

DUAL DECISION PROCESSES:

RETRIEVING PREFERENCES WHEN SOME CHOICES ARE AUTOMATIC

Francesco Cerigioni*
Universitat Pompeu Fabra and Barcelona GSE

January 13, 2017

ABSTRACT

Evidence from cognitive sciences suggests that some choices are conscious and reflect individual volition while others tend to be automatic, driven by analogies with past experiences. Under these circumstances, usual economic modeling might not be apt because not all choices are the consequence of individual tastes. We propose a behavioral model that can be used in standard economic analysis that formalizes how conscious and automatic choices arise by presenting a decision maker composed of two selves. One self compares past decision problems with the one the decision maker faces and, when the problems are similar enough, it replicates past behavior (Automatic choices). Otherwise, a second self is activated and preferences are maximized (Conscious choices). We then present a novel method capable of identifying a set of conscious choices from observed behavior and finally, we discuss its importance as a framework to study empirical puzzles in different settings.

(JEL D01, D03, D60)

KEYWORDS: Dual Processes, Fluency, Similarity, Revealed Preferences, Memory, Automatic Choice

*Department of Economics and Business, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005, Barcelona, Spain. E-mail: francesco.cerigioni@upf.edu. The paper has previously circulated under the title “Dual Decision Processes: Retrieving Preferences when some Choices are Intuitive”. I am grateful to the Ministerio de Ciencia y Innovacion for the FPI scholarship connected with the project ECO2010-17943. I thank Miguel A. Ballester for his insightful comments that were essential for the development of this paper and the editor Joel Sobel and four anonymous referees for their critiques and suggestions that greatly improved this work. Furthermore, I want to acknowledge the incredible environment I found in The Eitan Berglas School of Economics (TAU). In particular, I am in debt to Kfir Eliaz, Itzhak Gilboa, Ariel Rubinstein and Rani Spiegel for their help and suggestions. I am also grateful to David Ahn, Larbi Alaoui, Jose Apesteguia, B. Douglas Bernheim, Caterina Calsamiglia, Andrew Caplin, Laurens Cherchye, Vincent Crawford, Tugce Cuhadaroglu, Jan Eeckhout, Matthew Eelman, Fabrizio Germano, Johannes Gierlinger, Benjamin Golub, Edgardo Lara-Cordova, David Laibson, Marco Mariotti, Daniel Martin, Yusufcan Masatlioglu, Antonio Miralles, Isabel Melguizo, Muriel Nierderle, Efe Ok, Dilan Okcuoglu, Marco Pelliccia, Pedro Rey-Biel, Tomás Rodríguez-Barraquer, Michael Richter, Guillem Roig-Roig, John Tsoukalis and participants at SITE and BRIC for their comments and suggestions. All mistakes are mine.

1 Introduction

Behavioral economics has posed a serious challenge to standard economic theory since it has documented and studied numerous behaviors inconsistent with preference maximization. Inconsistencies tend to arise when individuals do not exert the effort to consciously analyze the problem at hand, e.g., Carroll et al. (2009). A key question arises then, when are people choosing consciously?

In this paper we provide a first theoretical framework to think about the coexistence of conscious and automatic, i.e., unconscious, behavior and we provide a method to understand the nature of individual decisions in different situations once this duality is taken into account. Studying this dichotomy is crucial to understand market outcomes. In fact, as highlighted by Simon (1987) and Kahneman (2002), *experts* such as managers, doctors, traders or policy makers often make automatic decisions. Nevertheless, the sources of automatic choices are still understudied in economics.¹

Evidence from cognitive sciences suggests that the *familiarity* of the choice environment is the key determinant of the relationship between conscious and automatic choices.² Unconscious, automatic choices are made in familiar environments. Thus, if we want to understand the role of automatic decisions in markets, we need (i) to have a model of automatic choices that takes into account the role of familiarity of the decision environment and (ii) to analyze whether it is possible from observed behavior to distinguish between conscious and automatic choices to understand the preferences that drive observed behavior. These are the research questions we address in the paper.

To answer the first question we propose a simple formalization of evidence coming from cognitive sciences whose main contribution is to provide a first model describing when and how choices should be conscious or automatic. In Section 3, we present a decision maker described by a simple procedure. Whenever the decision environment is familiar, i.e., whenever its similarity with what has been experienced, measured by

¹Nonetheless, in economics there is a growing attention to conscious and intuitive reasoning. See for example Rubinstein (2007) and Rubinstein (2016) for a distinction between conscious and intuitive strategic choices for players participating in a game or the distinction in Cunningham and de Quidt (2015) between implicit and explicit attitudes.

²See section 3.2 for a review of the literature discussing this dichotomy.

a similarity function, passes a certain threshold, past behavior is replicated. This is the source of *automatic* choices. Otherwise, the best option is chosen by maximizing a rational preference relation. This is the source of *conscious* choices. Think for example of a consumer that buys a bundle of products from the shelves of a supermarket. The first time he faces the shelves, he tries to find the bundle that best fits his preferences. Afterwards, if the price and arrangement of products do not change *too much*, he will perceive the decision environment as familiar and so he will automatically stick to the past chosen bundle. If on the contrary, the change in price and arrangement of the products is evident to him, he will choose by maximizing his preferences again.

Even in such a simple framework, there is no trivial way to distinguish which choices are made automatically and which ones are made consciously. Following the example, suppose our consumer faces again the same problem but this time a new bundle is available and he sticks with the old choice. Is it because the old bundle is preferred to the new one? Or is it because he is choosing automatically?

We show how to find conscious choices and hence restore standard revealed preference analysis by understanding which environments were not familiar. Section 4 assumes that (i) the decision maker behaves according to our model and (ii) the similarity function is known while the threshold is not.³ We then show that, for every sequence of decision problems, it is possible to identify by means of an algorithm a set of conscious observations and an interval in which the similarity threshold should lie. That is, we provide a novel method to restore revealed preference analysis.

First notice that *new* observations, i.e., those in which the choice is an alternative that had never been chosen before, must be conscious. No past behavior could have been replicated. Starting from these observations, the algorithm iterates the following idea. If an observation is consciously generated, any other *less familiar* observation, that is any problem which is less similar to those decision problems that preceded it, must be also consciously generated. Returning to our consumer, if we know that after a change in the price of the products on the shelf, the consumer chose consciously, then he must have done so also in all those periods where the change was even more

³See Sections 3 and 4 for a justification of the latter hypothesis.

evident.

The algorithm identifies a set of automatic decisions in a similar fashion, that is, first it highlights some decisions that have to be automatic and then finds *more familiar* observations to reveal other automatic decisions. Notice that understanding if some decisions were made automatically is very important to understand how familiarity of environments is determined. Even if automatic choices do not reveal individual preferences, they tell us what problems are considered familiar, i.e., similar enough, by the decision maker hence allowing for the identification of the interval in which the similarity threshold should lie.

The algorithm assumes that the decision maker behaves following our model, hence falsifiability of the model becomes a central concern. In Section 5 we propose a testable condition that is a weakening of the Strong Axiom of Revealed Preference that characterizes our model and thus renders it falsifiable. Section 6 studies the model as a general framework formalizing the coexistence of behavioral inertia and adaptive behavior that might give new insights into physicians behavior, pricing behavior of retailers and portfolio decisions of traders in financial markets. Section 7 concludes while the Appendix contains an extension of the identification algorithm assuming only partial knowledge of similarity comparisons.

2 Related Literature

In our model the presence of similarity comparisons makes behavior automatic, that is, if two environments are similar enough then behavior is replicated. This is a different approach with respect to the theory for decisions under uncertainty proposed in Gilboa and Schmeidler (1995) and summarized in Gilboa and Schmeidler (2001). In case-based decision theory, as in our model, a decision maker uses a similarity function in order to assess how much alike are the problem he is facing and the ones he has in his memory. In that model the decision maker tends to choose the action that performed better in past similar cases. There are two main differences with the approach we propose here. First, from a conceptual standpoint, our model relies on the idea of two selves interacting during the decision making process. Second, from a

technical point of view, our model uses the similarity in combination with a threshold to determine whether the individual replicates past behavior or maximizes preferences while in Gilboa and Schmeidler (1995) preferences are always maximized. Thus, as section 7 suggests, case-based decision theory can be ingrained in the more general structure proposed here. The model in Gilboa and Schmeidler (1995) can be seen as a particular way of making conscious decisions. Nevertheless, both models agree on the importance of analogies for human behavior. In Gilboa and Schmeidler (1995) analogies are used to find the action that maximizes individual utility in a world without priors, here analogies can be potentially dangerous because they determine whether the decision maker thinks about his choices or just chooses automatically without consciously analyzing the problem.

We would like to stress that even if the behavioral model we propose is new, the idea that observed behavior can be the outcome of the interaction between two different selves is not novel and it dates back at least to Strotz (1955). Strotz kind of models, such as Laibson (1997), Gul and Pesendorfer (2001) or Fudenberg and Levine (2006), are different from the behavioral model we introduce here, since they represent the two selves as two agents with different and conflicting preferences, usually *long-run* vs *short-run* preferences.⁴ In our approach however, the two selves are *inherently* different one from the other. One uses analogies to deal with the environment in which the decision maker acts, while the other one uses a preference relation to consciously choose among the alternatives available to the decision maker. Furthermore, which self drives a particular decision problem depends on the problems that have been experienced and how similar they are with the present one and thus, it does not depend on whether the decision is affecting the present or the future. Nevertheless, we do not exclude the possibility that the fact that a decision affects the present or the future has some kind of influence on how analogies are made.

Finally, it is important to notice that the preference revelation strategy we use in

⁴In some models the difference between the two selves comes from the fact that they have different information. See for example Cunningham (2015) that proposes a model of decision making where the two selves hierarchically aggregate information before choosing an alternative. Other papers use more than two selves to rationalize individual choices, but still all selves are represented by different preference relations, see for example Kalai et al. (2002), Manzini and Mariotti (2007) and Ambrus and Rozen (2015).

the paper agrees with the one used in Bernheim and Rangel (2009). They analyze the same problem of eliciting individual preferences from behavioral datasets, and they do this in two stages. In a first stage they take as given the *welfare relevant domain*, that is the set of observations from which individual preferences can be revealed, and then in a second stage they analyze the welfare relevant observations and propose a criterion for the revelation of preferences that does not assume any particular choice procedure to make welfare judgments.⁵ Even if similar, our approach differs in two important aspects. First, by modeling conscious and intuitive choices, we propose a particular method to find the welfare relevant domain, i.e. the algorithm highlighting a set of conscious choices. Second, by proposing a specific choice procedure, we use standard revealed preference analysis on the relevant domain, thus our method, by being behaviorally based, is less conservative for the elicitation of individual preferences. In this sense, our stance is also similar to the one proposed in Rubinstein and Salant (2012), Masatlioglu et al. (2012) and Manzini and Mariotti (2014) that make the case for welfare analysis based on the understanding of the behavioral process that generated the data.

3 Dual Decision Processes

3.1 The Model

Let X and E be two sets. The decision maker (DM) faces at time t a *decision problem* (A_t, e_t) with $A_t \subseteq X$ and $e_t \in E$. The set of alternatives A_t that is available at time t , and from which the DM has to make a choice, is called the *menu*. An alternative is any element of choice like consumption bundles, lotteries or streams of consumption. The *environment* e_t is a description of the possible characteristics of the problem that the DM faces at time t . As highlighted below, an environment can be the menu, the set of attributes of alternatives in the menu, the framing, etc. We denote by $a_t \in A_t$ the chosen alternative at time t . With a little abuse of the notation, we refer to the couple

⁵Notice that Apesteguía and Ballester (2015) propose an approach to measure the welfare of an individual from a given dataset that is also choice-procedure free. They do so by providing a model-free method to measure how close actual behavior is to the preference that best reflects the choices in the dataset.

formed by the decision problem (A_t, e_t) and the chosen alternative a_t as *observation* t . We denote the collection of observations in the sequence $\{(A_t, e_t, a_t)\}_{t=1}^T$ as D , i.e. $D = \{1, \dots, T\}$. Notice that the same menu or environment can appear more than once in the sequence.

A chosen alternative is the outcome of a two stage choice procedure that describes the DM and that formalizes the duality of automatic and conscious choices. Formally, let $\sigma : E \times E \rightarrow [0, 1]$ be the *similarity function*. The value $\sigma(e, e')$ measures how similar environment e is with respect to e' . The *automatic self* is endowed with a *similarity threshold* $\alpha \in [0, 1]$ that delimits which pairs of environments are similar enough. Whenever $\sigma(e, e') > \alpha$ the individual considers e to be similar enough to e' . At time t and facing the decision problem (A_t, e_t) , the automatic self executes a choice if it can replicate the choice of a previous period $s < t$ such that $\sigma(e_t, e_s) > \alpha$. The choice is the alternative a_s chosen in one such period. The *maximizing self* is endowed with a *preference relation* \succ over the set of alternatives. For ease of exposition, we assume that \succ is a strict order, i.e. an asymmetric, transitive and complete binary relation, defined over X . At time t , if the maximizing self is activated it chooses the alternative a_t that maximizes \succ in A_t .⁶ Summarizing:

$$a_t = \begin{cases} a_s \text{ for some } s < t \text{ such that } \sigma(e_t, e_s) > \alpha \text{ and } a_s \in A_t, \\ \text{the maximal element in } A_t \text{ with respect to } \succ, \text{ otherwise.} \end{cases}$$

Three remarks are useful here. First, notice that automatic and conscious decisions are separated by the behavioral parameter α . In some sense α is summarizing the cost of using the cognitive demanding system. The higher the cost, the lower the threshold. Thus, parameter α captures individual heterogeneity as preferences do. Notice that while the similarity function has been widely studied in cognitive sciences, e.g. Tversky (1977), Medin et al. (1993) and Hahn (2014), the cognitive costs of activating conscious decisions processes are still an unknown, thus the method we propose in section 4 can be seen as a first attempt of identifying from observed

⁶The idea that conscious behavior comes from the maximization of a preference relation is a simplification we use to focus the analysis on the main novelties of the framework presented in this paper. For a more detailed discussion regarding this point, see section 7.

behavior the interval in which such costs should lie, given the similarity function.

As a second remark, notice that we are describing a class of models because we do not impose any particular structure on replicating behavior. We do not specify which alternative would be chosen when more than one past choice can be replicated. Many different behaviors can be part of this class, e.g. choosing the alternative that was chosen in the most similar past environment or choosing the alternative that maximizes the preference relation over the ones chosen in similar enough past environments, etc. All the analysis that follows is valid for the class as a whole.

As a final remark, for illustrative purposes, we propose here some examples of environments that are relevant for economic applications and a possible similarity function that can be used in such cases.

Environments as Menus: *In many economic applications it seems sensible to see the whole menu of alternatives, e.g. the budget set, as the main driver of analogies. That is, E could be the set of all possible menus and two decision problems are perceived as similar as their menus are. In this framework, $E = 2^X$.*

Environments as Attributes: *Decision makers many times face alternatives that are bundles of attributes. In those contexts, it is reasonable to assume that the attributes of the available alternatives determine the decision environment. If A is the set containing all possible attributes, then $E = 2^A$.*

Environments as Frames: *We can think of the set E as the set of frames or ancillary conditions as described in Salant and Rubinstein (2008) and Bernheim and Rangel (2009). Frames are observable information that is irrelevant for the rational assessment of alternatives, for example how the products are disposed on a shelf. Every frame can be seen as a set of irrelevant features of the decision environment. Thus, if the set containing all possible irrelevant features is F , we have $E = 2^F$.*

In all the previous examples it is natural to assume that the similarity function relates different environments depending on their commonalities and differences. For example, $\sigma(e, e') = \frac{|e \cap e'|}{|e \cup e'|}$, that is, two environments are more similar the more characteristics they share relative to all the characteristics they have. Such function is just a symmetric specification of the more general class considered in Tversky (1977).

Although, it is sometimes not possible to have all the information regarding the similarity function, a case we analyze in the appendix, from now on we take E and σ as given.

3.2 Cognitive Foundations of the Model

This section maps recent findings in cognitive sciences into the model and shows which simplifying assumptions were made. The model tries to formalize the evidence concerning the activation of non-deliberative processes. With this objective in mind, first we present which mental state makes unconscious processes play an important role in individual decision making and then we study how such state arises.

Cognitive sciences have discovered that unconscious processes are extremely context dependent. The activation of non deliberative cognition depends on the characteristics of the environment. In particular, the studies on *fluency*, i.e., the subjective experience of ease associated with a mental process or a particular situation, highlight that the influence of unconscious processes is stronger in those contexts that trigger a higher sensation of cognitive ease.⁷

The studies surveyed by Oppenheimer (2008), one of the leading scholars in fluency research, show that in fluent, familiar, situations individuals decisions tend not only to be less conscious, but also less coherent and more prone to errors and biases.⁸ Here we propose a model of fluency that is based on the formalization of what is usually considered its source, i.e., priming.

Paraphrasing Oppenheimer (2008), the interpretation of a fluent experience relies

⁷Fluency of a situation is seen as the main determinant of automatic or fast thinking as defined in Dual Process Theory. See Schneider and Shiffrin (1977), Evans (1977), McAndrews et al. (1987), Evans (1989), Reber (1989), Epstein (1994) and Evans and Over (1996) for the development of Dual Process Theory as described in Evans and Frankish (2009) and Kahneman (2011) that uses much of the evidence presented here and that is related to the model proposed.

⁸For example, Alter et al. (2007) show that individuals choosing in a fluent environment tend to do worst in the Cognitive Reflection Test and in solving logical syllogisms than individuals in a disfluent situation. They also show that individuals in a fluent setting tend to evaluate consumer products by placing less weight on their objective characteristics. In another study, Hernandez and Preston (2013) show that putting people in a disfluent environment disrupts the confirmation bias. Another example comes from the study conducted by Costa et al. (2014). They show that students put in a fluent environment in which assignments are written in their mother tongue tend to make more biased decisions than children put in a disfluent one where assignments are written in a foreign language.

on the unconsciously perceived relationship between past experiences and the current context. Such relationship is analyzed by *priming* that studies the influence of unconscious or implicit memory on behavior through environmental cues. As defined in Tulving and Schacter (1990), priming is an unconscious change in the ability to identify or produce an item as a result of a specific prior encounter with the item.

Priming creates a sense of familiarity and ease that makes the environment fluent and behavior automatic and coherent with the context. The behavioral model proposed here, even in its simplicity, captures some of the main characteristics of priming in a stylized way.

Bargh (2005), one of the major contributors to the study of priming, highlights the following characteristics of the phenomenon:

1. It is unconscious.
2. It is based on the conceptual representation of the environment.
3. It determines behavior.
4. It is costly to control.
5. It has a long term effect on behavior and it is not affected by age.

The remainder of the section explains how the model incorporates these points.

Regarding point 1, the evidence presented in Bargh (2005) shows that priming is completely unconscious. Quoting Bargh:

“...individual’s behavior is being controlled by external stimuli not by his or her consciously accessible intentions or acts of will... action tendencies can be activated and triggered independently and in the absence of individual’s conscious choice or awareness of those causal triggers.”

Neurological evidence strengthens even more this finding. Many studies have shown a decrease in neural activity when an individual faces an environment that primes behavior, a phenomenon known as *repetition suppression*. The decrease of activity affects in particular neurons that are usually associated with conscious recognition of

the stimuli, while neurons associated with unconscious recognition are not affected as much. This physiological reaction points to the fact that recognition of environmental cues precedes any conscious control of behavior. The model captures this idea by assuming that the automatic self acts in the first stage of the choice procedure.

The evidence presented by Bargh (2005) shows that priming is effective even if the environment slightly changes with respect to the reference one, that is, even if some features have changed.⁹ As stated in point 2, priming is based on a conceptual representation of experienced environments. In the model, an environment can influence present behavior because it is similar enough to, i.e., it shares the same conceptual representation of, the environment that primes behavior.

Point 3 is the most important from a behavioral perspective and it is a defining characteristic of priming. Again, quoting Bargh:

“... an individual’s behavior can be directly caused by the current environment... this behavior can and will unfold without the person being aware of its external determinant.”

This idea is the basis of priming. The most radical examples in this direction come from subjects with prefrontal cortex dysfunctions, a critical region for conscious and controlled behavior. Many studies, e.g., Lhermitte (1986), Ghika et al. (1995) and Bisio et al. (2012), have shown that these individuals can be made to drastically change their behavior just by manipulating the environment. The reaction is coherent with the context priming behavior. Such extreme dependency on the environment, called *environmental dependency syndrome*, is a radical manifestation of priming, as Bargh (2005) states, given it activates the same areas of the brain.¹⁰ This complex idea is simplified in the paper by means of replication of past behavior from familiar environments.

⁹See also Higgins and Bargh (1987) and Bargh et al. (2001) for similar discussions.

¹⁰It is interesting to note that healthy individuals unconsciously experiencing priming try to consciously justify their behavior ex-post. That is, the influence of priming on behavior is much more subtle than what it seems at a first glance. Individuals can consciously justify their behavior ex-post without acknowledging the influence of the environment in the moment preceding the action. See Bargh and Morsella (2008) for a defense of the idea that action precedes reflection.

Finally, points 4 and 5 are important to understand why we assume that (i) α is an individual characteristic and (ii) memory is perfect. Regarding the first point, the similarity threshold tries to catch the idea that overcoming fluency to take control of one's own actions is costly and depends on individual cognitive characteristics. With respect to the second point, the assumption of perfect memory is a simplification of the fact that, as highlighted by the evidence in Bargh (2005) and also in Duhigg (2012), behavior can be primed by experiences stored in implicit memory that can be forgotten by explicit, or conscious, memory. For example, there is evidence that patients with amnesia are affected by priming as strongly as any other individual. Moreover, it seems that priming effects do not deteriorate with age, hence pointing to the fact that implicit memory is not affected by age as explicit memory is. Notice however, that this assumption can be weakened without affecting the results.

To conclude, the model captures important aspects of fluency and priming in a simple way while keeping the main insights of environmental influence on behavior.

3.3 An example

This section provides an example to illustrate the behavioral process we are modeling. Even if the example is abstract, it shows in a very direct way how the model works. For more concrete applications of the model to interesting economic settings, please refer to section 6.

Let $X = \{1, 2, 3, \dots, 10\}$. We assume that environments are menus, i.e. E is the set of all subsets of X and we assume that $\sigma(A, A') = \frac{|A \cap A'|}{|A \cup A'|}$. Suppose that the automatic self is described by $\alpha = .55$ and that the preference $1 \succ 2 \succ 3 \succ \dots \succ 10$ describes the maximizing one. We now explain how our DM makes choices from the following list of ordered menus:

$t = 1$	$t = 2$	$t = 3$	$t = 4$	$t = 5$	$t = 6$	$t = 7$	$t = 8$
3, 4, 5, 6	1, 2, 3, 4, 5, 6	1, 3, 4, 7	2, 4, 7	1, 3, 6	1, 2, 3, 4	2, 4, 8	2, 4, 8, 9, 10

In the first period, given the DM has no prior experiences, the maximizing self is going to be active. Thus, the choice comes from the maximization of preferences,

that is, $a_1 = 3$. Then, in the following period, given that $\sigma(A_2, A_1) = \frac{4}{6} > .55$, the automatic self is active and so we would observe a replication of past choices, that is, $a_2 = 3$. Period 1's environment makes period 2's one fluent, hence priming behavior. Now, in period 3, notice that the similarity between A_3 and A_2 or A_1 is always below the similarity threshold and this makes the maximizing self active. The preference relation is maximized and so $a_3 = 1$. A similar reasoning can be applied for the fourth and fifth periods to see that $a_4 = 2$ and $a_5 = 1$. Then, in period six, the automatic self is active given that $\sigma(A_6, A_3) = \frac{3}{5} > .55$, leading to $a_6 = 1$. In period seven, given no past environment is similar enough, the maximizing self is active and so $a_7 = 2$. Finally, in the last period the automatic self is active again given that $\sigma(A_8, A_7) = \frac{3}{5} > .55$ and so behavior will be replicated, i.e. $a_8 = 2$.

One may wonder what an external observer would understand from this choice behavior. Section 4 addresses this point

4 The Revealed Preference Analysis of Dual Decisions

In this section, we discuss how to recognize which observations were automatically or consciously generated in a DD process. This information is crucial to elicit the unobservables in the model that are the sources of individual heterogeneity, that is the preference relation and the similarity threshold. As previously discussed in Section 3, the similarity function is taken as given, while the similarity threshold α is elicited from observed behavior. Notice that, the similarity function and the similarity threshold define a binary similarity function that is individual specific. Thus, by separating the similarity function from the similarity threshold we are able to associate individual heterogeneity to a parameter related with individual cognitive costs without having too many degrees of freedom to properly run the technical analysis.

It is easy to recognize a set of observations that is undoubtedly generated by the maximizing self. Notice that all those observations in which the chosen alternative was never chosen in the past must belong to this category. This is so because, as no

past behavior has been replicated, the automatic self could not be active. We call these observations *new* observations.

In order to identify a first set of observations automatically generated, notice that the maximizing self being rational, it cannot generate cycles of revealed preference. As it is standard, a set of observations t_1, t_2, \dots, t_k forms a *cycle* if $a_{t_{i+1}} \in A_{t_i}$, $i = 1, \dots, k - 1$ and $a_{t_1} \in A_{t_k}$, where all chosen alternatives are different. Given the previous reasoning, for every cycle there must be at least one observation that is automatically generated. Intuitively, the one corresponding to the most familiar environment should be a decision mimicking past behavior.

The *unconditional familiarity* of observation t is

$$f(t) = \max_{s < t, a_s \in A_t} \sigma(e_t, e_s).$$

Whenever there is no $s < t$ such that $a_s \in A_t$, we say $f(t) = 0$.

That is, unconditional familiarity measures how similar observation t is to past observations from which behavior *could be* replicated, i.e., those past decision problems for which the chosen alternative is present at t . Then, we say that observation t is a *most familiar in a cycle* if it is part of a cycle of observations, and within it, it maximizes the value of the unconditional familiarity.

The major challenge is to relate pairs of observations in a way that allows to transfer the knowledge of which self generated one of them to the other. In order to do so, we introduce a second measure of familiarity of an observation t , that we call *conditional familiarity*. Formally,

$$f(t|a_t) = \max_{s < t, a_s = a_t} \sigma(e_t, e_s).$$

Whenever there is no $s < t$ such that $a_s = a_t$, we say $f(t|a_t) = 0$.

That is, conditional familiarity measures how similar observation t is with past observations from which behavior *could have been* replicated, i.e., those past decision problems for which the choice is the same as the one at t . The main difference between $f(t)$ and $f(t|a_t)$ is that the first one is an ex-ante concept, i.e., before considering the choice, while the second one is an ex-post concept, i.e., after considering the

choice. Our key definition uses these two measures of familiarity to relate pairs of observations.

Definition 1 (Linked Observations) *We say that observation t is linked to observation s , and we write $t \in L(s)$, whenever $f(t|a_t) \leq f(s)$. We say that observation t is indirectly linked to observation s if there exists a sequence of observations t_1, \dots, t_k such that $t = t_1$, $t_k = s$ and $t_i \in L(t_{i+1})$ for every $i = 1, 2, \dots, k - 1$.*

The definition formalizes the key idea behind the algorithm. Observation t is linked to observation s if its conditional familiarity is below the unconditional familiarity of s . As explained below, once an observation is categorized as consciously or automatically generated, it is through its link to other observations that such knowledge can be extended.

Denote by D^N the set of all observations that are indirectly linked to new observations and by D^C the set of all observations to which most familiar observations in a cycle are indirectly linked. The binary relation determined by the concept of linked observations is clearly reflexive, thus, new observations and most familiar observations in a cycle are contained in D^N and D^C respectively. We are ready to present the main result of this section. It establishes that observations in D^N are generated by the maximizing self, while observations in D^C are generated by the automatic self. As a consequence, it guarantees that the revealed preference of observations in D^N is informative with respect to the preferences of the individual and, moreover, it identifies an interval in which the similarity threshold has to lie. Before stating the proposition, it is useful to highlight that x is revealed preferred to y for a set of observations O , i.e., $xR(O)y$, if there is a sequence of different alternatives x_1, x_2, \dots, x_k such that $x_1 = x, x_k = y$ and for every $i = 1, 2, \dots, k - 1 \in O$, it is $x_i = a_t$ and $x_{i+1} \in A_t$ for some t .

Proposition 1 *For every collection of observations D generated by a DD process:*

1. *all observations in D^N are generated by the maximizing self while all observations in D^C are generated by the automatic self,*
2. *if x is revealed preferred to y for the set of observations D^N , then $x \succ y$,*

$$3. \max_{t \in D^N} f(t) \leq \alpha < \min_{t \in D^C} f(t|a_t).$$

Proof. We start by proving the statement regarding conscious observations. Trivially, new observations must be generated by the maximizing self since they cannot replicate any past behavior. Consider an observation $t \in D^N$. By definition, there exists a sequence of observations t_1, t_2, \dots, t_n with $t_1 = t$, $f(t_i|a_{t_i}) \leq f(t_{i+1})$ for all $i = 1, 2, \dots, n - 1$ and t_n being new. We prove that t is generated by the maximizing self recursively. We know that t_n is consciously generated. Now assume that t_k is generated by the maximizing self and suppose by contradiction that t_{k-1} is generated by the automatic one. From the assumption on t_k , we know that $f(t_k) \leq \alpha$. From the assumption on t_{k-1} , we know that $f(t_{k-1}|a_{t_{k-1}}) > \alpha$, which implies $f(t_{k-1}|a_{t_{k-1}}) > f(t_k)$, a contradiction with the hypothesis. Hence, t_{k-1} is also generated by the maximizing self, and the recursive analysis proves that observation t is also consciously generated.

We now prove the statement regarding automatic observations. Consider first an observation t which is a most familiar in a cycle and assume by contradiction that it is generated by the maximizing self. Then, $f(t) \leq \alpha$. By definition of most familiar in a cycle, it must be $f(s) \leq \alpha$ for every s in the cycle, making all decisions in the cycle being generated by the maximizing self. This is a contradiction with the maximization of a preference relation. Consider now an observation $t \in D^C$. By definition, there exists a sequence of observations t_1, t_2, \dots, t_n with $t_n = t$, $f(t_i|a_{t_i}) \leq f(t_{i+1})$ for all $i = 1, 2, \dots, n - 1$ and t_1 being a most familiar in a cycle. We proceed recursively again. Since t_1 is generated by the automatic self, we have $f(t_1|a_{t_1}) > \alpha$. Now assume that t_k is generated by the automatic self and suppose by contradiction that t_{k+1} is generated by the maximizing one. We then know that $f(t_k|a_{t_k}) > \alpha \geq f(t_{k+1})$, which is a contradiction concluding the recursive argument.

For the revelation of preferences part, since D^N can only contain observations generated by the maximizing self, it is straightforward to see that the revealed information from such a set must respond to the preferences of the DM. Regarding α , notice that since observations in D^N are generated by the maximizing self, we know that $\max_{t \in D^N} f(t) \leq \alpha$ and also that, since observations in D^C are generated by the automatic self, $\alpha < \min_{t \in D^C} f(t|a_t)$, which concludes the proof. ■

To understand the reasoning behind Proposition 1, consider first an observation t that we know is new, and hence consciously generated. This implies that its corresponding environment is not similar enough to any other previous environment. In other words, $f(t) \leq \alpha$. Then, any observation s for which the conditional familiarity is less than $f(t)$ must be generated by the maximizing self too. In fact, $f(s|a_s) \leq f(t) \leq \alpha$ implies that no past behavior that could have been replicated in s comes from an environment that is similar enough to the one in s . Thus, any observation linked with a new observation must be generated by the maximizing self. It is easy to see that this reasoning can be iterated, in fact, any observation linked with an observation generated by the maximizing self must be generated by the maximizing self too.

Similarly, consider a most familiar observation in a cycle t , that we know is automatically generated. Any observation s for which the unconditional familiarity is greater than the conditional familiarity of t must be generated by the automatic self too. In fact, we know that $\alpha < f(t|a_t)$ because t is generated by the automatic self. Then, any observation s to which t is linked has an unconditional familiarity above α , which implies that some past behavior *could be* replicated by the automatic self, and so such observation must be generated by the automatic self too. Again, the reasoning can be iterated. Thus, we can start from a small subset of observations undoubtedly generated either automatically or consciously, inferring from there which other observations are of the same type. We use the example of section 3.3 to show how the algorithm works.

Algorithm: Example of Section 3.3

Suppose that we observe the decisions made by the DM in the example of section 3.3, without any knowledge on his preferences \succ or similarity threshold α . The following table summarizes the different observations.

$t = 1$	$t = 2$	$t = 3$	$t = 4$	$t = 5$	$t = 6$	$t = 7$	$t = 8$
3, 4, 5, 6	1, 2, 3, 4, 5, 6	1, 3, 4, 7	2, 4, 7	1, 3, 6	1, 2, 3, 4	2, 4, 8	2, 4, 8, 9, 10
$a_1 = 3$	$a_2 = 3$	$a_3 = 1$	$a_4 = 2$	$a_5 = 1$	$a_6 = 1$	$a_7 = 2$	$a_8 = 2$

We can easily see that the only new observations are observations 1, 3 and 4, and hence we can directly infer that the corresponding choices were conscious.

We can go one step further and consider observation 5. From observed behavior we cannot understand whether the choice comes from maximization of preferences or the replication of past behavior in period 3. Nevertheless, the choice was conscious in period 3 and one can easily see that $f(5|a_5) = \frac{2}{5} \leq \frac{3}{7} = f(3)$, making observation 5 linked with observation 3 and according to Proposition 1, revealing it was conscious too.

Consider now observation 7. We cannot directly link observation 7 to either observations 1, 3 or 4, because $f(7|a_7) = \frac{1}{2} > \max\{f(1), f(3), f(4)\}$. However, we can indirectly link observation 7 to observation 3 through observation 5, because $f(7|a_7) = \frac{1}{2} \leq \frac{1}{2} = f(5)$, thus making 7 an element of D^N . No other observation is indirectly linked to observations 1, 3 or 4 and hence, $D^N = \{1, 3, 4, 5, 7\}$. The method rightfully excludes all automatic choices from D^N .

The example presents inconsistencies in the revealed preference. Observation 3 and 6 are both in conflict with observation 2. That is, observations 2 and 3 and 2 and 6 form cycles. Then, noticing that $\max\{f(2), f(3)\} = f(2)$ and that $\max\{f(2), f(6)\} = f(2) = f(6)$ we can say that observations 2 and 6 are generated by the automatic self thanks to Proposition 1, given they are most familiar in a cycle.

But then, notice that observation 6 is linked to observation 8 given that $f(6|a_6) = \frac{3}{5} \leq f(8) = \frac{3}{5}$ revealing that the latter must have been automatically generated too. Thus, we get $D^C = \{2, 6, 8\}$ that were the observations rightfully excluded from D^N . No decision made by the maximizing self has been cataloged as automatic.

The modified revealed preference exercise helps us determine that alternative 1 is better than any alternative from 2 to 7, alternative 3 is better than any alternative from 4 to 6, and alternative 2 is better than alternatives 4, 7 and 8 as it is indeed the case. The value of the similarity threshold α by Proposition 1 can be correctly determined to be in the interval $[0.5, 0.6)$ thanks to the information retrieved from

observations 7 and 8 respectively.

D^N contains new observations and those indirectly linked to them.¹¹ It may be the case that some observations consciously generated are not linked to any observation in D^N , hence making D^N a proper subset of the set of all consciously generated observations. For this reason, nothing guarantees that $D \setminus D^N$ are observations automatically generated and hence, Proposition 1 needs to show how to dually construct a set of automatic decisions D^C . Nonetheless, if the observations are *rich enough*, it is possible to guarantee that D^N and D^C contain all conscious and automatic decisions, respectively.¹² More importantly, notice that Proposition 1 relies on one important assumption, that is, the collection of observations is generated by a DD process. The following section addresses this issue.

5 A Characterization of Dual Decision Processes

In Section 4 we showed how to elicit the preferences and the similarity threshold of an individual that follows a DD process. Here, building upon the results of that section, we provide a necessary and sufficient condition for a set of observations to be characterized as a DD process with a known similarity function. In other words, we provide a condition that can be used to falsify our model.

From the construction of the set D^N , we understand that a necessary condition for a dataset to be generated by a DD process is that the indirect revealed preference we obtain from observations in D^N , i.e. $R(D^N)$, must be asymmetric. It turns out that this condition is not only necessary but also sufficient to represent a sequence of decision problems as if generated by a DD process. One simple postulate of choice characterizes the whole class of DD processes.

Axiom 1 (Link-Consistency) *A sequence of observations $\{(A_t, e_t, a_t)\}_{t=1}^T$ satisfies Link-Consistency if, in every cycle there is at least one observation not indirectly linked to a new observation.*

¹¹ D^N is never empty because it always contains the first observation.

¹²Material available upon request.

This is a weakening of the Strong Axiom of Revealed Preference. In fact it allows for cycles but only of a particular kind. Preferential information gathered from observations in D^N cannot be inconsistent. The next theorem shows that such condition is indeed necessary and sufficient to characterize DD processes with known similarity.

Theorem 1 *A sequence of observations $\{(A_t, e_t, a_t)\}_{t=1}^T$ satisfies Link-Consistency if and only if there exist a preference relation \succ and a similarity threshold α that characterize a DD process.*

Proof. Necessity: If D is generated by a DD process, then it satisfies Link-Consistency as explained in the text.

Sufficiency: Now suppose that D satisfies Link-Consistency. We need to show that it can be explained as if generated by a DD process. Notice that Link-Consistency implies that the revealed preference relation defined over D^N , i.e. $R(D^N)$, is asymmetric. In fact, asymmetry of $R(D^N)$ means that it is not possible to construct cycles composed by observations in D^N . By standard mathematical results, we can find a transitive completion of $R(D^N)$, call it \succ . By construction, all decisions in D^N can be seen as the result of maximizing \succ over the corresponding menu.

Define $\alpha = \max_{t \in D^N} f(t)$. Notice that by definition of D^N , there is no observation $s \notin D^N$ such that $f(s|a_s) \leq f(t)$ for some $t \in D^N$. This implies that for all $s \notin D^N$, $f(s|a_s) > \alpha$, so, for all of them, it is possible to find a preceding observation they would seem to replicate. In particular, the one defining $f(s|a_s)$.

Thus, we can represent the choices as if generated by an individual with preference relation \succ and similarity threshold α . ■

The theorem is saying that Link-Consistency makes possible to determine whether the DM is following a DD process or not. In particular, when the property is satisfied, we can characterize the preferences of the DM with a completion of $R(D^N)$ which is asymmetric thanks to Link-Consistency and use the lower bound of α as described in Proposition 1 to characterize the similarity threshold. In fact, by construction, for any observation t outside D^N it is possible to find a preceding observation that can be replicated, i.e., the one defining $f(t|a_t)$.¹³ In the Appendix we show how the algorithm

¹³Notice that Link-Consistency implies that there have to be no cycles among new observation,

would work if the knowledge of similarity comparisons is partial. Finally, notice that we do not assume any particular structure for the sequence of observations we use as data and hence, the characterization of preferences does not have to be unique, even when the similarity is known.¹⁴

6 Dual Decisions Processes: General Implications

This section presents some applications of the model for various settings. The aim is to show that the model provides a simple possible explanation for some puzzling empirical regularities that have been highlighted in different fields. It does so very directly, without the need of being tailored too much to a particular application. In particular, only the interpretation of E and X will change. In fact, for all settings, we assume that E and X are subsets of some real space. Furthermore, for any $e \in E$, we assume there exist an optimal action $x(e) \in X$, that is different for different environments, that is, for any $e, e' \in E$ we have $x(e) \neq x(e')$. Moreover, let the similarity function be $\sigma(e, e') = \frac{1}{1+d(e, e')}$ where d is the euclidean distance between vectors e and e' .

The key idea on which the remainder of the section builds is that the model formalizes a novel way of understanding the coexistence of sticky and adaptive behavior. In section 6.1 the model is used to reconcile some puzzling evidence on doctors' behavior, in section 6.2 the model helps understanding new evidence on sticky pricing and finally in section 6.3 the model provides a possible interpretation of some empirical regularities highlighted in financial markets.

6.1 Dual Decision Processes and Treatments Prescription

Suppose doctors are described by a DD process. Let $E \subset \mathbb{R}^n$ be the compact set of symptom profiles. That is, $e \in E$ describes a unique list of symptoms that a patient could display. Furthermore, let $X \subset \mathbb{R}^n$ be the set of possible treatments.

a concept independent of the particular similarity function chosen, thus making the model always falsifiable.

¹⁴Material providing a full characterization of the procedure with rich enough data is available upon request.

The automatic self makes analogies between the symptom profile of a patient visiting the doctor and the symptoms profiles the doctor has seen in the past. If there is one that is similar enough, the doctor automatically assigns to the new patient the same treatment he gave in the past. Otherwise, the maximizing self is activated and the doctor prescribes the optimal treatment.

Our doctor has a memory composed by past cases, $(e, x) \in M \subseteq E \times X$, with the interpretation that the doctor has seen a patient with symptoms e and has given him treatment x . Let E^M be the projection of M in E . This means that E^M contains all those symptom profiles the doctor has experienced.

This model clearly oversimplifies the complexities of treatment prescription. First, it does not consider the uncertainty inherent to clinical decision making. Second, it does not include the cost of diagnosis nor the possibility of mistakes when using the maximizing self, hence abstracting from the positive effect of experience through learning. These assumptions would worsen the exposition without adding to the general results.

The first implication of the model is immediate.

Remark 1 *For every new patient, doctors described by a DD process, either categorize the patient with past similar ones or they give him a tailored treatment.*

This is in line with the evidence shown in Frank and Zeckhauser (2007). Their results highlight the fact that doctors cluster some patients giving them the same treatment while individually studying other cases. Moreover, when prescribing categorized treatments doctors are faster, i.e., the visits are shorter. This might be an indication that the cognitive process underlying treatment prescription when patients are categorized is automatic, that is, it is less cognitively costly for the doctor. Moreover, Frank and Zeckhauser (2007) show that such behavior is far from optimal and call for a behavioral explanation of the phenomenon.

A second, direct, implication is that the automatic self causes repetition of prescription behavior, or more simply habits, to emerge.

Remark 2 *Doctors described by a DD process show habits in treatment prescription.*

Any time the automatic self is activated past treatments are prescribed again to new patients. This implies that for diseases our doctor has already treated, even if new treatments are discovered, new patients affected by the same diseases will receive the outdated and suboptimal treatment. This is in line with the evidence. For instance Hamann et al. (2004) show that some psychiatrists continue to prescribe first generation anti-psychotics even though last generation ones are available.

Obviously, this phenomenon could be due to learning or uncertainty aversion on the side of doctors, that is, doctors might prescribe more those drugs they have previously prescribed and for which they know the effects. This kind of explanation, plausible in many instances, is in contrast with the evidence shown in Hellerstein (1998). She examines prescription behavior of physicians when deciding between generics and branded drugs. Even if the U.S. Food and Drug Administration refers to generic drugs as “chemical carbon copies” of branded drugs, doctors that have prescribed branded drugs tend to continue to do so with new patients. This behavior, at odds with learning, might be due to trust. Doctors and/or patients tend not to trust generics. Nevertheless, the latter interpretation cannot be easily reconciled with the evidence provided in Hamann et al. (2004) and Shrank et al. (2011) that show that this kind of *inconsistent* behavior seems to be correlated with the age of the doctor. Older doctors tend to prescribe older, e.g., branded, drugs. The next remark shows that this empirical regularity can be reproduced in our model.

Remark 3 *Suppose that patients arrive to the doctor following a continuous time independent distribution F with full support on E . As the memory of the doctor increases, the probability of using the automatic self increases.*

Proof. First, call $\mathcal{N}_{\frac{1-\alpha}{\alpha}}(e)$ the $\frac{1-\alpha}{\alpha}$ -neighborhood of $e \in E$, that is

$$\mathcal{N}_{\frac{1-\alpha}{\alpha}}(e) = \{e' \in E \mid \sigma(e, e') > \alpha\} = \left\{ e' \in E \mid d(e, e') < \frac{1-\alpha}{\alpha} \right\}$$

Second, let $N(E^M)$ be the union of $\frac{1-\alpha}{\alpha}$ -neighborhoods, i.e., $\mathcal{N}_{\frac{1-\alpha}{\alpha}}$ of all the cases

in E^M intersected with E . Formally,

$$N(E^M) = \bigcup_{e \in E^M} \mathcal{N}_{\frac{1-\alpha}{\alpha}}(e) \cap E,$$

that is $N(E^M)$ is the set containing all the cases in the doctor's memory plus all those that would be considered similar enough by the doctor. Notice that this set is increasing in E^M .

Finally, notice that the probability that a new patient is considered similar enough to one of the cases in the memory, i.e., the probability of activating the automatic self, is $F(N(E^M))$ which is increasing in its argument and the result follows. ■

At first glance Remark 3 seems to suggest that older doctors should incur stronger habits, in line with the evidence. Nevertheless, the proposition is really saying that this behavior is not due to age *per se* but to the experience of doctors of which age is a proxy. Thus, in theory, two doctors of different ages but with the same experience over one category of diseases should have the same prescription behavior. This is something to test in the data that we leave for further research.

The experience of the doctor is clearly central in the model. Doctors' behavior is path dependent, in the sense that, whenever the automatic self is activated, which treatment the doctor prescribes heavily depends on what he has seen in the past. This creates heterogeneity in behavior in line with the evidence that Chandra et al. (2011) present about variation in treatment prescription among different regions in the U.S.. They show that such variation cannot be completely explained by demand and supply factors, and so they call for a behavioral theory to account for the remaining variation. The model proposed here might be a possible explanation to consider. This problem is central in health economics because it translates in an impossibility for the government to control the expenditures in medical services across regions.

Finally, notice that the kind of implications analyzed in this section are not limited to doctors' behavior. Experts in general, like managers, military officers or political analysts, can be described in this way. The trade-off between positive and negative effects of expertise have to be explored more deeply. As Kahneman (2002) highlights, experts show pattern of behavior that are inconsistent and heavily affected by biases

like confirmation bias and availability heuristics that are intertwined with the way analogical thinking is formalized here.

6.2 Dual Decision Processes and Prices

Suppose retailers are described by a DD process. Let $E \subset \mathbb{R}$ be a compact set of perfect signals about the demand of the product sold. That is, $e \in E$ is a signal that perfectly describes the state of the market for the retailer. Moreover, let $X \subseteq \mathbb{R}$ be the set of possible prices.

Given this structure it is immediate to see that pricing can be *sticky*. Suppose retailers receive in period t the signal e_t about the state of the market and that signals are drawn from a continuous and time invariant distribution F with full support on E . The automatic self compares e_t with any market environment previously encountered. If the market is similar enough with what has been experienced, then the automatic self is active and pricing is replicated. On the other hand, if environments are dissimilar enough, prices are optimally adapted to the market. This is highlighted in the following remark.

Remark 4 *If retailers are described by a DD process, prices are not always optimally adapted to the market environment.*

There is ample evidence that prices are sticky. The usual explanation first proposed by Sheshinski and Weiss (1977) is that adapting prices to the market has a cost, called *menu cost*, making it suboptimal to constantly optimally modify prices. Prices are changed only when the benefits of the update compensate the costs. In a dynamic environment, the longer the inaction the higher the benefits of updating the price because the old pricing rule gets farther away from the optimum. Thus in general, the probability that a price adapts to the market situation should be higher the longer the period of inaction.

Here we propose a different channel that gives opposite predictions.

Remark 5 *Suppose a retailer described by a DD process has not changed the price for T periods, then, ceteris paribus, the probability of changing the price in period $T + 1$ is decreasing in T .*

Proof. The mechanism is quite straightforward and it is similar to the one behind Remark 3. In fact, we use the same notation adopted there. Let M_T be the set containing all signals the retailer has received for the T periods in which the price did not change. Notice that $N(M_T)$ is increasing with respect to the cardinality of M_T given that E is compact. Moreover, M_T is increasing in T . Given that the retailer did not change the price for T periods, we have only two possible scenarios due to the assumption that to each signal corresponds a unique and different optimal price.

- M_T is a singleton. That is, the retailer received always the same signal.
- M_T is not a singleton. Given the retailer stucked to one price this implies that he received a sequence of signals such that every new signal was similar enough with some other signal already in M_T .

Notice that the first case has probability zero, so it must be that the second one applies which implies that M_T is increasing in T . Thus, given the probability of changing the price in $T + 1$ is equal to $1 - F(N(M_T))$, the result follows. ■

The intuition behind the result is quite straightforward. If a retailer described by a DD process does not adapt his price to the market, it can only be because he is replicating past choices. Thus, the longer he sticks to a price, the more signals he has experienced that are similar enough to the ones he previously saw. The more time passes, the less is the probability of receiving a signal not similar enough to one of the signals already experienced.

Interestingly, this prediction of the model is in line with recent empirical evidence. Using weekly scanner data from two small US cities, Campbell and Eden (2014) find that the probability of a nominal adjustment of a price decreases with the *age* of the price change. That is, the older is the price change the less probable it is that it will change again. This result, confirmed with Mexican CPI data by Baley et al. (2016), is at odds with standard economic modeling of price rigidities as underlined before, but it can be accommodated with the model proposed here.¹⁵

Another empirical regularity that Campbell and Eden (2014) find is that the probability of a nominal adjustment of a price is highest when a store's price substantially

¹⁵See Baley and Blanco (2016) for an alternative explanation of this phenomenon in which menu costs are combined with learning.

differs from the average of other stores prices and that this happens with relatively *young* prices. This fact can be easily explained in our framework when certain conditions hold. If all retailers receive independent signals from the same distribution, an *extreme* signal, in the sense of it being unlikely, has a higher probability of being dealt with consciously thus leading to a price adjustment. Furthermore, it is more probable that it will be followed by another price adjustment because the successive signal is more likely to be less extreme hence activating the conscious self or the automatic self replicating a *more familiar* pricing decision.

A subtle implication of the model is that the more volatile is the market environment, the more frequently prices should be adjusted because it is higher the probability of facing environments different one from the other. This is in line with the evidence shown in Bils and Klenow (2004). The authors find that stickiness of prices varies greatly between markets and that in particular oil and commodities markets, that are usually quite volatile, tend to be the ones showing lower levels of stickiness.

Finally, the mechanism described here might also be used to study bank managers reacting in a similar way to changes in monetary policy or, as the next section highlights, traders reacting to new information in financial markets.

6.3 Dual Decision Processes and Financial Markets

Suppose traders are described by a DD process. Let $E \subseteq \mathbb{R}$ be the set containing perfect and public signals regarding the state of the market. For example, in a simple economy with one risky asset and one riskless asset, E might be the set containing the perfect signals that inform on the risky asset, e.g., the expected dividends.¹⁶ Moreover, let X contain portfolio compositions. In the example, $x \in X$ would say how much to invest on the risky and on the riskless asset.

Suppose traders in every period t receive a public signal e_t drawn from a publicly known distribution G . It is evident that not every trader will consider the signal when choosing the optimal portfolio. Some traders will use the signal while others will stick to past decisions. Clearly, depending on the distribution of α in the population and on

¹⁶See Cerigioni (2016) for a deeper analysis of this framework.

how demand is determined, signals can have different impact on the price of assets. To be more precise, let α be distributed in the population following a continuous distribution F with full support on $[0, 1]$, then we can state the following:

Remark 6 *Given a history of signals M , for every signal e there exists an α_e^M such that:*

- $1 - F(\alpha_e^M)$ traders will fully incorporate the signal into their trading decisions.
- $F(\alpha_e^M)$ traders will not incorporate the new information and stick to past trading decisions.

Proof. The reasoning behind the remark is quite straightforward. Trader i considers the signal e only if

$$\max_{e' \in M} \sigma(e, e') \leq \alpha_i.$$

Thus, let $\alpha_e^M = \max_{e' \in M} \sigma(e, e')$ and the result follows. In fact, given signals are public α_e^M is uniquely determined for the whole population. ■

Simply put, traders that have a similarity threshold high enough to consciously perceive the new signal will adapt their choices, the others will replicate past choices. This is in line with the evidence presented in the literature related to *underreaction* of prices in financial markets. Since Cutler et al. (1991) and Bernard (1993) much empirical evidence has been accumulated that shows that markets tend to underreact to news. That is, information is not fully and immediately incorporated into prices as the model proposed here predicts.

There is another feature of the model that is worth underlining.

Remark 7 *If the market is composed by traders described by a DD process, the random process determining the realization of present and future signals creates two sources of uncertainty for the demand of risky assets:*

- *Uncertainty regarding future movements of the fundamentals.*
- *Uncertainty regarding the fractions of traders that will consciously use new information.*

While the first source is standard in the literature, the second one is due to the model analyzed here and it has the potential of increasing the level of uncertainty in a given market thus increasing its volatility. This is in line with many empirical studies, e.g. Shiller (1992), and it can give new insights into the *equity premium puzzle* presented in Mehra and Prescott (1985). In fact, the authors show that traders, given the historical return on equities, tend to underinvest in risky assets with respect to what standard models would predict. To accommodate for the observed trading decisions, usual models should assume unrealistic levels of risk aversion. Here we propose a different framework that, by increasing the risk of the economy, might help explain this kind of evidence.

Notice that the behavior of traders described in this section can also be a good description of households and consumption decisions. Households described by a DD process do not always adapt consumption to information concerning permanent income thus creating, on the population level, the aggregate smoothness of consumption evident in the data, e.g. Carroll et al. (2011), but not predicted by classical models. Moreover, the model would produce smoothness of aggregate consumption without abandoning heterogeneity of household behavior. In fact, following Remark 6, some households optimally adapt to shocks as highlighted by Carroll and Slacalek (2006) and in contrast with models of habit formation.

7 Final Remarks

Cognitive sciences have highlighted the fact that choices can be divided into two categories; conscious choices and automatic ones. This hinders greatly the use of standard economic models that generally do not take into account the consequences of there being a difference between conscious and automatic choices.

In this paper, we proposed a possible formalization of this duality that allows us to restore standard revealed preference analysis. In order to do this, we have assumed that conscious behavior is the maximization of a given preference relation. In some cases this assumption can be too strong. Inconsistent choices might arise also when choosing consciously. Fortunately, the framework and analysis we developed here do

not depend on the particular conscious behavior that is assumed. In fact, for any given decision environment analogies determine a partition with two components, one containing those problems that are similar enough to the reference environment and another one containing those that are not. Such a partition is based only on one and simple assumption, automatic choices must come from the replication of past behavior.

Alternative conscious behaviors are possible. The only element of the formal analysis that has to be changed is the consistency requirement that needs to be tested on those problems in D^N . Obviously, dually, to find automatic choices we should analyze violations of such requirements. Similarly, the assumption that automatic behavior comes from the replication of past behavior is much less restrictive than what would appear at a first glance. Automatic choices have to be *familiar* choices. For example, we can consider cases where the DM stores in his memory not only his experiences, but also his parents', siblings' or friends' ones thus making the concept more general than what the literal interpretation of the model would suggest.

Furthermore, the model can be of use for welfare considerations. Obviously, restoring a revealed preferences approach is crucial for welfare analysis. Nevertheless, the model proposed here highlights something that is usually not considered, i.e., cognitive costs. The upper and lower bound of the interval in which the similarity threshold lies are defining a lower and upper bound of the cognitive costs of activation of the maximizing self, respectively. This can help welfare analysis but opens the question of how to relate such costs to utility that we leave to future research.

Finally, another point deserves further attention. Here we made some simplifying assumptions in order to develop a new framework encompassing automatic choices. In particular, we assumed that the similarity function is known and the environments are given. The former assumption can be weakened as the Appendix shows, nonetheless it is essential to know something about similarity comparisons. Automatic behavior can be consistent with the maximization of a given preference relation making the observation of inconsistent behavior not enough to correctly categorize decisions. Analogies add a layer of complexity to the problem that demands richer data. A first step in this direction is to address the second assumption previously mentioned,

i.e., that environments are known. Understanding what are the important elements in a decision problem for similarity comparisons, is something crucial for the study of individual decision making. The model proposed here provides a setting to think about this issue in a structured way but leaves the question open for future research.

A Appendix

A.1 Partial Information on the Similarity

In this section we show that our algorithmic analysis is robust to weaker assumptions concerning the knowledge of the similarity function. In particular, we study the case in which only a partial preorder over pairs of environments is known, denoted by S . For example, the Symmetric Difference between sets satisfies this assumption. Such extension can be relevant in many contexts where it is not possible to estimate the similarity function. In such cases, it is sensible to assume that at least some binary comparisons between pairs of environments are known. Coming back to the example we use in the introduction of the paper, we might not know how the DM compares different prices and dispositions of the products on the shelf, but we might know that for any combination of prices, a small change in just one price, results in a more similar environment than a big change in all prices.

We show here that, even if the information regarding similarity comparisons is partial, it is still possible to construct two sets that contain only conscious and automatic observations respectively, and that one consistency requirement characterizes all DD processes. In order to do so, we assume that, if the individual follows a DD process, the similarity σ cardinally represents a completion of such partial order. Thus, for any $e, e', g, g' \in E$, $(e, e')S(g, g')$ implies $\sigma(e, e') \geq \sigma(g, g')$ and we say that (e, e') dominates (g, g') . We first adapt the key concepts on which the algorithmic analysis is based in order to encompass this new assumption.

The two concepts of familiarity need to be changed. In particular, given that it is not always possible to define the *most* familiar past environment, the new familiarity definitions will be sets containing undominated pairs of environments. Let $F(t)$ and

$F(t|a_t)$ be defined as follows:

$$F(t) = \{(e_t, e_s) | s < t, a_s \in A_t \text{ and there is no } w < t \text{ such that } (e_t, e_w)S(e_t, e_s) \text{ and } a_w \in A_t\},$$

$$F(t|a_t) = \{(e_t, e_s) | s < t, a_s = a_t \text{ and there is no } w < t \text{ such that } (e_t, e_w)S(e_t, e_s) \text{ and } a_w = a_t\}.$$

That is, $F(t)$ and $F(t|a_t)$ generalize the idea behind $f(t)$ and $f(t|a_t)$, respectively. In fact, $F(t)$ contains all those undominated pairs of environments where e_t is compared with past observations which choice *could be* replicated. Similarly, $F(t|a_t)$ contains all those undominated pairs of environments where e_t is compared with past observations which choice *could have been* replicated. We can easily redefine the concept of link. We say that observation t is *linked to the set* of observations O whenever either $F(t|a_t) = \emptyset$ or for all $(e_t, e) \in F(t|a_t)$ there exists $s \in O$ such that $(e_s, e')S(e_t, e)$, for some $(e_s, e') \in F(s)$. Two things are worth underlining. First, notice that $F(t|a_t) = \emptyset$ only if t is new, thus, as in the main analysis, new observations are linked with any other observation. Second notice that this time we defined the link between an observation t and a *set* of observations O . This helps to understand whether an observation is generated by the maximizing self once we know that another observation is. If all observations in O are generated consciously and for each one of them there exists a pair of environments that dominates a pair in $F(t|a_t)$ then it must be that the maximizing self generated t too. This is because for all observation s in O , the similarity of *all* pairs of environments contained in $F(s)$ must be below the similarity threshold.

Then, we say that observation t is *Consciously-indirectly linked* to the set of observations O if there exists a sequence of observations t_1, \dots, t_k such that $t = t_1, t_k$ is linked to O and t_i is linked to $\{t_{i+1}, t_{i+2}, \dots, t_k\} \cup O$ for every $i = 1, 2, \dots, k-1$. Define $D^{\hat{N}}$ as the set containing all new observations and all those observations indirectly linked to the *set* of new observations. Proposition 2 shows that $D^{\hat{N}}$ contains only observations generated by the maximizing self.

What about automatic choices? As in the main analysis, whenever a cycle is present in the data, we know that at least one of the observations in the cycle must be generated by the automatic self. This time, given that we assume only a partial

knowledge of the similarity comparisons, it is not always possible to define a most familiar observation in a cycle.¹⁷ Nevertheless, notice that whenever an observation is inconsistent with the revealed preference constructed from $D^{\hat{N}}$, it must be that such observation is automatically generated. Thus, say that observation t is *cloned* if it is either a most familiar in a cycle or $xR(t)y$ while $yR(D^{\hat{N}})x$.

Say that observation t is *Automatically-indirectly linked* to observation s if there exists a sequence of observations t_1, \dots, t_k such that $t = t_1$, $t_k = s$ and t_i is linked to t_{i+1} for every $i = 1, 2, \dots, k-1$. Whenever we know that observation t is automatically generated, we can infer that observation s is generated by the automatic self too, only if for all pairs of environments in $F(t|A_t)$ there exists a pair of environments in $F(s)$ that dominates it. In fact, in general, only the similarity of *some* pairs of environments contained in $F(t|a_t)$ is above the similarity threshold. As before, let $D^{\hat{C}}$ be the set containing all cloned observations and the observations to which they are indirectly linked. Proposition 2 below shows that $D^{\hat{C}}$ contains only observations generated by the automatic self.

Proposition 2 *For every collection of observations D generated by a DD process where only a partial preorder over pairs of environments is known:*

1. *all decisions in $D^{\hat{N}}$ are generated by the maximizing self and all decisions in $D^{\hat{C}}$ are generated by the automatic self,*
2. *if x is revealed preferred to y for the set of observations $D^{\hat{N}}$, then $x \succ y$.*

Proof. First we show that any observation t linked to a set O of conscious observations must be generated by the maximizing self too. In fact, notice that for any $s \in O$ we know that $\sigma(e_s, e') \leq \alpha$ for all $(e_s, e') \in F(s)$. Then, given t is linked to O we know that for any $(e_t, e) \in F(t|a_t)$ there exists $s \in O$ such that $(e_s, e')S(e_t, e)$, for some $(e_s, e') \in F(s)$. Now, given the definition of $F(t|a_t)$ this implies that $\sigma(e_t, e_w) \leq \alpha$ for all $w < t$ such that $a_w = a_t$ and the result follows. Then, by Proposition 1 in the paper we know that new observations are consciously generated and applying the

¹⁷Obviously, in this context, a most familiar observation in a cycle would be an observation t belonging to a cycle such that for any other observation s in the cycle, $F(t)$ dominates $F(s)$. That is, for any $(e_s, e) \in F(s)$ there exists $(e_t, e') \in F(t)$ such that $(e_t, e')S(e_s, e)$.

previous reasoning iteratively it is shown that $D^{\hat{N}}$ must contain only observations generated by the maximizing self.

As a second step, we show that any observation s to which an observation t generated by the automatic self is linked, must be automatically generated too. Given t is generated by the automatic self it means that there exists a $w < t$ such that $\sigma(e_t, e_w) > \alpha$ and $a_w = a_t$. Then, either $(e_t, e_w) \in F(t|a_t)$ or $(e_t, e_w) \notin F(t|a_t)$.

- Let $(e_t, e_w) \in F(t|a_t)$. Then, given t is linked to s , there exists a pair $(e_s, e') \in F(s)$ such that $(e_s, e')S(e_t, e_w)$. This implies $\sigma(e_s, e') \geq \sigma(e_t, e_w) > \alpha$, and the result follows.
- Let $(e_t, e_w) \notin F(t|a_t)$. Then, there exists a $w' < t$ such that $(e_t, e_{w'})S(e_t, e_w)$ and $(e_t, e_{w'}) \in F(t|a_t)$. This implies that $\sigma(e_t, e_{w'}) > \sigma(e_t, e_w) > \alpha$. Then, given t is linked to s we know that for all $(e_t, e) \in F(t|a_t)$, there exists a $(e_s, e') \in F(s)$ such that $(e_s, e')S(e_t, e)$. In particular, there exists a $(e_s, e') \in F(s)$ such that $(e_s, e')S(e_t, e_{w'})$. This implies $\sigma(e_s, e') \geq \sigma(e_t, e_{w'}) > \alpha$, and the result follows.

Then, given that cloned observations are automatically generated, we can apply the previous reasoning iteratively to show that $D^{\hat{C}}$ must contain only observations generated by the automatic self.

Finally, by a reasoning similar to the one developed in the proof of Proposition 1 in the main text, given all observations in $D^{\hat{N}}$ must be consciously generated, $R(D^{\hat{N}})$ reveals the preference of the DM. ■

Thus, we see that knowing only a partial preorder does not heavily affect the structure of the algorithm and the main logical steps behind it. What is of interest is that even with this assumption it is possible to characterize a DD process with just one single condition, that is $D^{\hat{N}}$ -Consistency.

Axiom 2 ($D^{\hat{N}}$ -Consistency) *A sequence of observations $\{(A_t, e_t, a_t)\}_{t=1}^T$ satisfies $D^{\hat{N}}$ -Consistency whenever $xR(D^{\hat{N}})y$ implies not $yR(D^{\hat{N}})x$.*

$D^{\hat{N}}$ -Consistency imposes asymmetry on the revealed preference obtained from $D^{\hat{N}}$. If a sequence of decision problems satisfies $D^{\hat{N}}$ -Consistency when only a partial preorder

is known, then we are able to characterize the preferences of the individual, the similarity threshold and, more importantly, a similarity function that respects such preorder. This is what the next theorem states. Notice that S is assumed to be known.

Theorem 2 *A sequence of observations $\{(A_t, e_t, a_t)\}_{t=1}^T$ satisfies $D^{\hat{N}}$ -Consistency if and only if there exist a preference relation \succ , a similarity function σ representing S and a similarity threshold α that characterize a DD process.*

Proof. Necessity: Suppose that the sequence $\{(A_t, e_t, a_t)\}_{t=1}^T$ is generated by a DD process. Then it satisfies $D^{\hat{N}}$ -Consistency given that, according to Proposition 2, $D^{\hat{N}}$ contains only conscious observations and \succ is a linear order.

Sufficiency: Suppose that the sequence $\{(A_t, e_t, a_t)\}_{t=1}^T$ satisfies $D^{\hat{N}}$ -Consistency. We need to show that it can be explained as if generated by a DD process. Notice that $D^{\hat{N}}$ -Consistency implies that the revealed preference relation defined over $D^{\hat{N}}$, i.e., $R(D^{\hat{N}})$, is asymmetric. Thus, by standard mathematical results, we can find a transitive completion of $R(D^{\hat{N}})$, call it \succ . By construction, all decisions in $D^{\hat{N}}$ can be seen as the result of maximizing \succ over the corresponding menu.

We now define σ . We first complete S . Notice that by construction of $D^{\hat{N}}$, for all $t \notin D^{\hat{N}}$ there exists a pair $(e_t, e) \in F(t|a_t)$ such that there is no $s \in D^{\hat{N}}$ for which $(e_s, e')S(e_t, e)$, for some $(e_s, e') \in F(s)$. That is, for all observations not in $D^{\hat{N}}$ there exists a pair of environments that is not dominated by any pair of environments of observations in $D^{\hat{N}}$, a pair that we call *undominated*. Then, let S' be the following reflexive binary relation. For any undominated pair $(e_t, e) \in F(t|a_t)$ with $t \notin D^{\hat{N}}$, let for all $s \in D^{\hat{N}}$ and for all $(e_s, e') \in F(s)$, $(e_t, e)S'(e_s, e')$ and not $(e_s, e')S'(e_t, e)$. Let S'' be the transitive closure of $S \cup S'$. Notice that S'' is an extension of S that preserves its reflexivity and transitivity. Thus we can find a completion S^* of S'' and a similarity function $\sigma : E \times E \rightarrow [0, 1]$ that represents S^* .

Finally, we can define α . For any observation t , let $f^*(t)$ be as follows:

$$f^*(t) = \max_{s < t, a_s \in A_t} \sigma(e_t, e_s),$$

Then let $\alpha = \max_{t \in D^{\hat{N}}} f^*(t)$. Notice that by construction of σ for all $t \notin D^{\hat{N}}$ there

exists a pair of environments $(e_t, e) \in F(t|a_t)$ such that for all $s \in D^{\hat{N}}$, $\sigma(e_t, e) > f^*(s)$, hence $\sigma(e_t, e) > \alpha$. So, for every observation not in $D^{\hat{N}}$ we can find a preceding observation to imitate.

Thus, we can represent the choices as if generated by an individual with preference relation \succ , similarity function σ and similarity threshold α . ■

Intuitively, the observations in $D^{\hat{N}}$ are used to construct the preference relation of the individual. The similarity function represents a possible extension of the partial preorder that respects the absence of links between observations in $D^{\hat{N}}$ and the ones outside that set. This is possible thanks to how $D^{\hat{N}}$ has been constructed and it allows for the definition of the similarity threshold in a similar fashion as before.

References

- Alter, A. L., D. M. Oppenheimer, N. Epley, and R. N. Eyre (2007). Overcoming intuition: Metacognitive difficulty activates analytic reasoning. *Journal of Experimental Psychology: General* 136(4), 569.
- Ambrus, A. and K. Rozen (2015). Rationalising choice with multi-self models. *The Economic Journal* 125(585), 1136–1156.
- Apestequia, J. and M. A. Ballester (2015). A measure of rationality and welfare. *Journal of Political Economy* 123(6), 1278–1310.
- Baley, I. and J. A. Blanco (2016). Menu costs, uncertainty cycles, and the propagation of nominal shocks. *Working Paper*.
- Baley, I., F. Kochen, and D. Sámano (2016). Price duration and pass-through. *Mimeo-Banco de Mexico*.
- Bargh, J. A. (2005). Bypassing the will: Toward demystifying the nonconscious control of social behavior. In R. R. Hassin, J. S. Uleman, and J. A. Bargh (Eds.), *The New Unconscious*, pp. 37–58. Oxford: Oxford University Press New York.

- Bargh, J. A., P. M. Gollwitzer, A. Lee-Chai, K. Barndollar, and R. Trötschel (2001). The automated will: Nonconscious activation and pursuit of behavioral goals. *Journal of Personality and Social Psychology* 81(6), 1014.
- Bargh, J. A. and E. Morsella (2008). The unconscious mind. *Perspectives on Psychological Science* 3(1), 73–79.
- Bernard, V. L. (1993). *Stock Price Reactions to Earnings Announcements*, Volume 1. New York: Russell Sage Foundation.
- Bernheim, B. and A. Rangel (2009). Beyond revealed preference: Choice-theoretic foundations for behavioral welfare economics. *The Quarterly Journal of Economics* 124(1), 51–104.
- Bils, M. and P. J. Klenow (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy* 112(5), 947.
- Bisio, A., M. Casteran, Y. Ballay, P. Manckoundia, F. Mourey, and T. Pozzo (2012). Motor resonance mechanisms are preserved in alzheimer’s disease patients. *Neuroscience* 222, 58–68.
- Campbell, J. R. and B. Eden (2014). Rigid prices: Evidence from us scanner data. *International Economic Review* 55(2), 423–442.
- Carroll, C. D. and J. Slacalek (2006). Sticky expectations and consumption dynamics. *Mimeo-John Hopkins University*.
- Carroll, C. D., J. Slacalek, and M. Sommer (2011). International evidence on sticky consumption growth. *Review of Economics and Statistics* 93(4), 1135–1145.
- Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick (2009). Optimal defaults and active decisions. *The Quarterly Journal of Economics* 124(4), 1639–1674.
- Cerigioni, F. (2016). Dual decision processes and noise trading. *Working Paper*.

- Chandra, A., D. Cutler, and Z. Song (2011). Who Ordered That? The Economics of Treatment Choices in Medical Care. In *Handbook of Health Economics*, Volume 2, pp. 397–432. Elsevier B.V.
- Costa, A., A. Foucart, I. Arnon, M. Aparici, and J. Apesteguia (2014). “piensa” twice: On the foreign language effect in decision making. *Cognition* 130(2), 236–254.
- Cunningham, T. (2015). Hierarchical aggregation of information and decision-making. https://dl.dropboxusercontent.com/u/13046545/paper_heuristics.pdf.
- Cunningham, T. and J. de Quidt (2015). Implicit preferences inferred from choice. *Unpublished*.
- Cutler, D. M., J. M. Poterba, L. H. Summers, et al. (1991). Speculative dynamics. *Review of Economic Studies* 58(3), 529–46.
- Duhigg, C. (2012). *The Power of Habit: Why We Do What We Do in Life and Business*. New York: Random House.
- Epstein, S. (1994). Integration of the cognitive and the psychodynamic unconscious. *American Psychologist* 49(8), 709.
- Evans, J. S. (1977). Toward a statistical theory of reasoning. *Quarterly Journal of Experimental Psychology* 29, 297–306.
- Evans, J. S. (1989). *Bias in Human Reasoning: Causes and Consequences*. Hove and London: Lawrence Erlbaum Associates, Inc.
- Evans, J. S. and K. E. Frankish (2009). *In Two Minds: Dual Processes and Beyond*. Oxford: Oxford University Press.
- Evans, J. S. and D. E. Over (1996). *Reasoning and Rationality*. New York: Psychology Press.
- Frank, R. G. and R. J. Zeckhauser (2007). Custom-Made versus Ready-to-Wear Treatments: Behavioral Propensities in Physicians’ Choices. *Journal of Health Economics* 26(6), 1101–1127.

- Fudenberg, D. and D. K. Levine (2006). A dual-self model of impulse control. *American Economic Review* 96(5), 1449–1476.
- Ghika, J., M. Tennis, J. Growdon, E. Hoffman, and K. Johnson (1995). Environment-driven responses in progressive supranuclear palsy. *Journal of the Neurological Sciences* 130(1), 104–111.
- Gilboa, I. and D. Schmeidler (1995). Case-based decision theory. *The Quarterly Journal of Economics* 110(3), 605–639.
- Gilboa, I. and D. Schmeidler (2001). *A Theory of Case-Based Decisions*. Cambridge: Cambridge University Press.
- Gul, F. and W. Pesendorfer (2001). Temptation and self-control. *Econometrica* 69(6), 1403–1435.
- Hahn, U. (2014). Similarity. *Wiley Interdisciplinary Reviews: Cognitive Science* 5(3), 271–280.
- Hamann, J., B. Langer, S. Leucht, R. Busch, and W. Kissling (2004). Medical Decision Making in Antipsychotic Drug Choice for Schizophrenia. *The American Journal of Psychiatry* 161(7), 1301–1304.
- Hellerstein, J. K. (1998). The importance of the physician in the generic versus trade-name prescription decision. *RAND Journal of Economics* 29(1), 108–136.
- Hernandez, I. and J. L. Preston (2013). Disfluency disrupts the confirmation bias. *Journal of Experimental Social Psychology* 49(1), 178–182.
- Higgins, E. T. and J. A. Bargh (1987). Social cognition and social perception. *Annual Review of Psychology* 38(1), 369–425.
- Kahneman, D. (2002). Maps of bounded rationality: A perspective on intuitive judgment and choice. *Nobel Prize Lecture* 8, 351–401.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

- Kalai, G., A. Rubinstein, and R. Spiegel (2002). Rationalizing choice functions by multiple rationales. *Econometrica* 70(6), 2481–2488.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics* 112(2), 443–478.
- Lhermitte, F. (1986). Human autonomy and the frontal lobes. part ii: Patient behavior in complex and social situations: the “environmental dependency syndrome”. *Annals of Neurology* 19(4), 335–343.
- Manzini, P. and M. Mariotti (2007). Sequentially rationalizable choice. *American Economic Review* 97(5), 1824–1839.
- Manzini, P. and M. Mariotti (2014). Welfare economics and bounded rationality: The case for model-based approaches. *Journal of Economic Methodology* 21(4), 343–360.
- Masatlioglu, Y., D. Nakajima, and E. Y. Ozbay (2012). Revealed attention. *American Economic Review* 102(5), 2183–2205.
- McAndrews, M. P., E. L. Glisky, and D. L. Schacter (1987). When priming persists: Long-lasting implicit memory for a single episode in amnesic patients. *Neuropsychologia* 25(3), 497–506.
- Medin, D. L., R. L. Goldstone, and D. Gentner (1993). Respects for similarity. *Psychological Review* 100(2), 254.
- Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics* 15(2), 145–161.
- Oppenheimer, D. M. (2008). The secret life of fluency. *Trends in Cognitive Sciences* 12(6), 237–241.
- Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General* 118(3), 219.

- Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *The Economic Journal* 117(523), 1243–1259.
- Rubinstein, A. (2016). A typology of players: Between instinctive and contemplative. *The Quarterly Journal of Economics* 131(2), 859–890.
- Rubinstein, A. and Y. Salant (2012). Eliciting welfare preferences from behavioral dataset. *Review of Economic Studies* 79, 375–387.
- Salant, Y. and A. Rubinstein (2008). (a, f): Choice with frames. *The Review of Economic Studies* 75(4), 1287–1296.
- Schneider, W. and R. M. Shiffrin (1977). Controlled and automatic human information processing i: Detection, search and attention. *Psychological Review* 84(1), 1–66.
- Sheshinski, E. and Y. Weiss (1977). Inflation and costs of price adjustment. *Review of Economic Studies* 44(2), 287–303.
- Shiller, R. (1992). *Market Volatility*. The MIT Press.
- Shrank, W. H., J. N. Liberman, M. A. Fischer, C. Girdish, T. A. Brennan, and N. K. Choudhry (2011). Physician perceptions about generic drugs. *Annals of Pharmacotherapy* 45(1), 31–38.
- Simon, H. A. (1987). Making management decisions: The role of intuition and emotion. *The Academy of Management Executive* (1987-1989), 57–64.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies* 23(3), 165–180.
- Tulving, E. and D. L. Schacter (1990). Priming and human memory systems. *Science* 247(4940), 301–306.
- Tversky, A. (1977). Features of similarity. *Psychological Review* 84(4), 327.