

Experiments on Decisions under Uncertainty: A Theoretical Framework[†]

By ERAN SHMAYA AND LEEAT YARIV*

The analysis of lab data entails a joint test of the underlying theory and of subjects' conjectures regarding the experimental design itself, how subjects frame the experiment. We provide a theoretical framework for analyzing such conjectures. We use experiments of decision making under uncertainty as a case study. Absent restrictions on subjects' framing of the experiment, we show that any behavior is consistent with standard updating ("anything goes"), including those suggestive of anomalies such as overconfidence, excess belief stickiness, etc. When the experimental protocol restricts subjects' conjectures (plausibly, by generating information during the experiment), standard updating has nontrivial testable implications. (JEL C91, D11, D81, D83)

Experiments studying behavior under uncertainty (be it regarding some underlying state, such as income in consumption-saving problems, or regarding others' behavior, as is the case in practically all strategic interactions) usually consist of three stages. First, the uncertainty is realized; second, subjects are provided with partial information about that realization; third, subjects choose an action. When analyzing data of such experiments, one essentially tests joint hypotheses regarding *responses* to the experimental design (namely, the realization of uncertainty and the information generation procedure) and subjects' *beliefs* about the design itself. The focus of this paper is the analysis of the potential consequences of such joint tests. We allow subjects to hold arbitrary conjectures about the experimental design and identify links between classes of conjectures subjects hold and the testable implications of theoretical predictions in the lab.

The idea that subjects may form (potentially inaccurate) beliefs about an experiment's design to which they respond is long standing in the social sciences. It falls under the general umbrella of *experimenter demand*, suggesting that the design itself makes subjects (consciously or subconsciously) frame the problem at hand

*Shmaya: Managerial Economics and Decision Sciences (MEDS), Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208 (e-mail: e-shmaya@kellogg.northwestern.edu); Yariv: Division of the Humanities and Social Sciences, Caltech, 1200 E. California Boulevard, MC 228-77, Pasadena, CA 91125 (e-mail: lyariv@hss.caltech.edu). We thank Alessandro Lizzeri for numerous comments on an earlier version of the paper. Christopher Chambers, Federico Echenique, and Matthew O. Jackson provided very helpful suggestions. Financial support from the National Science Foundation (SES 0963583) and the Gordon and Betty Moore Foundation (grant 1158) is gratefully acknowledged.

[†]Go to <http://dx.doi.org/10.1257/aer.20120978> to visit the article page for additional materials and author disclosure statement(s).

in a particular way that makes them believe certain responses are more appropriate than others. It has not only been a concern for economists (Zizzo 2010, and references therein); it is also considered a potential source of distortion in responses to psychological surveys (see Paulhus 1991) and a channel through which subjects in sociological experiments choose their actions (see Gillespie 1991). Experimenter demand may even play a role in the generation of placebo effects in medicine (see Beecher 1955 and work that followed). Despite its importance for experimental deductions, the literature has offered little in the form of a theoretical framework for inspecting the testable implications of experiments accounting for experimenter demand effects. Our goal is to take a step in that direction.

We concentrate on a general case study of laboratory decision making under uncertainty. Most experiments in this class involve some form of updating. Consequently, a natural first step in the theoretical analysis of such experiments pertains to experiments having to do with elicitation of beliefs, which are the experiments we study in this paper.

We consider experiments in which payoffs depend on some unknown state. After the state is realized, a subject is provided with some information regarding the state in the form of a sequence of potentially informative signals. Ultimately, the subject chooses one of several alternatives. Many extant experiments fall under this category (particularly when thinking of other subjects' actions as part of the state). For instance, consumption-saving experiments, individual experimentation and learning experiments, herding, and sequential voting experiments all naturally fall under this rubric (see, for example, Kagel and Roth 1997 for an overview of experiments in different fields of economics).

A Motivating Example.—Before describing our formal results, consider the following example, providing a simple caricature of the structure of experiments of decision making under uncertainty.

There are two states of nature, a and b , realized at the outset of the experiment. For example, a and b can stand for whether or not going on vacation is a good idea, as in Tversky and Shafir's (1992) disjunction effect experiments. The subject's goal is to report which of the states is most likely given different amounts of information.¹ For simplicity, assume that initially, she is given no information. Then she receives some information, say in the form of a signal taking the values u or d . For example, in Tversky and Shafir's (1992) experiments, subjects were asked to contemplate either failing or passing an important exam. Figure 1 captures, up to relabeling, the three types of responses one could observe in such an experiment. The actions at the roots of each tree correspond to the reports absent any signals, while actions at the other two nodes correspond to the subject's responses following the respective signals.

In panel A, the subject does not change her mind regardless of which signal she observes. This is consistent with Bayesian updating. For instance, the subject may have a strong prior favoring the state a and view the signals as uninformative. In panel B, the subject changes her mind when the signal is d . This can also be

¹For instance, she may be paid a fixed amount only when her response matches the realized state.




FIGURE 1. SIMPLE REVERSAL

explained with Bayesian updating by assuming, e.g., that the prior for a is $p > 1/2$ while the signal s_1 follows the following distribution:

$$P(s_1 = u | a) = P(s_1 = d | b) = q \quad \text{and} \quad P(s_1 = d | a) = P(s_1 = u | b) = 1 - q,$$

where $q > p$. Thus, upon observing the signal u , the subject views a as more likely and upon observing the signal d , the subject views the state b as more likely.²

Panel C describes a situation in which regardless of what the signal is, the subject changes her view of what is the most likely state. These observations cannot be explained with simple Bayesian updating. Indeed, if the subject updates in a standard manner, absent any information, she must put some weight on the realization of s_1 being u and complementary weight on the realization of s_1 being d . In particular, her belief that the state is a when she does not observe a signal must be a convex combination of her beliefs following the realization of the signal. Thus, unless the subject is indifferent, a belief that puts weight lower than $1/2$ on the state a must be a convex combination of two beliefs that put weight greater than $1/2$ on the state a , which is not possible.

In fact, many experiments that are used to illustrate updating anomalies exhibit observations of the form of panel C: Tversky and Shafir's (1992) disjunction effect experiments take this format,³ as do an assortment of experiments on

²In fact, one can come up with many belief structures for the subject that would be consistent with the observed reports and Bayesian updating for either panel A or panel B.

³Specifically, Tversky and Shafir (1992) asked students whether they would be interested in a vacation package that would take place after the results of an exam would be revealed. Thirty-two percent of the subjects were interested, 7 percent of the subjects were not interested, and the remaining 61 percent were willing to pay \$5 to have the option to cancel the vacation after realizing whether they had passed or failed. Nonetheless, when asked whether they would buy the same vacation package *after* discovering they had failed or passed the exam, 57 percent or 54 percent, respectively, of the subjects reported affirmatively.

overconfidence, belief polarization, etc. (for an overview, see, e.g., Camerer, Loewenstein, and Rabin 2006 or Nisbett and Ross 1980).

To summarize, even without placing any restrictions on the subject's belief about the initial distribution of states and the signal generation process, Bayesian updating has clear testable implications. Panels A and B can be explained, while panel C cannot.

In the current paper we also entertain the idea that the subject conjectures the *amount* of information revealed to her by the experimenter is a function of the realized uncertainty: the state and the signals. Then, panel C can easily be explained with Bayesian updating with such richer conjectures. For instance, the subject could conjecture that a priori the states are equally likely, but that when the state is b (or when the appropriate response is b), the experimenter does not reveal any information to her, while when the state is a , the experimenter reveals either signal with equal probability. In particular, behavior that seems to suggest anomalous updating procedures may be attributed to the subject's framing of the experiment instead.

Nonetheless, note that in explaining panel C, we only need the subject to conjecture the amount of information depends on the realized *state* (and not necessarily on the realized signals). Such restricted conjectures are relevant for experiments of decisions under uncertainty in which there is transparency regarding how signals are generated (for example, signals may be generated during the experiment, *after* their number had been determined, as is common practice in, e.g., voting experiments). Interestingly, they do entail testable implications.

Consider Figure 2, that illustrates reports from an experiment as above in which there are two signal realizations, each taking the value of u or d . Note that the sub-tree corresponding to $s_1 = u$ is of the form appearing in panel C of Figure 1. In particular, in order to explain the reversal following either signal s_2 , we need to assume the subject conjectures that the mere observation of two signals is more likely when the state is b . However, switch to the other sub-tree corresponding to $s_1 = d$. Here, we see a similar reversal, but now the subject reports the state a regardless of s_2 . The only way to explain this would be with a conjecture that makes the observation of two signals more likely when the state is a . It is therefore impossible to explain these observations with such restricted conjectures. Thus, transparency of the signal generation procedure assures that Bayesian updating has clear testable implications, even without restricting the subject's conjecture on the distribution of the underlying state, signal generation process, *and* the experimenter's strategy of information revelation (that may depend on the realized state).

This example highlights two important methodological messages of the paper. First, when no restrictions are placed on subjects' framing of the experiment, any behavior is consistent with standard updating ("anything goes"). Second, when the experimental protocol restricts subjects' conjectures (say, by generating information during the experiment), standard updating entails nontrivial testable implications. Namely, such protocols restrict the type of action reversals that Bayesian subjects may exhibit when they are provided with additional information.

Description of Results.—In general, the design of an experiment dictates how much information, i.e., the number of signals, subjects observe prior to making choices. Subjects may hold a wide range of beliefs regarding the link between the




FIGURE 2. REVERSAL NOT EXPLAINED BY A PARTIALLY RESTRICTED CONJECTURED EXPERIMENT

amount of information they observe and the underlying uncertainty. We use the term *conjectured experiment* to denote the subject's beliefs about the design of the experiment. We study the impacts of two types of conjectured experiments: one in which the amount of information is perceived independent of any realized uncertainty and another in which the amount of information can be correlated (in an arbitrary manner) with the realized uncertainty.

We emphasize that in experiments of decision making under uncertainty, the amount of information per se is usually independent of the actual state and the signals' realizations. There are many reasons why a subject's conjecture may not coincide with the actual experimental design. For instance, subjects may arrive at the lab with real-world recipes of behavior that affect their interpretation of experimental instructions; they may also misunderstand the notion of independence, or theorize about the experimenter's goals and selective choice of information (see Friedman and Sunder 1994; Kagel and Roth 1997; and Zizzo 2010). Subjects' conjectures do not necessarily reflect suspicions regarding how the experiment is carried out.⁴

We assume subjects report to the experimenter the state they view as most likely. This is a natural starting point for our analysis since looking at reports of the most likely state requires the subject to go through very simple calculations that do not require expected utility maximization. The entire analysis here is based on ordinal utilities (a monotonic transformation of the subject's utility will not change her optimal response). In particular, the analysis is not sensitive to subjects' risk preferences. In contrast, in order to elicit the full beliefs a subject holds, choosing the proper utility would require certain behavioral assumptions, e.g., a quadratic scoring rule combined with maximization of expected utility.

Suppose first that subjects' conjectured experiments are restricted so that the *amount* of information revealed to them is independent of both the realized state and

⁴Indeed, most economic experiments are preceded by a set of instructions, and the prevailing norm in the profession banning deception would presumably trickle down to subjects.

the realized signals. In that case, there is a natural restriction on subjects' reports for them to be consistent with Bayesian updating (generalizing our introductory example). Indeed, Theorem 1 shows that a subject is Bayesian if and only if the following condition holds: the subject chooses the action a for signal sequence s whenever all continuation signal sequences lead her to choose a .⁵

The other polar case of subjects' conjectures occurs when no restrictions are put and subjects' conjectures are allowed to entail correlations between the realization of all uncertainty (on the state and signals) and the number of signals observed; subjects are free to *frame* the experiment however way they like. In that case, in line with the example above, Theorem 2 illustrates an "anything goes" result: *any* responses of subjects can be explained as arising from a particular conjectured experiment and Bayesian updating of the released information. In other words, in that case the theoretical predictions have no bite once jointly tested with the conjectured experiment.

From a positive and prescriptive point of view, certain experimental protocols are likely to place restrictions on the conjectures subjects hold. For instance, in many voting experiments, the number of signals is determined at the outset, but subjects produce their own signals, e.g., by drawing a ball from a physical urn (the composition of which represents the underlying state), during the experiment.⁶ In terms of the subjects' conjectures, such designs are likely to entail independence between the realized signals and the number of signals provided.

In Section IV we study the intermediate case in which subjects' conjectures allow for correlations between the amount of information revealed and the realized state, but the amount of information is independent of the signal realizations. Such partially restricted conjectures allow for the volume of information to be associated with the state realization. Theorem 3 generalizes the example corresponding to Figure 2 and provides a characterization of the class of reversals that violate Bayesian updating when subjects hold partially restricted conjectures. Much as in the example, the reversals that are consistent with Bayesian updating are roughly described as going "in the same direction" (or tending to a specific report) for a particular length of signal sequence. The theorem provides the tools for deducing departures from Bayesian updating when the signal generation process is transparent in the lab.

The underlying assumption in our analysis is that there is a natural sequencing of signals. This is why the experimental design is captured by the *number* of signals reported to the subject. This assumption is applicable in many situations (indeed, any context in which signals are tied with time), and eases the presentation. In the online Appendix we provide a similar analysis for environments in which there is no natural ordering of signals and a general conjectured experiment pertains to the *dimensions* of information that are reported (or which elements of the set of signals are reported) and their correlation with the underlying uncertainty.⁷

⁵ This is reminiscent of the "Sure Thing Principle" à la Savage (1954) and dynamic consistency. However, the experimental observations are mappings from sequences of signals. Thus, they contain a subset of elements considered by the Savage (1954) type of analysis. In particular, the sufficiency of the condition cannot be deduced from the work on foundations of expected utility. We elaborate on this point in Section II.

⁶ See Palfrey (2006) for a description of various experiments on voting behavior.

⁷ Our "anything goes" result still holds when conjectured experiments are unrestricted. However, the implications imposed by the natural analogue of restricted conjectured experiments do not carry through. In fact, we identify stronger observational restrictions that are in the spirit of "Dutch books." Experimental observations can be

As mentioned, our analysis is very closely related to the informalized notion of *experimenter demand*, the idea that the design itself makes subjects (consciously or subconsciously) frame the problem at hand in a particular way that makes them believe certain responses are more appropriate than others (see, e.g., Friedman and Sunder 1994; Kagel and Roth 1997; and Zizzo 2010). Experimenter demand type of arguments broadly take one of two forms. The first suggests that the way the experimenter phrases problems indicates something to the subject about the realized uncertainty and the correct response. The second refers to the subject trying to “help” or “satisfy” the experimenter by aiming at the answers the experimenter is seeking. Our approach provides a first step in formalizing the former manifestation of experimenter demand, by contemplating a large class of conjectures, or ways to *frame* the experimental design. It is a necessary step for future models of the second demand manifestation as well, which potentially requires more behavioral and strategic assumptions on both experimenter and subject.

The idea that individuals may exhibit a variety of nonstandard behaviors in the presence of uncertainty, such as excess stickiness of beliefs, overconfidence, etc., has received much attention in both psychology and economics (see, for example, Part 1 in Brocas and Carrillo 2004; Kahneman, Slovic, and Tversky 1982; and Nisbett and Ross 1980).⁸ The economics literature has taken two approaches for utilizing these observations:

- (i) Suggest alternative models of belief updating to that prescribed by Bayes’ rule (e.g., Rabin and Schrag 1999 and Compte and Postlewaite 2004);
- (ii) Modify subjects’ utilities to account for arguments going beyond pure monetary rewards and directly depending on held beliefs (e.g., Bénabou and Tirole 2002; Part 3 in Brocas and Carrillo 2003; Köszegi 2006; Yariv 2005, and references therein).

Our approach is very different in that we fix utilities and the belief updating algorithm, but allow for a wide range of theories subjects hold on the design of the experiment itself.⁹ The analysis of this paper does not propose a particular model describing subjects’ behavior. Rather, it focuses on the *testable implications* of existing (standard) models. In particular, it suggests the class of circumstances under which new descriptive models may be useful.

explained if and only if, after observing the subjects’ responses, the experimenter cannot design a sequence of bets that would lead the agent to lose money for sure.

⁸In fact, some work suggests that certain biases in probability assessments are associated with mental health: see Taylor and Brown (1988).

⁹There is an old and wide literature going back to Savage (1954) and Anscombe and Aumann (1963) that considers the question of when experimental observations match some form of optimization (allowing utilities to be arbitrary) *together with* standard probabilistic assessments/updates, essentially restricting the conjectured experiment to coincide with the actual design (for recent references, see Green and Park 1996).

I. Description of the Model

A. Setup

Let A be a finite set of *alternatives* or *states* and S a finite set of *signals*. Let N be the number of available signals and $\mathcal{N} = \{0, 1, \dots, N\}$. We denote by $S^{\leq N}$ the set of all *instances*, i.e., sequences of elements of S of length no greater than N , including the empty sequence e :

$$S^{\leq N} = \bigcup_{0 \leq n \leq N} S^n.$$

Experimental observations are summarized by a mapping $\sigma : S^{\leq N} \rightarrow A$. For every signal sequence \mathbf{s} , $\sigma(\mathbf{s})$ is the subject's report of the most probable alternative given \mathbf{s} .¹⁰ The reason for focusing on these sorts of reports is that we want to identify behaviors that cannot be rationalized by means that are not related to experimental demand (for instance, via risk aversion). Reports of the most likely state are not sensitive to the shape of a subject's utility (in particular, his/her risk attitudes).

For $x = (s_1, \dots, s_N) \in S^N$ and $1 \leq n \leq N$, let $x|_n$ be the *truncation* of x to the first n signals: $x|_n = (s_1, \dots, s_n)$.

If $\mathbf{s} = (s_1, \dots, s_n)$ is an instance and $s \in S$ we denote by $\mathbf{s}\hat{s}$ the *concatenation* of \mathbf{s} with s : $\mathbf{s}\hat{s} = (s_1, \dots, s_n, s)$.

The set $S^{\leq N}$ of instances has a natural rooted tree structure: the *root* is the empty sequence, the *depth* $d(\mathbf{s})$ of an instance $\mathbf{s} = (s_1, \dots, s_n)$ is $d(\mathbf{s}) = n$ and the *children* of \mathbf{s} are instances of the form $\mathbf{s}\hat{s}$ for $s \in S$. The *n*th layer of the tree is the set of nodes of depth n . Technically, all our results and proofs depend only on the rooted tree structure of the set of instances. In particular, the results remain true if the number of available signals is infinite.

B. Conjectured Experiments

We consider subjects who hold beliefs regarding the experimental procedures, namely about the link between the signals they observe and the underlying realized alternative (and their preferred action). We call such a belief a *conjectured experiment*, defined formally as:

DEFINITION 1 (Conjectured Experiment): A *conjectured experiment* is given by a triplet $(\alpha, \tau, \zeta = \{\zeta_n\}_{1 \leq n \leq N})$ of random variables over some probability space (Ω, \mathcal{A}, P) with values in $A, \mathcal{N}, S^{\leq N}$, respectively.

The random variable α denotes the conjecture about the set of alternatives A , τ stands for the length of the observed signal sequence that, therefore, takes values

¹⁰Our primitive, the experimental observations, or the mapping from all possible signal collections to alternatives chosen, is reminiscent of what is observed in experiments utilizing the *strategy method*, an experimental procedure dating back to Selten (1967) under which subjects report *contingent choices*, thereby eliciting their preferred alternatives for any observed realization of uncertainty in the lab.

in \mathcal{N} , and ζ captures the random variables corresponding to the realization of any signal sequence.¹¹

Our results inspect the impacts of different restrictions on the conjectured experiment (that are consequences of either the transparency of the experimental design, or the subjective interpretations of subjects).¹²

In Definition 1 we pose no restrictions on the dependence between the variables α , τ , and ζ . To emphasize this fact we will often refer to such a conjectured experiment as an *unrestricted conjectured experiment*.

In particular, the opposite polar case to that of an unrestricted conjectured experiment is a *restricted conjectured experiment* in which the subject believes the number of signals she sees is uncorrelated with neither the realized alternative nor with the realization of the signals themselves. This is the case, for example, when the subject believes that the experimenter does not know both the realized alternative and realized signals when determining how many signals to provide. Formally,

DEFINITION 2 (Restricted Conjectured Experiment): A *restricted conjectured experiment* is a conjectured experiment (α, τ, ζ) such that τ is independent of the pair (α, ζ) .

We are interested in conjectured experiments—structures of information with which the subject interprets observations—that explain a subject's behavior as arising from Bayesian updating given the number of signals available to her and their realizations. That is,

DEFINITION 3 (Explaining Observations): A *conjectured experiment* $(\alpha, \tau, \zeta = \{\zeta_n\}_{1 \leq n \leq N})$ *explains the experimental observations* σ if:

(i) for every $n \in \mathcal{N}$ and every $s_1, \dots, s_n \in S$,

$$P(\tau = n, \zeta_i = s_i \text{ for } 1 \leq i \leq n) > 0.$$

(ii) For every instance $\mathbf{s} = (s_1, \dots, s_n)$,

$$(1) \quad \sigma(\mathbf{s}) = \arg \max_{a \in A} P(\alpha = a | \tau = n, \zeta_i = s_i \text{ for } 1 \leq i \leq n).$$

Here and elsewhere, when we use the arg max notation, we implicitly assert the uniqueness of the maximizer.¹³

¹¹Note that the set A of alternatives, the number N of signals, and the set S from which signals are drawn are fixed. The only element of the experimental design about which the subject conjectures regards the *distribution* over A , N , and S^N .

¹²Throughout, we assume a natural sequencing of signals. This assumption is applicable in most economics experiments (for instance, it holds whenever signals are provided in sequence) and eases the presentation. In the online Appendix we analyze environments in which there is no natural ordering of signals and a general conjectured experiment pertains to the *dimensions* of information reported.

¹³In principle, we could also assume that σ is a correspondence. In that case, $\sigma(\mathbf{s})$ would stand for the *set* of all alternatives the agent views as most likely after a sequence \mathbf{s} of observations. If we assume that the set $\sigma(\mathbf{s})$ is observed after every sequence \mathbf{s} and modify Definition 3 accordingly then our main results continue to hold, but notation becomes cumbersome. If, on the other hand, we assume that only one alternative in $\sigma(\mathbf{s})$ is observed (which is the case in an experiment, where the subject has to select one alternative), then any observations can be trivially

The first condition in the definition requires that every finite sequence of signals the subject is faced with is indeed conceivable (has positive probability) under her conjectured experiment.¹⁴

The second condition requires that under the conjectured experiment, the subject (Bayesian) updates on which alternative is most likely, conditioning on the signals she receives, and selects that alternative.

We note that the signals in our framework have no content per se. We chose this setting since we want to focus on how different assumptions about the role of the experimenter in providing those signals might affect a subject's behavior, rather than inspect subjects' beliefs about how signals are generated for any alternative. For this reason, we give the subject as much freedom as possible in interpreting the signals, and we show that, even with this freedom, there is a gap between the behaviors that can be explained under different assumptions on the experimenter's role. Nonetheless, often signals are associated with a natural interpretation and it is reasonable to restrict subjects' conjectures to particular plausible ones. In the online Appendix we detail instances of two types of such plausible conjectures: conditionally i.i.d. and conditionally exchangeable.

II. Restricted Conjectured Experiments

We start with the case in which subjects' conjectures are restricted. That is, subjects believe that the *amount* of information they observe is independent of the realized alternative and the realized signals. In that case, conjectured experiments can be written in a simplified manner.

REMARK 1 (Restricted Conjectures, Simplified Notation): *If τ is independent of the pair (ζ, α) then (1) becomes*

$$(2) \quad \sigma(\mathbf{s}) = \arg \max_{a \in A} P(\alpha = a | \zeta_i = s_i \text{ for } 1 \leq i \leq n).$$

Thus, a restricted conjectured experiment that explains the experimental observations is identified by a pair (α, ζ) of random variables with values in A, S^N such that (2) is satisfied.

Let \mathbf{s} be an instance and suppose that no matter what additional signal $s \in S$ is observed, the subject reports the alternative a^* to be the most likely. That is, $\sigma(\mathbf{s} \hat{s}) = a^*$ for every $s \in S$. If the subject is updating using Bayes rule, with a restricted conjecture, she cannot deduce anything regarding the realized state by the sheer volume of signals she observes. Thus, her assessment of each realized alternative under \mathbf{s} is a convex combination of the corresponding assessments over all

rationalized if we assume that the subject believes that all alternatives are equally likely and that the signals are independent of the alternatives.

¹⁴The Bayesian paradigm is silent on what should happen when a Bayesian agent observes a sequence of signals to which she attaches probability zero. Since we want to perform our analysis completely within the Bayesian paradigm, we assume that this cannot happen. Without imposing this assumption, we would have to pose an ad hoc alternative assumption about what a Bayesian agent should do when facing a zero-probability event.

continuations $\hat{s}s$. In particular, the most likely alternative must be a^* . This suggests a clear necessary requirement for observations to be explained. Theorem 1 illustrates that this requirement is, in fact, also sufficient. That is,

THEOREM 1 (Restricted Conjectured Experiments): *The experimental observations σ can be explained by a restricted conjectured experiment if and only if the following condition is satisfied:*

Let \mathbf{s} be an instance. If, for some $a^ \in A$, one has $\sigma(\hat{s}s) = a^*$ for every $s \in S$, then $\sigma(\mathbf{s}) = a^*$.*

The idea underlying the proof of the theorem is instructive for some of the future analysis in the paper. While some details appear in the Appendix, we stress that illustrating the sufficiency of the requirement is done in two steps. First, we consider an auxiliary (and hypothetical) experiment in which subjects report full posteriors over states for every instance (generating a mapping from $S^{\leq N}$ to $\Delta(A)$) and identify the responses consistent with Bayes' rule. Second, for any experimental observations satisfying the requirement of Theorem 1, we construct recursively a set of posteriors satisfying these identified restrictions that are consistent with the experimental observations.¹⁵

The first step, then, is captured by the following lemma, which addresses the question of which assignments of posterior distributions over states of nature can be explained.

LEMMA 1 (Explainable Posteriors, Restricted): *For every instance $\mathbf{s} \in S^{\leq N}$, let $p_{\mathbf{s}} \in \Delta(A)$. Then there exist random variables $\{\alpha, \zeta\}$ over some probability space (Ω, \mathcal{A}, P) with values in A, S^N respectively such that:*

- (i) $P(\zeta_1 = s_1, \dots, \zeta_N = s_N) > 0$ for every $(s_1, \dots, s_N) \in S^N$; and
 - (ii) For every instance $\mathbf{s} = (s_1, \dots, s_n)$ and every $a \in A$
- $$(3) \quad P\{\alpha = a \mid \zeta_i = s_i \text{ for } 1 \leq i \leq n\} = p_{\mathbf{s}}[a]$$

if and only if for every $n < N$ and every $\mathbf{s} = (s_1, \dots, s_n)$ one has¹⁶

$$(4) \quad p_{\mathbf{s}} \in ri(\text{Conv}\{p_{\hat{\mathbf{s}}s} \mid s \in S\}).$$

The proof of the lemma appears in the Appendix. The lemma effectively provides the testable implications of Bayesian updating. It suggests that an individual is Bayesian if and only if the following holds: the probability distribution that

¹⁵The assumption that the subject is Bayesian is important for Theorem 1. For instance, were we to use a multiple prior updating rule à la Gilboa and Schmeidler (1993), one can show that any experimental observations could be rationalized.

¹⁶The relative interior $ri(C)$ of a convex set C in finite dimensional real vector space is the interior of C with respect to the smallest affine space that contains C . If F is a finite set of points then $ri(\text{Conv}F)$ is given by the set of all convex combinations of elements of F with strictly positive coefficients,

$$ri(\text{Conv}F) = \left\{ \sum_{v \in F} \lambda_v v \mid \lambda_v > 0 \ \forall v \in F \text{ and } \sum_{v \in F} \lambda_v = 1 \right\}.$$

represents the agent's belief after observing any finite sequence of signals is a convex combination of the probability distributions that represent her beliefs conditional on observing sequences of signals that are the possible continuations of the original sequence.

It is interesting to note that the lemma implies that a large class of updating rules that exhibit stickiness to prior beliefs (i.e., rules in which the reported beliefs are always tilted toward previous reports) is observationally equivalent to Bayesian updating. Indeed, assume for simplicity that the subject believes that signals are conditionally independent and identically distributed given the state, and calculates the posteriors p_s after observing signals $\mathbf{s} = (s_1, \dots, s_n)$ as follows:

$$p_s[a] = \alpha \cdot \frac{p_{s|n-1}[a]\rho(s_n|a)}{\sum_{b \in A} p_{s|n-1}[b]\rho(s_n|b)} + (1 - \alpha)p_{s|n-1}[a],$$

where $0 < \alpha < 1$ and $0 < \rho(s|a) < 1$ for every state a and signal s , and $s|_{n-1} = (s_1, \dots, s_{n-1})$.

That is, the agent weighs the correct posterior with her previous prediction. It is straightforward to show that p_s would, in fact, satisfy the lemma's condition. In particular, it would be indistinguishable from Bayesian updating (with a prior different than μ).

As a technical note, the relative interior ri appearing in condition (4) assures that there is positive weight on any instance observed, which will be useful for constructing conjectured experiments that entail this type of restriction. Simply requiring $p_s \in \text{Conv}\{p_{s^*s} | s \in S\}$ would correspond to consistent posterior reports, but ones that may place zero probability on certain instances.

We are now ready to present the proof of Theorem 1.

PROOF OF THEOREM 1:

It follows directly from Lemma 1 that the experimental observations $\sigma : S^{\leq N} \rightarrow A$ that can be explained by restricted conjectured experiments must satisfy the condition in Theorem 1. Indeed, let (α, τ, ζ) be a restricted conjectured experiment that explains σ , and let

$$p_s[a] = P(\alpha = a | \zeta_i = s_i \text{ for } 1 \leq i \leq n).$$

From (2), $\sigma(\mathbf{s}) = \arg \max_a p_s[a]$, and from Lemma 1 we get that p_s satisfies (4). Assume that for some instance \mathbf{s} and some $a^* \in A$ one has $\sigma(\mathbf{s}^*s) = a^*$ for every $s \in S$. Then it follows that $\arg \max_a p_{s^*s}[a] = a^*$ for every s . Therefore, $\sigma(\mathbf{s}) = \arg \max_a p_s[a] = a^*$, where the last equality follows from the previous equality and (4).

To prove the other direction, we start with experimental observations that satisfy the condition of Theorem 1 and construct posterior distributions p_s for every instance \mathbf{s} such that (4) is satisfied and, in addition, $\sigma(\mathbf{s}) = \arg \max_a p_s[a]$ for every \mathbf{s} .

We use the following additional lemma:

LEMMA 2 (One Period Construction): *Let $\Delta^*(A) = \left\{ p \in \Delta(A) \mid p[a] < \frac{1}{|A|-1} \right.$ for every $a \in A \left. \right\}$ and let $p \in \Delta^*(A)$. Suppose $a_1, \dots, a_m \in A$ are such that either $a_i \neq a_j$ for some $1 \leq i, j \leq m$, or $a_1 = \dots = a_m = \arg \max p$. Then there exists $p_1, \dots, p_m \in \Delta^*(A)$ such that $p \in \text{ri}(\text{Conv}\{p_1, \dots, p_m\})$ and $\arg \max p_i = a_i$.*

The proof of Lemma 2 is technical and is deferred to the Appendix. It follows from the lemma that we can inductively assign posterior probabilities $p_s \in \Delta^*(A)$ for every instance s such that (4) is satisfied and, in addition,

$$(5) \quad \arg \max_a p_s[a] = \sigma(a)$$

for every instance s (indeed, assuming we already defined p_s , we define $p_{s \cdot S}$ using the lemma, with $m = |S|$). By Lemma 1 there exists random variables α, ζ over some probability space with values in A, S^N , respectively, such that ζ has full support and (3) is satisfied. Let τ be some random variable with values in \mathcal{N} , independent of (α, ζ) , whose distribution has full support. Then for every $a \in A$ and $s = (s_1, \dots, s_n) \in S^{\leq N}$,

$$\begin{aligned} P(\alpha = a \mid \tau = n, \zeta_1 = s_1, \dots, \zeta_n = s_n) \\ = P(\alpha = a \mid \zeta_1 = s_1, \dots, \zeta_n = s_n) = p_s[a]. \end{aligned}$$

By (5) and the last equation it follows that (α, ζ, τ) explains σ , as desired. ■

Since the proof that the condition in Theorem 1 is necessary was based solely on (2), it bears on a larger set of conjectures that allow for some correlation between the volume of information and the realization of signals. Indeed, let us say that a conjectured experiment (α, τ, ζ) is *adapted* if

$$\begin{aligned} P(\tau = n \mid \alpha = a, \zeta_1 = s_1, \dots, \zeta_N = s_N) \\ = P(\tau = n \mid \zeta_1 = s_1, \dots, \zeta_n = s_n), \end{aligned}$$

for every $n \in \mathcal{N}$, $a \in A$, and $s_1, \dots, s_N \in S$ (see Prokhorov and Shiryaev 1997, p. 113 for a general definition of adaptable stopping times). This means that when the experimenter decides whether to continue providing signals after stage n , he can base his decision upon the already released signals s_1, \dots, s_n (which were observed by the subject). Importantly, every restricted conjectured experiment is adapted.

It is easy to verify that (2) is still valid for adapted conjectured experiments. Therefore, even though the set of adapted conjectured experiments is larger than the

set of restricted conjectured experiments, both sets of conjectured experiments have the same explanatory power and we get the following corollary:

COROLLARY 1 (Adapted Conjectured Experiments): *If experimental observations can be explained by an adapted conjectured experiment then they can also be explained by a restricted conjectured experiment.*

In particular, the Corollary implies that allowing subjects to conjecture that the experimenter provides information using a sequential procedure, one in which the experimenter decides whether or not to provide an additional signal depending on the signals that have been revealed thus far, leads to the same testable implications described in Theorem 1.

The Sure Thing Principle and Dynamic Consistency.—The condition of Theorem 1 is reminiscent of the “Sure Thing Principle” and the notion of dynamic consistency. Dynamic consistency suggests that if the decision maker knows that given additional information tomorrow she will prefer action a to b regardless of what that information will be, then she should prefer a to b today (see, e.g., Epstein and Le Breton 1993; Ghirardato 2002, and references therein; as well as Savage 1954). Note, however, that while the condition in Theorem 1 is expressed solely in terms of the maximal element within the agent’s conditional beliefs, dynamic consistency is a property of a preference relation over contingent plans. Moreover, our primitive σ determines only the subject’s action conditioned on a certain class of events, i.e., events that are given by the assignment of values to certain sets of signals. The condition of the theorem is therefore weaker than dynamic consistency: if the order induced by observations σ over actions can be extended to an order over contingent plans that satisfies dynamic consistency, then σ must satisfy the condition. The Theorem implies that the converse is also true: if σ satisfies the condition, then it can be explained by Bayesian updating, and therefore it can be extended to an order over contingent plans that satisfies dynamic consistency. We emphasize that the sufficiency part of Theorem 1 cannot be derived from results regarding dynamic consistency. In fact, without imposing the fixed ordering of signals (thereby allowing information to be provided via arbitrary sets of signals), the condition of the theorem is not sufficient for the observations to be explained by a restricted conjectured experiment. As mentioned, we provide an analysis of this more general setup in the online Appendix.

Generating Restricted Conjectures in the Lab.—While it is difficult to control the subjective beliefs of subjects in the lab, some experiments are more likely to induce restricted conjectures by specifying the number of signals subjects will observe at the outset and later determining both the state and the signals in a transparent manner. We mention two examples that match our setting particularly well.

The first has to do with the formation of herd behavior in the lab. For instance, in the classic experiments of Anderson and Holt (1997), the original cascade model of Bikhchandani, Hirshleifer, and Welch (1992) was tested. In the main treatment, there were six subjects in each session, one of which was randomly designated as the “monitor” and helped the experimenters generate the uncertainty and signals in

the experiment. In each of 15 periods, the monitor used a die to determine which of two urns would be used for that period. Urn A contained two “a” balls and one “b” ball, and urn B contained two “b” balls and one “a” ball. Urn A was used if the throw of the die was 1, 2, or 3; urn B was used otherwise. Subjects were randomly ordered and each saw one private draw from the selected urn. After seeing their private draw, each subject had to guess which urn was selected. This guess was then announced to other subjects. In this way, each subject knew his or her own private draw and the prior decisions of others, if any, before making a guess. This process continued until all subjects had made guesses. At the end of each period, the monitor announced which urn had been selected and each subject was paid \$2 if their guess was correct and nothing otherwise. In this design, the number of signals was determined at the outset and one of the subjects, the “monitor,” helped assure subjects the experimenters did not release information selectively. Furthermore, as in our theoretical setting, it was in the best interest of the subjects to report the alternative, which urn had been selected, they thought was more likely.¹⁷

The results suggested about one-half as many reverse cascades as there were normal cascades that theory predicts with Bayesian updating based on the underlying uncertainty and equilibrium behavior. The authors suggested a model of errors, which is consistent with subjects relying more on their own private information. However, they also noted that about one-third of their subjects seem to rely on a simple count of signals, which is consistent with Bayesian updating with a different information generation process (e.g., one in which each preceding subject’s guess reflects their private draw and is not impacted by previous guesses).

The second example of an experiment that fits our restricted conjectures setting is provided by the recent work of Möbius et al. (2014), who look at the evolution of subjects’ beliefs about being among the top half of performers in an IQ quiz. At the outset, subjects responded to an IQ quiz. Their beliefs about being among the top half of performers were elicited after the quiz. There were then four stages, which were announced at the outset. At each stage, each subject initially received a binary signal that indicated whether the subject was among the top half of performers; that signal was correct with 75 percent probability. Subjects then reported their belief about being among the top half of performers. That is, subjects received four conditionally independent signals that depended on their own performance and others’. Since both the realized state and the signals were generated by the subjects themselves, this design was likely to induce restricted conjectures.

The results illustrate that subjects over-weigh positive feedback relative to negative feedback and are conservative—they have sticky beliefs in that they update too little in response to both positive and negative feedback. The authors suggest, in line with our results, that subjects update as if they misinterpret signal distributions in a self-serving manner, but process these modified signals using Bayes’ rule.

We now turn to analyze the testable implications derived when conjectured experiments are fully unrestricted.

¹⁷The use of subjects as “monitors” may be a useful way to restrict subjects’ conjectures and appears in several other experimental designs. For another experiment in the context of Bayesian updating, see El-Gamal and Grether (1995).

III. Unrestricted Conjectured Experiments

When conjectured experiments are unrestricted, the amount of information observed can be perceived as correlated with the realization of uncertainty. The requirement appearing in Theorem 1 is then too strong. As an example, consider the classical experiment of Tversky and Kahneman (1983). In their experiment, subjects were told that “Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.” Subjects were then asked to determine which is more likely: 1. Linda is a bank teller; or 2. Linda is a bank teller and is active in the feminist movement. Eighty-five percent of the subjects chose the second option. This was taken as evidence for the failure of conjunction of probabilities. However, a different interpretation is that the mere fact the experimenter chose to expose Linda’s participation in demonstrations suggests that 2 is the appropriate answer. In other words, the information revealed (captured by the details appearing in the blurb describing Linda) could be perceived to be correlated with the underlying state (which answer the experimenter expects). In fact, many experiments falling under the general rubric of the *Representative Heuristic* (see Tversky and Kahneman 1974) have this feature of a scenario, or a person’s profile, that is described in some detail and followed by a classification query.

The following is an example in the same spirit utilizing our notation:

Example 1 (Explanatory Power of Unrestricted Conjectures): Assume that $N = 1$, $S = \{u, d\}$, $A = \{a, b\}$, and consider the experimental observations σ given by

$$\sigma(e) = a \quad \text{and} \quad \sigma(u) = \sigma(d) = b,$$

as depicted in tree form in panel C of Figure 1. Note that the condition of Theorem 1 is not satisfied, and therefore σ cannot be explained by a restricted conjectured experiment. However, σ can be explained by an unrestricted conjectured experiment (α, τ, ζ) such that τ and ζ are independent conditional on the realized alternative α and, in addition, characterized by the following:

$$P(\alpha = a) = P(\alpha = b) = \frac{1}{2}$$

$$P(\tau = 0 | \alpha = a) = 1 - P(\tau = 1 | \alpha = a) = 0.9$$

$$P(\tau = 0 | \alpha = b) = 1 - P(\tau = 1 | \alpha = b) = 0.1$$

$$P(\zeta = u | \alpha = \tilde{a}) = P(\zeta = d | \alpha = \tilde{a}) = 0.5 \quad \text{for every } \tilde{a} \in \{a, b\}.$$

As explained in the paper’s introductory example, the conjectured experiment is such that when the realized alternative is $\alpha = a$, the subject perceives receiving no signals as extremely likely, whereas when the realized alternative is $\alpha = b$, the

subject perceives receiving a signal as very likely. However, the determination of the number of signals observed is done independently of the generation of the signals themselves. In fact, the signals in and of themselves are uninformative.

It turns out that *any* experimental observations can be explained with an unrestricted conjectured experiment, formally captured in the following “anything goes” result:

THEOREM 2 (Unrestricted Conjectured Experiments): *For every $\sigma : S^{\leq N} \rightarrow A$ the experimental observations σ admit an explanation by an unrestricted conjectured experiment.*

From the theorem, when subjects’ conjectures are unconstrained, anomalous updating behavior is indistinguishable from particular framing of the experimental design itself.

The proof of Theorem 2, as the proof of Theorem 1, follows two steps. We first show that any set of posteriors can be explained with an unrestricted conjectured experiment. This is the crucial step in the proof, since it is then immediate to choose a sequence of posteriors that is consistent with the observations, and therefore rationalizes them with an unrestricted conjecture. Formally, the analogue of Lemma 1 is the following:

LEMMA 3 (Explainable Posteriors, Unrestricted): *Let $p_s \in \Delta(A)$ be an assignment of probability distribution over A for every instance $s \in S^{\leq N}$. Then there exist random variables (α, τ, ζ) over some probability space (Ω, \mathcal{A}, P) , with values in A, \mathcal{N}, S^N respectively, such that*

(i) *for every $n \in \mathcal{N}$ and every $s_1, \dots, s_n \in S$,*

$$P(\tau = n, \zeta_i = s_i \text{ for } 1 \leq i \leq n) > 0;$$

(ii) *For every instance $s = (s_1, \dots, s_n)$ and every $a \in A$*

$$(6) \quad P(\alpha = a | \tau = n, \zeta_i = s_i \text{ for } 1 \leq i \leq n) = p_s[a].$$

The lemma is conceptually important in that it highlights how robust the message of Theorem 2 is to what is elicited by the experimenter when subjects’ conjectures are unrestricted. Indeed, even if for each instance the *full belief* were elicited (instead of the most likely state), Lemma 3 assures that an analogous “anything goes” result would still hold and any experimental observations (now entailing posterior beliefs for all instances) could be explained with an unrestricted conjecture.

The proof of Lemma 3, appearing in the Appendix, is at the root of our “anything goes” result. Intuitively, an unrestricted conjectured experiment ultimately corresponds to a joint distribution μ over $A \times \mathcal{N} \times S^N$. For any assignment of posteriors $\{p_s\}_{s \in S^{\leq N}}$, $p_s \in \Delta(A)$, pick an arbitrary distribution ν over $\mathcal{N} \times S^N$ and define μ as

$$\mu(k, n, s_1, \dots, s_N) = \nu(n, s_1, \dots, s_N) p_{(s_1, \dots, s_N)}[a].$$

Then, the conditional distribution of μ given the number of signals n and full realization (s_1, \dots, s_N) is $p_{(s_1, \dots, s_N)}$. It follows that the conditional distribution of μ given n and (s_1, \dots, s_n) is p_s , where $\mathbf{s} = (s_1, \dots, s_n)$, as required. The proof of Theorem 2 now follows immediately:

PROOF OF THEOREM 2:

For experimental observations $\sigma : S^{\leq N} \rightarrow A$ we choose arbitrarily, for every instance $\mathbf{s} \in S^{\leq N}$, a probability distribution $p_s \in \Delta(A)$ such that $\arg \max_a p_s[a] = \sigma(\mathbf{s})$. The corresponding random variables α , τ , and ζ , whose existence is asserted in Lemma 3, explain σ . ■

Aggregate Data.—Theorems 1 and 2 pin down the testable implications of restricted and unrestricted conjectured experiments when we observe the response of a single subject after every sequence of signals. In practice, experimenters usually work with aggregate data from responses of a random sample of subjects, each of whom has responded only to some sequences of signals.

Suppose each subject is asked to make a decision after a single signal sequence \mathbf{s} . Theorem 1 can be used to deduce the testable implications of such data, assuming the hypothetical (unobserved) responses of every subject are rationalized by some conjectured experiment (though not necessarily the same conjecture for every subject). In the online Appendix we provide a detailed model of such a setting. Here, assume, for example, that $A = \{a, b\}$ and $S = \{u, d\}$ and, for every sequence \mathbf{s} of signals, let $\theta(\mathbf{s})$ be the proportion of subjects who pick a when provided with the signal sequence \mathbf{s} . Assume further that each of these subjects responds to the experimental queries using a function $\sigma : S^{\leq N} \rightarrow A$ that can be rationalized by some restricted conjectured experiment (i.e., satisfies the condition of Theorem 1). Then, it must be the case that $\theta(\mathbf{s}) \leq \theta(\mathbf{s}^{\wedge}u) + \theta(\mathbf{s}^{\wedge}d) \leq \theta(\mathbf{s}) + 1$ for every $\mathbf{s} \in S^N$. Indeed, $\theta(\mathbf{s}) \leq \theta(\mathbf{s}^{\wedge}u) + \theta(\mathbf{s}^{\wedge}d)$ since every subject who chooses a after \mathbf{s} must be either among the subjects who would choose a after $\mathbf{s}^{\wedge}u$ or among the subjects who would choose a after $\mathbf{s}^{\wedge}d$. Similarly, $1 - \theta(\mathbf{s}) \leq (1 - \theta(\mathbf{s}^{\wedge}u)) + (1 - \theta(\mathbf{s}^{\wedge}d))$ (or, equivalently, $\theta(\mathbf{s}^{\wedge}u) + \theta(\mathbf{s}^{\wedge}d) \leq \theta(\mathbf{s}) + 1$) since every subject who chooses b after \mathbf{s} must be either among the subjects who would choose b after $\mathbf{s}^{\wedge}u$ or among the subjects who would choose b after $\mathbf{s}^{\wedge}d$.

We can estimate $\theta(\mathbf{s})$ from the sample. Statistically significant violations of this inequality will therefore refute the hypothesis that subjects respond according to some restricted conjectured experiment. On the other hand, it follows from Theorem 2 that every distribution of subjects' responses can be rationalized as the aggregation of responses of subjects associated with some unrestricted conjectured experiment.

To summarize our results up to now, we have characterized the testable implications of Bayesian updating for two polar cases pertaining to subjects' freedom regarding the framing of the experimental design. When conjectures are fully restricted, standard updating is equivalent to a type of dynamic consistency condition. When conjectures are unrestricted, standard updating has no testable implications. In what follows, we analyze testable implications derived from intermediate restrictions.

IV. Partially Restricted Conjectured Experiments

In this section, we consider intermediate restrictions that may correspond to experimental designs in which signals (but not states) are determined in the lab in front of the subjects (and the number of signals is a design choice). Indeed, many of the recent social learning and voting experiments follow protocols of this nature. In these experiments, two uncertain alternatives are often manifested in two jars that contain a majority of either, say, red or blue balls. Subjects, not knowing which jar had been selected can then draw a *predetermined* number of balls, often with replacement, from the chosen jar and observe their color¹⁸ (for a survey of such experiments, see Palfrey 2006, forthcoming).

With such designs, it is natural to assume the subject believes the number of signals she receives is uncorrelated with the *realizations* of the signals (but may be correlated with the realized alternative). We formalize such conjectures as follows:

DEFINITION 4 (Partially Restricted Conjectured Experiment): *A partially restricted conjectured experiment is a conjectured experiment (α, τ, ζ) such that τ and ζ are conditionally independent given α , i.e.,*

$$P(\tau = n, \zeta = x | \alpha = a) = P(\tau = n | \alpha = a) \cdot P(\zeta = x | \alpha = a),$$

for every $n \in \mathcal{N}$, $x \in S^N$, and $a \in A$.

We assume hereafter a binary state space, $A = \{a, b\}$.¹⁹ Note that the conjecture proposed in Example 1 is partially restricted and so that example illustrates that the explanatory power of partially restricted conjectured experiments is larger than that of restricted conjectured experiments. However, not all experimental observations can be explained by a partially restricted conjectured experiment, as illustrated by the following example, formalizing the example we discussed in our introduction:

Example 2 (Unexplainable Reversals): Assume that $N = 2$, $S = \{u, d\}$, and consider the experimental observations σ depicted in Figure 2. Let \mathbf{s} and \mathbf{t} be the signal sequences (u) and (d) respectively (so $d(\mathbf{s}) = d(\mathbf{t}) = 1$). On the one hand, controlling for any learning from the sheer amount of information released, \mathbf{s} is more supportive of the alternative a than \mathbf{t} (from the subject's point of view) since $\sigma(\mathbf{s}) = a$ and $\sigma(\mathbf{t}) = b$. On the other hand, the same can be said in reverse, since conditioning on depth 2 of the tree, $\sigma(\hat{\mathbf{s}}s) = b$ and $\sigma(\hat{\mathbf{t}}t) = a$ for every $s, t \in \{u, d\}$. This inconsistency implies that the experimental observations cannot be explained by a partially restricted conjectured experiment.

Examples 1 and 2 illustrate that the explanatory power of partially restricted conjectured experiments is strictly between that of unrestricted and restricted conjectured experiments.

¹⁸Our framework encompasses such strategic interactions since we can always rename the alternatives appropriately. Namely, in strategic interaction experiments, subjects select the action (which serves as our alternative) that is more likely to deliver higher payoffs.

¹⁹Our characterization in this section relies on techniques tailored for a binary set of alternatives. While many experiments satisfy this restriction (for instance, much of the experimental literature on voting focuses on binary elections), we view the extension to richer sets of alternatives as an interesting direction for future research.

Our characterization of the class of observations that are explainable with partially restricted conjectures entails ruling out a generalized set of reversals of the type appearing in Example 2. Roughly speaking, reversals that are consistent with Bayesian updating are described by “tending to a specific report” for a particular length of signal sequence.

Formally, fix experimental observations σ . For a pair of instances \mathbf{s}, \mathbf{t} of the same depth we define recursively what we mean by “the conditional probability of state ‘ a ’ given \mathbf{s} is behaviorally revealed higher than the conditional probability of state ‘ a ’ given \mathbf{t} ,” where the conditioning is on the value of signals received, but not their number. We say shortly that \mathbf{s} is *revealed higher* than \mathbf{t} .²⁰

DEFINITION 5 (Revealed Higher Relation): *Let $\sigma : S^{\leq N} \rightarrow \{a, b\}$ be experimental observations. For a pair of instances $\mathbf{s}, \mathbf{t} \in S^{\leq N}$ of the same depth the relation ‘ \mathbf{s} is revealed higher than \mathbf{t} under σ ’ is recursively defined using the following rules:*

- (i) *If $\sigma(\mathbf{s}) = a$ and $\sigma(\mathbf{t}) = b$, then \mathbf{s} is revealed higher than \mathbf{t} .*
- (ii) *If $\mathbf{s} \hat{s}$ is revealed higher than $\mathbf{t} \hat{t}$ for every $s, t \in S$, then \mathbf{s} is revealed higher than \mathbf{t} .*

The following lemma illustrates the implication of one instance being revealed higher than another in terms of probabilistic assessments:

LEMMA 4 (Revealed Higher Probabilities): *Let $\sigma : S^{\leq N} \rightarrow \{a, b\}$ be experimental observations, and let (α, τ, ζ) be a partially restricted conjectured experiment that explains σ . If $\mathbf{s} = (s_1, \dots, s_n)$ and $\mathbf{t} = (t_1, \dots, t_n) \in S^{\leq N}$ are a pair of instances such that \mathbf{s} is revealed higher than \mathbf{t} under σ , then*

$$P(\alpha = a | \zeta_1 = s_1, \dots, \zeta_n = s_n) > P(\alpha = a | \zeta_1 = t_1, \dots, \zeta_n = t_n).$$

Lemma 4 assures that if the experimental observations can be explained by a partially restricted conjectured experiment, it cannot be the case that there are reversals of the form \mathbf{s} is revealed higher than \mathbf{t} and \mathbf{t} is revealed higher than \mathbf{s} for some instances \mathbf{s} and \mathbf{t} . As it turns out, the reverse is also true. Indeed, the following theorem characterizes the set of experimental observations that can be explained by a partially restricted conjectured experiment.

THEOREM 3 (Partially Restricted Conjectured Experiments): *The experimental observations $\sigma : S^{\leq N} \rightarrow \{a, b\}$ can be explained by a partially restricted conjectured experiment if and only if the order revealed higher is antisymmetric. That is, there exists no pair \mathbf{s}, \mathbf{t} of instances such that \mathbf{s} is revealed higher than \mathbf{t} and \mathbf{t} is revealed higher than \mathbf{s} .*

²⁰Note that “revealed higher” need not be a complete relation, nor is it necessarily antisymmetric.

The formal proof of the sufficiency of the condition is intricate and appears in the online Appendix, together with the formal proof of Lemma 4.

In order to provide the reader with some intuition, we present here a sketch of the proof of the theorem. Assume that the experimental observations $\sigma : S^{\leq N} \rightarrow \{a, b\}$ satisfy the theorem's condition. We construct the partially restricted conjectured experiment (α, τ, ζ) that explains σ in two steps. First, we construct random variables (α, ζ) such that

$$P(\alpha = a | \zeta_1 = s_1, \dots, \zeta_n = s_n) > P(\alpha = a | \zeta_1 = t_1, \dots, \zeta_n = t_n)$$

for every pair $\mathbf{s} = (s_1, \dots, s_n)$ and $\mathbf{t} = (t_1, \dots, t_n)$ of instances of the same length such that $\sigma(\mathbf{s}) = a$ and $\sigma(\mathbf{t}) = b$. In the second step, we add a random variable τ that is independent of ζ given α and, for every n , we choose the probabilities $P(\alpha = a | \tau = n)$ so that

$$P(\alpha = a | \tau = n, \zeta_1 = s_1, \dots, \zeta_n = s_n) > 1/2$$

for all instances $\mathbf{s} = (s_1, \dots, s_n)$ such that $\sigma(\mathbf{s}) = a$ and

$$P(\alpha = a | \tau = n, \zeta_1 = t_1, \dots, \zeta_n = t_n) < 1/2$$

for all instances $\mathbf{t} = (t_1, \dots, t_n)$ such that $\sigma(\mathbf{t}) = b$.

In words, we first construct conjectures that are consistent within each layer of the tree of observations (that is, for a *fixed* number of signals). We then construct the correlation between the number of signals observed and underlying state so that the assessments *across* layers are consistent.

The main difficulty is in the first stage. The proof makes use of the tree structure over the set of instances. For every instance $\mathbf{s} = (s_1, \dots, s_n) \in S^{\leq N}$ we assign a number $p_{\mathbf{s}} \in [0, 1]$ such that $p_{\mathbf{s}} \in \text{ri}(\text{Conv}\{p_{\mathbf{s}\zeta_s} | s \in S\})$ and $p_{\mathbf{s}} > p_{\mathbf{t}}$ whenever \mathbf{s}, \mathbf{t} are two instances of the same length and \mathbf{s} is revealed higher than \mathbf{t} . Then, by Lemma 1 we can construct random variables α, ζ such that

$$p_{\mathbf{s}} = P(\alpha = a | \zeta_1 = s_1, \dots, \zeta_n = s_n)$$

for every instance $\mathbf{s} = (s_1, \dots, s_n)$.

We assign the numbers $p_{\mathbf{s}}$ to nodes \mathbf{s} by induction over the depth of \mathbf{s} , working from the root of the tree to the leafs. Assume that we have already defined the values $p_{\mathbf{s}}$ for nodes of depth $n - 1$ and consider now the set X of nodes of depth n . The relation revealed higher induces a partial order over X , which we denote by \leq . In addition, there is a natural partition \mathcal{A} over the set X , whose atoms are the set of instances with a common parent in the tree. Thus, the atoms of \mathcal{A} correspond to nodes of depth $n - 1$. Since we have already defined $p_{\mathbf{s}}$ over these nodes, we need to “lift” p to a function over X that is monotone with respect to \leq . Much of the

technical aspects of the proof are dedicated to showing that this is indeed possible, using the fact that the order “revealed higher” is an interval order.²¹

We note that Theorem 3 gives rise to a simple algorithm for checking whether experimental observations in a finite tree can be explained by a partially restricted conjectured experiment: go over all the layers of the tree of instances, from the layer of depth $d = N$ to the layer of depth $d = 0$ (i.e., from the leaves to the root of the corresponding tree). For every layer d , construct the relation “ s is higher than t ” for instances s, t of that layer (using Definition 5 and the already constructed relation over layer $d + 1$) and check that the condition is satisfied over nodes in that layer. Using this algorithm, the reader can verify that the experimental observation in the following example cannot be explained by a partially restricted conjectured experiment.

Example 3 (Explainability Algorithm): Consider the Example depicted in Figure 3 (notation following that of our previous examples). Note that (u) is revealed higher than (d) . However, algorithmically proceeding from the leaves to layer 1 we see that, from layer 3, (d, u) and (d, d) are revealed higher than (u, u) , and from layer 2, (d, u) and (d, d) are also revealed higher than (u, d) , which therefore implies that (d) is revealed higher than (u) . In particular, the condition of Theorem 3 is not satisfied and these observations cannot be explained by a partially restricted conjectured experiment.

Example 3 also demonstrates that lack of “simple reversals” of the type of Example 2 is not sufficient for experimental observations to admit an explanation by a partially restricted conjectured experiment.

V. Conclusions

This paper provides a theoretical framework for analyzing experimental data accounting for subjects’ conjectures regarding the experimental design itself. When subjects’ conjectures are unrestricted, we illustrate an “anything goes” result: any experimental observations can be explained with standard updating and the appropriate choice of a conjectured experiment (in fact, generically, multiple conjectured experiments would explain the observations). When subjects’ conjectures are restricted, in terms of the perceived correlation between the amount of information revealed to them in the lab and the underlying realized uncertainty, our results provide a full characterization of the testable implications standard updating entails. To the extent that experimental transparency (say, regarding the way information is generated in the lab) yields more restricted conjectures, the overwhelming message from our results is that transparency is crucial for meaningful testable implications of theoretical predictions pertaining to decision making under uncertainty.

²¹In terms of technical generality, we note that Theorems 1, 2, and 3 all easily extend to environments in which the number of signals is not bounded.

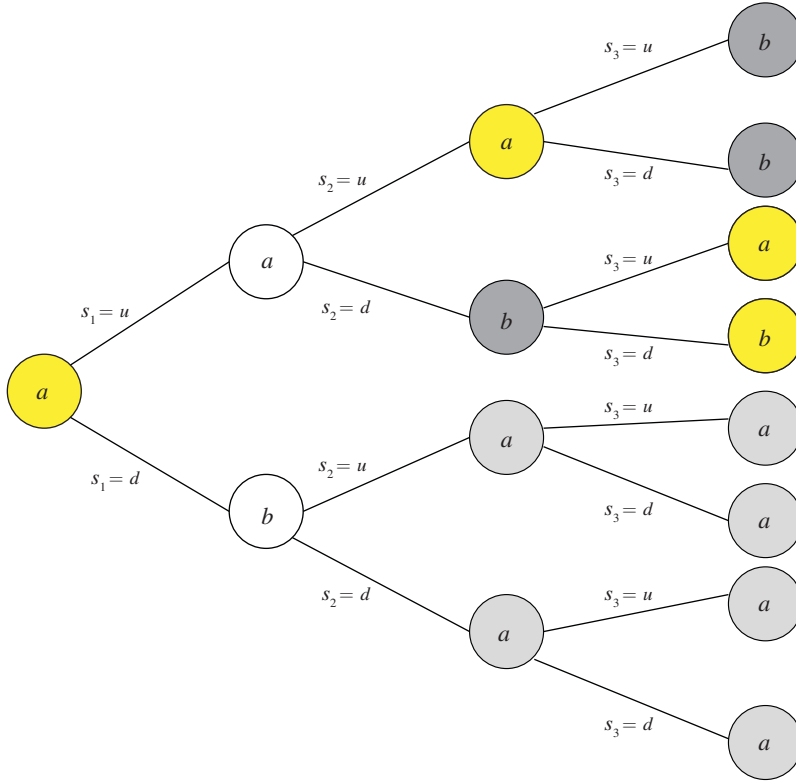


FIGURE 3. ALGORITHMICALLY CHECKING CONSISTENCY

APPENDIX

A. Restricted Conjectured Experiments

PROOF OF LEMMA 1:

Let $\{\alpha, \zeta\}$ be random variables that satisfy the condition of the Lemma. Then, for every $n < N$ and every $\mathbf{s} = (s_1, \dots, s_n) \in S^n$ it follows from condition 2 and the law of total probability that

$$\begin{aligned} p_{\mathbf{s}}[a] &= P(\alpha = a \mid \zeta_i = s_i \text{ for } 1 \leq i \leq n) \\ &= \sum_{s \in S} P(\zeta_{n+1} = s \mid \zeta_i = s_i \text{ for } 1 \leq i \leq n) \\ &\quad \cdot P(\alpha = a \mid \zeta_i = s_i \text{ for } 1 \leq i \leq n, \zeta_{n+1} = s) = \sum_{s \in S} \lambda_s p_{\mathbf{s} \hat{\mathbf{s}}} [a], \end{aligned}$$

where $\lambda_s = P(\zeta_{n+1} = s \mid \zeta_i = s_i \text{ for } 1 \leq i \leq n) > 0$ by condition 1. Therefore, $p_{\mathbf{s}} \in \text{ri}(\text{Conv}\{p_{\mathbf{s} \hat{\mathbf{s}}} \mid s \in S\})$ as desired.

Assume now that for every $\mathbf{s} = (s_1, \dots, s_n)$ one has $p_{\mathbf{s}} \in \text{ri}(\text{Conv}\{p_{\mathbf{s} \hat{\mathbf{s}}} \mid s \in S\})$, so that

$$(7) \quad p_{\mathbf{s}} = \sum_{s \in S} \lambda_{\mathbf{s} \hat{\mathbf{s}}} p_{\mathbf{s} \hat{\mathbf{s}}}$$

for some $\lambda_{\mathbf{s}^*} > 0$ such that $\sum_{\mathbf{s}} \lambda_{\mathbf{s}^*} = 1$. Let $\zeta_1, \dots, \zeta_n, \alpha$ be random variables such that

$$(8) \quad P(\zeta_{n+1} = s | \zeta_1 = s_1, \dots, \zeta_n = s_n) = \lambda_{\mathbf{s}^*}$$

for every $n < N$, $\mathbf{s} = (s_1, \dots, s_n)$, $s \in S$, and

$$(9) \quad P(\alpha = a | \zeta_1 = s_1, \dots, \zeta_N = s_N) = p_{\mathbf{s}}[a]$$

for every $\mathbf{s} = (s_1, \dots, s_N)$. Since $\lambda_{\mathbf{s}^*} > 0$, it follows from (8) that

$$P(\zeta_1 = s_1, \dots, \zeta_N = s_N) = \prod_{n < N} P(\zeta_{n+1} = s_{n+1} | \zeta_1 = s_1, \dots, \zeta_n = s_n) > 0,$$

so that the first condition in Lemma 1 is satisfied. We now prove the second condition by backward induction over n . For $n = N$ the condition follows from (9). Let $n < N$, $\mathbf{s} = (s_1, \dots, s_n) \in S$, and assume the condition is true for $n + 1$. Then,

$$\begin{aligned} & P(\alpha = a | \zeta_i = s_i \text{ for } 1 \leq i \leq n) \\ &= \sum_s P(\zeta_{n+1} = s | \zeta_i = s_i \text{ for } 1 \leq i \leq n) \\ & \cdot P(\alpha = a | \zeta_i = s_i \text{ for } 1 \leq i \leq n+1) = \sum_s \lambda_{\mathbf{s}^*} p_{\mathbf{s}^*}[a] = p_{\mathbf{s}}[a], \end{aligned}$$

where the first equality follows from the law of total probability, the second from (8) and the induction hypothesis, and the third from (7). ■

PROOF OF LEMMA 2:

If $a_1 = \dots = a_m = \arg \max_a p$ then we can choose $p_i = p$ for $1 \leq i \leq m$.

Assume now that $a_1 \neq a_2$. We choose $p_3, \dots, p_m \in \Delta^*(A)$ arbitrarily such that $\arg \max_a p_i[a] = a_i$. Let $\varepsilon > 0$ be sufficiently small so that $(m-2)\varepsilon < 1$ and $p' \in \Delta^*(A)$, where

$$(10) \quad p' = \frac{1}{1 - (m-2)\varepsilon} (p - \varepsilon(p_3 + \dots + p_m)).$$

The existence of such ε follows from the fact that $\Delta^*(A)$ is an open set and the right-hand side of (10) converges to $p \in \Delta^*(A)$ as ε goes to 0.

We now define p_1 and p_2 . Choose q such that

$$(11) \quad \max_a p[a] < q < \frac{1}{|A| - 1},$$

and let $r = 1 - \sum_{a \neq a_1, a_2} p'[a] - q$. Note that from the fact that $p'[a] < \frac{1}{|A| - 1}$ for every $a \in A$ and $q < \frac{1}{|A| - 1}$ it follows that $r > 0$. Moreover, since $q > p'[a_1]$ it follows that

$$(12) \quad r = 1 - \sum_{a \neq a_1, a_2} p'[a] - q < 1 - \sum_{a \neq a_2} p'[a] = p'[a_2] < q < \frac{1}{|A| - 1}.$$

For $i = 1, 2$ and $j = 3 - i$, let

$$(13) \quad p_i[a] = p'[a] \quad \text{for } a \neq a_1, a_2$$

$$(14) \quad p_i[a_i] = q$$

$$(15) \quad p_i[a_j] = r.$$

Then, it follows from (12) that $p_i \in \Delta^*(A)$ and $\arg \max_a p_i[a] = a_i$.

In addition, it follows from (11) and (12) that $r < p'[a_2] < q$. Therefore,

$$(16) \quad p'[a_2] = \lambda r + (1 - \lambda)q$$

for some $0 < \lambda < 1$. We claim that

$$(17) \quad p' = \lambda p_1 + (1 - \lambda)p_2$$

$$(18) \quad \text{i.e., } p'[a] = \lambda p_1[a] + (1 - \lambda)p_2[a] \quad \text{for every } a \in A.$$

Indeed, for $a \neq a_1, a_2$, the equality follows from the fact that, by (13), $p_1[a] = p_2[a] = p'[a]$. For $a = a_2$, the equality follows from (14), (15), and (16). For $a = a_1$, the equality follows from the equality in all other coordinates and the fact that both sides of (18) sum to 1 (over $a \in A$). Finally, it follows from (10) and (17) that

$$\begin{aligned} p &= (1 - (m - 2)\varepsilon)p' + \varepsilon(p_3 + \cdots + p_m) \\ &= \lambda_1 p_1 + \lambda_2 p_2 + \lambda_3 p_3 + \cdots + \lambda_m p_m, \end{aligned}$$

where

$$\begin{aligned} \lambda_1 &= (1 - (m - 2)\varepsilon) \cdot \lambda, \\ \lambda_2 &= (1 - (m - 2)\varepsilon) \cdot (1 - \lambda), \quad \text{and} \\ \lambda_3 &= \cdots = \lambda_m = \varepsilon. \end{aligned}$$

Therefore, $p \in \text{ri}(\text{Conv}\{p_1, \dots, p_m\})$, as desired. ■

B. Unrestricted Conjectured Experiments

PROOF OF LEMMA 3:

Choose arbitrarily a distribution ν over $\mathcal{N} \times S^N$ with full support. Let α, τ, ζ be random variables over some probability space (Ω, \mathcal{A}, P) with values in A, \mathcal{N}, S^N such that:

- (i) The joint distribution of τ and ζ is ν .

- (ii) The conditional distribution of α given $\tau = n$ and $\zeta = x$ is p_s , where $s = x|_n$,

$$(19) \quad P(\alpha = a | \tau = n, \zeta = x) = p_s[a].$$

We proceed to prove that (6) is satisfied for every instance $s = (s_1, \dots, s_n)$. The event $\{\tau = n, \zeta_i = s_i \text{ for every } 1 \leq i \leq n\}$ is the disjoint union of the events $\{\tau = n, \zeta = x\}$, ranging over all $x \in S^N$ such that $x|_n = s$. Therefore, by the law of total probability,

$$\begin{aligned} P(\alpha = a | \tau = n, \zeta_i = s_i \forall 1 \leq i \leq n) \\ &= \sum_{x \in S^N \text{ and } x|_n = s} P(\zeta = x | \tau = n, \zeta_i = s_i \forall 1 \leq i \leq n) \\ &\quad \times P(\alpha = a | \tau = n, \zeta = x) = p_s[a], \end{aligned}$$

where the last equality follows from (19). ■

REFERENCES

- Anderson, Lisa R., and Charles A. Holt. 1997. "Information Cascades in the Laboratory." *American Economic Review* 87 (5): 847–62.
- Anscombe, Francis J., and Robert J. Aumann. 1963. "A Definition of Subjective Probability." *Annals of Mathematical Statistics* 34 (1): 199–205.
- Beecher, Henry K. 1955. "The Powerful Placebo." *Journal of the American Medical Association* 159 (17): 1602–06.
- Bénabou, Roland, and Jean Tirole. 2002. "Self-Confidence and Personal Motivation." *Quarterly Journal of Economics* 117 (3): 871–915.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100 (5): 992–1026.
- Brocas, Isabelle, and Juan D. Carrillo, eds. 2003. *The Psychology of Economic Decisions, Volume 1: Rationality and Well-Being*. Oxford: Oxford University Press.
- Brocas, Isabelle, and Juan D. Carrillo, eds. 2004. *The Psychology of Economic Decisions, Volume 2: Reasons and Choices*. Oxford: Oxford University Press.
- Camerer, Colin F., George Loewenstein, and Matthew Rabin, eds. 2006. *Advances in Behavioral Economics*. Princeton, NJ: Princeton University Press.
- Compte, Olivier, and Andrew Postlewaite. 2004. "Confidence-Enhanced Performance." *American Economic Review* 94 (5): 1536–57.
- El-Gamal, Mahmoud A., and David M. Grether. 1995. "Are People Bayesian? Uncovering Behavioral Strategies." *Journal of the American Statistical Association* 90 (432): 1137–45.
- Epstein, Larry G., and Michel Le Breton. 1993. "Dynamically Consistent Beliefs Must Be Bayesian." *Journal of Economic Theory* 61 (1): 1–22.
- Friedman, Daniel, and Shyam Sunder. 1994. *Experimental Methods: A Primer for Economists*. New York: Cambridge University Press.
- Ghirardato, Paolo. 2002. "Revisiting Savage in a Conditional World." *Economic Theory* 20 (1): 83–92.
- Gilboa, Itzhak, and David Schmeidler. 1993. "Updating Ambiguous Beliefs." *Journal of Economic Theory* 59 (1): 33–49.
- Gillespie, Richard. 1991. *Manufacturing Knowledge: A History of the Hawthorne Experiments*. Cambridge: Cambridge University Press.
- Green, Edward J., and In-Uck Park. 1996. "Bayes Contingent Plans." *Journal of Economic Behavior and Organization* 31 (2): 225–36.
- Kagel, John H., and Alvin E. Roth. 1997. *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.

- Kahneman, Daniel P., Paul Slovic, and Amos Tversky, eds.** 1982. *Judgment under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.
- Köszegi, Botond.** 2006. "Ego Utility, Overconfidence, and Task Choice." *Journal of the European Economic Association* 4 (4): 673–707.
- Möbius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat.** 2014. "Managing Self-Confidence." <http://www.nber.org/papers/w17014>.
- Nisbett, Richard E., and Lee Ross.** 1980. *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- Palfrey, Thomas R.** 2006. "Laboratory Experiments." In *Handbook of Political Economy*, edited by Barry R. Weingast and Donald A. Wittman, 915–36. Oxford: Oxford University Press.
- Palfrey, Thomas R.** Forthcoming. "Experiments in Political Economy." In *The Handbook of Experimental Economics*. Vol. II, edited by J. Kagel and A. Roth. Princeton, NJ: Princeton University Press.
- Paulhus, Delroy L.** 1991. "Measurement and Control of Response Bias." In *Measures of Personality and Social Psychological Attitudes*, edited by John P. Robinson, Phillip R. Shaver, and Lawrence S. Wrightsman, 17–60. San Diego: Academic Press.
- Prokhorov, Yurij V., and Albert N. Shiryaev, eds.** 1997. *Probability Theory III: Stochastic Calculus*. Encyclopaedia of Mathematical Sciences. New York: Springer.
- Rabin, Matthew, and Joel L. Schrag.** 1999. "First Impressions Matter: A Model of Confirmatory Bias." *Quarterly Journal of Economics* 114 (1): 37–82.
- Savage, Leonard J.** 1954. *The Foundations of Statistics*. New York: Wiley.
- Selten, Reinhard.** 1967. "Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopol-experiments." In *Beiträge zur Experimentellen Wirtschaftsforschung*, edited by Heinz Sauer mann, 136–68.
- Taylor, Shelley E., and Jonathan D. Brown.** 1988. "Illusion and Well-Being: A Social Psychological Perspective on Mental Health." *Psychological Bulletin* 103 (2): 193–210.
- Tversky, Amos, and Daniel Kahneman.** 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185 (4157): 1124–31.
- Tversky, Amos, and Daniel Kahneman.** 1983. "Extension versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment." *Psychological Review* 90 (4): 293–315.
- Tversky, Amos, and Eldar Shafir.** 1992. "The Disjunction Effect in Choice under Uncertainty." *Psychological Science* 3 (5): 305–309.
- Yariv, Leeat.** 2005. "I'll See It when I Believe It." <http://people.hss.caltech.edu/~lyariv/papers/Believe.pdf>.
- Zizzo, Daniel John.** 2010. "Experimenter Demand Effects in Economic Experiments." *Experimental Economics* 13 (1): 75–98.