that feelings of regret are separable in this way. Sarver basically focuses on regret feelings that are triggered by *actually observing* the counterfactual outcomes ("I could have eaten this juicy steak that I now see, but instead I chose fish"). In other words, he models regret as if it were a temptation-driven emotion. This definitely captures an aspect of regret, and it is understandable given that Sarver's starting point must have been the Gul-Pesendorfer framework. However, by doing so, he misses the element of responsibility that I find essential to the experience of regret.

This argument suggests that the preferences-over-menus approach to ex-post regret suffers from an infinite-regress problem. Our ability to use observed preferences over menus to elicit ex-post regret relied on the "identifying assumption" that these preferences do not generate ex-post regret themselves (or that the regret they generate is "separable"). However, if we reject this assumption, this means that any effect of ex-post regret that can be elicited from observed choices *between* menus should also be possible to elicit by observing choices *from* menus. Thus, if choices from menus are unable to reveal ex-post regret, then so are choices between menus (or between menus of menus, and so forth).

Suppose we accept the idea that decision makers regret an outcome if they feel *responsible* for it, as well as the idea that responsibility arises from *any* action in the course of a dynamic decision process that rules out some choice alternatives - how *can* we elicit this psychological force from observed behavior? The answer, in my opinion,

[1] In reality, decision makers' sense that an action they took in the ancient past is responsible for the final outcome is often weakened by intervening "moves of nature" that lie beyond their control. It may be interested to introduce such moves of nature into Sarver's model.

has to do with preference incompleteness. If we see that the decision maker refuses to choose between alternatives a and b, but he is happy to let another agent choose for him, this suggests that the decision maker is partly motivated by regret avoidance.

Incompleteness of revealed preferences is essential for this argument. We cannot reformulate this behavior in terms of complete preferences over some other domain (such as menus), because whenever the decision maker is asked to provide a clear ranking between an action that eventually induces a and an action that eventually induces b, he will refuse to state this ranking. Thus, the attempt to model regret-driven behavior with a utility maximization model over any domain will necessarily fail to capture the crucial element of responsibility that lies behind feelings of ex-post regret.

Indecisiveness

The preference incompleteness generated by ex-post regret is an example of the more general phenomenon of indecision. What are the behavioral manifestations of this phenomenon? When decision makers face decisions that involve intense value judgments or deep underlying uncertainty, they often defer choices, cling to default options, change their minds constantly ("flip-flopping"), seek additional information before making a decision (even when the information is irrelevant in the sense that it never changes the ultimate decision), and look for new options even if the search cost seems prohibitively large. Moreover, indecisiveness seems to be an aspect the decision maker's personality: some people tend to be more indecisive than others.

If one wants to study these choice effects, the natural domain of choice objects seems to consist of terminal histories in extensive-form decision problems. However, observed choices in this domain are typically inconsistent with utility maximization over this domain. Therefore, one might think it possible to capture the effects in question by looking at ex-ante preferences over dynamic choice problems. The idea is that the way the decision maker ranks choice problems with and without deadlines, defaults or reversible actions may reveal the forces of indecision that cause him to display the effects listed above.

(In some sense, Kreps' original model of preferences for flexibility is such a model. However, the language of preferences over menus is insufficiently rich for describing deadline effects, for instance. Thus, one would have to extend the model from preferences over menus to preferences over extensive-form decision problems. Also, since Kreps rationalizes the taste for flexibility by expected-utility maximization with a subjective state space, he rules out the seeking-irrelevant-information effect: if the decision maker found the same action optimal in all subjective states, he would exhibit no taste

for flexibility.)

To see why this approach suffers from the infinite-regress problem, consider preferences over choice problems with and without deadlines. Why would a decision maker who displays indecision *within* a choice problem be able to decide ex ante whether he wants to have a deadline or not? If indecisiveness is a personality trait, it should be more-or-less fixed throughout all stages of a dynamic decision problem. Again, it is hard to think of a concrete choice setting that would give rise to preferences that anticipate the behavioral effect without being plagued by it themselves.

2.2 Anticipation of Choice Effects can Annul Them

Ex-ante anticipation of behavioral effects having to do with bounded rationality can eliminate these effects altogether, thus rendering preferences over menus incapable of eliciting them. For instance, consider the case of framing effects. A framing effect is a type of preference inconsistency, where an "irrelevant" change in the formulation of the choice problem affects choice. For example, if the alternatives x (y) and x' (y') are logically equivalent and differ only in the units of measurement employed in their description, then a rational decision maker should prefer x to y if and only if he prefers x' to y'. When the decision maker violates this consistency requirement, we say that his behavior is sensitive to framing.[2]

In reality, framing effects seem to work when the decision maker is not aware that the effect exists. If two ways of framing an alternative are simultaneously presented to the decision maker, he is likely to recognize the conceit and the framing effect is likely to disappear as a result. This is why psychologists use inter-subject experimental methodology to elicit framing effects. What are the implications of this observation for preferences over menus of framed alternatives?

First, we need to ask whether the decision maker is frame-sensitive at the stage of choosing between menus. Assuming that he is immune to framing effects raises the infinite-regress criticism we have already discussed. Putting this criticism aside, suppose that the decision maker faces a choice between the menus $\{x, y\}$ and $\{x', y'\}$. Since both ways of framing each alternative are presented to the decision maker at that choice stage, he can recognize the effect and become immune to it throughout the decision process. In particular, he will be indifferent between the two menus, and so choices between menus cannot reveal the framing effects.

[2] Eliaz and Spiegler (2010) and Piccione and Spiegler (2010) analyze market models in which consumers choose between framed alternatives in a way that exhibits sensitivity to framing. However, in these models, framing causes preference incompleteness rather than preference reversal.

In contrast, suppose that the decision maker is vulnerable to framing at the stage of choosing between menus. Consider how he ranks singleton menus. When he faces a choice between $\{x\}$ and $\{x'\}$, or between $\{y\}$ and $\{y'\}$, he recognizes the conceit and becomes indifferent. However, when he faces a choice between $\{x\}$ and $\{y\}$, or a choice between $\{x'\}$ and $\{y'\}$, he falls victim to the framing effect and exhibits a preference reversal, thereby violating transitivity of the preference relation over menus. Therefore, a utility function over the set of all menus of framed alternatives cannot capture framing effects.

To conclude, if we accept the idea that framing effects are annulled by their anticipation, we have to conclude that the preferences-over-menus approach is incapable of capturing them, whether or not we assume that the decision maker is vulnerable to framing effects at the stage of choosing between menus.

2.3 Ex-Ante Attitudes are Sometimes Elusive

I have stated repeatedly that ex-ante preferences over choice problems reflect subsequent behavior within the choice problem through the prism of the decision maker's anticipation of his future choices. In fact, there is a double prism, because preferences over choice problem also reflect the decision maker's ex-ante *attitudes* to future choices. But is it always easy to develop strong intuitions about what these ex-ante attitudes might be?

Recall the case of indecisiveness. We noted that indecision often causes delayed action. Using the preferences-over-choice-problems approach to elicit this effect, we can examine choice problems with or without deadlines. If the decision maker anticipates his future indecisiveness, he is unlikely to be indifferent to the existence of tight deadlines. But would the decision maker prefer choice problems with deadlines or without them? If we assume that he shares the taste for delay revealed by his behavior within the choice problem, he should prefer looser deadlines. If, however, he is opposed to the delay, he should prefer tighter deadlines. I do not have strong intuitions on this matter. (As already mentioned, my strongest intuition is that the decision maker will be unable to decide whether or not he wants a tight deadline to be imposed.) The lack of a strong intuition suggests that we may have picked the wrong domain of choice objects for eliciting the effects of indecision.

3 Profiles of Preferences over Acts

In this section I discuss another modeling framework, which follows the general impulse of staying as close as possible to a utility maximization model, but gives up on the quest for a single preference relation over some "Krepsian" domain. Instead, it takes as primitive a collection of complete and transitive preference relations over the original domain of choice objects, where each preference relation is indexed by some aspect of the choice problem.[3]

In a highly stimulating recent paper, Ahn and Ergin (2010) take up this approach in the context of decision under uncertainty. In the Ahn-Ergin model, a frame of an act is a partition of the state space. A "framed act" is a pair (x, r), such that r is a partition and x is an Anscombe-Aumann act that is measurable with respect to r. The choice objects in the Ahn-Ergin model are thus framed acts. One could study choices in this domain that are consistent with maximizing a utility function over it. However, that would effectively rule out the framing effects Ahn and Ergin are interested in.

Not willing to give up on utility maximization as a modeling tool, Ahn and Ergin define a *profile* of complete and transitive preference relation over subsets of framed acts. Specifically, for each frame r, they consider a preference relation over all r-measurable Anscombe-Aumann acts. The framing effect that Ahn and Ergin's model highlights is the dependence of the decision maker's subjective probability of events on the partitions that are used to frame the various available acts. For example, consider the way a life insurance policy specifies payouts in case of a work accident. If the policy describes every type of accident in great detail, this can lead a consumer who contemplates purchasing the policy to assign a higher probability to work accidents than if the policy does not go into such detail. In Ahn and Ergin's elegant utility representation, each preference relation $\succsim_r$ is represented by an expected utility functional, where the utility numbers are independent of r but the subjective probability is a function of r that captures the framing effect Ahn and Ergin are interested in.

However, even when we know everything about this profile of preference relations and their utility representation, how much do we know about the way they map into observed choices? In order for us to be able to talk about choice, we must define it in terms of the preference-profile primitive. Ahn and Ergin offer the following definition, which is an attractive extension of the usual revealed preference definition: when confronted with the pair of framed acts (x, r) and (y, s), the decision maker chooses

[3] Kalai, Rubinstein and Spiegler (2002) formalize this type of rationalization and classify choice functions according to the number of preference relations needed to rationalize them in this manner.

(x, r) if $x \succ_{r \vee s} y$, where $r \vee s$ is the join of the two partitions - namely, their coarsest refinement.

For a fixed frame, the Ahn-Ergin model coincides with standard expected utility maximization. However, as noted above, when we allow the frame to vary, observed choices are inconsistent with maximization of a utility function defined over framed acts. Binary choice cycles may arise because it is possible that $x \succ_{r \vee s} y$, $y \succ_{s \vee t} z$ and $z \succ_{t \vee r} x$. In particular, this means that we cannot use our definition of binary choices as a basis for consistently determining choices from larger choice sets. In fact, Ahn and Ergin do not provide a definition of choices from arbitrary sets. The natural extension of their definition to larger choice sets is as follows: given a set $\{x_i, r_i\}_{i=1,\ldots,n}$, the decision maker chooses (x_i, r_i) over all alternatives if $x_i \succ_{(\vee_{k=1,\ldots,n} r_k)} x_j$ for all $j \neq i$. It is now possible to see that violations of the Independence-of-Irrelevant-Alternatives axiom are consistent with the model. For example, for any given menu of acts, add an act that is manifestly dominated by all the acts in the original menu yet defined in terms of an extremely fine partition, and this will change the frame that the decision maker employs to rank the other acts. The change in frame may reverse the decision maker's preference ranking.

Thus, when we translate the preference-profile model into observed choices, we see that: *(i)* the definition of choices in terms of pairs is incomplete and needs to be extended; and *(ii)* choice behavior violate fundamental rationality properties. Obviously, the latter point is not meant as a criticism of the Ahn-Ergin model, because - as mentioned above - the whole point of modeling framing effects is to account for the violations of rationality they cause. My point is that the mapping from the preference-profile model to observed choices is opaque. It is not clear which properties of observed choices characterize the Ahn-Ergin model. The axiomatization of the profile of frame-indexed preferences gives us very few clues about what we are ultimately interested in, namely observed choices. Maybe it is possible to axiomatize the Ahn-Ergin model in terms of simple properties of observed choices among framed acts; maybe it is not. We just don't know.

This limitation is not special to the Ahn-Ergin model. Any attempt to cast frame-sensitive behavior in terms of a profile of frame-dependent preferences is likely to meet the same fate: unless the model leads to fully rational choice behavior, it will be hard to come up with a transparent characterization of the preference-profile model in terms of properties of observed choice. Again, we see that insisting on using the language of utility maximization to model choice behavior that departs from narrow specifications of rationality can take the choice phenomenon of interest out of focus.

4 Reversing the Priorities

I have critically assessed a popular decision-theoretic methodology that is based on two principles: first, a search for a domain of choice objects that is most effective in eliciting a psychological force of interest; and second, an insistence on utility maximization as a modeling device. But what happens when the two principles are in conflict - i.e., when choices in the most "natural" domain are inconsistent with utility maximization (because revealed preference are incomplete or inconsistent)? My impression is that many decision theorists in the literature that followed Gul and Pesendorfer (2001) have tended to prioritize the two principles such that utility maximization is retained at all costs, even if this means tampering with the domain of choice objects.

An alternative approach, which I would like to advocate here, is to reverse these priorities, and insist on analyzing observed choices in a "natural" domain that appears most strongly connected with the psychological phenomenon of interest, or most relevant for economic applications - even if this means sacrificing utility maximization as a modeling tool. This approach is "Krepsian" in the sense that it takes choices as primitive, and it allows us to expand the notion of a choice problem beyond the mere set of feasible alternatives, to incorporate things like the framing of alternatives, default options, the order in which alternatives are presented to the decision maker, and so forth.[4] Choice functions are defined over such an extended domain, and the main task is - as usual in decision theory - to connect properties of this function to decision models that involve unobserved mental entities. The crucial difference is that the decision model is not necessarily a utility maximization model.

This path has been taken by Ok and Masatlioglu (2005) and Ortoleva (2010) in the case of choices that display status-quo bias (Ortoleva's paper explicitly addresses choice under Knightian uncertainty), Rubinstein and Salant (2006,2008) in the case of choices that may depend on the order in which alternatives are presented and other types of frames, and Mariotti and Manzini (2007) and de Clippel and Eliaz (2009) in the case of choices that result from some procedure for aggregating multiple preference criteria - to name a few examples.

Models in this extended choice-theoretic tradition cannot be justified on normative grounds, because they do not lead to a conventional utility representation theorem. Therefore, they only admit descriptive interpretations. This does not mean that these models preclude the elicitation of a notion of welfare from choice data. However, one

[4] Rubinstein (2008) also proposes adding non-choice observable data as the *output* of such an extended choice function.

needs to modify the conventional definition of revealed preferences and make "identifying assumptions" to this end. For example, Masatlioglu, Nakajima and Ozbay (2009) explore the elicitation of welfare rankings when the decision maker has limited attention, Rubinstein and Salant (2010) discuss this question when the decision maker is prone to small decision errors, while Bernheim and Rangel (2009) explore the limitations of non-paternalistic welfare analysis when decision makers make mistakes. Because of this departure from conventional revealed preference methodology, I expect this to be among the interesting and controversial directions in decision theory in the foreseeable future.

References

[1] Ahn, D. and H. Ergin (2010), "Framing Contingencies," Econometrica 78, 655-695.

[2] Bernheim, D. and A. Rangel (2009), "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," Quarterly Journal of Economics 124, 51-104.

[3] de Clippel, G. and E. Eliaz (2009), "Reason-Based Choice: A Bargaining Rationale for the Attraction and Compromise Effects," working paper, Brown University.

[4] Dillenberger, D. and P. Sadowski (2010), "Ashamed to be Selfish," working paper, University of Pennsylvania.

[5] Eliaz, K. and R. Spiegler (2010), "Consideration Sets and Competitive Marketing," Review of Economic Studies, forthcoming.

[6] Gul, F. and W. Pesendorfer (2001), "Temptation and Self-Control," Econometrica 69, 1403-1436.

[7] Kalai, G., A. Rubinstein and R. Spiegler (2002), "Rationalizing Choice Functions by Multiple Rationales," Econometrica 70, 2481-2488.

[8] Kreps, D. (1979), "A Representation Theorem for 'Preference for Flexibility'," Econometrica 47, 565-576.

[9] Kreps, D. and E. L. Porteus (1978), "Temporal Resolution of Uncertainty and Dynamic Choice Theory," Econometrica 46, 185-200.

[10] Manzini, P. and M. Mariotti (2007), "Sequentially Rationalizable Choice," American Economic Review 97, 1824-1839.

[11] Noor, J. (2009), "Temptation and Revealed Preference," Econometrica, forthcoming.

[12] Masatlioglu, Y. and E. Ok (2005), "Rational Choice with Status Quo Bias," Journal of Economic Theory 121, 1-29.

[13] Masatlioglu, Y., D. Nakajima and E. Ozbay (2009), "Revealed Attention," mimeo.

[14] Ortoleva, P. (2010), "Status Quo Bias, Multiple Priors and Uncertainty Aversion," Games and Economic Behavior 69, 411-424.

[15] Piccione, M. and R. Spiegler (2010), "Price Competition under Limited Comparability," mimeo.

[16] Rubinstein. A. (2008), "Comments on Neuroeconomics," Economics and Philosophy 24, 485-494.

[17] Rubinstein, A. and Y. Salant (2006), "A Model of Choice from Lists," Theoretical Economics 1, 3-17.

[18] Rubinstein, A. and Y. Salant (2008), "(A,f), Choice with Frames," Review of Economic Studies 75, 1287-1296.

[19] Rubinstein, A. and Y. Salant (2010), "Eliciting Welfare Preferences from Behavioral Dataset," mimeo.

[20] Sarver, T. (2008), "Anticipating Regret: Why Fewer Options May Be Better," Econometrica 76, 263-305.