Rational Status Quo*

Itzhak Gilboa[†] and Fan Wang [‡]

July 2018

Abstract

Decision makers often stick to a status quo without explicitly reconsidering it at every period. Yet, their observed choices might be compatible with a fully rational model. We ask when such a rationalization is possible. We assume as observable only the choice of sticking to the status quo vs. changing it, as a function of a database of cases. We state conditions on the set of databases that would make the decision maker change the status quo, which are equivalent to the following representation: the decision maker entertains a set of theories, of which one is that her current choice is the best; she is inert as long as that theory beats any alternative theory according to a maximum likelihood criterion.

1 Introduction

Some decisions are taken, and some just happen. Consider, for example, Mary, who wakes up in the morning and goes to work. She could have quit her job on this particular day. The fact that she didn't would typically be

*This paper subsumes "On Deciding When to Decide". Gilboa gratefully acknowledges ISF Grants 704/15 and 1077/17, the Sapir Center for Research in Economics, the Foerder Institute for Economic Research, the Investissements d'Avenir ANR -11- IDEX-0003 / Labex ECODEC No. ANR - 11-LABX-0047, AXA Chair and the Chair for Economic Theory. We thank the editor and three anonymous referees for comments and suggestions that greatly influenced the paper.

†HEC, Paris-Saclay, and Tel-Aviv University. tzachigilboa@gmail.com

‡HEC, Paris-Saclay. fan.wang1@hec.edu

modeled as a *decision* to stick to her job. If asked, Mary might say that no explicit, conscious decision process was involved; the thought simply hasn't crossed her mind that morning. Indeed, one might conjecture that most people, on most mornings, do not explicitly consider some of their weightier life choices, including the jobs they have, the partners they live with, and the countries they reside in. Still, while the observed choices to stick to the status quo are not explicitly or even consciously made, we tend to model them as decisions.

The fact that people do not consciously consider all their choices obviously does not mean that they behave irrationally. As economists, we are used to thinking of observed choice *as if* it were a result of conscious deliberation. Should Mary write down her options, list outcomes and possible scenarios, estimate utilities and probabilities, and calculate expected utility, she would probably find that it does not make sense to quit her job. Often she would figure out that the decision problem she faces has not changed dramatically since she last deliberated it. Thus, Mary can economize on the cognitive and emotional costs of decision making. A vague sense that the decision problem has not changed very much might be sufficient to make an implicit decision not to engage in explicit decision making. Further, it is possible that this implicit decision does not require awareness, or, differently put, that consciousness is summoned only when a sufficiently strong external impetus seems to require it.

Economists who espouse the rational choice paradigm can therefore explain the status quo bias as a rational decision to stick to an alternative that seems to be a best choice, until the arrival of information that convinces the decision maker that it isn't. Further, a wide range of behavior patterns can be rationalized if one takes into account sufficiently rich sources of information. Consider the famous studies on pension contributions (401(K)). Choi et al. (2002) and Carroll et al. (2009) showed the dramatic effect of changing a default decision, indicating that people may make sub-optimal decisions due to insufficient consideration of the options available. This phenomenon is justifiably considered to be an example in which behavioral economics had important policy implications. Moreover, there is no doubt that economists have not been sensitive to such phenomena before Tversky and Kahneman (1981) raised the issue of framing effects. At the same time, economists might argue that changing one's decision as a result of a change in the default option isn't irrational in any way. The way information is imparted to the decision maker is often informative in and of itself.[1] Indeed, rational players in a game are assumed to derive conclusions from the fact that they received certain signals and not others, from the incentives of those who chose which signals to send, and so forth.[2] Along similar lines, one might speculate that forms are generally designed by benevolent administration in such a way that the default option is the modal one. With this working assumption, an employee who does not have sufficient information about her optimal pension contribution might do wisely in imitating the modal choice. This form of social learning, or "cognitive free-riding" may or may not lead to an optimal decision. In particular, it can lead to information cascades and sub-optimal herd behavior. Yet, the very fact that a certain option was chosen as the default might carry information, and it might be "more rational" to consider this information than to ignore it.

In this paper we ask which status quo decisions can be rationalized along similar lines. We assume as observable the decision maker's choice between sticking to the status quo (S) and changing it (D), as a function of a *database*, which is a collection of *cases*. In the pension contribution problem, we would conceptualize the default as a proxy for many cases in which that choice was selected by others. Thus, we would interpret the choice of the default option as a result of social learning, and expect it to be made given databases that include relatively many such choices. Specifically, a database is modeled as

[1] The earliest reference to this point we are aware of, in the context of classical framing effects, was in personal communication by Ehud Kalai, in the early 1990s.

[2] This is often the point of Monty Hall 3-door puzzle when taught in classes.

a "counter vector", attaching an integer $I(c) \geq 0$ to each "case type" c, signifying the number of times cases of this type have been encountered.[3] The set of all conceivable databases is divided into two – the set in which the decision maker sticks to the status quo, $\mathcal{S}$, and its complement, $\mathcal{D}$, in which a different decision is made. We state conditions on $\mathcal{S}$ and $\mathcal{D}$, which are equivalent to the following representation: the decision maker entertains a set of "theories", T, each of which is a probability distribution over the case types. That is, for $t \in T$, $t(c)$ is the probability theory t assigns to a case of type c. A database I is in $\mathcal{S}$ if an only if

$$\prod_c [t_0(c)]^{I(c)} \geq \prod_c [t(c)]^{I(c)} \tag{1}$$

for all $t \in T$.

In this representation, theory t_0 can be interpreted as stating "the status quo choice is the best one available to me" whereas other theories t suggest that alternative courses of action are more promising. The decision maker is assumed to select a theory by the maximum likelihood criterion, and use it for decision making. Should (1) hold, theory t_0 can be selected as the most likely theory, and if it suggests the status quo decision, it makes sense to stick to the status quo.[4] In the default choice example discussed above, t_0 could be thought of as "choice a is optimal", while cases involve choices made by other, similar individuals. Assume that these individuals made the choices that were best for them, and that they have done so independently. Thus, the accumulation of many cases in which a was chosen by others would bolster the hypothesis that a is indeed optimal. If we further assume that the default choice is the modal choice in a group of similar individuals, one can rationalize the default effect as a form of social learning.[5]

[3] This is the set-up used in Gilboa and Schmeidler (2003, 2010).

[4] See a discussion of this mode of decision making, and its compatibility with the Bayesian one, in subsection 4.1.

[5] Clearly, if the default is based on modal choice by people who have not given the problem much thought, or who followed their colleagues, the implicit assumptions of optimality

Cases in our model are abstract entities that may correspond to various pieces of information. In the job search example, cases may involve information about the decision maker, such as positive or negative reviews she received; about others, such as the fact that another person changed her job; or about no one in particular, such as the news that there is higher demand for certain skills in certain markets. In an emigration decision example, a case might be the decision by another person to emigrate; a hate crime in one's country; a statement by a politician; and so forth. Thus, cases may or may not be accounts of similar decision problems (faced by the self or by others), and if they are, they may or may not specify outcomes and payoffs.

The representation (1) clearly implies that, for every I, $I \in \mathcal{D}$ if and only if $kI \in \mathcal{D}$ for all $k > 1$. Thus, rationalization of status quo decisions requires that the status quo decision be followed for a database I depending only on the relative frequencies of cases therein. Does it rule out anything else? Because in (1) the status quo is rationalized by a single theory t_0, it follows that for two databases $I, J \in \mathcal{S}$ we have to have also $I + J \in \mathcal{S}$. This condition suggests another test of a rational status quo: if each of two databases suggests that the current option is the best one, we cannot observe a deviation from the status quo given their union. Our main result is that these two conditions basically suffice for the representation we seek. Theorem 1 proves that, coupled with so-called "technical" assumptions, status quo decisions can be rationalized in our sense if and only if they satisfy these two conditions.

Thus, Theorem 1 suggests conditions on data that allow a rationalization of status quo choices, and characterizes those patterns of choice that are incompatible with optimal decision making. For example, assume that we observe people's decision to emigrate from their home country as a result of the arrival of information – about others emigrating, about political developments in the home country, and so forth. For this application it makes sense

and independence would be unwarranted.

to assume that only the binary decision to emigrate or not is observable, as data on final destinations may be unavailable (and perhaps only loosely related to the immigrants' plans). Are people who reside in a country just subject to a status quo bias, or do they rationally choose to live where they do? When we see a wave of emigration out of a country, is this a result of "dormant" decision maker who finally "wake up" and start thinking about their lives, or is it no more than rational information processing? Our result suggests an answer.

Observe that, as stated, our result is a bit simplistic as it does not allow for learning with noisy information. In the immigration example, suppose that database I contains a single case of a hate crime against an ethnic group. A single such case could be dismissed as an outlier. However, for a large k, kI would suggest that the country isn't safe. Thus, we could see $I \in \mathcal{S}$ but $kI \in \mathcal{D}$. While this is incompatible with the representation (1), it is hardly irrational. We therefore provide a more general result (Theorem 2) that allows for such patterns of behavior. It characterizes the sets $\mathcal{S}$ that can be represented by a set of theories as above, where each theory has a coefficient $d_t > 0$ such that I is in $\mathcal{S}$ if and only if

$$\prod_c [t_0(c)]^{I(c)} \geq d_t \cdot \prod_c [t(c)]^{I(c)} \qquad (2)$$

for all $t \in T$. This representation allows more freedom, and can captures situations in which a rational decision maker decides to switch from the status quo decision only with the accumulation of sufficient evidence. At the same time, as Theorem 2 shows, the representation is not vacuous. In particular, if a collection of news induces an immigration decision, it cannot happen that replicating these news would lead the decision maker back to the status quo decision.

To relate these results to the discussion above, assume that a behavioral economist says, "People are highly irrational. Some 50% of them fall prey to framing effects in their most important economic decisions. The 401(K) example shows that half of the employees make default-dependent choices." A more classically-oriented economist might respond, "This might be a wild overestimate. Some of the people you talk about might have been reacting very rationally to the information encapsulated in the definition of the default." The former might wonder is there anything that isn't compatible with such an account, and, relatedly, whether the rational approach has any empirical content. Possible answers are given by our characterization theorems. As Theorem 2 shows, even the more general representation (2) has refutable predictions. For example, if a database $I \in \mathcal{S}$ denotes a person's original information, and if she changes her behavior as a result of new information J (so that $I + J \in \mathcal{D}$), it is impossible that she would change her decision back as a result of additional information of the same type (that is, it is impossible that $I + kJ \in \mathcal{S}$ for $k > 1$). Thus, the classical economist can challenge the behavioral one to find such violations of the representation (2). We do not attempt to design such experiments, let alone to speculate about the scope of deviations from the rational choice paradigm. Our goal here is only to sharpen the debate, and provide testable conditions that are not easily explained by the rational information processing account.[6]

Our model can also be interpreted as a decision whether to make a decision. In this interpretation, the set $\mathcal{S}$ stands for the databases for which the decision maker makes an implicit decision not to make an explicit one, whereas in $\mathcal{D}$ she makes an explicit, conscious, decision. The latter may end up being identical to the status quo. In that case it might be hard to observe whether the decision maker sticks to the status quo without reconsidering it, or whether she deliberated the problem and decided not to change her choice. However, when applied to organization decisions, the decision whether to decide will typically be observable.[7] Consider, for example, an

[6]Observe that our representations retain the assumption of conditional independence of the cases given the theories. Clearly, one may relax this assumption and obtain more general accounts of rational behavior. However, at a certain level of generality the rational theory might indeed become vacuous. We discuss these issues in subsections 4.1 and 4.3.

[7]We thank an anonymous referee for suggesting this interpretation to us. Gilboa and

academic department that is running business as usual until a proposal for a new policy is made. In order to adopt the new policy the proposal will have to be discussed in a department meeting. In this case the decision whether to decide is a conscious, rational decision –but of a single individual, rather than the organization itself: the department chair has to decide whether to convene a meeting. Our question will then be, under which circumstances will a meeting be convened, and can we rationalize the decision to call for a meeting?[8]

The rest of this paper is organized as follows. The next sub-section is devoted to a discussion of related literature. Section 2 presents the model and the main result. Some extensions are presented in Section 3, while Section 4 contains a discussion.

1.1 Related Literature

Simon (1957) and March and Simon (1958) suggested the model of "satisficing", according to which a decision maker doesn't engage in conscious optimization, but rather sticks to her past choice as long as her objective performance is above a certain threshold, viewed as an aspiration level. It is well known that such patterns of behavior can also result from optimal Bayesian optimization in a multi-armed bandit problem, as in Gittins (1979). Our paper is similar in its motivation: starting with a commonplace observation, that people do not engage in active deliberation of each and every choice, we ask, under which conditions can an economist still assume that they do? However, as opposed to Simon's, our model does not describe the choices made once the status quo has changed. It is therefore better suited to deal with problems in which there are no data on alternative choices as (arguably)

Schmeidler (2012) and Cres, Gilboa, and Vieille (2012) model organization decisions using tools of case-based decision theory, but they do not deal with the decision when to decide.

[8] In the organizational interpretation, the decision whether to convene a meeting will typically be observable. Specifically, the organization's records would look differently if a meeting were not held or if a meeting were held and resulted in a decision to stick to the status quo.

in the emigration decision. Another distinction is that the cases in our model need not be past choices by the same decision maker in the same, or even similar problem. In particular, our model more readily captures phenomena such as social learning, contagion, etc.

Samuelson and Zeckhauser (1988) documented the "status quo bias", and also discussed its potential rationalizations. (See also Kahneman, Knetsch, and Thaler, 1991, discussing the status quo bias in comparison with loss aversion and the endowment effect.) Our main focus are status quo behaviors that may not involve an explicit decision process. In most experiments it appears obvious that participants were made to consider their decisions, and the status quo bias refers to an explicit decision to stick with the status quo. However, in many applications it might be difficult to tell whether a decision process was set into motion, resulting in a decision to stick with the status quo choice, or whether no such process took place to begin with. Our results, characterizing the status quo behaviors that can be rationalized in our sense, can thus apply to a wider set of applications, including daily choices as not-quitting one's job or not-emigrating from one's country. At the same time, they can also be applied to experiments where participants are practically forced to make explicit decisions.

Sims (2003) discusses "rational inattention", explicitly taking into account the cost of attention and its impact on consumption decisions. Along similar lines, Gabaix and Laibson (2002), Moscarini (2004), and Reis (2006) show that the cost of attention and computation can explain patterns in the data that are not readily explained otherwise. These models assume much more structure on the information that the decision maker can obtain than does our model, and they explicitly model the cost of information processing. By contrast, we merely ask which patterns of status quo bias can be rationalized by a particular model of decision making.

This line of research focuses on the choice of a continuous variable, such as a consumption bundle or a financial portfolio, and studies how frequently the choice is re-optimized in light of new information about the relevant economic variables.

Kahneman (2010) suggested the distinction between the fast, intuitive, “System 1 (S1)”, and the slow, more rational, “System 2 (S2)”. If the set $\mathcal{S}$ is interpreted as “deciding (implicitly) not to decide (explicitly)”, the distinction between $\mathcal{S}$ and its complement, $\mathcal{D}$, might overlap the S1/S2 distinction. Following habits and sticking with status quo decisions could be viewed as belonging to S1, especially because the rational decision process of S2 isn’t invoked in these cases. However, many of the examples of decisions made by S1 aren’t status quo decisions. Correspondingly, S1 might lead to decisions that would be considered mistaken for various reasons, such as biases and heuristics. Of these, our focus is only on the status quo bias, that is, on situation in which the only possible mistake is refraining from thinking.

Cerigioni (2017) offers a model of sequential decision making, in which, as long as a decision problem is similar to past ones, the decision maker makes the same choice she made in the past, and she engages in conscious optimization only for novel problems. This is a behavioral model of the two systems, S1, S2, and their interaction. Like Cerigioni (2017), we are interested in a decision maker who sometimes makes a conscious decision and sometimes doesn’t. Moreover, choices that retain the status quo, such as keeping one’s job or staying in one’s country, can be viewed as making the same choice that has been made in similar cases in the past. However, there are several differences between the models. First, our model allows the status quo choice to be a result of a conscious and explicit decision making process. Second, it is silent on the decision making process once it is set into motion. Third, our model allows for various cases to invoke conscious decision making, and not only the appearance of novel problems.[9] In sum, our model is not

[9] This is similar to the distinction between our model and satisficing: while Simon (1957) suggests that a decision maker would be prompted into making a decision by low payoffs, Cerigioni (2017) assumes that it is the similarity, rather than the payoff function, that is key to deliberation. Both, however, focus on past choices by the same decision

designed to identify the existence of S1 thinking from data, but to clarify what is actually assumed when status quo decisions are modeled as if they were conscious and rational.

2 Model and Representation

The set of case-types is C, assumed finite, with $|C| = n$. A database is a function $I : C \to \mathbb{Z}_+$, with $\mathcal{I} = \mathbb{Z}_+^n$ being the set of all databases. For $I \in \mathcal{I}$ define $|I| = \sum_{j=1}^{n} I(j)$. $0 \in \mathcal{I}$ denotes the origin, and summation, subtraction, multiplication by a number, and inequality on $\mathcal{I}$ are interpreted pointwise.

A set $\mathcal{S} \subset \mathcal{I}$ is given, interpreted as the set of databases for which the status quo is retained. We assume throughout that $0 \in \mathcal{S}$, that is, that in the absence of any data, the status quo will be retained. (This can be viewed as part of the definition of a "status quo".) The set $\mathcal{S}$ is considered to be the set of databases for which the status quo decision is chosen, whereas $\mathcal{D} = \mathcal{I} \backslash \mathcal{S}$ is interpreted as the set of databases for which a different decision is made.

We assume the Combination condition of Gilboa and Schmeidler (2003) for the set $\mathcal{S}$, but not for $\mathcal{D}$. This condition states that, if two disjoint databases lead the decision maker to keep the status quo, so should their union. In our formulation, the union is modeled by the sum of the two counter vectors that represent the databases:

C**1.** S**-Combination:** If $I, J \in \mathcal{S}$, then $I + J \in \mathcal{S}$.

The logic of this condition is the following: if a database I supports that conclusion that the current decision is the optimal one, and so does a database J, when they are taken together the conclusion can only be further reinforced. We do not assume a similar condition for the set $\mathcal{D}$, because it is possible that one database, I, gives reason to believe that an act a is better than the status quo, while another, J, that an act $b \neq a$ is better than the

maker.

status quo, but given the union, $I + J$, neither a nor b seems very promising, and the status quo decision can still appear to be a reasonable choice, or a safe compromise.

By contrast, if we add a database I to itself, it stands to reason that, if we start in $\mathcal{D}$ we will also end up in it. Thus we have

C**2.** **D-Replicability:** If $I \in \mathcal{D}$, then $kI \in \mathcal{D}$ for $k \in \mathbb{N}$.

Notice that $\mathcal{D}$-Replicability could follow from a condition similar to the Combination one, if I were restricted to be a replication of J. Clearly, given C1 and C2, for every $I \in \mathcal{I}$ and every $k \in \mathbb{N}$, $I \in \mathcal{S}$ iff $kI \in \mathcal{S}$ (and $I \in \mathcal{D}$ iff $kI \in \mathcal{D}$).

Both conditions C1 and C2 have a flavor of exchangeability: the accumulation of observations of a given type is assumed to strengthen evidence. Indeed, the representation we obtain, namely a likelihood function that is the product of conditional probabilities suggests that the cases are independent conditional on each of the possible theories. Some notion of conditional independence is certainly inherent both to our axioms and to the representation we obtain. It may be less restrictive than it may appear if the definition of a "case" is informative enough to include some summary statistics of its history. Clearly, including too much such information in the description of a case would render many databases logically incoherent.[10]

We also adapt an Archimedean condition from Gilboa and Schmeidler (2003):

C**3.** **Archimedeanity:** If $I \in \mathcal{D}$, then for every $J \in \mathcal{S}$ there exists a positive integer k such that $kI + J \in \mathcal{D}$.

Archimedeanity states that, if database I, on its own, contains sufficient information to dethrone the status quo decision, then, even if another data-

[10]For example, if a case is "driving a car without having to fill up its gas tank", many occurrences of the case make it less likely to occur again. But if we include the amount of gas in the tank as a state variable, independence becomes a reasonable assumption. At the same time, this state variable would make certain sequences of cases incoherent or hard to imagine.

base J does not, sufficiently many replications of I, when added to J, would do the same. In other words, if there is evidence in I that some other act, say a, is strictly better than the status quo, then, even if there is initial evidence J that favors the status quo, with sufficient accumulation of I-type evidence, the decision maker will favor a.

Notice that we do not require a similar condition for $I \in \mathcal{S}$. Basically, the Archimedean condition is an openness condition, and only one of the two sets would be a natural candidate to be open. As we think of $\mathcal{S}$ as defined by weak inequalities, it is natural to require Archimedeanity of its complement, $\mathcal{D}$, but not of $\mathcal{S}$ itself. (To consider an extreme example, note that for $I = 0 \in \mathcal{S}$ the corresponding condition would imply $\mathcal{D} = \varnothing$.)

A *theory* is a probability distribution of the case types. That is, it is a function $t : C \to [0,1]$ with $\sum_{c \in C} t(c) = 1$ where $t(c)$ is interpreted as the probability of case-type c occurring were t true. It is implicitly assumed that the probabilistic assessments are independent of past cases.

Theorem 1 *The following two statements are equivalent:*

(i) C1, C2, and C3 hold;

(ii) There exists a theory t_0 and a set of alternative theories T such that, for every $I \in \mathcal{I}$, $I \in \mathcal{S}$ iff

$$\prod_c [t_0(c)]^{I(c)} \geq \prod_c [t(c)]^{I(c)} \qquad \forall t \in T.$$

Thus, conditions C1 to C3 are equivalent to the representation we started out with, that is, to the claim that the decision maker behaves as if she were explicitly considering a set of theories about alternative paths of action, but dismisses them based on the data. Theory t_0 is interpreted as stating that the status quo choice is the best one, and theories $t \neq t_0$ – as identifying alternative choices as optimal.

Note that the theorem says nothing about uniqueness. It will be clear from the proof of the theorem that the theories correspond to hyperplanes, so that $\mathcal{S}$ is a convex cone which is obtained as the intersection of half-spaces defined by these hyperplanes. Generally, sets of such hyperplanes (defining a given set) will not be unique. If the set is a polytope, one may seek a minimal set of hyperplanes that is unique. But more generally there need not be a minimal set of such hyperplanes. [11]

3 Extensions

In the previous section we used two assumptions that, combined, implied that for all I and all natural k, $I \in \mathcal{S}$ iff $kI \in \mathcal{S}$. This property simplifies the analysis and, in particular, allows a relatively simple condition such as $\mathcal{S}$-Combination to imply the convexity of the set $\mathcal{S}$. However, the assumption that all that matters is the relative frequency of case types in a database, is somewhat restrictive. First, theories can differ in their complexity, the number of parameters they employ, and so forth. For example, the widely-used Akaike Information Criterion ranks theories by a linear function of their log-likelihood scores and the number of degrees of freedom they allow. Second (and relatedly), if there is an a priori preference for a given choice, it is possible that only a database of a certain size might outweigh it. For example, a person might like the country she lives in, and, in the absence of any information to the contrary (say, for $I = 0$), she may think that the status quo is the best choice for her. Next, assume that a hate crime against her ethnic group is performed. The case does not support the status quo decision;

[11] For example, suppose that $C = \{1, 2, 3\}$ and

$$S = \{0\} \cup \left\{ (I(1), I(2), I(3)) \neq 0 \mid \sum_{r=1}^{3} \left(\frac{I(r)}{I(1) + I(2) + I(3)} - \frac{1}{3} \right)^2 \leq \frac{1}{12} \right\}$$

that is, $\mathcal{S}$ consists of the vectors whose projection on the two-dimensional simplex is within a disk of the simplex center. Since a disk can be defined as the intersection of half-spaces defined by the supporting hyperplanes at any dense subset of its boundary, the set above can also be defined by any such collection of half-spaces. However, the set of such sets has no minimal elements.

in fact, it suggests that it might be better to emigrate. However, a single such case may be shrugged off as an aberration, so that for the database J containing this case alone, we still have $J \in \mathcal{S}$. At the same time, ten repetitions of such cases may be too much to ignore, and may convince the person it is time to leave. That is, we may have $10J \in \mathcal{D}$ even if $J \in \mathcal{S}$, and this should be expected when large databases are needed as in hypotheses testing.

To generalize our representation, we first ask a question that is asymptotic in nature: for $I \in \mathcal{S}$, is it the case that I supports the status quo only because of the paucity of data? To find out what is the information contained in I "in principle", we consider replicas of I. If, for some $k \geq 1$ we find that $kI \notin \mathcal{S}$, we would think of I as a database that, should essentially make the decision maker prefer another choice over the status quo. If, on the contrary, $kI \in \mathcal{S}$ for all $k \geq 1$, we may think of I as a database that is truly supports the status quo. Our first condition requires that at least one such database, which is strictly positive, exists:

$\mathcal{S}$-Conceivability: There is a database $I \in \mathbb{Z}^n_{++}$ such that $kI \in \mathcal{S}$ for all $k \geq 1$.

This condition assumes the existence of a database $I \in \mathbb{Z}^n_{++}$ that "truly" supports the decision maker's status quo decision. To be specific, if the decision maker deems every case-type possible, databases with the empirical frequency that exactly matches the decision maker's probabilistic assessment should be in $\mathcal{S}$. [12] Since a decision maker who abides by $\mathcal{S}$-Conceivability has to deem every case-type possible, we later discuss a relaxation of this condition that allows the decision maker to take some case-types as being impossible.

We say that a database I is *approximated by* $\mathcal{S}$ if for any $\epsilon > 0$, there is a finite collection of databases $\{I_i\}_{i=1}^m$ in $\mathcal{S}$ such that $\sum_{c=1}^n |I(c) - \sum_{i=1}^m \alpha_i I_i(c)| <$

[12] This implicitly assumes that the decision maker's probability assessments are rational numbers.

ϵ, where $\alpha_i > 0$ and $\sum_{i=1}^{m} \alpha_i = 1$.

$\mathcal{S}$-Closedness: Any database that is approximated by $\mathcal{S}$ is in $\mathcal{S}$.

$\mathcal{S}$-Closedness is assumed to make sure that there is no database in $\mathcal{D}$ that is also on the "boundary" of $\mathcal{S}$. In Theorem 1 this was guaranteed by the assumptions that implied that $\mathcal{S}$ and $\mathcal{D}$ are cones, combined with the Archimedeanity condition. In the present setup $\mathcal{S}$ and $\mathcal{D}$ need not be cones, and this requires a considerable strengthening of the Archimedeanity condition. The reason has to do with the fact that sets of integer vectors are messier objects than their continuous counterparts. We discuss this issue and our choice of the closedness condition in the appendix.

Theorem 2 *The following two statements are equivalent:*

1. $\mathcal{S}$ satisfies Conceivability and Closedness.

2. There exists a theory t_0 that matches the empirical frequencies of some database $I_0 \in \mathbb{Z}^n_{++}$, a set of alternative theories T, and constants $(d_t)_{t \in T}$ in $(0, 1]$ such that, for every $I \in \mathcal{I}$, $I \in \mathcal{S}$ iff

$$\prod_c [t_0(c)]^{I(c)} \geq d_t \cdot \prod_c [t(c)]^{I(c)} \qquad \forall t \in T.$$

This representation uses not only likelihood functions, but also a constant, d_t, for each theory t. This constant can be thought of as an a priori bias in favor of theory t_0, which is the "status quo theory". Observe that, in the log representation, the constant $\log(d_t) < 0$ can be interpreted as a theory-specific cost, as in Akaike Information Criterion. Along the tradition of classical statistics, d_t can be understood as a threshold chosen to control the chance of committing type-I error, as in the Neyman-Pearson lemma. In a Bayesian analysis, one can take d_t to be the ratio of the prior probability put on theory t to that of t_0. Geometrically, the fact that one can have these coefficients that need not all be 1 means that the supporting hyperplanes used to characterize $\mathcal{S}$ do not all go through the origin.

4 Discussion

4.1 Compatibility with the Bayesian Account

In our model the status quo decision is rationalized by selecting a single theory, interpreted as "the status quo is a best choice", according to the maximum likelihood criterion. This mode of decision-making differs from the standard subjective expected utility model, in that the decision maker is assumed to select a single theory rather than update her prior beliefs over all theories. It is a mode of decision making that is closer to classical than to Bayesian statistics: the decision maker can be thought of as using a "null hypothesis" t_0 as long as it is a maximum likelihood theory, and use only that theory.

However, this decision making procedure is compatible with subjective expected utility maximization. A simple way to embed our status-quo decision procedure in a Bayesian one is the following. Let there be given a set of theories T as in our model, with $t_0 \in T$. Assume for simplicity that T is finite. For each $t \in T$, let there be a possible act a_t and a state of the world, ω_t. Let the decision matrix have payoff 1 on the diagonal and 0 off the diagonal (that is, $u(a_t)(\omega_{t'}) = 1_{\{t=t'\}}$ for any two theories t, t'). In essence, the decision maker has to guess the correct theory t and she obtains a payoff of 1 iff she does so correctly (otherwise – 0). Assume that the prior probability over that states is uniform, and that the decision maker believes that, in state ω_t the probability of case c is $t(c)$ at each period, independently of history. Given the observations in I, the decision maker's posterior over $\{\omega_t\}_t$ is proportional to the likelihood function, and t_0 is the best guess if and only if it is a maximizer of that function.[13]

In general, the Bayesian approach can explain patterns of behavior that much richer than our model. First, it can allow for decision matrices that are

[13] A similar account would hold for an infinite set T, with the caveat that one needs to deal with an improper prior.

not diagonal as above. Indeed, a Bayesian decision maker need not choose the most likely theory and behave as if it were true; she should generally take into account all possible theories, weigh them by their posterior probabilities and calculate her subjective expected utility for each possible act. Second, the decision maker need not assume that the cases are independently drawn, conditional on the theory. Much more involved patterns can be imagined, including the possibility that cases are causally related (as in the example of herding).

In this sense, our rationalization of status quo biases is a "lower bound" of sorts: patterns of behavior that do not satisfy our conditions can still be rationalized by more involved models. However, it appears that one needs to impose some discipline on the Bayesian account in order to make sure it is meaningful. Green and Park (1996) show that, if we observe a decision maker's optimal choice given any possible (finite) sequence of payoff-irrelevant signals (corresponding to cases in our model), the only condition that is needed for compatibility with the Bayesian account is a version of the Sure Thing Principle: if we consider the tree of all possible sequences of signals, then an act that is optimal in every child of a given node should also be optimal at that node. Along similar lines, Shmaya and Yariv (2016) studied the evolution of beliefs over a state space Ω, as a function of such signals. Their conclusion is similar: for compatibility with the Bayesian account, all one needs is the (obviously necessary) condition that the posterior at a node be in the convex hull of the posteriors at its children.

Our model does not specify which acts are available to the decision maker, and certainly not which beliefs she might entertain. It thus allows many degrees of freedom, and renders the question of compatibility with the Bayesian approach, in its full generality, close to vacuous. To make the question more meaningful, we assume here a particular model in which cases are independent, given each theory, and where the most likely theory is selected. Other specific assumptions are certainly possible.

We observe in passing that in the presence of social learning, such as information cascades and herding, even if an agent is implicitly aware of the fact that cases are not conditionally independent, our conditions might still hold. Specifically, if a database I supports the status quo, but a database $I+J$ does not, it is not obvious that one can find a database $I+kJ$ (for $k>1$) that brings the agent back to the status quo decision. Such a pattern is, however, possible, with models in which one learns simultaneously both about the choices and about the information sources available. For example, assume that a case c is an endorsement of a product by a pop star. Repetition of this case, as in an advertisement campaign, would be expected to prod economic agents to quit the status quo in favor of the advertised product. However, a campaign that is too aggressive might seem suspicious, and a pop star who becomes known mostly through advertisements might lose popularity. In these cases we would expect agents to violate our conditions, though their choices can still be accounted for by a rational Bayesian model.

4.2 Multiple Choices

In our model we assume as observable only the binary choice of sticking with the status quo or not. In the immigration decision, this would mean that we know whether a person decided to emigrate from her country, but not where she decided to immigrate to. Similarly, we might have data on couples getting divorced, without any data on the personal lives of the divorcees.

Assume, however, that one can observe what people choices are instead of the status quo. For example, we might know where people chose to immigrate to, how many of the divorcees are in new relationships, and so forth. Clearly, uncertainty is involved. People might leave their country and end up in a different country than they one they envisioned. Similarly, people might get a divorce with the hope of getting into a new relationship, but end up being single. But, if we ignore this problem for the time being, we can imagine a model in which multiple alternatives to the status quo are observable, and ask, which databases would be consistent with rational learning, and which would indicate the presence of a status quo bias.

The simplest rational model would divide the set of databases $\mathcal{I}$ into convex cones, in each of which one act is considered optimal. The model would be very similar to those of Young (1975) and Myerson (1995) (see also Gilboa and Schmeidler, 2003). In such a model, it would be natural to assume that people tend to imitate the decisions of others, and "waking up" to change the status quo in the same way that others did would be perfectly compatible with rational learning. For example, if we see a wave of immigration from a given country to country A, it makes sense that individuals use the signal provided by others' choices and conclude it is a smart move. By contrast, if immigration to country A seems to result also in immigration to a different country B, which is very different from A, one might wonder whether this was a result of rational learning, or an example of an agent who was woken up to consider her status quo decision. Similarly, if we observe a rise in divorce rate in a given society at a given period, we can explain it by the fact that people learn from the experience of others, perhaps finding out that divorce isn't such a bad option. But suppose that we learn that John decided to divorce, quit his job, and join a spiritual sect, and, following that, his close friend Jim also divorced, left his job, and started to sail around the globe. Clearly, one could still come up with some rational account of Bayesian learning for Jim's decision, but such an account might seem awkward. It would probably be more reasonable to argue that such decisions are an indication of decision makers exhibiting a status quo bias simply because they are not making a conscious decision, and the fact that others do make such decisions prods them to consider their own choices.

Back to the example of pension contribution, one can argue that switching to the default option on a form is not a proof of irrational behavior, but switching to another option might be. To be concrete, suppose that employees have to select among three options, $\{a, b, c\}$, and that the status quo changed from a to b. If, as a result, some employees switch from a to c, this would be stronger evidence for an irrational status quo bias than would switching to b. Obviously, some Bayesian rationalization can still be provided for such a choice. Perhaps the fact that b is the default tells us something about c. In some situations, this could be rather reasonable indeed: assume that a, b, c are health plans, and that a offers no coverage for a certain disease, while b and c do, and the latter provides more extensive coverage than the former. Finding that the default changed from a to b, a rational employee could infer that people find this disease to be more likely than they used to. Given this signal about the uncertainty everybody faces, an individual might apply her own risk attitude and choose a more extensive coverage than do others. Thus, with sufficient information about the options a, b, c, a switch from a to c due to a change of the status quo changing from a to b need not be irrational. Yet, in many problems such a rational learning account would be dubious.

4.3 The Definition of Cases

Our definition of a "case" is rather flexible, allowing for instances of choices made by the decision maker in question, by other decision makers, as well as for pieces of information that do not involve decision making at all. The question naturally arises, isn't this notion too general? What is *not* a case? How does a modeler know what are the relevant pieces of information that should be included in the model?

The simple answer is that any piece of information might be included in the model, and if it turns out to be irrelevant, the model will be equivalent to a reduced one without that piece of information. More precisely, suppose that the modeler introduces into the set C a case c that is redundant. This will be reflected in the decision maker's set $\mathcal{S}$: for any two databases, I, I', such that $I(d) = I'(d)$ for all $d \neq c$, we will have $I \in \mathcal{S}$ iff $I' \in \mathcal{S}$. This means that the representation we obtain can be chosen to have all theories

$t \in T$ constant on c.[14]

This answer is obviously somewhat theoretical, as including anything imaginable in the model is hardly practical. Indeed, the practice of modeling would require some judgment of which cases might be relevant to the problem at hand. In this sense, modeling is more an art than a science, and the situation is similar to standard models of decision making. For example, when constructing a Savage-type model, one is faced with modeling choices, such as determining which states of the world to include (or which outcomes to consider). The general model is a "conceptual framework" allowing for specific theories to be developed within it. Theoretically, one may consider all conceivable states of the world, and let the decision maker's preferences identify null states, partitions that are too fine, etc. In practice, economists who use the conceptual framework use their judgment for determining a reasonable set of states that would leave the problem tractable, yet capture the essence of the phenomenon in question. Similarly, common sense and economic insight will be helpful in defining a set of cases in our framework, in a way that results in a tractable theory that captures the essence of the phenomenon in question.

4.4 Multiple Theories

Our representation results focus on a single theory t_0 that is supposed to be at least as likely as any other in T, for the status quo to be maintained. One may ask whether the status quo might be maintained as long as one of several theories is at least as likely as any other. Specifically, if the current choice is optimal given any theory in $T_0 \subset T$, the decision maker need not know which of the theories in T_0 is the maximizer of the likelihood function in order to conclude that she is doing the best she can. Such a generalization would replace the set $\mathcal{S}$ by a union of several (or perhaps infinitely many)

[14] In principle, the model can also be used to describe behavior given false information. Clearly "fake news" need not be irrelevant, and one may then consider a database I that becomes shorter with time, as a result of "deletion" of cases describing false information.

such sets, and would allow to consider status-quo sets that are not convex. Clearly, with too much freedom, the theory may become vacuous, say, if for every $I \in \mathcal{S}$ there is an element $t_0(I) \in T_0$ that is the most likely theory given I. Yet, putting some structure on the set T_0 or limiting its cardinality might lead to interesting generalizations.

5 Appendix: Proofs and Related Analysis

In the appendix, we refer to case-types by $c \in C$ or $i \in \{1, 2, ..., n\}$ depending on the context. A bar over a set notation denotes the closure of that set.

5.1 Proof of Theorem 1:

We start by presenting a lemma which may be of interest on its own.

$E \subset \mathbb{R}^n$ is an *algebraic neighborhood* of $x \in \mathbb{R}^n$, if for all $y \in \mathbb{R}^n$, there is an $\alpha \in (0, 1)$ such that $x + r(y - x) \in E$ for all $r \in [0, \alpha)$. $E \subset \mathbb{R}^n$ is *algebraically open*, if it is an algebraic neighborhood of every $x \in E$.

E is an algebraic neighborhood of x if and only if we can associate each point e in the unit sphere a positive number ϵ_e such that $\{y \in \mathbb{R}^n : ||y-x|| < \epsilon_{(y-x)/||y-x||}\} \subset E$. Notice that if the ϵ_e's can be chosen to be the same, E contains an usual open neighborhood around x.

Lemma 1 a). If $E \subset \mathbb{R}^n$ is an algebraic neighborhood around x, then x is not in the closure of any convex subset of E^c.

b). If $E \subset \mathbb{R}^n$ is algebraically open and E^c is convex, then E is open.

Proof. a). Due to the linear nature of the usual topology on $\mathbb{R}^n$, we can assume $x = 0$ without loss of generality. If 0 is in the closure of a convex subset F of E^c, there is a sequence $\{x_l\}$ in F converging to 0. Let m denote the dimension of the subspace spanned by F. We know $m > 1$, because $m = 1$ will imply that x_l are proportional, which contradict E being an algebraic neighborhood around 0. Pick $m + 1$ elements $\{y_i\}_{i=1}^{m+1}$ in F such that $conv(\{y_i\}_{i=1}^{m+1})$ has nonempty interior relative to this subspace spanned by F. This can be done, because otherwise F has to be on a hyperplane in this subspace that contains 0 (since $x_l \to 0$), which means the dimension of F has to be smaller than m. Pick an element y_0 and an open ball $B(y_0, \epsilon)$ of radius $\epsilon > 0$ around it in the interior of $conv(\{y_i\}_{i=1}^{m+1})$. By the convexity of F, we have $B(y_0, \epsilon) \subset F$. And, there is an $\alpha \in (0, 1)$ such that $\alpha y_0 \in E$.

For every l, the point $z_l = y_0 - \frac{1-\alpha}{\alpha} x_l$ satisfies $(1-\alpha)x_l + \alpha z_l = \alpha y_0$. Since $x_l \to x$, it must be the case that $z_l \in B(y_0, \epsilon)$ for large l. Then, for large l, both z_l and x_l are in F. This along with the convexity of F contradict $\alpha y_0 \in E$.

b). By a), every point $x \in E$ is not in the closure of E^c, indicating E is an open set. ∎

Corollary 1 *Suppose set $F \subset \mathbb{R}^n$ is convex and $x \notin F$. If for every $y \in F$ there is an $\alpha \in (0,1)$ such that $x + r(y-x) \notin F$ for all $r \in [0, \alpha)$, then x is not in the closure of F.*

Proof. The argument used in the lemma only requires the existence of α for directions coming from set F. ∎

Two observations about the Archimedeanity condition are due here. First, the Archimedeanity property also holds for $J \in \mathcal{D}$. Because, for any $J \in \mathcal{D}$, if $J + I \in \mathcal{D}$, we are done. If $J + I \in \mathcal{S}$, we apply the Archimedeanity condition to $J + I$ and get $(J+I) + kI \in \mathcal{D}$ for some $k \in \mathbb{N}$.

Second, for any $I \in \mathcal{D}$ and any $J \in \mathcal{I}$, there exists positive integer K such that for all $k > K$, $kI + J \in \mathcal{D}$. If $k_1 < k_2 < k_3$ and both $k_1 I + J$ and $k_3 I + J$ are in $\mathcal{S}$, we have $k_2 I + J = \alpha(k_1 I + J) + (1-\alpha)(k_3 I + J)$ for some $\alpha \in \mathbb{Q} \cap (0,1)$. Since there exists positive integers n and m such that $\alpha = \frac{m}{n}$, we have $n(k_2 I + J) = m(k_1 I + J) + (n-m)(k_3 I + J)$. Since $\mathcal{S}$-Combination implies $m(k_1 I + J) + (n-m)(k_3 I + J) \in \mathcal{S}$, it must be the case that $k_2 I + J \in \mathcal{S}$. Successive applications of Archimedeanity indicate that there are infinitely many ks such that $kI + J \in \mathcal{D}$. So, by the above argument, we can't have infinitely many ks such that $kI + J \in \mathcal{S}$.

We now proceed to prove Theorem 1. Since it is obvious that (ii) implies (i), we focus on (i) implying (ii). To begin with, we assume $\mathcal{S} \neq \{0\}$. The special situations $\mathcal{S} = \mathcal{I}$ and $\mathcal{S} = \{0\}$ are analyzed in the end.

We extend the $\mathcal{D}-\mathcal{S}$ partition from $\mathcal{I}=\mathbb{Z}_+^n$ to $\mathbb{Q}_+^n$. For any $I\in\mathbb{Q}_+^n$, find a $q\in\mathbb{Z}_+$ such that $qI\in\mathcal{I}$ and let $I\in\mathcal{D}$ if and only if $qI\in\mathcal{D}$, and $I\in\mathcal{S}$ if and only if $qI\in\mathcal{S}$. This is well defined because both $\mathcal{D}$ and $\mathcal{S}$ have the replicability property. We still call these two extended sets $\mathcal{D}$ and $\mathcal{S}$ respectively.

By $\mathcal{D}$-replicability and $\mathcal{S}$-combination, we have "$I\in\mathcal{D}$ implies $qI\in\mathcal{D}$, $\forall q\in\mathbb{Q}_+$", and "$I,J\in\mathcal{S}$ implies $pI+qJ\in\mathcal{S}$, $\forall p,q\in\mathbb{Q}_+$". This indicates $\mathcal{S}=conv(\mathcal{S})\cap\mathbb{Q}_+^n$, where $conv(\mathcal{S})$ is the convex hull of $\mathcal{S}$ in $\mathbb{R}_+^n$. To see the reason, notice that $\mathbb{Q}$ being a field indicates every point in $conv(\mathcal{S})\cap\mathbb{Q}_+^n$ can be obtained as a convex combination of points in $\mathcal{S}$ using only rational coefficients.

The Archimedeanity condition implies that given any $I\in\mathcal{D}$ and $J\in\mathbb{Q}_+^n$, there is a $\alpha\in(0,1)$ such that $I+r(J-I)\in\mathcal{D}$ for all $r\in[0,\alpha)\cap\mathbb{Q}$. This along with Corollary 1 and the fact that $\mathbb{Q}$ is a field together imply $\mathcal{D}\cap\overline{\mathcal{S}}=\emptyset$, where $\overline{\mathcal{S}}$ is the closure of $\mathcal{S}$ in $\mathbb{R}_+^n$.

$\overline{\mathcal{S}}$ is a closed convex cone. It is closed by definition and it is convex because it is also the closure of the convex set $conv(\mathcal{S})$. To see it is a cone, assume that $I\in\overline{\mathcal{S}}$ is the limit of a sequence $\{I_n\}$ in $\mathcal{S}$ and let q be in $\mathbb{Q}_+$. Since $\{qI_n\}$ must converge to qI, we have $I\in\overline{\mathcal{S}}$ if and only if $qI\in\overline{\mathcal{S}}$ for all $q\in\mathbb{Q}_+$. By the closedness of $\overline{\mathcal{S}}$, it must be the case that $qI\in\overline{\mathcal{S}}$ for all $q\in\mathbb{R}_+$. Hence, $\overline{\mathcal{S}}$ is a closed convex cone.

Given any $\beta\in\mathbb{R}^n$, let $c_\beta=inf_{x\in\overline{\mathcal{S}}}\sum_{i=1}^n x(i)\cdot\beta(i)$. Because $\overline{\mathcal{S}}$ is a cone, it must be the case that $c_\beta\in\{-\infty,0\}$. Let $B=\{\beta\in\mathbb{R}^n:||\beta||_1=1,c_\beta=0\}$, where $||\beta||_1=\sum_{i=1}^n|\beta(i)|$. The (supporting) half-space associated with a vector $\beta\in B$ is denoted by $H_\beta=\{x\in\mathbb{R}_+^n:\sum_{i=1}^n x(i)\cdot\beta(i)\geq 0\}$. It's well known that $\overline{\mathcal{S}}=\cap_{\beta\in B}H_\beta$. Thus, we conclude that, for all $I\in\mathcal{I}$, I is in $\mathcal{S}$ if and only if, for every $\beta\in B$ we have $\sum_{i=1}^n I(i)\beta(i)\geq 0$.

For the purpose above of characterizing $\overline{\mathcal{S}}$ by intersecting half-spaces, we only need to keep those $\beta\in B$ with the property that it has at least one strictly positive entry and at least one strictly negative entry. To see that, notice first that since we are looking at $\mathbb{Z}_+^n$, $\mathbb{Q}_+^n$, and $\mathbb{R}_+^n$, a half-space H_β associated with a non-negative β does not have any bite. Second, if β is non-positive, it can be approximated by other elements of B that have the property we want. To be explicit, if β is non-positive, it must have some entry(ies) equal to 0 since $\mathcal{S} \neq \{0\}$. For every $m \geq 2$, construct vector β_m by multiplying the strictly negative entries of β by the factor $\frac{m-1}{m}$ and allocating the positive weight $\frac{1}{m}$ uniformly to β's zero entries. Then $||\beta_m||_1 = 1$ for all $m \geq 2$, $||\beta_m - \beta||_1 \to 0$, and $\sum_{i=1}^n x(i)\beta_m(i) \geq \sum_{i=1}^n x(i)\beta(i)$ for any $x \in \mathbb{R}_+^n$, and thus we have $\beta_m \in B$ and $\cap_{m\geq 2} H_{\beta_m} = H_\beta$. With some abuse of notation, B is from now on understood as the set of the elements of the original B with the prescribed property.

The proof is completed if we can find a probability distribution $t_0 \in \Delta_n \equiv \Delta\{1, 2, ..., n\}$ such that for each $\beta \in B$, there is a probability distribution $t_\beta \in \Delta_n$ and

$$\ln(t_0) - \ln(t_\beta) \equiv \{\ln(t_0(c)) - \ln(t_\beta(c))\}_{c\in C} = k_\beta \cdot \beta$$

for some $k_\beta > 0$. This is because for each $\beta \in B$,

$$\sum_{i=1}^n I(i)\beta(i) \geq 0 \Leftrightarrow \prod_{i=1}^n [exp(k_\beta \cdot \beta(i)]^{I(i)} \geq 1 \Leftrightarrow \prod_{c\in C} [exp(\ln \frac{t_0(c)}{t_\beta(c)})]^{I(c)} \geq 1,$$

indicating that

$$I \in \mathcal{S} \Leftrightarrow \prod_{c\in C} t_0(c)^{I(c)} \geq \prod_{c\in C} t(c)^{I(c)}, \forall t \in T \equiv \{t_\beta\}_{\beta\in B}.$$

Notice that if we represent the set Δ_n as the following manifold P of log-likelihood

$$P = \{x \in \mathbb{R}_{--}^n : \sum_{i=1}^n exp(x^i) = 1\},$$

then our question can be reformulated as the following: is there a $\ln(t_0) \in P$ (and thus $t_0 \in \Delta_n$) such that for every β, $\ln(t_0) - k_\beta \cdot \beta \in P$ for some $k_\beta > 0$?

For any point $x = \{x^i\}_{i=1}^n$ on P, the hyperplane passing through x and tangent to P is the set $H_x = \{z \in \mathbb{R}^n : \sum_{i=1}^n z^i \cdot exp(x^i) = \sum_{i=1}^n x^i \cdot exp(x^i)\}$. Our question has a positive answer if we can find a $\ln(t_0) \in P$ such that: 1. $H_{\ln(t_0)}$ separates P and $\ln(t_0)+B$, and 2. the only intersection between $H_{\ln(t_0)}$ and $\ln(t_0) + B$ is the singleton $\{\ln(t_0)\}$. To see why, notice that for every $\beta \in B$, $\ln(t_0) - k_\beta \cdot \beta \in P$ for some $k_\beta > 0$ is equivalent to saying that the ray starting from $\ln(t_0) \in P$ with direction $-\beta$ will hit P again. Since each β has at least one strictly positive entry and at least one strictly negative entry, the ray starting from $\ln(t_0) \in P$ with direction $-\beta$ will eventually go to an orthant other than $\mathbb{R}^n_-$, in which process it will hit the strictly convex manifold P again if and only if its "opposite ray" will not, that is $\ln(t_0) + \beta$ and P are on different sides of $H_{\ln(t_0)}$.

Now we are ready to construct t_0 and $T \equiv \{t_\beta\}_{\beta \in B}$. Consider the following three situations.

Situation 1: suppose $\overline{\mathcal{S}}$ has full dimensionality, which happens when $\mathcal{S}$ has full dimensionality. Let $L = \{x \in \mathbb{R}^n_{++} : \sum_{i=1}^n x(i) = 1\}$ be the set of points in $\mathbb{R}^n_{++}$ whose L_1 norm is equal to 1. Since $\overline{\mathcal{S}}$ has full dimension, the intersection between its interior and L is nonempty, that is $int(\bar{\mathcal{S}}) \cap L \neq \emptyset$. Pick a point $I_0 \in int(\bar{\mathcal{S}}) \cap L$, we must have $\sum_{i=1}^n I_0(i) \cdot \beta(i) > 0$ for all $\beta \in B$. Notice that the tangent hyperplane $H_{\ln(I_0)}$ at $\ln(I_0) \equiv \{\ln(I_0(i))\}_{i=1}^n \in P$ is characterized by the gradient $\{exp(\ln(I_0(i)))\}_{i=1}^n = I_0$, indicating $H_{ln(I_0)}$ separates P and $(\ln(I_0) + B)$ and the only intersection between $H_{\ln(I_0)}$ and $(\ln(I_0)+B)$ is the singleton $\{\ln(I_0)\}$. Since these two conditions are satisfied, we can take I_0 as our t_0, and for every $\beta \in B$, we take the unique $t_\beta \in \Delta_n$ whose log-likelihood corresponds to $\ln(t_0) - k_\beta \cdot \beta \in P$ for some (which is also unique) $k_\beta > 0$.

Situation 2: suppose $\overline{\mathcal{S}}$ does not have full dimensionality but $\overline{\mathcal{S}} \cap L \neq \emptyset$. This happens when $\mathcal{S}$ does not have full dimensionality but every case-type is observed in $\mathcal{S}$, that is (by $\mathcal{S}$-convexity) $\mathcal{S}$ includes a database that contains a non-zero observation of each case-type. Fix a point $I_0 \in \overline{\mathcal{S}} \cap L$, and for every $m \geq 10$ define

$$\mathcal{S}_m \equiv \{x \in \mathbb{R}^n_+ : \exists y \in \overline{\mathcal{S}}, ||\frac{x}{||x||_1} - \frac{y}{||y||_1}||_1 < \frac{1}{m}\}.$$

Notice that $\overline{\mathcal{S}} = \cap_{m\geq 10}\overline{\mathcal{S}_m}$ and each $\overline{\mathcal{S}_m}$ fits in the first situation because it is a full-dimensional closed convex cone. And, since $I_0 \in int(\overline{\mathcal{S}_m}) \cap L$ for all $m \geq 10$, by the analysis of the first situation we have for all m a common theory $t_0 = I_0$ and a set of theories T_m that represents $\overline{\mathcal{S}_m}$. If we use the same $t_0 = I_0$ and take $T = \cup_{m,}T_m$, then we have for any database I, $\prod_c t_0(c)^{I(c)} \geq \prod_c t(c)^{I(c)}$ for all $t \in T$ if and only if $I \in \overline{\mathcal{S}_m}$ for all $m \geq 10$. Since $\overline{\mathcal{S}} = \cap_{m\geq 10}\overline{\mathcal{S}_m}$, t_0 and T represent $\overline{\mathcal{S}}$.

Situation 3: suppose $\overline{\mathcal{S}}$ does not have full dimensionality and $\overline{\mathcal{S}} \cap L = \emptyset$. This happens when some case-type is never observed in $\mathcal{S}$. Let $C_1 \subset C$ be the set of case-types that appear in $\mathcal{S}$ and $C_2 \equiv C \backslash C_1$ be the complement. The construction here takes 2 steps. For step one: a). if $\mathcal{S}$ restricted on the "subspace" $\mathbb{Z}^{C_1}_+$ coincides with $\mathbb{Z}^{C_1}_+$, then we take both t_0 and (the only) t_β to be the probability that assigns 0 probability to C_2 and distributes uniformly on C_1; b). if $\mathcal{S}$ restricted on the "subspace" $\mathbb{Z}^{C_1}_+$ does not coincide with $\mathbb{Z}^{C_1}_+$, we use the results from the above two situations to get t_0 and T in the "subspace" $\mathbb{Z}^{C_1}_+$, and extend it to $\mathcal{I} = \mathbb{Z}^n_+$ by setting the probability on C_2 to be 0 for t_0 and all $t_\beta \in T$. Here we need the custom that $0^0 = 1$. For step two, we add to T another theory t' constructed based on t_0 by bringing down the probability numbers t_0 assigns to C_1 by a factor of $\frac{1}{2}$ and allocating the left-over $\frac{1}{2}$ probability uniformly on C_2. This newly added t' will beat t_0 when cases in C_2 appear in a database, yet it won't have any impact on $\mathcal{S}$ restricted on $\mathbb{Z}^{C_1}_+$ because it can never fit better than t_0 on $\mathbb{Z}^{C_1}_+$.

For the situation $S = \mathcal{I}$, let us take t_0 be the uniform distribution over C and let the only alternative theory t be t_0 itself.

Finally, for the situation $\mathcal{S} = \{0\}$, let us take $t_0 \in \Delta_n$ such that $t_0(i) \in \mathbb{R}\backslash\mathbb{Q}$ for all $i \in \{1, 2, ..., n\}$, and take $T = \Delta_n \cap \mathbb{Q}^n$. Notice that for any $I \in \mathbb{Z}^n_+ \backslash \{0\}$ the maximum likelihood is uniquely achieved at the probability that is equal to the frequency, which is a vector of rational numbers. So, for any $I \in \mathbb{Z}_+^n \backslash \{0\}$, there is always a theory in T that beats t_0. ∎

5.2 Proof of Theorem 2

Proof: We first show the necessity of the conditions. Since t_0 matches the empirical frequency of some database $I_0 \in \mathbb{Z}_{++}^n$, the highest likelihood of I_0 is achieved at t_0, indicating $k \cdot I_0 \in \mathcal{S}$ for all $k \geq 1$. S-conceivability is thus valid. S-closedness is also valid because the system of inequalities in statement 2 characterizes $\{x \in \mathbb{R}_+^n : \forall t \in T, \sum_{i=1}^n x(i) \cdot \ln\left(\frac{t_0(i)}{t(i)}\right) \geq \ln(d_t)\}$, which is a closed convex set in $\mathbb{R}_+^n$.

Now we turn to the sufficiency of the conditions. Let us begin by defining $\mathcal{S}' \equiv \{I \in \mathcal{I} : \forall k \in \mathbb{Z}_+, k \cdot I \in \mathcal{S}\}$. Apparently, $\mathcal{S}' \subset \mathcal{S}$ and we have $\overline{conv(\mathcal{S}')} \cap \mathcal{D} = \emptyset$ by the closedness of $\mathcal{S}$. Depending on the property of $\mathcal{S}'$, we consider two situations.

First, suppose $\mathcal{S}'$ has full dimension. Take a dataset $I_0 \in \mathbb{Z}_{++}^n$ in the interior of $\overline{conv(\mathcal{S}')}$ and let $B' = \{\beta \in \mathbb{R}^n : ||\beta||_1 = 1, inf_{I \in \mathcal{S}'} \sum_{i=1}^n \beta(i)I(i) \neq -\infty\}$ represent the supporting hyperplanes of $\mathcal{S}'$. As in Theorem 1, for the rest of the proof we only need to consider those hyperplanes β that have at least one strictly positive entry and at least one strictly negative entry. To characterize $\mathcal{S}'$, by Theorem 1 we can take $t_0 = \frac{I_0}{||I_0||_1}$ and for each $\beta \in B'$ there is a $t_\beta \in \Delta_n$ such that $ln\,(t_0) - ln\,(t_\beta) = k_\beta \cdot \beta$ for some $k_\beta > 0$.

Let $B = \{\beta \in \mathbb{R}^n : ||\beta||_1 = 1, inf_{I \in \mathcal{S}} \sum_{i=1}^n \beta(i)I(i) \neq -\infty\}$ represent the supporting hyperplanes of $\mathcal{S}$ and define $c_\beta = inf_{I \in \mathcal{S}} \sum_{i=1}^n \beta(i)I(i) \leq 0$ for all $\beta \in B$. Then, we have $\overline{conv(\mathcal{S})} = \cap_{\beta \in B}\{x \in \mathbb{R}_+^n : \sum_{i=1}^n x(i)\beta(i) \geq c_\beta\}$. Notice that by $\mathcal{S}$-closedness, we have $\overline{conv(\mathcal{S})} \cap D = \emptyset$. Because $\mathcal{S}' \subset \mathcal{S}$, we must have $B \subset B'$. For each $\beta \in B$, let $d_\beta = exp(k_\beta \cdot c_\beta) \in (0, 1]$. Since for any $x \in \mathbb{R}^n_+$ and any $\beta \in B$, we have

$$\begin{aligned}\sum_{i=1}^{n} x(i)\beta(i) \geq c_\beta \quad &\Leftrightarrow \quad \prod_{i=1}^{n}[exp(k_\beta \cdot \beta(i)]^{I(i)} \geq exp(k_\beta \cdot c_\beta) \\ &\Leftrightarrow \quad \prod_{c \in C}[t_0(c)]^{x(c)} \geq d_\beta \prod_{c \in C}[t_\beta(c)]^{x(c)},\end{aligned}$$

the construction is thus valid.

Second, suppose $\mathcal{S}'$ does not have full dimension. For every $m \geq 10$, define

$$\mathcal{S}'_m \equiv \{z \in \mathbb{R}^n_+ : \exists z' \in conv(\mathcal{S}'), ||\frac{z}{||z||_1} - \frac{z'}{||z'||_1}||_1 < \frac{1}{m}\}$$

and let $\mathcal{S}_m = \overline{conv(\mathcal{S} \cup \mathcal{S}'_m)}$. Now we prove $\overline{conv(\mathcal{S})} = \cap_{m=10}^{\infty}\mathcal{S}_m$. Fix a point $x \notin \overline{conv(\mathcal{S})}$ and define $\eta \equiv inf_{y \in \overline{conv(\mathcal{S})}}||x-y||_1 > 0$ as the distance (in L_1 norm here and later) between x and $\overline{conv(\mathcal{S})}$. We want to show the distance between x and $conv(\mathcal{S} \cup \mathcal{S}'_m)$ is eventually positive, indicating x is eventually outside of $\mathcal{S}_m$. We only need to consider points in $conv(\mathcal{S} \cup \mathcal{S}'_m)$ that have L_1 norm no larger than $2||x||_1$, because other points are further away from x than $0 \in \mathcal{S}$. Any such point can be represented as $\alpha y + (1-\alpha)z$, where $\alpha \in [0, 1]$, y is a convex combination of points in $\mathcal{S}$, and z is a convex combination of points in $\mathcal{S}'_m \backslash conv(\mathcal{S})$. Because $||\alpha y + (1-\alpha)z||_1 \leq 2||x||_1$, we must have $1 - \alpha \leq \frac{2||x||_1}{||z||_1}$. Notice that by the definition of $\mathcal{S}'_m$ we can always decompose z as $z' + \epsilon$, where $z' \in conv(\mathcal{S}')$ and $||\epsilon||_1 < \frac{||z||_1}{m}$. This leads to

$$\begin{aligned}||x - \alpha y - (1-\alpha)(z' + \epsilon)||_1 \quad &\geq \quad ||x - (\alpha y + (1-\alpha)z')||_1 - (1-\alpha)||\epsilon||_1 \\ &\geq \quad \eta - (1-\alpha)\frac{||z||_1}{m} \geq \eta - \frac{2||x||_1}{m}\end{aligned}$$

where the first inequality is due to the triangle inequality, the second is due to $\alpha y + (1-\alpha)z' \in conv(\mathcal{S})$ and $||\epsilon||_1 < \frac{||z||_1}{m}$, and the third is due to $1-\alpha \leq \frac{2||x||_1}{||z||_1}$. Thus, x is eventually outside of $\mathcal{S}_m$ and we must have $\overline{conv(\mathcal{S})} = \cap_{m=10}^{\infty}\mathcal{S}_m$.

For each $\mathcal{S}_m$, let $B_m = \{\beta \in \mathbb{R}^n : ||\beta||_1 = 1, inf_{x \in \mathcal{S}_m} \sum_{i=1}^{n} \beta(i)x(i) \neq -\infty\}$ represent its supporting hyperplanes and let $c^m_\beta = inf_{x \in \mathcal{S}_m} \sum_{i=1}^{n} \beta(i)x(i) \leq$

0 denotes the "residual" of $\beta \in B_m$ on $\mathcal{S}_m$. Because $\mathcal{S}_{m+1} \subset \mathcal{S}_m$, we have $B_m \subset B_{m+1}$, and $c_\beta^m \leq c_\beta^{m+1}$ for any $\beta \in B_m \cap B_{m+1}$. Define $B \equiv \cup_{m=10}^{\infty} B_m$ and for every $\beta \in B$ define $c_\beta \equiv lim_{m\to\infty} c_\beta^m \leq 0$. Combined with the fact that $\overline{conv(\mathcal{S})} = \cap_{m=10}^{\infty} \mathcal{S}_m$, we have

$$\begin{aligned} x \in \overline{conv(\mathcal{S})} \quad &\Leftrightarrow \quad \sum x(i) \cdot \beta(i) \geq c_\beta^m, \forall \beta \in B, \forall m \geq 10 \\ &\Leftrightarrow \quad \sum x(i) \cdot \beta(i) \geq c_\beta, \forall \beta \in B \end{aligned}$$

Now fix a dataset $I_0 \in \mathbb{Z}_{++}^n$ in $\mathcal{S}'$. Since each $\mathcal{S}'_m$ has full dimension, by earlier result we have for each $\mathcal{S}_m$ a common $t_0 = \frac{I_0}{||I_0||_1}$ and a set of alternative theories $T_m = \{t_\beta\}_{\beta \in B_m}$. Apparently we have $T_m \subset T_{m+1}$ for all $m \geq 10$. For the unified representation, we still use this $t_0 = \frac{I_0}{||I_0||_1}$ and take $T = \{t_\beta\}_{\beta \in B} = \cup_{m=10}^{\infty} T_m$. And, for every t_β, let $d_\beta = exp(k_\beta \cdot c_\beta) \in (0, 1]$ where k_β is the positive real number such that $lnt_0 - lnt_\beta = k_\beta \cdot \beta$. This completes the whole construction. ∎

5.3 A Discussion of S-Closedness

As we explained in the main text, $\mathcal{S}$-closedness is used to make sure that $\mathcal{D} \cap \overline{conv(\mathcal{S})} = \emptyset$. It is arguable that the behaviorally nicer approach to this end is to follow the style of Theorem 1 by invoking some Archimedeanity condition. However, if $\mathcal{S}$ and $\mathcal{D}$ are not cones, doing this requires a notion of Archimedeanity that might be difficult for our decision maker to understand. This problem remains even if $\mathcal{S}$ is *integer-convex* in the sense that a database must be in $\mathcal{S}$ if it is equal to a convex combination of some databases in $\mathcal{S}$. We illustrate the problem through the following two examples.

Example 1: $\mathcal{I} = \mathbb{Z}_+^2$, $\mathcal{S} = \{I \in \mathcal{I} : I(1) = I(2)\} \cup \{(0, 1)\}$. Here we have $\overline{conv(\mathcal{S})} \cap \mathcal{I} = \{I \in \mathcal{I} : I(1) = I(2) \ or \ I(1) + 1 = I(2)\}$, which incorporates many databases in $\mathcal{D}$ into $\overline{conv(\mathcal{S})}$.

In this example, we have $(0, 1) \in \mathcal{S}$ and $(0, 1) + (1, 1) = (1, 2) \in \mathcal{D}$. Adding $(1, 1)$ to a database in $\mathcal{S}$ wakes up the decision maker, indicating $(1, 1)$ contains information that is surprising to the decision maker. But the fact that $k \cdot (1,1) \in \mathcal{S}$ for all $k \geq 1$ suggests otherwise. This example thus points to a natural strengthening of the Archimedeanity condition: if $J \in \mathcal{S}$ and $J + I \in \mathcal{D}$ for some database $I \in \mathcal{I}$, then for every $J' \in \mathcal{S}$ there exists a positive integer k such that $J' + kI \in \mathcal{D}$.

Unfortunately, the situations ruled out by this condition are not necessarily generic. Databases in $\mathcal{D}$ may be incorporated into $\overline{conv(\mathcal{S})}$ in subtler ways, as illustrated by the next example.

Example 2: Suppose $\mathcal{I} = \mathbb{Z}_+^3$ and consider three rays in $\mathbb{R}_+^3$: $l_1 = \{x \in \mathbb{R}_+^3 : x_2 = x_1, x_3 = 0\}$, $l_2 = \{x \in \mathbb{R}_+^3 : x_2 = 2x_1, x_3 = 0\}$, and $l_3 = \{x \in \mathbb{R}_+^3 : x_3 = \sqrt{2}x_1, x_2 = 0\}$. Now let $\mathcal{O} = conv(l_1, l_2, l_3) \cap \mathbb{Z}_+^3$, $\mathcal{O}' = (1,0,0) + \mathcal{O}$, and $\mathcal{S} = conv(\mathcal{O} \cup \mathcal{O}') \cap \mathbb{Z}_+^3$. Then, notice that $(1,0,1) \in \overline{conv(\mathcal{S})}$ while $(1,0,1) \in \mathcal{D}$.

In this example, we have $\frac{\sqrt{2}}{2}(0,0,0) + (1 - \frac{\sqrt{2}}{2})(1,0,0) = (1 - \frac{\sqrt{2}}{2}, 0, 0) \in conv(\mathcal{S})$ because $(0,0,0)$ and $(1,0,0)$ are in $\mathcal{S}$, and we have $(1 - \frac{\sqrt{2}}{2}, 0, 0) + (\frac{\sqrt{2}}{2}, 0, 1) = (1,0,1) \in \mathcal{D}$. So, adding $(\frac{\sqrt{2}}{2}, 0, 1)$ to a point in $conv(\mathcal{S})$ ends up in $\mathcal{D}$. This seems to suggest that databases with empirical frequency very close to $(\frac{\sqrt{2}}{2}, 0, 1)$ should be surprising to the decision maker in principle. However, the ray $\{k \cdot (\frac{\sqrt{2}}{2}, 0, 1) : k \in \mathbb{R}_+\}$ coincides with l_3, which is on the boundary of $conv(\mathcal{S})$.

Example 2 suggests a further strengthening of the Archimedeanity condition. To state it, we define two new notions. A *limiting frequency* is any vector $\sigma \in \mathbb{R}_+^n$ such that $||\sigma||_1 = 1$. We say a limiting frequency σ is *approximated* by $\mathcal{S}$, if every point on the ray $\{k \cdot \sigma : k \in \mathbb{R}_+\}$ is approximated by $\mathcal{S}$.

Strong Archimedeanity: If $x \in conv(\mathcal{S})$ and $x + y \in \mathcal{D}$ for some $y \in \mathbb{R}_+^n$, then the limiting frequency $\frac{y}{||y||_1}$ is not approximated by $\mathcal{S}$.

We believe that this condition is probably too complicated for a decision maker to understand, and we opt for the more direct $\mathcal{S}$-closedness condition to retain the rhetorical nature of this paper.

5.4 Zero Probability and Theorem 2

For the case-types that the decision maker deems possible, databases with the empirical frequency that exactly matches the decision maker's probabilistic assessment should be in $\mathcal{S}$. Indeed, these databases exactly fit the decision maker's expectation. And, if I is such a database, then for any $k \geq 1$ the database $k \cdot I$ should also be in $\mathcal{S}$ because the empirical frequency is the same. From this understanding, it is reasonable to say that the set of case-types that the decision maker deems possible is a subset of $C_0 = \{c \in C : \exists I \in \mathcal{S}, I(c) > 0, and\ \forall k \geq 1, kI \in \mathcal{S}\}$. For any case type in $C \backslash C_0$, the decision maker deems it impossible and any observation of it should be able to alert the decision maker. The following condition formalize this intuition.

$\mathcal{S}$-Consistency: $C_0 = \emptyset$ and a database has to be in $\mathcal{D}$ if it has a non-zero observation of some case-type in $C \backslash C_0$.

Notice that if we have a database $I \in \mathbb{Z}^n_{++}$ such that $k \cdot I \in \mathcal{S}$ for all $k \geq 1$, $\mathcal{S}$-consistency automatically holds because $C_0 = C$.

Corollary 2 *The following two statements are equivalent:*

1. $\mathcal{S}$ satisfies consistency and closedness.

2. There exists a theory $t_0 \in T$ that matches the empirical frequency of a nonzero database, a set of alternative theories T in which at least one has full support, and constants $(d_t)_{t \in T}$ in $(0,1]$ such that, for every $I \in \mathcal{I}$, $I \in \mathcal{S}$ iff

$$\prod_c [t_0(c)]^{I(c)} \geq d_t \cdot \prod_c [t(c)]^{I(c)} \qquad \forall t \in T.$$

Proof: Since the necessity of the conditions is obvious, we only prove their sufficiency.

Theorem 2 holds on the subspace $\{I \in \mathcal{I} : I(c) = 0, \forall c \in C \backslash C_0\}$ by setting the probability on $C \backslash C_0$ to be 0 for t_0 and all $t \in T$. If $C_0 = C$, we are done (notice that all $t \in T$ in Theorem 2 has full support). If not, we add to T a theory t_0': $t_0'(c) = \frac{1}{2} t_0(c)$ for all $c \in C_0$ and $t_0'(c) = \frac{1}{2(n - \#C_0)}$ for all $c \notin C_0$, where $\#C_0$ denotes the number of case-types in C_0. ∎

6 References

Akaike, H. (1974), "A new look at the statistical model identification". *IEEE Transactions on Automatic Control* **19** (6), 716–723.

Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick (2009), "Optimal Defaults and Active Decisions", *Quarterly Journal of Economics*, **124**: 1639-1674.

Cerigioni, F. (2017), "Dual Decision Processes: Retrieving Preferences when some Choices are Automatic".

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2002), "Defined Contribution Pensions: Plan Rules, Participant Choices, and the Path of Least Resistance", *Tax Policy and the Economy*, **16**: 67-113.

Cres, H., I. Gilboa, and N. Vieille, "Bureaucracy in Quest for Feasibility", 2012.

Gabaix, X. and D. Laibson (2002), "The 6D Bias and the Equity Premium Puzzle", *NBER Macroeconomics Annual 2001*, 257-311.

Gilboa, I. and D. Schmeidler (2003), "Inductive Inference: An Axiomatic Approach", *Econometrica*, **71**: 1-26.

——— (2010), "Likelihood and Simplicity: An Axiomatic Approach", *Journal of Economics Theory*, **145**: 1757-1775.

——— (2012), "Authorization Decisions".

Gittins, J. C. (1979), "Bandit Processes and Dynamic Allocation Indices", *Journal of the Royal Statistical Society B*, **41**: 148-177.

Green, E. J. and Park, I.-U. (1996), "Bayes Contingent Plans", *Journal of Economic Behavior and Organizations*, **31**: 225-236.

Kahneman, D. (2011), *Thinking, Fast and Slow.* Farrar, Straus, and Giroux.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1991), "The Endowment Effect, Loss Aversion, and the Status Quo Bias", *Journal of Economic Perspectives*, **5**: 193-206.

March, J. G. and H. A. Simon (1958), *Organizations.* New York: Wiley.

Moscarini, G. (2004), "Limited Information Capacity as a Source of Inertia", *Journal of Economic Dynamics and Control* **28**: 2003-2035.

Myerson, R. B. (1995), "Axiomatic Derivation of Scoring Rules Without the Ordering Assumption", *Social Choice and Welfare*, 12, 59-74.

Reis, R. (2006), "Inattentive Consumers", *Journal of Monetary Economics*, **53**: 1761-1800.

Samuelson, W. and R. Zeckhauser (1988), "Status Quo Bias in Decision Making", *Journal of Risk and Uncertainty*, **1**: 7-59.

Savage, L. J. (1954), *The Foundations of Statistics.* New York: John Wiley and Sons. (Second addition in 1972, Dover)

Shmaya, E., and L. Yariv (2016), "Experiments on Decisions Under Uncertainty: A Theoretical Framework", *American Economic Review*, **106**: 1775-1801.

Simon, H. A. (1957), *Models of Man.* New York: John Wiley and Sons.

Sims, C. A. (2003), "Implications of Rational Inattention", *Journal of Monetary Economics*, **50**: 665-690.

Tversky, A. and D. Kahneman (1981), "The Framing of Decisions and the Psychology of Choice". *Science*, **211**: 453–458.

von Neumann, J. and O. Morgenstern (1947), *Theory of Games and Economic Behavior.* Princeton: Princeton University Press.

Young, H. P. (1975), "Social Choice Scoring Functions", *SIAM Journal of Applied Mathematics*, 28: 824-838.