Essays in Search and Learning

by

Collin Buckley Raymond

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2012

Doctoral Committee:

Assistant Professor Yusuf Can Masatlioglu, Chair
Professor Yan Chen
Associate Professor Daniel Susman Silverman
Assistant Professor Daisuke Nakajima
Professor Lones A. Smith, University of Wisconsin — Madison

To my parents, for their continual encouragement and support.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

CHAPTER I

Search and Temptation[1]

1.1 Introduction

In recent years there has been a growing interest in modeling agents who experience self-control problems. A large literature has identified agents who experience self-control problems as those who would prefer to commit to smaller choice sets, in order to avoid items they find tempting. In the real world, consumers often construct choice sets through a search process. Search provides consumers with a way to commit through their actions: they can stop searching and choose among the currently available items. In fact, the decision to restrict the available choice set, through stopping, is an essential component of search. This suggests that search is a natural environment in which to observe the effects of preferences for commitment. Furthermore, understanding how consumers search is an active area of both theoretical and empirical research. More accurate models of the search decisions of consumers that care not just about the item that they purchase, but also unchosen, tempting options, can help better explain behavior in search markets.

The canonical result in sequential search theory, dating back to McCall (1970), is that consumers follow a simple threshold strategy. If the utility of an item falls above some constant threshold, the consumer stops and purchases it. Otherwise, the consumer continues searching. The canonical model has provided a tractable way of modeling sequential search, but recent research has shown that these theoretical predictions are often inconsistent with observed behavior. In particular, decision-makers' choices can be affected by options that have been seen during the search process but remain unchosen.

[1] I thank seminar participants at the XI Summer School on Economics and Philosophy, the Michigan Economics Theory Lunch, the Michigan Economics Labor Lunch, the Midwest Economic Theory Meetings, the North American Winter Meetings of the Econometric Society 2012, Notre Dame, National University of Singapore, Simon Frasier, Oxford, Pompeu Fabra, the Federal Trade Commission, and Charles River Associates, as well as Daniel Benjamin, Tilman Borgers, Yan Chen, Josh Cherry, Jesse Gregory, Matt Hall, Dmitry Lubensky, Matthew Rabin, Dan Silverman, Lones Smith, and Sarah Taylor for helpful comments. Particular thanks goes to Yusufcan Masatlioglu and Daisuke Nakajima for their support. I also would like to thank the University of Michigan School of Information's Socio-Technical Infrastructure for Electronic Transactions Multidisciplinary Doctoral Fellowship funded by NSF IGERT grant 0654014 for financial support. All remaining errors are my own.

For example, consumers often engage in a search process while purchasing a new car. Consumers typically want a car that is affordable, reliable and gets good gas mileage. However, during the search process, consumers will observe cars that are 'tempting', with larger engines, additional features, and luxury materials. Anecdotally, consumers often seem to engage in a path-dependent choices. If consumers observe a car at a dealership that is very tempting, such as a high end sedan, they end up purchasing cars that are more tempting (e.g. have leather seats) relative to what they would have purchased had they not seen the sedan. This paper shows that path-dependent choice is a natural outcome of search when the consumer is influenced by temptation.

In this paper, agents engage in sequential search from a known distribution with perfect recall. Agents have temptation-based preferences (introduced by Gul and Pesendorfer, 2001) and so experience a cost of self-control at the time when they stop searching, and choose among the already observed items. Agents' payoffs depend not only on the chosen item but also the most tempting item already seen. Agents discount the future, and are uncertain about what they will observe next. In order to focus on the effects of context-dependent preferences, I assume agents know the distribution of outcomes.

I demonstrate that the addition of temptation to an otherwise standard search model delivers implications starkly different from the standard model of search, as well as search models with quasi-hyperbolic discounters, but that match anecdotal evidence and intuition about choice behavior in the real world. These results help underscore the importance of context-dependent preferences in economic settings. Although economists have recently made great strides at modeling behavior when utility is affected by unchosen options from choice sets, there has been limited work understanding how these preferences affect behavior in basic economic settings. Search is a particularly important application of temptation for two reasons. First, search provides a natural environment to observe preferences for commitment, and so can serve as a testing ground for models that capture temptation. This is because search is a natural process for agents to endogenously form choice sets, and so we should expect forward-looking agents who are concerned about not just the item that they choose, but also unchosen observed alternatives, to behave differently than agents with classical preferences. For example, agents have a simple way of committing to a choice set — by simply refusing to continue search. Therefore, we should expect that agents will more often stop when they are more concerned about self-control. Second, temptation can help us better understand real-world search behavior. Search is an important economic phenomenon, occurring in a wide variety of settings, including housing, retail, and labor markets. Understanding how preferences that reflect important, although non-classical, considerations change behavior in search settings can lead to a clearer understanding of market outcomes, including path-dependent choice, advertising, and product ordering.

Section 2 provides a brief summary of related literature, looking at related applications of models of self-control. It also reviews empirical studies that have reported behavior consistent with, and suggestive of, context-dependent preferences in dynamic settings. Although no empirical work directly speaks to the model of self-control developed in this paper, there is direct evidence in psychology on temptation and self-control, as well as the fact that unchosen options but observed can affect choice in search settings in ways inconsistent with the classical model.

Section 3 lays out the basic model of search with temptation. Due to its popularity and wide applicability, I focus on the model of temptation-based preferences introduced by Gul and Pesendorfer (2001) (hereafter GP).[2] Each outcome is characterized by two values: a commitment value and a temptation value. The commitment value reflects how the agent would like to choose, if she could commit to having a single item available in her choice set. The temptation value reflects the self-control problems the agent experiences.

The agent would like to have only the item with the highest commitment value available. However, when the agent stops and chooses an item, she chooses the item that has the highest sum of temptation and commitment values, but suffers a utility loss equal to the highest temptation value of any item in the choice set. This means that search now carries additional risk; the agent could come across a more tempting item that she would prefer to avoid. If the agent chooses to search, she receives a payoff of zero.[3]

Section 4 examines the behavioral implications of search with temptation. Because search now has a cost of a psychological nature (observing a more tempting item), the agent's incentives to search changes relative to the standard model. The optimal stopping rule takes the form of a threshold: the agent stops and chooses an item if and only if the sum of its commitment and temptation values exceeds a reservation value. The reservation value is a non-decreasing function of the most tempting item already observed in the search process. This is because the implicit psychological cost of additional search is falling over time. Accordingly, her expected future length of search is rising in the current length of search. I refer to this behavior as the 'threshold effect'. As the threshold changes, the set of items that an agent is willing to stop and choose changes in a particular way. As

[2] Appendix B generalizes the results to include other context-dependent preferences. In this case agents' utility can depend on the entire history of observations, and items are characterized by a vector of values, which are aggregated over the available set of options. Of course, even this type of context-dependence has its limitation. One can imagine situations where agents are tempted not only by what they have seen, but also by what they expect to see during search. Although not addressed in this paper, the model can be modified to accommodate specifications where utility depends on expected outcomes as well.

[3] This assumption, although restrictive, seems natural in many contexts where the agent is simply acquiring options. The alternative assumption, that agents experience temptation as a flow cost while searching, is discussed as an extension in Appendix B.

the agent observes more tempting items, she rules out items that are relatively low in temptation value, but is now able to purchase items relatively low in commitment value. I refer to the effect of observing a more tempting item on the trade-off between commitment and temptation values of items that the agent is willing to choose as the 'compromise effect'. It matches the intuition that a consumer, upon observing a very tempting item, does not necessarily choose it, but becomes more likely to purchase something relatively more tempting, but which is worse in terms of ex-ante preferences. Changes in the environment, or in the susceptibility of the agent to temptation can impact behavior. As an agent becomes more susceptible to temptation, modeled as her costs of self-control rising, she reduces her threshold. Because search can carry the risk of encountering more tempting items, agents who experience self-control problems want to search less, relative to agents who does not experience temptation. Changes in the distribution of commitment value that satisfy first order stochastic dominance and mean preserving spreads accord with the effects seen in the standard model — the threshold increases.

Section 5 compares the observable behavior of search under temptation to that generated by other models of search, including search with other models of self-control, and provide easy behavioral tests for assessing the appropriateness of each model. Section 6 summarizes the results and concludes. The appendices contain proofs and extend the techniques and results to other types of context-dependent preferences.

1.2 Related Literature

Although many types of context-dependent preferences have been provided decision-theoretic characterizations in both static and dynamic settings, there have been relatively few applications of them. Temptation and self-control based preferences have been widely applied, especially models of quasi-hyperbolic discounting. There are several models applying GP to consumption contexts including those developed by Amador, Angeletos and Werning (2006) and Krusell, Kuruscu and Smith (2011). Esteban, Miyagawa and Shum(2007), Esteban and Miyagawa (2006a, 2006b) solve for a firm's optimal menu of choice sets when consumers experience heterogeneous costs of self-control. O'Donoghue and Rabin (1999), DellaVigna and Paserman (2005), Fudenberg and Levine (2006), and Miao (2008) consider optimal stopping problems when agents have self-control problems. However, these papers typically model self-control as an intertemporal phenomenon; agents are tempted to delay costs and push forward benefits. A very similar paper is Irons and Hepburn (2007) , which considers a search setting with a finite number of periods when consumers are subject to regret.

Unlike previous models of optimal stopping with agents that suffer from self-control problems, I assume temptation is not a purely intertemporal phenomenon. This paper allow for temptation to occur for reasons that are not directly related to intertemporal concerns. Instead, individuals can simply find certain attributes of items tempting. Considering temptation and self-control as non-temporal phenomenon allows for novel of results, such as non-time stationary reservation values.

Although there is little direct evidence that temptation affects behavior in search, there is indirect evidence pointing to the importance of self-control in other settings, as well as evidence indicating the search behavior is effected by unchosen options.[4] Baumeister, Sparks, Stillman and Vohs (2008) report a large body of evidence showing that subjects who first faced a task meant to deplete willpower were significantly more likely to choose the more tempting option in a subsequent choice, generating a similar compromise effect to the one in this paper.[5] There is also evidence from psychology of the effects of context on behavior in search contexts, albeit evidence that focuses on regret, not temptation. Simonson (1992) and Cooke, Meyvis, and Schwartz (2001) find that either regret about options already observed, or options that will be passed up in the future, can induce changes in the agent's search behavior. In a closely related setting, Levav, Heitmann, Herrmann and Iyengar (2010) manipulate the order of attribute presentation in the product customization of suits and automobiles. Although the paper focuses not on temptation, but willingness to accept a default option, they find that consumers' purchase prices are sensitive to both what attributes are in the sequence of choice, and the order in which the attributes appear. They suggest that this means that consumers can be exploited by adjusting the placement of its high-priced attributes (which are oftentimes tempting).

Within economics there is a large empirical literature on search models. However, isolating the effects of temptation or other context-dependent preferences from other aspects of search, such as learning or boundedly rational heuristics, can be difficult in field data. Furthermore, as mentioned, the experimental work on search tends to provide outcomes that are monetary, and so do not possess different attributes which may be more or less tempting. Therefore, neither body of work provides a pre-existing data-set with which to test the model against.[6] However, as Section 5 makes clear,

[4]Most of the experiment evidence on search focuses on payoffs that are purely monetary in value, and so it is difficult to conceive of the subjects ranking them in the multi-dimensional fashion that temptation requires.

[5]Relatedly, Fudenberg and Levine (2006) and Noor and Takeoka (2010) are among a set of papers in economics that attempts to capture the compromise effect in static settings

[6]Although the empirical studies may not pose as direct tests of the model, insofar as context-dependent preferences introduce a path-dependent cost of search, they can help shed light on why individuals in settings may behave in an fashion that is inconsistent with a constant cost of search. For example, using field data collected from consumers who are searching for items among internet retailers, De los Santos, Hortacsu, and Wildenbeest (2011) report that even in an setting where items are observed sequentially agents behavior is not adequately described by the standard search model, even when accounting for learning. Similarly, Brown, Flinn and Schotter (2011) find that even when experimental subjects know that the distribution they are sampling from is stationary, they have a non-stationary

there are notable difference between the behavior predicted by this model of search and self-control and previous models of search, providing a way to distinguish between the alternative models.

1.3 Model

1.3.1 Preferences

In order to capture the effects of temptation, GP considered preferences not over outcomes, but over choice sets. They capture temptation as a preference for commitment: agents prefer to avoid choice sets that require self-control. [7] In this case, GP prove that all items are ranked according to two orders. First, a commitment value, the value of item if the agent could commit to only having that item in their choice set, denoted $u(x)$, and second, the temptation value: $v(x)$. Not choosing the most tempting item in a choice set imposes a cost of self-control, which is the difference in the temptation value of the chosen item and the most tempting item:

$$V(A) = \max_{x\in A}\{ \underbrace{u(x)}_{\substack{\text{commitment value} \\ \text{of chosen item}}} - \underbrace{(\max_{y\in A} v(y) - v(x))\}}_{\substack{\text{cost of} \\ \text{self-control}}}$$

The natural interpretation of second period choice is the agent chooses the lottery $\text{argmax}_{x\in A}\{u(x)+v(x)\}$, while facing a cost that depends on the choice set, specifically the temptation value of the most tempting item in the choice set $\max_{x'\in A} v(x')$. This can be more clearly seen by rewriting the value of choice set A as the difference between the highest total value (i.e. the sum of commitment and temptation values) of the items in A and the temptation value of the most tempting item in A:

$$V(A) = \overbrace{\max_{x\in A}\{u(x)+v(x)\}}^{\substack{\text{commitment and temptation} \\ \text{value of chosen item}}} - \overbrace{\max_{y\in A} v(y)}^{\substack{\text{temptation value} \\ \text{of most tempting item}}}$$

There is a gap between choice within a choice set and choice between choice sets. Singleton choice

threshold policy.

[7] Formally, GP capture temptation in the following fashion. Denote X as the set of outcomes, and 2^X as the set of choice sets, with arbitrary choice sets A and B. GP capture self-control in the following way: if $A \succsim B$ then $A \succsim A \cup B \succsim B$. This condition is called 'set-betweenness'. Using preferences over choice sets of lotteries, GP characterized temptation-based preferences as those that satisfy not only standard axioms (weak order, continuity, and independence), but also 'set-betweenness'. Although I do not explicitly maintain the fact that outcomes are lotteries, it can be justified by the assumption that item has having uncertain attributes. For example, one can never be sure of how much maintenance a car will require.

set $\{x\}$ has a value of $u(x)$. However, the item the agent chooses from any given choice set A is

$$c(A) = \operatorname*{argmax}_{x \in A} u(x) + v(x)$$

which may not be the item in the highest valued singleton choice set. The value of choice set A is always less than or equal to $u(x_c)$. Additional items beyond the one chosen can only harm the individual, and so there can be welfare gains from restricting the size of choice sets.

Preferences over choice sets can be represented graphically as in Figure 1.1. Items are points in (v, u) space. Choice sets are sets of points. For example, choice set A is a set of items. The indifference curves for $c(A)$, the first term of $V(A)$ (indifference curves over the item chosen, fixing the maximum amount of temptation experienced) are the set of points that give the same value of $u(x) + v(x)$, and so are lines of slope -1. The indifference curves for the second term of $V(A)$ (indifference curves over the most tempting item, fixing the item chosen from a choice set) are the set of points that give the same value for $v(x)$, and so can be represented as vertical lines.

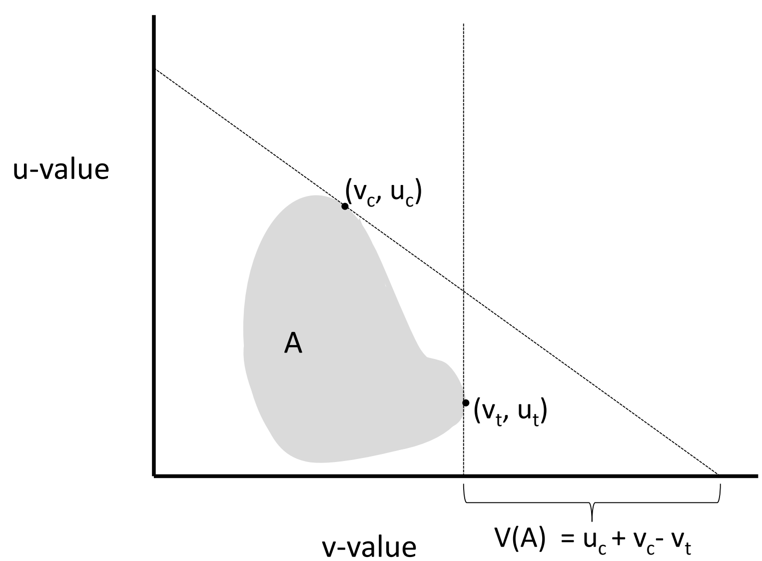

Figure 1.1: Graphical representation of temptation-based preferences

Graphically, the value of $c(A)$ is the value $u(x) + v(x)$ of the first type of indifference curve that passes through an element of A, and is farthest to the northeast in Figure 1.1. The value of the second term in $V(A)$ is the value $v(x)$ of the second type of indifference curve that passes through an element of A, and is farthest to the east in Figure 1.1. The utility of choice set A is the difference between the v-intercepts of the highest valued positive indifference curve and the highest valued negative indifference curve.

1.3.2 The Agent's Problem

The environment the agent faces is relatively simple. There is a set of outcomes X, for example, cars, available for purchase. Each outcome $x \in X$ has an associated temptation and commitment values $v(x)$ and $u(x)$ values, and so each outcome is uniquely determined by these values. Goods are jointly distributed according to a distribution $F(v, u)$, which has support on $M = [\underline{v}, \bar{v}] \times [\underline{u}, \bar{u}] \subset \mathbb{R}^2_+$.[8]

The agent engages in sequential search. In any given period t, the agent starts with choice set $A_{t-1} = \{x_1, x_2, ..., x_{t-1}\} = \{(v_1, u_1), ...(v_{t-1}, u_{t-1})\}$. At the beginning of each period she observes an outcome, denoted $x_t = (v_t, u_t)$, drawn from F.[9] At this time the agent must decide between two actions: stop or continue. If she stops, she must choose a good from the choice set $A_t = \{x_1, x_2, ..., x_t\} = A \cup x_t$.[10] At this time she receives a payoff[11]

$$U(A_t) = \max_{x \in A_t}\{u(x) + v(x)\} - \max_{y \in A_t} v(y).$$

If instead, the agent continues searching, she receives a payoff of 0 in period t, and the next period begins. The agent discounts the future by $\beta < 1$.[12] The agent faces no temptation until she makes a choice among the items she has observed during the search process.

As with other search problems, the agent's problem can be defined recursively. First, there are only two goods in any choice set that are payoff relevant — the good that would be chosen if the agent stopped, $\text{argmax}_{x \in A}\{u(x) + v(x)\}$, and the most tempting good, $\text{argmax}_{x \in A} v(x)$. Denote the former as $x_c = (v_c, u_c)$, and the latter as $x_t = (v_t, u_t)$. However, because u_t is not payoff relevant, the state space of the problem can be reduced to three variables (v_c, u_c, v_t). Using the newly defined state variables, the payoff from stopping is $v_c + u_c - v_t$. The payoff from continuing to search depends on what good will be observed next period, and therefore added to the choice set of options.

[8] The only result that depends on assumption of bounded support is the exact functional form of the threshold

[9] For expositional ease, I assume A_0 is the null set, so that in period 1 the agent faces a singleton choice set. No results in the derivation of the optimal policy depend on this.

[10] Note that the assumption of perfect recall helps motivate the fact previous items are tempting. In fact, if the agent is not tempted by items they have previously seen, but cannot recall, the agent can be made better off. It will also generate very different search behavior. For example, individuals will now have an incentive to continue searching at times in order to rid themselves of particularly tempting items. Of course, context-dependence still has bite if individuals' payoffs are affected by items that cannot be recalled. This makes particular sense in contexts such as regret. See Raymond (2010) for details on identification of recall and its interaction with preferences.

[11] Although this paper uses what seems to be the static model payoff of GP in an explicitly dynamic setting, Raymond (2010) provides a decision-theoretic characterization of preferences in search contexts, including the preferences used in this paper

[12] In certain situations one might imagine that the cost of search is not time, but rather a utility cost of c per period searched. In this case the qualitative results concerning the threshold remain unchanged.

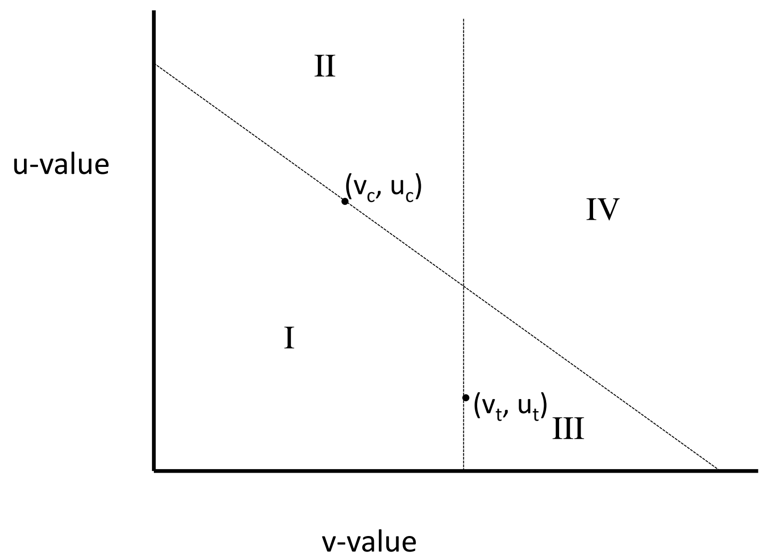

Figure 1.2: Outcome of additional observation

Given that there are two state points (x_c and x_t), there are four possible outcomes next period, depending on whether one or both of the state points are replaced by new draw. In Figure 1.2 these four outcomes are shown graphically. The current choice set is A, and the good that is observed next period is denoted $x' = (v', u')$. In the first case (Region I), the newly observed good will not change either state point, i.e. it would neither be chosen from $A \cup x'$ if the agent stopped nor be the most tempting item in $A \cup x'$. This means that $v' + u' \leq v_c + u_c$ and $v' \leq v_t$. In the second case (Region II), the newly observed good would chosen from $A \cup x'$ if the agent stopped, but it is not the most tempting item, so $v' + u' > v_c + u_c$ but $v' \leq v_t$. The third case (Region III) is where the newly observed good would not be chosen from $A \cup x'$ if the agent stopped, but it is the most tempting item, so $v' + u' \leq v_c + u_c$ but $v' > v_t$. In the last case (Region IV), the newly observed item is both the item that would be chosen from from $A \cup x'$ if the agent stopped and is the most tempting item: $v' + u' > v_c + u_c$ and $v' > v_t$.

Table 1.1 summarizes the results of search for the agent. Note that the payoffs to stopping change depending on the new draw. In particular, in Region IV, since the newly observed item is both the item that would be chosen from $A \cup x'$ if the agent stopped, and is the most tempting item, the change in payoff to stopping can either increase or decrease. This is because the newly drawn item changes both state variables, and if the increase in $v_c + u_c$ more than offsets the increase in v_t, then the payoff to stopping rises. Otherwise it falls.

Region	Change (v_c, u_c)	Change (v_t, u_t)	Stopping Payoff
I	No	No	Same
II	Yes	No	Better
III	No	Yes	Worse
IV	Yes	Yes	Varies

Table 1.1: Summary of Outcomes from Search

Given these state variables, and the four regions, the agent's problem can then be recursively defined as max $W(v_c, u_c, v_t)$, where

$$\begin{aligned} W(v_c,u_c,v_t) &= \max_{stop,search}\{\overbrace{v_c+u_c-v_t}^{\text{Payoff from Stopping}}, \\ & \qquad \beta(\underbrace{\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{v_c+u_c-v'} W(v_c,u_c,v_t)dF(v',u')}_{\text{Region I}} \\ &+ \underbrace{\int_{\underline{v}}^{v_t}\int_{v_c+u_c-v'}^{\bar{u}} W(v',u',v_t)dF(v',u')}_{\text{Region II}} \\ &+ \underbrace{\int_{v_t}^{\bar{v}}\int_{\underline{u}}^{v_c+u_c-v'} W(v_c,u_c,v')dF(v',u')}_{\text{Region III}} \\ &+ \underbrace{\int_{v_t}^{\bar{v}}\int_{v_c+u_c-v'}^{\bar{u}} W(v',u',v')dF(v',u'))\}}_{\text{Region IV}} \end{aligned}$$

The first element of the maximization is the payoff from stopping, while the second element is the payoff from continuing to search (the payoff from each case discussed above presented in their respective order). I refer to the first term inside the max operator as the value of stopping, and the second as the value of search. Although the state space has three elements, by construction if $v_c + u_c = v_c' + u_c'$, it must be the case that $W(v_c, u_c, v_t) = W(v_c', u_c', v_t)$ since the payoff to stopping is the same, and the four regions that define the payoff to searching are also the same. Therefore, I will sometimes say that $W(v_c, u_c, v_t) \equiv W(v_c + u_c, v_t)$.

1.4 Results

The recursively defined agent's problem has a unique continuous solution because of the contraction mapping theorem.

Proposition 1 *There exists a unique continuous function $W(v_c, u_c, v_t)$ that solves the agent's problem.*

This result is standard in the search literature. In this model, as in the standard model, search generates a benefit: the ability to see additional options. Unlike the standard model, search also imposes a costs (in addition to the implicit cost of discounting): the risk of observing a more tempting item. Fixing the most tempting item already observed (fixing v_t), an increase in either the commitment or temptation value of the item that would be chosen if the agent stopped (an increase in $u_c + v_c$) raises the value of stopping. Furthermore, it decreases the value of search, because the agent is less likely to find a better option. Fixing the item that would be chosen if the agent stopped (fixing $u_c + v_c$), an increase in the temptation value of the most tempting item (an increase in v_t), reduces the value of stopping today, but also increases the value of search, because the agent is less likely to come across an item that increases the level of temptation currently faced.

The agent must balance these costs and benefits when deciding when to search or stop. In fact, the agent's optimal decision rule implies that the value function follows a threshold rule: the agent's value is equal to the value of stopping if and only if the best item in the choice set has a total value above a threshold that is conditional on the amount of temptation currently experienced. This means that if the agent is better off stopping when she comes across a good item, then she is also better off stopping if she comes across a better item (conditional on this item not changing the maximum amount of temptation experienced).

Intuitively, the threshold policy is optimal because fixing v_t, the value of stopping and the value of search can cross only once. This is because the value of search is always (weakly) rising less than the value of stopping in $u_c + v_c$. To see why this is true, note what happens in each region in Figure 2.1. If the new item observed is in Region I, because search does not change any state variables, it leaves the decision-maker unaffected. In Region II, search is solely beneficial, as it increases the value of the item that would be chosen, while leaving the amount of temptation faced unchanged. In Region III, search is solely harmful, because it only increases the amount of temptation faced. In Region IV, there are both benefits and costs to search (as the new item is both the one that will be chosen, and the most tempting item), and their relative value depends on the actual commitment and temptation value.

Formally, suppose the agent stops searching with state variables (v, u, v_t) and consider (v', u', v_t) such that $u + v \leq u' + v'$. This change represents a shift of mass from Region II to Region I, and from Region IV to Region III. Recall that shifting mass from Region II to Region I shifting mass from a region where search is beneficial to the value of stopping to a region where it doesn't change the value of stopping. Similarly shifting mass from Region IV to Region III shifts mass from a region where search could be beneficial to the value of stopping to where it is always harmful. Both these shifts represent a reduction in the value of search.

Therefore, the value of stopping has risen (since the payoff is higher), while the benefits of search have fallen (since the agent is less likely to come across an even higher payoff), and the costs of search have not changed (since the agent is as likely to come across a more tempting item).

Proposition 2 *The value function follows a threshold: there exists a function* $\overline{w}(v_t) : \mathbb{R} \to \mathbb{R}$ *such that* $v_c + u_c \geq \overline{w}(v_t)$ *if and only if* $W(v_c, u_c, v_t) = v_c + u_c - v_t$

In the standard model, the value function also follows a threshold rule. However, in the standard model this threshold only depends on the agent's beliefs about the distribution she is facing. In this model, the beliefs about the distribution are stationary. However, the threshold depends on v_t. This is because v_t is a measure of the current cost the agent will pay upon stopping, and as such, is also a measure of the marginal expected cost the agent expects to face from an additional draw (since an additional draw can increase v_t).

The fact the value function follows a threshold rule implies that the policy function also follows a conditional threshold.

Corollary 1 *The optimal policy is a threshold: the agent stops if* $v_c + u_c > \overline{w}(v_t)$*, is indifferent between stopping and continuing searching if* $v_c + u_c = \overline{w}(v_t)$*, and continues searching if* $v_c + u_c < \overline{w}(v_t)$*.*

1.4.1 Threshold Effect

Over the course of the search process, the agent can expect the most tempting item she observes to weakly increase in the level of temptation. When the agent observes a more tempting item, the implicit psychological cost of additional search falls. This is because she is less likely to draw an item that will increase the maximum value temptation in her choice set. Fixing the item that would be chosen if the agent stopped, an increase in v_t causes the cost of stopping to rise (because the temptation cost must then be paid) and the cost of searching to decline (since more tempting items are less likely to be encountered). Therefore, the agent will find search relatively more attractive, and will search for higher values of $u_c + v_c$. Formally, consider state variables (v_c, u_c, v_t) where $v_c + u_c = \overline{w}(v_t)$, and then let v_t increase to v_t'. First, this directly reduces the value of stopping. Second, this also causes a shift in mass from Region III, a region where search makes the agent worse off, to Region I, region where search does not change the payoff. Third, this shifts mass from Region IV, a region where search confers both benefits and costs, to Region II, a region where search is solely beneficial. The latter two effects mean that the costs of search fall, since it is less likely that the agent will observe a more tempting item. Since the benefits to search are rising, and the payoff to stopping is falling, in order to keep stopping indifferent to searching, the benefits of stopping must rise, while the benefits of search must fall. An increase in the threshold is the only way to accomplish both.

Proposition 3 *The threshold payoff $\overline{w}(v_t)$ is increasing in the level of temptation v_t.*

The fact that the threshold is increasing in the maximum temptation value in a choice set has an immediate implication for observable choice — any item chosen by agent will be the item most recently observed. Since the threshold $\overline{w}(v_t)$ weakly rises after each additional draw, any option previously unacceptable is still unacceptable

Corollary 2 *The agent, when stopping, chooses the item last observed.*

Because the threshold is rising as more temptation is experienced, continuing to search after observing a more tempting item means that the agent will expect to search longer, as there are fewer items that fall above the new, higher threshold. Furthermore, since v_t is weakly increasing after every draw, then the threshold is weakly rising after every draw. Therefore, the probability that the next good observed falls above the threshold falls after every draw.

Corollary 3 *Conditional on still searching at time t, the expected length of future search is growing.*

Since the agent will never go back and choose an item that she previously passed over, and because the items are drawn independently from the distribution F, if the agent searches the value function depends only on v_t.

Lemma 1 *If $u_c + v_c < \overline{w}(v_t)$, then the value function is a function only of v_t: $W(v_c, u_c, v_t) = \overline{w}(v_t) - v_t$*

For choice sets where the item that the agent would choose if she stopped is below the threshold, the value function is equal to $\overline{w}(v_t) - v_t$. For choice sets where the item that the agent would choose if she stopped is above the threshold the value function rises one-for-one with the total value of this item: $u_c + v_c$.

1.4.2 Compromise Effect

Since the threshold changes over time, the set of items an agent is willing to stop and choose changes with v_t. Denote the set of items that an agent could choose, given her optimal policy, and conditional on v_t, as $\mathbb{S}(v_t)$. Because the threshold rises with v_t, $\mathbb{S}(v_t)$ shifts in (v, u) space with v_t. Characterizing $\mathbb{S}(v_t)$ will allow us to better predict individual behavior in search settings.

A useful way of thinking about the threshold is to restrict consideration to choice sets where the item that would be chosen if the agent stopped is the most tempting item. This is equivalent to restricting consideration to choice sets where the state space can be represented as (v_t, u_c, v_t). In this case the agent's decision to stop is solely determined by comparing the value of u_c to a function of v_t: $\overline{w}(v_t) - v_t$. Denoting this value as $\mu(v) \equiv \overline{w}(v) - v$, then the agent stops if and only if $u_c \geq \mu(v)$. In other words, given v_t, $\mu(v_t)$ is the minimum commitment value that the agent requires in order to stop. Therefore $\mathbb{S}(v_t) = \{(v, u) | u + v \geq \mu(v_t) + v_t \text{ and } v \leq v_t\}$.

To understand how $\mu(v_t)$ changes with v_t recall that value function W does not depend on u_c or v_c if $u_c + v_c < \overline{w}(v_t)$. Furthermore, conditional on $u_c + v_c < \overline{w}(v_t)$ the value function W is falling in v_t at a rate less than one-for-one. Consider a choice set A containing item $x = (v, u)$, such that $v_t = v_c = v$ and $u_c = u$; x is the most tempting item in A as well as the item that the agent would choose if she stopped. Furthermore, suppose that the agent is indifferent between searching and stopping, and so $u_c = \mu(v_t)$. Now consider some other choice set A' with item $x' = (v', u)$ such that $v'_t = v'_c = v' > v$ and $u_c = u$. Therefore x' is the item that would be chosen from A' if the agent stopped, and is also the most tempting item in A'. Relative to A, the value of stopping has not changed at A' since the payoff is u in both cases. However, the value of stopping at the threshold $\overline{w}(v'_t)$ must fall relative to $\overline{w}(v_t)$. Therefore the agent must strictly prefer to stop, and so $\mu(v_t) > \mu(v'_t)$. This directly implies that the threshold is rising less than one-for-one with v_t.

Proposition 4 $\mu(v_t) \equiv \overline{w}(v_t) - v_t$ *is decreasing in v_t.*

Proposition 4 has a direct implication about $\mathbb{S}(v_t)$. First, the minimum temptation value that the agent is willing to stop at is rising in v_t. Second, conditional on the stopping with v_t, the minimum commitment value that an agent is willing to stop at is falling in v_t. When an agent observes a more tempting item she rules out choosing some alternatives. Items that are no longer acceptable to stop at are those that are relatively high in commitment value but low in temptation value. However, because v_t has increased, the agent could also stop and choose some items that she previously could not. These items are relatively high in temptation value but low in commitment value. This compromise effect, whereby the agent shifts toward choosing more tempting items, is part of the optimal strategy of the agent. It captures an intuition about the world that as consumers observe more tempting outcomes, they end up compromising. Although they do not necessarily choose the most tempting item that have seen so far, such as an Italian sports car, they do end up choosing something that has more tempting attributes than what they initially would have preferred to choose.

Proposition 5 $\min_{(v,u)\in\mathbb{S}(v_t)} v$ *is rising* v_t*, while* $\min_{(v,u)\in\mathbb{S}(v_t)} u$ *is falling* v_t*.*

Figure 1.3 summarizes the stopping region $\mathbb{S}(v_t)$ for an agent. $\mu(v_t)$ is drawn with a dotted line. Conditional on currently facing a most tempting item with temptation level v_t, first find the value $\mu(v_t)$, and this characterizes $\mathbb{S}(v_t)$. The agent stops during the current period if any item in the choice set falls above the line with slope -1 that intersects $(v, \mu(v))$ (drawn as a solid line), and to the left of v_t.

1.4.3 Characterizing the Threshold

When an agent is deciding whether to stop she must not only consider what the likelihood of drawing additional items above her current threshold is, but she must also think about what her new threshold will be if the most tempting item she has observed during search changes. Therefore, the threshold for any particular v_t depends on threshold for all $v'_t > v_t$. This means that the functional form of the threshold can be characterized by two conditions. The first is a condition describing the threshold at the maximum level of temptation: $\overline{w}(\bar{v})$. Because there is no possibility of drawing a more tempting item, this condition is analogous to the threshold in the standard model (and can be solved for in the same manner). Given this "final" threshold, a differential equation then characterizes the threshold for all other levels of temptation, which captures the agent's forward looking behavior.

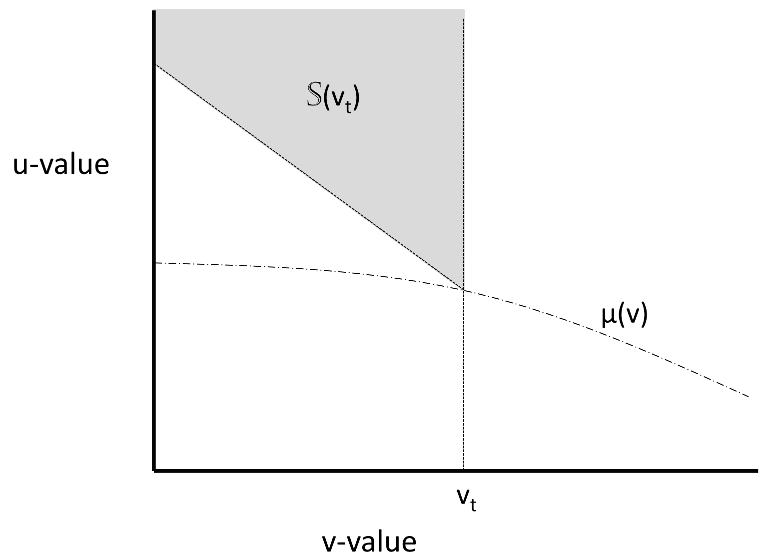

Figure 1.3: Stopping Regions

Proposition 6 *The threshold function $\overline{w}(v_t)$ is defined by two conditions. The first is the threshold at $\overline{w}(\bar{v})$:*

$$\overline{w}(\bar{v}) - \bar{v} = \beta\Big(\int_{\underline{v}}^{\bar{v}} \int_{\underline{u}}^{\overline{w}(\bar{v})-v'} (\overline{w}(\bar{v}) - v_t) dF(v', u') + \int_{\underline{v}}^{v_t} \int_{\overline{w}(\bar{v})-v'}^{\bar{u}} (u' + v' - \bar{v}) dF(v', u')\Big)$$

The second is a differential equation describing $\frac{d\overline{w}(v_t)}{dv_t}$:

$$\frac{d\overline{w}(v_t)}{dv_t} - 1 \quad = \quad \beta\Big(\int_{\underline{v}}^{v_t} \int_{\underline{u}}^{\overline{w}(v_t)-v'} \frac{\partial \overline{w}(v_t)}{\partial v_t} dF(v', u') - \int_{\underline{v}}^{v_t} \int_{\underline{u}}^{\bar{u}} (1) dF(v', u')\Big)$$

1.4.4 Comparative Statics

Despite the fact there is no general closed analytic solution to the differential equation, useful comparative statics can still be derived. These can help better predict behavior of agents across different situations. I will consider two types of comparative statics. The first is when agents become more or less susceptible to temptation. The second is when the environment changes — the distribution of u shifts with respect to first order stochastic dominance and mean preserving spreads.

It is quite plausible that different consumers will have heterogeneous reactions to temptation. In this paper, I capture this heterogeneity by by allowing the the cost of self-control to vary by individual. Recall that previously, the agent's payoff upon stopping with state variables (v_c, u_c, v_t) was $u_c - (v_t - v_c)$. I will now introduce heterogeneity by parameterizing the cost of self-control. The agent's payoff upon stopping, with state variables (v_c, u_c, v_t), will be $u_c - \alpha(v_t - v_c)$, where $\alpha \geq 0$ (since the cost of self-control always reduces the utility of the consumer). α parameterizes the agent's cost of self-control. If this increases, α rises.[13] When $\alpha = 0$ the agent simply maximizes u, and so behaves as if she does not experience any temptation at all, but is a classical consumer engaged in sequential search from a distribution F_u. In contrast, as $\alpha \to \infty$, for any set of state variables (v_c, u_c, v_t), where $v_c \neq v_t$, the cost of self-control also goes to infinity. Therefore, the consumer will simply choose the most tempting item. In fact, as α rises, and the cost of self-control rises, the agent's threshold will fall. This is because the cost of self-control has risen and so the benefits to search have fallen.

Proposition 7 *Fix F and suppose there are two agents, with preferences (u, v, α_1) and (u, v, α_2), such that $\alpha_1 \geq \alpha_2$. If $u_c + \alpha_2 v_c \geq \overline{w}_{\alpha_2}(\alpha_2 v_t)$ then $u_c + \alpha_1 v_c \geq \overline{w}_{\alpha_1}(\alpha_1 v_t)$.*

Intuitively, this says that the threshold for Agent 1 is lower than for Agent 2 — whenever Agent 2 wants to stop, so does Agent 1. The intuition for this result can be obtained by considering a situation where the agents have observed only a single outcome, so the state variables are $(v_t, u_c v_t)$ — therefore the item that would be chosen if any agent stopped is also the most tempting item. First, consider Agent 1, and assume that $v_c + u_c = \overline{w}_1(v_t)$, so Agent 1 is indifferent between stopping and searching. The value of stopping today is simply u_c (since the agent pays no cost of self-control). This is equal to the value of search, which is $W_{\alpha_1}(u_c, v_t, v_t)$.

Next consider Agent 2, where $\alpha_2 \leq \alpha_1$, but in the exact same situations, with state variables (u_c, v_t, v_t). Agent 2 also pays no cost of self-control if she stops, and so her payoff from stopping is simply u_c. Rather than considering what Agent 2's payoff to searching is given her optimal policy, I will first consider her payoff under a (weakly) sub-optimal policy. In particular, suppose that Agent 2 uses Agent 1's optimal policy if she decides to search. In this case, Agent 2's probability of ending up actually stopping at any particular set of state variables (u_c', v_c', v_t') is exactly the same of Agent

[13] To understand the basic static implications of varying the cost of self-control, suppose that two agent's are the same other than differing in their cost of self-control, so that $\alpha_1 \geq \alpha_2$. Then it is the case the two agents' preferences agree over singleton sets. Furthermore, for any item x and choice set A if $A \succeq_1 x$ then $A \succeq_2 x$. However, this is a different comparative static than the ones discussed in GP. They allow the correlation between temptation and commitment value to vary. These comparative statics, while interesting, do not allow a comparison of the thresholds between individuals, because the rankings of items in one or both dimensions will vary across individuals

1's probability, because they are using the same policy and F is fixed. However, for any set of state variables (u_c', v_c', v_t') the payoff to stopping for Agent 2 is weakly higher than the payoff to Agent 1. This is because $u_c' - \alpha_2(v_t' - v_c') \geq u_c' - \alpha_1(v_t' - v_c')$. Therefore, the value function at any point must be weakly higher for Agent 2 relative to Agent 1 (since the value function is simply the average of the payoffs upon stopping). Therefore the value to searching at (v_t, u_c, v_t) must be weakly higher (since F has not changed) for Agent 2 than for Agent 1. Since the payoff to stopping is the same for both agents the threshold must weakly rise.

Next, I will consider shifts in the distribution of the commitment value of items.[14] First, consider a shift in the distribution F to F', such that for any v, the marginal distribution $F'_{u|v}$ first order stochastically dominates (FOSD) $F_{u|v}$. This means that the probability of drawing items with higher commitment values is higher. Conditional on any (v_c, u_c, v_t), weight is being shifted from Region I to Region II, and from Region III to Region IV. Furthermore, within each region, weight is being shifted from low u values to high u values. All of these make search more valuable. Therefore the threshold should rise under F' relative to F.

Similarly, consider a shift from F to F', such that such that for any v, the marginal $F'_{u|v}$ is a mean preserving spread (MPS) $F_{u|v}$. In the current model, like in other search models, search contains an option value for seeing a higher value of u, which is solely beneficial, and so a distribution that is more risky is beneficial. This result is entirely analogous to the result in the standard model.

Proposition 8 *Let $F'_{u|v}$ FOSD $F_{u|v}$ for all v or let $F'_{u|v}$ be a mean preserving spread of $F_{u|v}$ for all v. Then $\overline{w}_{F'}(v_t) \geq \overline{w}_F(v_t)$ for all v_t.*

1.5 Comparison to Alternative Models

Of course, there are a variety of other models of search with different implications for observable behavior. These include the basic model of search from a known distribution, search with learning (about an initially unknown distribution), and search with other types of biases. Key features that distinguish these models from one another include whether the threshold is stationary, how the threshold moves if it is non-stationary, and what item is chosen upon stopping.

The model discussed in this paper nests the standard model of search from a known distribution, where agents do not suffer from temptation. In this case $v = 0$ for all items. As shown in McCall

[14] I do not consider equivalent shifts in the temptation value of items. Although changes in the distribution of v which do not satisfy the conditions of Proposition 7 are mathematically meaningful, they are behaviorally difficult to interpret. For example, if v is a point mass, then the agent experiences no cost of self-control, regardless of the level of v.

Model	Stationary Threshold	Monotonic Threshold	Sophisticated Agents
Standard Model	Yes	N/A	Yes
Quasi-Hyperbolic	Yes	N/A	Yes and No
With Learning	No	No	Yes
With Temptation	No	Yes	Yes

Table 1.2: Summary of Search Models

(1970) this generates a threshold policy, but one that never varies (as would be expected from this paper's model when there is only a single value of v). This results in the agent, when stopping, always choosing the item last observed. Because the threshold never varies, the exit rate of agents from search (i.e. the hazard rate) is constant over from time.

When the agent does not suffer from temptation, but is uncertain about the distribution she is drawing from, there are far fewer general results. In fact, depending on the possible distributions, the optimal policy may not even be a threshold (for examples in this literature, see Rothschild, 1974, Burdett and Vishwanath, 1988, and Sharma and Bikhchandani, 1996). Even if the optimal policy is a threshold, the reservation value can rise or fall depending on the actual realization of outcomes, in contrast to the current paper. Although it is certainly the case that in the real-world learning is almost always occurring, it is still informative to think about how to compare the predictions of the model in this paper with those generated by models that accommodate learning about the distribution of outcomes. For example, it is possible to compare the effects of showing items at the beginning of search that are either relatively informative about the set of prices available and not tempting, or relatively tempting but uninformative.

As mentioned previously, several other papers have considered optimal stopping and search in the context of agent's with self-control problems. O'Donoghue and Rabin (1999), Miao (2008), and Fudenberg and Levine (2006) all consider optimal stopping when consumers have self-control problems. DellaVigna and Paserman (2005) model an agent engaging in wage search with quasi-hyperbolic discounting with variable search effort. In these models, there is no recall, so the choice is always between choosing from a singleton choice set today and continuing to search. Self-control problems arise because the agent faces intertemporal trade-offs of costs and benefits. As a result, these papers find that the optimal threshold policy is stationary. This is true regardless of whether the agents are sophisticated (they can accurately predict their self-control problems) or naive (they believe that although they suffer from self-control problems now, they will not in the future). In particular, Miao (2008), Fudenberg and Levine (2006), and DellaVigna and Paserman (2005) find a unique optimal policy with a stationary threshold. Furthermore, Miao (2008) and Fudenberg and

Levine (2006) find that an extension of O'Donaghue and Rabin (1999) to a stochastic environment with sophisticated agents produces multiple equilibria, but that at least one of them features a stationary threshold.[15] In contrast, the model in this paper allows self-control problems to arise because of temptation across items available at any given time, and can therefore accommodate behavior that is path-dependent.[16] Table 1.2 summarizes some of the differences between the current model (last line of the table) and the other models discussed in this section.

1.6 Conclusion

In this paper I develop a new application of temptation-based preferences, modeling an agent that first engages in search to construct an "optimal" choice set from which she will then choose. I show that the policy the agent follows is a threshold: conditional on experiencing a particular level of temptation, the agent will stop if and only if the best item in an agent's choice set exceeds a certain value. Otherwise the agent will continue to search. I also consider comparative statics of the threshold, and discuss some testable implications that distinguish this model from other models of search.

An important consideration in much of the behavioral literature is the difference between sophisticated agents and naive agents. This paper assumes that agents are sophisticates. One way to model naivete in the current settings is to assume that the agent knows that she will experience self-control problems if she stops today, but believes that in the future she will not experience any problems with self-control. In this case the threshold effect will still occur. This is because although the agent believes the value of search is constant, regardless of her draw, the value of stopping falls when she observes more tempting items, all else being equal. Furthermore, the agent will, upon observing a more tempting item, still rule out items relatively low in temptation value.

The results in this paper are independent of the distribution of alternatives. This raises questions about what types of market outcome should be expected when consumers are known to have self-control problems. In ongoing work, Raymond (2011), I consider market outcomes when there are rational profit maximizing firms and consumers who experience temptation. If temptation value is relatively cheaper to produce than commitment value for firms, then it might be expected that firms would offer product lines that induce consumers to shift their purchases toward items relatively high

[15] Although I do not provide explicit results for naifs in a stochastic environment, because there is no recall, a threshold strategy will be used because the agent believes that they will have no self-control problems

[16] Although Miao allows temptation to be generated by either the payoff from stopping or continuing, the structure of temptation is such that it can by captured by my model if the agent has no recall.

in temptation value but low in commitment value. However, in equilibrium both competition and the fact that consumers can avoid all options by not searching discipline firm behavior.

Although this paper focused on examining the implications of GP temptation-based preferences in search settings, the techniques used here are applicable to the analysis of other types of context-dependent preferences, including other models of temptation, regret and a preference for variety. These preferences also generate psychological costs and benefits of search that are path dependent. The related models of self-control by Dekel, Lipman and Rustichini (2009), Stovall (2010) and Masatlioglu, Nakajima and Ozdenoren can also be analyzed using the tools developed here, and generate similar results to those described in this paper. There are also other types of context-dependent preferences that have consequences for search behavior. For example, Irons and Hepburn (2007) develop a model of finite search over a set of n objects that results in a path dependent threshold. Their agents experience regret at having either passed up good choices in the past, or passing up possibly better choices by stopping now. However, they only solve for a three-period model of search. The techniques in this paper allow for a generalization of their approach to infinite horizon search problems, and can accommodate not just regret, but a wide variety of other preferences.

1.7 Appendix A: Proofs

I will endow the set of continuous real-valued functions on M with the sup-norm, which makes it a complete metric space.

Proof of Proposition 1

I use Blackwell's sufficiency conditions to show that the value function defines a contraction mapping, and therefore there exists a unique solution. Consider the set of bounded functions on M. Endowed with the sup norm metric this forms a complete metric space. Denote $x = (v_c, u_c, v_t)$, and let T be the functional

$$\begin{aligned} T(W)(x) &= \max\{v_c + u_c - v_t, \beta(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{v_c+u_c-v'} W(v_c, u_c, v_t)dF(v', u') + \int_{\underline{v}}^{v_t}\int_{v_c+u_c-v'}^{\bar{u}} W(v', u', v_t)dF(v', u') \\ &+ \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{v_c+u_c-v'} W(v_c, u_c, v')dF(v', u') + \int_{v_t}^{\bar{v}}\int_{v_c+u_c-v'}^{\bar{u}} W(v', u', v')dF(v', u'))\} \end{aligned}$$

that maps continuous bounded functions into continuous bounded functions. Therefore, by Theorem 3.3 of Stokey, Lucas and Prescott (1989), if T satisfies monotonicity and discounting, then a unique solution W^* exists to the decision maker's problem.

First I show monotonicity. Let W'' and W' be bounded functions on M, with $W''(x) < W'(x)$ for all $x \in M$.

It is clear that regardless of $W(x)$, the value of stopping is constant - it depends only on x. However, value of searching

$$\begin{aligned} S &= \beta(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{v_c+u_c-v'} W(v_c,u_c,v_t)dF(v',u') + \int_{\underline{v}}^{v_t}\int_{v_c+u_c-v'}^{\bar{u}} W(v',u',v_t)dF(v',u') \\ &+ \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{v_c+u_c-v'} W(v_c,u_c,v')dF(v',u') + \int_{v_t}^{\bar{v}}\int_{v_c+u_c-v'}^{\bar{u}} W(v',u',v')dF(v',u')) \end{aligned}$$

depends on the function $W(x)$. If $W\prime\prime(x) < W\prime(x)$ for all x, then each element of the value of search is less under $W\prime\prime$ than under $W\prime$. Therefore $T(W\prime\prime(x)) \leq T(W\prime(x))$.

Next I show discounting. Consider a function $W(x)$ on M. Define a series of functions on M: $(W+a)(x) = W(x)+a$ for all $a \geq 0$. Then:

$$\begin{aligned} T(W+a)(x) &= \max\{v_c+u_c-v_t, \beta(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{v_c+u_c-v'} (W(v_c,u_c,v_t)+a)dF(v',u') \\ &+ \int_{\underline{v}}^{v_t}\int_{v_c+u_c-v'}^{\bar{u}} (W(v',u',v_t)+a)dF(v',u') \\ &+ \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{v_c+u_c-v'} (W(v_c,u_c,v')+a)dF(v',u') \\ &+ \int_{v_t}^{\bar{v}}\int_{v_c+u_c-v'}^{\bar{u}} (W(v',u',v')+a)dF(v',u'))\} \\ &= \max\{v_c+u_c-v_t, S+\beta a\} \end{aligned}$$

Suppose that the maximization process in T chooses the first element (stopping) for $T(W+a)(x)$. Then it most be the case the maximization process chooses the first element (stopping) for $T(W(x)+0)$. Then $T(W+a)(x) = v_c+u_c-v_t < v_c+u_c-v_t+\delta a$ for all $\delta \in (0,1)$.

Now suppose the maximization process chooses the second element (continuing to search) for $T(W+a)(x)$. There are two cases.

- Assume that the second element (continuing to search) is also chosen under $T(W)(x)$. Then $T(W+a)(x) = S+\beta a < S+\delta a = T(W)(x)+\delta a$ for $\delta \in (\beta,1)$.

- Assume that the first element (stopping) is better under $T(W)(x)$. Then $T(W+a)(x) = S + \beta a < v_c + u_c - v_t < v_c + u_c - v_t + \delta a = T(W)(x) + \delta a$ for all $\delta \in (0,1)$.

Therefore T is a contraction mapping, and so there exists a unique continuous bounded W^* that solves the agent's problem.

□

Proof of Proposition 2 and Corollary 1

A short set of results will demonstrate that the value function follows a threshold that is conditional on the amount of temptation observed. To see this, first I will show that conditional on the amount of temptation, the value function is weakly increasing in $u_c + v_c$, but never more than at a rate one-for-one with $u_c + v_c$.

Lemma 2.1 *Fixing* v_t

1. $W(v_c, u_c, v_t)$ *is weakly increasing in* $u_c + v_c$ *and*

2. $\frac{\partial W(v_c,u_c,v_t)}{\partial(u_c+v_c)} \leq 1$ *wherever* $W(v_c, u_c, v_t)$ *is differentiable.*

Proof of Lemma 2.1 I prove these properties using the contraction mapping theorem. Note that the set of functions such that are weakly increasing, and the set of functions such that $\frac{\partial W(v_c,u_c,v_t)}{\partial(u_c+v_c)} \leq 1$ wherever W is differentiable are both closed sets. I assume that a function $W(x)$ possesses these two properties and show that they are preserved under $T(W)(x)$. Then by Corollary 1 of Stokey, Lucas and Prescott (1989) the solution to the decision maker's problem must possess these two properties.

1. Assume W is increasing in $u_c + v_c$. If $T(W)$ equals the value of stopping, then clearly it is increasing in $u_c + v_c$. If $T(W)$ equals the value of searching, then there are two effects. Assume $u_c + v_c$ increases. Within Regions I and III, W rises because $u_c + v_c$ is rising. W is constant within Regions II and II. Fixing the bounds of integration, this causes the value of $T(W)$ to rise. Furthermore, the bounds of integration shift. Region I grows at the expense of Region II. Assume some point (u', v') shifts from Region II to Region I. This increases the W at (u', v', v_t) since $u_c + v_c \geq u' + v'$. Similarly, Region III grows at the expense of Region IV. Assume some point (u', v') shifts from Region III to Region IV. This increases the W at (u', v', v_t) since $u_c + v_c \geq u' + v'$. Therefore $T(W)$ must be increasing. This immediately implies that $T(W)$ must be differentiable almost everywhere with respect to $u_c + v_c$.

2. Let $w = u_c + v_c$. Fix v_t, and assume $v < v_t$. By assumption $W(x)$ is increasing and $\frac{\partial T(w)(x)}{\partial w} \leq 1$ where $W(x)$ is differentiable (and it is differentiable almost everywhere by the first assumption).

   $T(W)(x)$ is either equal to the value of stopping or the value of search (or both). When it is equal to the value of stopping, then $T(W)(x) = w - v_t$, and so $\frac{\partial W(x)}{\partial w} = 1$. In addition, when T is equal to the value of stopping then using Leibniz's rule,

$$\begin{aligned}
\frac{\partial T(w)(x)}{\partial w} &= \beta(\int_{\underline{v}}^{\bar{v}} \int_{\underline{u}}^{\bar{u}} \frac{\partial W(v', u', v_t) dF(v', u')}{\partial w} \\
&+ \int_{\underline{v}}^{v_t} W(v_c + u_c, v_t) dF(v'|w = v' + u_c) - \int_{\underline{v}}^{v_t} W(v' - v' + v_c + u_c, v_t) dF(v'|w = v' + u_c) \\
&+ \int_{v_t}^{\bar{v}} W(v_c + u_c, v') dF(v'|w = v' + u_c) - \int_{v_t}^{\bar{v}} W(v' - v' + v_c + hatu, v') dF(v'|w = v' + u_c))
\end{aligned}$$

   which reduces to

$$\frac{\partial T(w)(x)}{\partial w} = \beta(\int_{\underline{v}}^{\bar{v}} \int_{\underline{u}}^{\bar{u}} \frac{\partial W(v', u', v_t) dF(v', u')}{\partial w})$$

   This is the average of the derivative of $W(x)$ over Regions I, II, III, IV. Since in all regions the slope of $W(x)$ is less than or equal to 1, the average slope must be less than 1 as well, and $T(W)(x)$ inherits the property of being almost everywhere differentiable. Obviously $T(W)(x)$ is not be differentiable at points where it switches from taking on the value of stopping to the value of searching, or vice-versa.

□

Next, I prove the proposition itself.

Condition on $v \leq v_t$. Assume that W follows a threshold. Consider $T(W)$. Then there are two options, either $T(W(u_c, v_c, v_t))$ is equal to the value of stopping or the value of search. Consider the properties of the value of stopping. It can take on all value between $\underline{v} + \underline{u} - (v_t)$ and $\bar{u}$, rising one-for-one with w. The previous lemma implies that the increase in value of search with respect to w is less than one-for-one. Therefore, there are two cases:

1. The value of search and the value of stopping never cross. This occurs when either the value of stopping starts off above the value of search ($\underline{v} + \underline{u} - v_t > W(\underline{v}, \underline{u}, v_t)$) or the the value

of stopping never reaches the value of search ($\bar{u} < W(v_t, \bar{u}, v_t)$). The first case means the threshold is is $\overline{w}(v_t) = \underline{v} + \underline{u}$, and the second $\overline{w}(v_t) = v_t + \bar{u}$.

2. The value of search and the value of stopping cross once. If the two cross once, then the threshold is interior. It's clear that since the value of stopping has a (weakly) larger slope, above the threshold the value function equals the value of stopping, and below it the value of search.

The corollary is simply an implication of the fact that

- The value function is unique
- The value function follows a threshold
- If $W(u_c, v_c, v_t) \neq u_c + v_c - v_t$ then the agent must strictly prefer to search, but if $W(u_c, v_c, v_t) = u_c + v_c - v_t$, and for all u_c', v_c' if $u_c' + v_c' < u_c + v_c$ then $W(u_c, v_c, v_t) \neq u_c + v_c - v_t$, then the agent must be indifferent between stopping and searching, but if $W(u_c, v_c, v_t) = u_c + v_c - v_t$ and there exists u_c', v_c' such that $u_c' + v_c' < u_c + v_c$ and $W(u_c, v_c, v_t) = u_c + v_c - v_t$ then the agent must strictly prefer to search.

□

Proof of Proposition 3

First I prove a lemma.

Lemma 3.1 *Fixing u_c and v_c*

1. *$W(v_c, u_c, v_t)$ is weakly decreasing in v_t and*
2. *$\frac{\partial W(v_c, u_c, v_t)}{\partial v_t} \geq -1$ wherever $W(v_c, u_c, v_t)$ is differentiable.*

Proof of Lemma 3.1 I prove these properties using the contraction mapping theorem. Note that the set of functions such that are weakly decreasing, and the set of functions such that $\frac{\partial W(v_c, u_c, v_t)}{\partial v_t} \geq -1$ are both closed sets. I assume that a function $W(x)$ possesses these two properties and show that they are preserved under $T(W)(x)$. Then by Corollary 1 (pg 52) of Stokey, Lucas and Prescott (1989) the solution to the decision maker's problem must possess these two properties.

1. Assume W is decreasing in v_t. If $T(W)$ equals the value of stopping, then clearly it is decreasing in v_t. If $T(W)$ equals the value of searching, then once again there are two effects. The first

is that within Regions I and II v_t rises, which implies within each of those regions W must fall. Within Regions III and IV the value of W is unchanged. The second effect that is the bounds of integration for each region shift, with mass shifted from Region III to Region I and from Region IV to Region II. Note that for any point (u_c, v_c, v') that is shifted from Region III to Region I (or from Region IV to II), W must fall, since $v_t \geq v'$. Therefore $T(W)$ must be decreasing in v_t. This implies that $T(W)$ must be differentiable almost everywhere with respect to v_t.

2. The proof is similar to the proof used in Lemma 2.1. Its clear that when T is equal to the value of stopping then the claim is true since the payoff to stopping is differentiable everywhere and the slope is -1 with respect to v_t

   When T is equal to the value of searching, then we simply reapply Leibniz's Rule as in the proof of Lemma 2.1; the derivatives with respect to the bounds of integration cancel out, and so the slope of $T(W)$ equals a weighted average of the slopes of V over all regions, discounted by β. Therefore $0 \geq \frac{\partial T(W)}{\partial v_t} \geq -1$.

□

Next, I prove the proposition. Using the results from the previous lemma, consider $\overline{w}(v_t)$ and $\overline{w}(v_t')$ where $v_t < v_t'$. Going from v_t to v_t', that the value of stopping must fall 1 for 1 with the increase in v_t, and the value of search falls by less (weakly). Therefore, the point where the value of stopping exceeds the value of search must increase. □

Proof of Corollary 2 Assume she does not. Then an item that was available before the last draw is chosen. Call this item x. Denote the highest value of temptation before the last draw as $\widetilde{v_0}$, and the highest value of temptation after the last draw $\widetilde{v_1}$. Therefore, $u(x) + v(x) < \overline{w}(\widetilde{v_0})$, $u(x) + v(x) \geq \overline{w}(\widetilde{v_1})$. But $\overline{w}(\widetilde{v_0}) \leq \overline{w}(\widetilde{v_1})$. This is a contradiction. □

Proof of Corollary 3 The value $E(v_t|\text{searching at time t}) = E(\max_{y \in A} v(y)|\text{searching at time t})$ is rising in t, and the the threshold $\overline{w}(v_t)$ is rising in t. Therefore, the probability that the next good observed falls above the threshold $(P(u(x) + v(x) > v_t))$ falls. □

Proof of Lemma 1 I prove this property using the contraction mapping theorem. Note that the set of functions that do not depend on (v_c, u_c) below some threshold is a closed set. I assume that a function $W(x)$ possesses this property and show that it is preserved under $T(W)(x)$. Then by

Corollary 1 (pg 52) of Stokey, Lucas and Prescott (1989) the solution to the decision maker's problem must possess these two properties.

Once again, fix v_t, and condition on $v_c + u_c < \overline{w}(v_t)$ and $v \leq v_t$. Consider the value function below the threshold, and denote it $W_-(v_c, u_c, v_t)$. Then:

$$\begin{aligned} W_-(v_c, u_c, v_t) &= \beta\Big(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{v_c+u_c-v'} W(v_c, u_c, v_t)dF(v', u') + \int_{\underline{v}}^{v_t}\int_{v_c+u_c-v'}^{\bar{u}} W(v', u', v_t)dF(v', u') \\ &+ \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{v_c+u_c-v'} W(v_c, u_c, v')dF(v', u') + \int_{v_t}^{\bar{v}}\int_{v_c+u_c-v'}^{\bar{u}} W(v', u', v')dF(v', u')\Big) \end{aligned}$$

This integral can be reformulated, using the threshold $\overline{w}(v_t)$ as the boundary for integration, and using the assumption that the value of search depends only on v_t.

$$\begin{aligned} W_-(v_c, u_c, v_t) &= \beta\Big(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}(v_t)-v'} W_-(v_t)dF(v', u') + \int_{\underline{v}}^{v_t}\int_{\overline{w}(v_t)-v'}^{\bar{u}} W(v', u', v_t)dF(v', u') \\ &+ \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}(v')} W_-(v')dF(v', u') + \int_{v_t}^{\bar{v}}\int_{\overline{w}(v')}^{\bar{u}} W(v', u', v')dF(v', u')\Big) \end{aligned}$$

Note that none of these terms depends on v_c or u_c, so therefore $W_-(v_c, u_c, v_t) = W_-(v_t)$. Since the value of stopping at the threshold is $\overline{w}(v_t) - v_t$, and at the threshold this is qual to the value of searching then $W_-(v_t) = \overline{w}(v_t) - v_t$. □

Proof of Proposition 4 Note that by construction $\overline{w}(v_t) - v_t = \mu(v_t) + v_t - v_t = \mu(v_t) = W(\mu(v_t), v_t, v_t)$. Because the value function is flat below the threshold, $\overline{w}(v_t) - v_t$ is equal to $W(\mu(v_t), v_t, v_t)$ as well as $W(u_c, v_c, v_t)$ for all u_c, v_c such that $u_c + v_c \leq \mu(v_t) + v_t$.

Consider any two sets of state variables (u_c, v_c, v_t) and (u_c, v_c, v_t') such that $v_t' \geq v_t$ and $u_c + v_c \leq \overline{w}(v_t)$. Therefore, $u_c + v_c \leq \overline{w}(v_t')$. Furthermore since the value function is falling in its third argument (as proved in Lemma 3.1), $W(u_c, v_c, v_t) \geq W(u_c, v_c, v_t')$. It also must be the case that $W(\mu(v_t), v_t, v_t) = W(u_c, v_c, v_t)$ and $W(\mu(v_t'), v_t', v_t) = W(u_c, v_c, v_t')$. Therefore $W(\mu(v_t), v_t, v_t) \geq W(\mu(v_t'), v_t', v_t)$. Therefore $\mu(v_t) \geq \mu(v_t')$.

In order to ensure that for all v_t and v_t' such a comparison can be made, extend the possible values of u that goods can be distributed on, changing the lower bound from $\underline{u}$ to $\underline{u} - (\bar{v} - \underline{v})$. Let the $f(v,u) = 0$ for all $(v,u) \in [\underline{v}, \bar{v}] \times [\underline{u} - (\bar{v} - \underline{v}), \underline{u}]$. This extension of the u values implies that regardless of state points (v_c, u_c) and (v_t, u_t) in M, I can construct Regions I and III. $\square$

Proof of Proposition 5 The follows directly from Proposition 4. $\square$

Proof of Proposition 6 Given the previous results, it is possible to characterize the threshold function using an endpoint condition and a differential equation. For any point (v_c, u_c, v_t) that is on the threshold the value of search must equal the value of stopping:

$$\begin{aligned} W(v_c, u_c, v_t) &= \overline{w}(v_t) - v_t = v_c + u_c - v_t \\ &= \beta\Big(\int_{\underline{v}}^{v_t} \int_{\underline{u}}^{v_c + u_c - v'} W(v_c, u_c, v_t) dF(v', u') + \int_{\underline{v}}^{v_t} \int_{v_c + u_c - v'}^{\bar{u}} W(v', u', v_t) dF(v', u') \\ &+ \int_{v_t}^{\bar{v}} \int_{\underline{u}}^{v_c + u_c - v'} W(v_c, u_c, v') dF(v', u') + \int_{v_t}^{\bar{v}} \int_{v_c + u_c - v'}^{\bar{u}} W(v', u', v') dF(v', u')\Big) \end{aligned}$$

It is possible to simplify this expression by using the fact that the policy is a threshold strategy, so any point above the threshold simply gives the payout of stopping.

$$\begin{aligned} W(v_c, u_c, v_t) &= \overline{w}(v_t) - v_t = v_c + u_c - v_t \\ &= \beta\Big(\int_{\underline{v}}^{v_t} \int_{\underline{u}}^{v_c + u_c - v'} (u + v - v_t) dF(v', u') + \int_{\underline{v}}^{v_t} \int_{v_c + u_c - v'}^{\bar{u}} (u' + v' - v_t) dF(v', u') \\ &+ \int_{v_t}^{\bar{v}} \int_{\underline{u}}^{v_c + u_c - v'} W(v_c, u_c, v') dF(v', u') + \int_{v_t}^{\bar{v}} \int_{v_c + u_c - v'}^{\bar{u}} W(v', u', v') dF(v', u')\Big) \end{aligned}$$

Now redefine Regions III and IV into III′ and IV′, where III′ is the region below the new threshold (for the new value of temptation), and IV′ is the region above the new threshold. See Figure 1.4 for a graphical representation of the new regions.

In addition, in Region III′ the value function depends only on the value of temptation, and that it must equal $\overline{w}(v') - v'$ (since this is the value of stopping, and the two are equal at the threshold). Rewriting everything in terms of the thresholds :

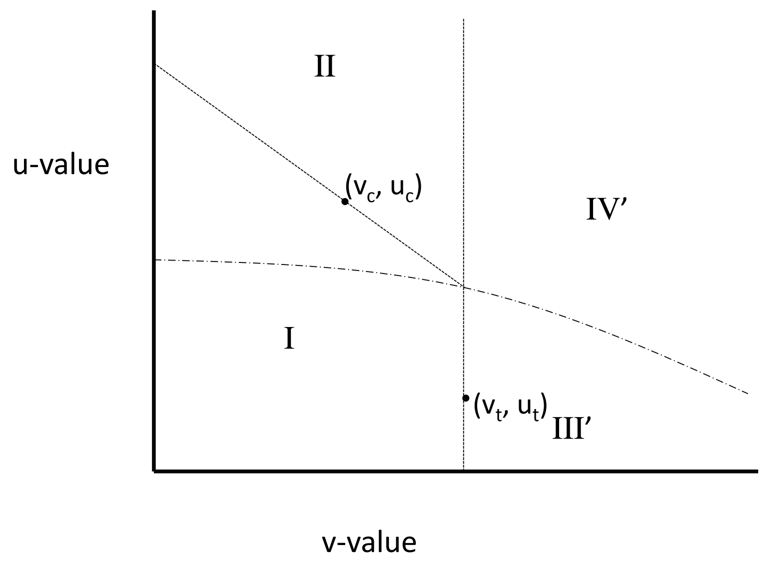

Figure 1.4: Reclassifying Search Outcomes

$$
\begin{aligned}
W(v_c, u_c, v_t) &= \overline{w}(v_t) - v_t \\
&= \beta\Big(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}(v_t)-v'} (\overline{w}(v_t) - v_t)dF(v', u') + \int_{\underline{v}}^{v_t}\int_{\overline{w}(v_t)-v'}^{\bar{u}} (u' + v' - v_t)dF(v', u') \\
&+ \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}(v')-v'} (\overline{w}(v') - v')dF(v', u') + \int_{v_t}^{\bar{v}}\int_{\overline{w}(v')-v'}^{\bar{u}} u'dF(v', u')\Big)
\end{aligned}
$$

However, the solution still depends on elements that are recursively defined - $\overline{w}(v')$ in Regions III and IV - the regions where the most tempting item is replaced with the new draw. These regions don't exist when $\bar{v}$ has been observed, because no more tempting item can be observed. Therefore, define the cutoff point at $v_t = \bar{v}$:

$$
\overline{w}(\bar{v}) - \bar{v} = \beta\Big(\int_{\underline{v}}^{\bar{v}}\int_{\underline{u}}^{\overline{w}(\bar{v})-v'} (\overline{w}(\bar{v}) - \bar{v})dF(v', u') + \int_{\underline{v}}^{v_t}\int_{\overline{w}(\bar{v})-v'}^{\bar{u}} (u' + v' - \bar{v})dF(v', u')\Big)
$$

This equation solves for the threshold when the most tempting point has been observed. Note that it is almost exactly equivalent to the solution to the McCall (1970) search model where the wage is equal to $u + v$. This describes the initial condition. Next, I solve for the form of the differential equation.

Letting the value of stopping be the LHS of an equation, and the value of search be the RHS of the equation, and moving along the threshold, take derivative of both sides, and applying Leibniz's rule repeatedly:

$$\begin{aligned}
\frac{d\overline{w}(v_t)}{dv_t} - 1 &= \beta(\int_{\underline{u}}^{\overline{w}(v_t)-v_t} (\overline{w}(v_t) - v_t) dF(v', u') \\
&+ \int_{\underline{v}}^{v_t} (\frac{\partial \overline{w}(v_t)}{\partial v_t})(\overline{w}(v_t) - v_t) dF(v', u') + \int_{\underline{v}}^{v_t} \int_{\underline{u}}^{\overline{w}(v_t)-v'} (\frac{\partial \overline{w}(v_t)}{\partial v_t} - 1) dF(v', u') \\
&+ \int_{\overline{w}(v_t)-v_t}^{\bar{u}} (u' + v_t - v_t) dF(v', u') \\
&+ \int_{\underline{v}}^{v_t} -\frac{\partial \overline{w}(v_t)}{\partial v_t} (\overline{w}(v_t) - v' + v' - v_t) dF(v', u') + \int_{\underline{v}}^{v_t} \int_{\overline{w}(v_t)-v'}^{\bar{u}} (-1) dF(v', u') \\
&- \int_{\underline{u}}^{\overline{w}(v_t)-v_t} (\overline{w}(v_t) - v_t) dF(v', u') \\
&+ \int_{v_t}^{\bar{v}} (0) dF(v', u') + \int_{v_t}^{\bar{v}} \int_{\underline{u}}^{\overline{w}(v')-v'} (0) dF(v', u') \\
&- \int_{\overline{w}(v_t)-v_t}^{\bar{u}} u' dF(v', u') \\
&+ \int_{v_t}^{\bar{v}} (0) dF(v', u') + \int_{v_t}^{\bar{v}} \int_{\overline{w}(v')-v'}^{\bar{u}} (0) dF(v', u')) \\
&= \beta(\int_{\underline{v}}^{v_t} (\frac{\partial \overline{w}(v_t)}{\partial v_t})(\overline{w}(v_t) - v_t) dF(v', u') + \int_{\underline{v}}^{v_t} \int_{\underline{u}}^{\overline{w}(v_t)-v'} (\frac{\partial \overline{w}(v_t)}{\partial v_t} - 1) dF(v', u') \\
&+ \int_{\underline{v}}^{v_t} -\frac{\partial \overline{w}(v_t)}{\partial v_t} (\overline{w}(v_t) - v_t) dF(v', u') - \int_{\underline{v}}^{v_t} \int_{\overline{w}(v_t)-v'}^{\bar{u}} (1) dF(v', u')) \\
&= \beta(\int_{\underline{v}}^{v_t} \int_{\underline{u}}^{\overline{w}(v_t)-v'} (\frac{\partial \overline{w}(v_t)}{\partial v_t} - 1) dF(v', u') - \int_{\underline{v}}^{v_t} \int_{\overline{w}(v_t)-v'}^{\bar{u}} (1) dF(v', u')) \\
&= \beta(\int_{\underline{v}}^{v_t} \int_{\underline{u}}^{\overline{w}(v_t)-v'} \frac{\partial \overline{w}(v_t)}{\partial v_t} dF(v', u') \\
&- \int_{\underline{v}}^{v_t} \int_{\underline{u}}^{\overline{w}(v_t)-v'} (1) dF(v', u') - \int_{\underline{v}}^{v_t} \int_{\overline{w}(v_t)-v'}^{\bar{u}} (1) dF(v', u')) \\
&= \beta(\int_{\underline{v}}^{v_t} \int_{\underline{u}}^{\overline{w}(v_t)-v'} \frac{\partial \overline{w}(v_t)}{\partial v_t} dF(v', u') - \int_{\underline{v}}^{v_t} \int_{\underline{u}}^{\bar{u}} (1) dF(v', u'))
\end{aligned}$$

□

Proof of Proposition 7 First, recall that by $\mu(v_t$ completely characterizes the threshold. If $\mu(v_t)$ increases, the threshold must increase. To see the effects of the change in α on μ consider an situation where the agents have observed only a single item, and so the state variables are (v_t, u_c, v_t), where

$v_c + u_c = \overline{w}(v_t)$. Therefore $u_c = \mu(v_t)$.

First, consider Agent 1. The value of stopping today is simply $u_c = W_{\alpha_1}(v_t, u_c, v_t)$, which is also the value of search. Now, consider Agent 2, where $\alpha_2 \leq \alpha_1$, but with the exact same state variables (v_t, u_c, v_t). Agent 2's payoff from stopping is simply u_c.

Assume that Agent 2 chooses to search today, and thereafter uses Agent 1's policy. Given Agent 1's optimal policy P_1, and current state variables (v_t, u_c, v_t) it is possible to rewrite the value of search as a function of the payoff from stopping each possible set of state variables (u_c', v_c', v_t'), denoted for Agent i as $\pi_i(v_c', u_c', v_t')$ and the probability of stopping given each possible set of future state variables (v_c', u_c', v_t') and current state variables (v_c, u_c, v_t) as $h((v_c', u_c', v_t')|(v_t, u_c, v_t))$. Under Agent 1's policy P_1 the value of search given (v_t, u_c, v_t) is then simply

$$\int_{\underline{v}}^{\bar{v}} \int_{\underline{u}}^{\bar{u}} h((v_c', u_c', v_t')|(v_t, u_c, v_t))\pi_i(v_c', u_c', v_t')dudv$$

It is the case that $\pi_2(v_c', u_c', v_t') = u_c' - \alpha_2(v_t' - v_c') \geq u_c' - \alpha_1(v_t' - v_c') = \pi_1(v_c', u_c', v_t') = \pi_1(v_c', u_c', v_t')$. Since h is same for Agent 1 and Agent 2 (while she is using P_1), the value function for Agent 1 must be weakly less than for Agent 2 using P_1. This immediately implies that the value of search, which is simply the expected value function given one additional observation, must be weakly higher for Agent 2, using P_1, than for Agent 1, since the two agents face the same distribution F of future draws. Note that P_1 must be weakly sub-optimal for Agent 2. Therefore, the value to search for Agent 2, under her optimal policy must be weakly higher than under P_1. Therefore, Agent 2 must (weakly) prefer to search, and so the threshold must have weakly risen. $\square$

Proof of Proposition 8 I will prove it for each stochastic dominance shift in turn.

1. First, I will prove the claim for FOSD. Denote the value function under distribution F as $W_F(v_c, u_c, v_t)$. I will show that the value function under distribution F', $W_{F'}(v_c, u_c, v_t)$ is always weakly treater than $W_F(v_c, u_c, v_t)$: $W_{F'}(v_c, u_c, v_t) \geq W_F(v_c, u_c, v_t)$ for all (v_c, u_c, v_t). I will prove the proposition using the contraction mapping theorem. Note that the set of function $\{f(v_c, u_c, v_t)|f(v_c, u_c, v_t) \geq W_F(v_c, u_c, v_t)$ for all $(v_c, u_c, v_t)\}$ is closed in the sup-norm metric. I assume that function $W_{F'}(x)$ is a member of the set, and show that membership is preserved under $T(W)(x)$. Then by Corollary 1 (pg 52) of Stokey, Lucas and Prescott (1989) the solution to the decision maker's problem must possess these two properties.

   First, assume that $W_{F'}(v_c, u_c, v_t) \geq W_F(v_c, u_c, v_t)$ for all (v_c, u_c, v_t). This implies that $\overline{w}F'(v_t) \geq \overline{w}F(v_t)$ for all v_t. Now consider $T(W_{F'})(v_c, u_c, v_t)$ and compare this to $W_F(v_c, u_c, v_t)$. Note that since W_F is a fixed point $W_F(v_c, u_c, v_t) = T(W_F)(v_c, u_c, v_t)$. The value to stopping at (v_c, u_c, v_t) under both F and F' is the same —- it is $u_c + v_c - v_t$.

Consider the value to searching under F.

$$
\begin{aligned}
&\beta\Big(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v_t)-v'}(\overline{w}_F(v_t)-v_t)dF(v',u') + \int_{\underline{v}}^{v_t}\int_{\overline{w}_F(v_t)-v'}^{\bar{u}}(u'+v'-v_t)dF(v',u') \\
+ \;& \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\overline{w}_F(v')-v')dF(v',u') + \int_{v_t}^{\bar{v}}\int_{\overline{w}_F(v')-v'}^{\bar{u}}u'dF(v',u')\Big)
\end{aligned}
$$

I will implement two changes to compare this to the value of searching under F'. First, keeping the distribution constant, I will shift the value function from W_F to W'_F. Recall that W is completely characterized by the threshold function. This changes the previous equation to

$$
\begin{aligned}
&\beta\Big(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_{F'}(v_t)-v'}(\overline{w}_{F'}(v_t)-v_t)dF(v',u') + \int_{\underline{v}}^{v_t}\int_{\overline{w}_{F'}(v_t)-v'}^{\bar{u}}(u'+v'-v_t)dF(v',u') \\
+ \;& \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}(v')-v'}(\overline{w}_{F'}(v')-v')dF(v',u') + \int_{v_t}^{\bar{v}}\int_{\overline{w}_{F'}(v')-v'}^{\bar{u}}u'dF(v',u')\Big)
\end{aligned}
$$

This change solely increases the value of search. This is because by assumption $\overline{w}_{F'}(v_t) \geq \overline{w}_F(v_t)$. Next, I change the distribution from F to F'.

$$
\begin{aligned}
&\beta\Big(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_{F'}(v_t)-v'}(\overline{w}_{F'}(v_t)-v_t)dF'(v',u') + \int_{\underline{v}}^{v_t}\int_{\overline{w}_{F'}(v_t)-v'}^{\bar{u}}(u'+v'-v_t)dF'(v',u') \\
+ \;& \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}(v')-v'}(\overline{w}_{F'}(v')-v')dF'(v',u') + \int_{v_t}^{\bar{v}}\int_{\overline{w}_{F'}(v')-v'}^{\bar{u}}u'dF'(v',u')\Big)
\end{aligned}
$$

This again increases the value of search. To see this, note that for any given (v_c, v_t) combination, mass is shifted from low values of u_c to higher values of u_c, and that as shown previously, the value function is increasing in u_c. If mass is shifted Region I, the value doesn't change, and similarly within Region III$'$. However, if mass is shifted from Region I to Region II, this increases the RHS, since by construction, fixing v_c and v_t, the value function is strictly increasing in u_c. Similarly the RHS is increased as mass shifts from Region III$'$ to Region IV$'$. Furthermore, within Regions II and IV$'$ a shift in FOSD of $F_{u|v}$ must increase the RHS since the payoff is increasing in u_c. Therefore, since in a pairwise comparison of the payoff from stopping and searching it is the case that the payoff under F' is greater than the payoff under

F. Therefore, $T(W_{F'})(v_c,u_c,v_t) \geq T(W_F)(v_c,u_c,v_t) = W_F(v_c,u_c,v_t)$ for all (v_c,u_c,v_t). This proves the claim.

2. Next I will prove the proposition for mean preserving spreads. Fix v_t. Consider a point$x =$ (v_c,u_c) such that v_c+u_c equals the threshold value $\overline{w}_F(v_t)$. At the threshold, the value of search is equal to the value of stopping. Using the equation derived previously:

$$\begin{aligned}\overline{w}_F(v_t)-v_t &= \beta(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v_t)-v'}(\overline{w}_F(v_t)-v_t)dF_{u'|v'}dF_{v'}+\int_{\underline{v}}^{v_t}\int_{\overline{w}_F(v_t)-v'}^{\bar{u}}(u'+v'-v_t)dF_{u'|v'}dF_{v'}\\ &+\int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\overline{w}_F(v')-v')dF_{u'|v'}dF_{v'}+\int_{v_t}^{\bar{v}}\int_{\overline{w}_F(v')-v'}^{\bar{u}}u'dF_{u'|v'}dF_{v'})\end{aligned}$$

Breaking up the LHS into Regions I, II, III′ and IV′ and then shifting terms around gives:

$$\begin{aligned}&(1-\beta)\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v_t)-v'}(\overline{w}_F(v_t)-v_t)dF_{u'|v'}dF_{v'}\\ =&\int_{\underline{v}}^{v_t}\int_{\overline{w}_F(v_t)-v'}^{\bar{u}}(\beta(u'+v')-\overline{w}_F(v_t)+(1-\beta)v_t)dF_{u'|v'}dF_{v'}\\ +&\int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\beta(\overline{w}_F(v')-v')-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'}\\ +&\int_{v_t}^{\bar{v}}\int_{\overline{w}_F(v')-v'}^{\bar{u}}(\beta u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'}\end{aligned}$$

Now, adding on the $1-\beta$ times the value under the integral on the LHS but for regions II, III′ and IV′ onto both sides gives:

$$\begin{aligned}(1-\beta)(\overline{w}_F(v_t)-v_t) &= \int_{\underline{v}}^{v_t}\int_{\overline{w}_F(v_t)-v'}^{\bar{u}}(\beta(u'+v')-\overline{w}_F(v_t)+(1-\beta)v_t+(1-\beta)(\overline{w}_F(v_t)-v_t))dF_{u'|v'}dF_{v'}\\ &+\int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\beta(\overline{w}_F(v')-v')-\overline{w}_F(v_t)+v_t+(1-\beta)(\overline{w}_F(v_t)-v_t))dF_{u'|v'}dF_{v'}\\ &+\int_{v_t}^{\bar{v}}\int_{\overline{w}_F(v')-v'}^{\bar{u}}(\beta u'-\overline{w}_F(v_t)+v_t+(1-\beta)(\overline{w}_F(v_t)-v_t))dF_{u'|v'}dF_{v'}\end{aligned}$$

Reducing gives:

$$
\begin{aligned}
(1-\beta)(\overline{w}_F(v_t)-v_t) &= \beta(\int_{\underline{v}}^{v_t}\int_{\overline{w}_F(v_t)-v'}^{\bar{u}}(u'+v'-\overline{w}_F(v_t))dF_{u'|v'}dF_{v'} \\
&+ \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\overline{w}_F(v')-v'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&+ \int_{v_t}^{\bar{v}}\int_{\overline{w}_F(v')-v'}^{\bar{u}}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'})
\end{aligned}
$$

The next step involves setting up the regions of integration on the RHS such that it is possible to to integrate each term from 'the bottom up' with the lower bounds of integration being $\underline{u}$ and $\underline{v}$. In addition, define the function $\overline{w}_{F,v}(v) = \overline{w}_F(v_t)$ if $v \leq v_t$ and $\overline{w}_F(v) = \overline{w}_F(v)$ if $v > v_t$. Then:

$$
\begin{aligned}
(1-\beta)(\overline{w}_F(v_t)-v_t) &= \beta(\int_{\underline{v}}^{v_t}\int_{\overline{w}_F(v')-v'}^{\bar{u}}(u'+v'-\overline{w}_F(v_t))dF_{u'|v'}dF_{v'} \\
&+ \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(u'+v'-\overline{w}_F(v_t))dF_{u'|v'}dF_{v'} \\
&- \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(u'+v'-\overline{w}_F(v_t))dF_{u'|v'}dF_{v'} \\
&+ \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\overline{w}_F(v')-v'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&+ \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\overline{w}_F(v')-v'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&- \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\overline{w}_F(v')-v'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&+ \int_{v_t}^{\bar{v}}\int_{\overline{w}_F(v')-v'}^{\bar{u}}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&+ \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&- \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&+ \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&- \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&+ \int_{\underline{v}}^{v_t}\int_{\overline{w}_F(v')-v'}^{\bar{u}}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&- \int_{\underline{v}}^{v_t}\int_{\overline{w}_F(v')-v'}^{\bar{u}}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'})
\end{aligned}
$$

Combining terms gives:

$$
\begin{aligned}
(1-\beta)(\overline{w}_F(v_t)-v_t) &= \beta(\int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\bar{u}}(u'+v'-\overline{w}_F(v_t))dF_{u'|v'}dF_{v'} \\
&- \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(u'+v'-\overline{w}_F(v_t))dF_{u'|v'}dF_{v'} \\
&+ \int_{\underline{v}}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\overline{w}_F(v')-v'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&- \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(\overline{w}_F(v')-v'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&+ \int_{\underline{v}}^{\bar{v}}\int_{\underline{u}}^{\bar{u}}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} - \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\bar{u}}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'} \\
&- \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}(u'-\overline{w}_F(v_t)+v_t)dF_{u'|v'}dF_{v'})
\end{aligned}
$$

Simplifying the terms and integrate by parts:

$$
\begin{aligned}
(1-\beta)(\overline{w}_F(v_t)-v_t) &= \beta((E_u(u')+v'-\overline{w}_F)F_{v'}|_{\underline{v}}^{v_t} - \int_{\underline{v}}^{v_t}F_{v'}dv' \\
&- \int_{\underline{v}}^{v_t}(u'+v'-\overline{w}_F(v_t))F_{u'|v'}|_{\underline{u}}^{\overline{w}_F(v')-v'}dF_{v'} + \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}F_{u'|v'}du'dF_{v'} \\
&+ \int_{\underline{v}}^{\bar{v}}(\overline{w}_F(v')-v'-\overline{w}_F(v_t)+v_t)F_{u'|v'}|_{\underline{u}}^{\overline{w}_F(v')-v'}dF_{v'} - 0 \\
&- \int_{\underline{v}}^{v_t}(\overline{w}_F(v')-v'-\overline{w}_F(v_t)+v_t)F_{u'|v'}|_{\underline{u}}^{\overline{w}_F(v')-v'}dF_{v'} + 0 \\
&+ E_u(u')-\overline{w}_F(v_t)+v_t-(E_u(u')-\overline{w}_F(v_t)+v_t)F_{v'}|_{\underline{v}}^{v_t} + 0 \\
&- \int_{v_t}^{\bar{v}}(u'-\overline{w}_F(v_t)+v_t)F_{u'|v'}|_{\underline{u}}^{\overline{w}_F(v')-v'}dF_{v'} + \int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}F_{u'|v'}du'dF_{v'})
\end{aligned}
$$

Canceling out terms leaves:

$$
\begin{aligned}
(1-\beta)(\overline{w}_F(v_t)-v_t) &= \beta(-\int_{\underline{v}}^{v_t}F_{v'}dv' + \int_{\underline{v}}^{v_t}\int_{\underline{u}}^{\overline{w}_F(v')-v'}F_{u'|v'}du'dF_{v'} + E_u(u') \\
&- \overline{w}_F(v_t)+v_t+\int_{v_t}^{\bar{v}}\int_{\underline{u}}^{\overline{w}_F(v')-v'}F_{u'|v'}du'dF_{v'})
\end{aligned}
$$

Combining and shifting terms gives:

$$\overline{w}_F(v_t) - v_t = \beta(E_u(u') - \int_{\underline{v}}^{v_t} F_{v'} dv' + \int_{\underline{v}}^{\bar{v}} \int_{\underline{u}}^{\overline{w}_F(v')-v'} F_{u'|v'} du' dF_{v'})$$

This equation holds for an arbitrary distribution. Since the LHS is simply the value of stopping at (v_c, u_c, v_t), then the RHS must be equal to the value of search at the point of being indifferent between searching and stopping. Furthermore, this is equal to the value function below the threshold.

Denote the value function under distribution F as $W_F(v_c, u_c, v_t)$. I will show that the value function under distribution F', $W_{F'}(v_c, u_c, v_t)$ is always weakly treater than $W_F(v_c, u_c, v_t)$: $W_{F'}(v_c, u_c, v_t) \geq W_F(v_c, u_c, v_t)$ for all (v_c, u_c, v_t).

I will prove the proposition using the contraction mapping theorem. Note that the set of function $\{f(v_c, u_c, v_t)|f(v_c, u_c, v_t) \geq W_F(v_c, u_c, v_t)$ for all $(v_c, u_c, v_t)\}$ is closed in the sup-norm metric. I assume that function $W_{F'}(x)$ is a member of the set, and show that membership is preserved under $T(W)(x)$. Then by Corollary 1 (pg 52) of Stokey, Lucas and Prescott (1989) the solution to the decision maker's problem must possess these two properties.

First, assume that $W_{F'}(v_c, u_c, v_t) \geq W_F(v_c, u_c, v_t)$ for all (v_c, u_c, v_t). This implies that $\overline{w}F'(v_t) \geq \overline{w}F(v_t)$ for all v_t.

Now consider $T(W_{F'})(v_c, u_c, v_t)$ and compare this to $W_F(v_c, u_c, v_t)$. Note that since W_F is a fixed point $W_F(v_c, u_c, v_t) = T(W_F)(v_c, u_c, v_t)$.

Assume that stopping is optimal under F at (v_c, u_c, v_t), then $W_F(v_c, u_c, v_t) = u_c + v_c - v_t$, which must be greater than the value of search. If search is optimal under F at (v_c, u_c, v_t), then

$$W_F(v_c, u_c, v_t) = \beta(E_u(u') - \int_{\underline{v}}^{v_t} F_{v'} dv' + \int_{\underline{v}}^{\bar{v}} \int_{\underline{u}}^{\overline{w}_F(v')-v'} F_{u'|v'} du' dF_{v'})$$

which must be greater than the value of stopping.

Next, consider the value function under distribution F'. Assume that stopping is optimal under F' at (v_c, u_c, v_t), then $W_F(v_c, u_c, v_t) = u_c + v_c - v_t$, which must be greater than the value of search. If search is optimal under F' at (v_c, u_c, v_t), then

$$W_{F'}(v_c, u_c, v_t) = \beta(E_u(u') - \int_{\underline{v}}^{v_t} F'_{v'} dv' + \int_{\underline{v}}^{\bar{v}} \int_{\underline{u}}^{\overline{w}_{F'}(v')-v'} F'_{u'|v'} du' dF_{v'})$$

which must be greater than the value of stopping. Furthermore, note that by assumption, since the value function under F is everywhere lower than the value function under F' then $\overline{w}_{F'}(v_t) - v_t \geq \overline{w}_F(v_t) - v_t$. Therefore

$$\beta(E_u(u') - \int_{\underline{v}}^{v_t} F'_{v'} dv' + \int_{\underline{v}}^{\bar{v}} \int_{\underline{u}}^{\overline{w}_{F'}(v')-v'} F'_{u'|v'} du' dF_{v'}) \geq \beta(E_u(u') - \int_{\underline{v}}^{v_t} F_{v'} dv' + \int_{\underline{v}}^{\bar{v}} \int_{\underline{u}}^{\overline{w}_F(v')-v'} F_{u'|v'} du' dF_{v'})$$

To see this, observe first that the only difference between the LHS and the RHS occurs in the last term. Furthermore, note that by the fact that $F'_{u|v}$ is a mean preserving spread of $F'_{u|v}$

$$\int_{\underline{u}}^{\overline{w}_F(v')-v'} F_{u'|v'} du' \leq \int_{\underline{u}}^{\overline{w}_F(v')-v'} F'_{u'|v'} du'$$

In addition, since $F'_{u|v}$ is increasing,

$$\int_{\underline{u}}^{\overline{w}_F(v')-v'} F'_{u'|v'} du' \leq \int_{\underline{u}}^{\overline{w}_{F'}(v')-v'} F'_{u'|v'} du'$$

This proves the inequality.

Given this, there are four cases:

- It is optimal to search under both F and F'. In this case $T(W_{F'})(v_c, u_c, v_t) = W_F(v_c, u_c, v_t) = u_c + v_c - v_t$.
- It is optimal to search under both F and F'. It was just proven above that in this case $T(W_{F'})(v_c, u_c, v_t) \geq W_F(v_c, u_c, v_t)$
- It is optimal to search under F' but to stop under F. In this case $T(W_{F'})(v_c, u_c, v_t) \geq u_c + v_c - v_t = W_F(v_c, u_c, v_t)$
- It is optimal to search under F but to stop under F'. This is a contradiction because the value of search under F' is higher than under F, and the payoff to stopping is the same under both distributions.

Therefore, $T(W_{F'})(v_c, u_c, v_t) \geq W_F(v_c, u_c, v_t)$ and the claim is proved.

$\square$

1.8 Appendix B: Extensions

The basic model presented in the previous sections presents a single stylized model that tries to capture the effects of a particular type of context-dependent preferences in search situations. This section will present two natural extensions of the model. The first is when the decision maker experiences temptation not only at the point of stopping, but also while continuing to search. The second presents a model that allows for more general forms of context dependence — the agent's preferences can depend on more than just two of the items observed in the search process, and can therefore be represented with the restriction of the model of Dekel, Lipman and Rustichini (2001) to finite subjective state spaces.

1.8.1 Ongoing Temptation

Although there are certain situations where temptation would be experienced only when purchasing a good, an agent might also experience temptation at all decision nodes. In these situations the cost of temptation is no longer born when the agent stops — instead it is a flow cost that is born every period during search. In this case, the agent maximizes:

$$\begin{aligned}
\max \mathsf{W}(v_c, u_c, v_t) \quad = & \max_{\text{stop, search}} \{v_c + u_c, \\
& \beta\Big(\int_{\underline{v}}^{v_t} \int_{\underline{u}}^{v_c+u_c-v'} \mathsf{W}(v_c, u_c, v_t) dF(v', u') \\
+ & \int_{\underline{v}}^{v_t} \int_{v_c+u_c-v'}^{\bar{u}} \mathsf{W}(v', u', v_t) dF(v', u') \\
+ & \int_{v_t}^{\bar{v}} \int_{\underline{u}}^{v_c+u_c-v'} \mathsf{W}(v_c, u_c, v') dF(v', u') \\
+ & \int_{v_t}^{\bar{v}} \int_{v_c+u_c-v'}^{\bar{u}} \mathsf{W}(v', u', v') dF(v', u')\Big)\} - v_t
\end{aligned}$$

A unique solution exists for the agent's problem, and it is a threshold function.

Proposition 10 *There exists a unique value function $\mathsf{W}(v_c, u_c, v_t)$ that solves the agents' problem and is continuous in v_c, u_c and v_t. Furthermore, The value function follows a threshold: there exists a function $\overline{w}(v_t) : \mathbb{R} \to \mathbb{R}$ such that $v_c + u_c \geq \overline{w}(v_t)$ if and only if $\mathsf{W}(v_c, u_c, v_t) = v_c + u_c - v_t$*

Proof of Proposition 10 The proof for the existence of a unique continuous value function

$\mathsf{W}(v_c, u_c, v_t)$ follows from Blackwell's sufficiency conditions, which are shown to hold using an analogous proof to that used in Proposition 1. Similarly, one can show that Lemma 2.1 still holds, and therefore the value function follows a threshold. □

However, unlike the model in the body of the paper, the threshold is now declining in the value of temptation. Observing an item that increases v_t decreases the value of stopping one-for-one (as in the base model). But now, it also decreases the value of continuing to search by more than one-for-one, causing the threshold to fall in the level of temptation experienced.

Intuitively, the cost of temptation is paid this period regardless of what action the agent takes. However, it will also decrease the agent's payoff in every future period as well. Therefore, the threshold falls.

Proposition 11 *The threshold payoff $\overline{w}(v_t)$ is decreasing in the level of temptation v_t.*

Proof of Proposition 11 First I shall prove a lemma.

Lemma 10.1 *Fixing u_c and v_c*

1. *$\mathsf{W}(v_c, u_c, v_t)$ is weakly decreasing in v_t and*
2. *$\frac{\partial \mathsf{W}(v_c,u_c,v_t)}{\partial v_t} \leq -1$ wherever $\mathsf{W}(v_c, u_c, v_t)$ is differentiable.*

Proof of Lemma 11.1 I prove these properties using the contraction mapping theorem. Note that the set of functions such that are weakly decreasing, and the set of functions such that $\frac{\partial \mathsf{W}(v_c,u_c,v_t)}{\partial v_t} \leq -1$ are both closed sets. I assume that a function $\mathsf{W}(x)$ possesses these two properties and show that they are preserved under $T(\mathsf{W})(x)$. Then by corollary 1 of Stokey, Lucas and Prescott (1989) the solution to the decision maker's problem must possess these two properties.

1. Assume W is decreasing in v_t. If $T(\mathsf{W})$ equals the value of stopping, then clearly it is decreasing in v_t. If $T(\mathsf{W})$ equals value of searching then by similar reasoning used in Lemma's 2.1 and 3.1 , $T(\mathsf{W})$ must be decreasing in v_t. This implies that $T(v_t)$ must be differentiable almost everywhere with respect to v_t.

2. The proof is similar to Lemma 2.1 and Lemma 3.1. Its clear that when T is equal to the value of stopping then the statement is true. When T is equal to the value of searching, then Leibniz's rule must be applied. The bounds of integration cancel out, and so the slope of T is

the average slope of W over all regions, less 1 (since v_t enters as an additively separable term). W is by assumption declining in all regions, so T must be declining.

Note that, fixing v_c and u_c, the value of searching must fall more than 1 for 1 with v_t. This is because if T equals the value of stopping, then the value falls 1 for 1 with v_t. If T equals the value of search, then its slope equals a weighted average of the slopes of W over all regions, discounted by β, less 1. Therefore $-1 \geq \frac{\partial T(\mathsf{W})}{\partial v_t}$.

$\square$

Now I prove the proposition itself. Consider the trade-off between stopping and search at a point on the threshold. Now increase the level of temptation v_t while fixing v_c and u_c. Consider only what occurs inside the maximization sign, since the $-v_t$ term doesn't change the trade-off between stopping and searching. In this case, the payoff to stopping (inside the maximization) doesn't change. But the payoff to stopping must fall. Therefore the threshold must fall.

$\square$

Because the threshold is falling, the agent no longer chooses the item most recently observed when stopped. When the amount of temptation faced increases, the agent may decide to stop, and go back and choose an item that previously was below reservation value, but is now above the (new) reservation value. This also reverses the unconditional (on the level of temptation experienced) expectation of the length of search conditional on still searching — now the agent will be expected to search less in the future the longer she has already searched.

Turning the the compromise effect, it is no longer the case that after observing a more tempting item, the agent rules out items that are lowest in commitment value. Instead, conditional on having observed a more tempting outcome (higher v_t) the lowest commitment valued item an agent is willing to choose is falling in v_t.

1.8.2 Other Context-Dependent Preferences

In addition to the model of temptation by Gul and Pesendorfer (2001) discussed at length in this paper, there is an extensive literature on characterizing context-dependent preferences (also referred to in the literature as menu-dependent or set-dependent preferences). Building on characterizations of preferences over choice sets by Kreps (1979) and Dekel, Lipman and Rustichini (2001) authors have modelled individuals whose utility depends on a range of non-classical concerns. [17] In fact, it

[17] Another key context-dependent concern in search is regret. There are multiple views of regret. One dates back to Loomes and Sugden (1982), while another is due to Sarver (2008). Irons and Hepburn (2007) model search with a

is possible to extend the analysis to incorporate a variety of preferences that can depend on more than just two of the previously observed items. In doing so, I will rely on preferences described by Dekel, Lipman and Rustichini (2001), restricted to finite subjective states. Then each item can be represented as a point in $\mathbb{R}^{\mathbb{P}+N}$ Euclidean space. There are $\mathbb{P}$ positive states, and N negative states. Positive state i has a payoff denoted u_i while negative state j has a payoff v_j. Items are distributed with joint probability F in $\left(\times_{i=1}^{\mathbb{P}}[\underline{u}_i, \overline{u}_i]\right) \times \left(\times_{j=1}^{N}[\underline{v}_j, \overline{v}_j]\right)$. The payoff from stopping with choice set A is $V(A) = \Sigma_{i=1}^{\mathbb{P}} \max_{x \in A} u_i(x) - \Sigma_{j=1}^{N} \max_{x' \in A} v_j(x')$. Otherwise, the agent's problem remains the same. I will denote $\hat{u}_i = \max_{x \in A} u_i(x')$ and $\hat{v}_j = \max_{x' \in A} v_j(x')$. The case of temptation modeled in the body of the paper is where there is a single positive state ($\mathbb{P} = 1$) and a single negative state ($N = 1$), with the additional restriction that the positive state is a function of the negative state. As a technical restriction, so that the effect of each state on the agent can be identified, no state can be solely a function of any other state. The agent's problem is then:

$$\begin{aligned} \max W(\hat{u}_1, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N) \quad = \quad & \max_{\text{stop, search}} \{\Sigma_{i=1}^{\mathbb{P}} \hat{u}_i - \Sigma_{j=1}^{N} \hat{v}_j, \\ & \beta\left(\int_{\underline{u}_1}^{\overline{u}_1} ... \int_{\underline{u}_{\mathbb{P}}}^{\overline{u}_{\mathbb{P}}} \int_{\underline{v}_1}^{\overline{v}_1} ... \int_{\underline{v}_N}^{\overline{v}_N} V(\hat{u}'_1, ..., \hat{u}'_{\mathbb{P}}, \hat{v}'_1, ..., \hat{v}'_N)\right)\} \end{aligned}$$

The inclusion of additional positive and negative dimension does not change the basic result — there still exists a continuous solution the agent's problem, and it is a threshold function. Because there are multiple positive states, the threshold can be defined for any one of the positive states (in terms of the values of the other positive states and negative states). The proposition defines the threshold in terms of the u_1, but an equivalent statement is true for any u_i.

Proposition 12 *There exists a unique continuous bounded value function $W(\hat{u}_1, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N)$ that solves the agent's problem. Furthermore, the value function follows a threshold: there exists a function $\overline{w}(\hat{u}_2, ..., \hat{u}_{\mathbb{P}}, \hat{v}_2, ..., \hat{v}_N) : \mathbb{R} \to \mathbb{R}$ such that $\hat{u}_1 \geq \overline{w}(\hat{u}_2, ..., \hat{u}_{\mathbb{P}}, \hat{v}_2, ..., \hat{v}_N)$ if and only if $W(\hat{u}_1, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N) = \Sigma_{i=1}^{\mathbb{P}} \hat{u}_i - \Sigma_{j=1}^{N} \hat{v}_j$*

Proof of Proposition 12 The existence and uniqueness of the value function follows from Blackwell's sufficiency conditions, which W can be shown to satisfy using the same proof as in Proposition 1. Furthermore, one can show that a modification of Lemma 2.1 still holds:

finite number of possible draws in a search context. However, their treatment is very different from this paper because they model search with a finite number of possible draws, inducing an additional source of non-stationarity. Regret may not only be backward looking (as the model of temptation in this paper), but can also be forward looking.

Lemma 12.1 *Fixing* $\hat{u}_2, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N$

1. $W(\hat{u}_1, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N)$ *is weakly increasing in* $\hat{u}_1$ *and*

2. $\frac{\partial W(\hat{u}_1,...,\hat{u}_{\mathbb{P}},\hat{v}_1,...,\hat{v}_N)}{\partial \hat{u_1}} \leq 1$ *wherever* $W(\hat{u}_1, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N)$ *is differentiable.*

This can be proven exactly as Lemma 2.1, and therefore, for the same reasons as in in the proof of Proposition 2, it implies the threshold property. □

Furthermore, just as in the case with temptation, the threshold will be rising in the values of the negative states. However, in addition, the threshold is falling in the value of the other positive states.

Proposition 13 *The threshold* $\overline{w}(\hat{u}_2, ..., \hat{u}_{\mathbb{P}}, \hat{v}_2, ..., \hat{v}_N)$ *is increasing in* $\hat{v}_j$ *for all* $j = 1...N$ *and decreasing in* $\hat{u}_i$ *for all* $i = 2...\mathbb{P}$.

Proof of Proposition 13 First, it is easy to prove, using the techniques employed previously, that the following two lemmas are true.

Lemma 13.1 *For each* $i = 1, ...\mathbb{P}$*, fixing all* $\hat{u}_{k\neq i}$ *and all* $\hat{v}_j$*,*

1. $W(\hat{u}_1, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N)$ *is weakly increasing in* u_i *and*

2. $\frac{\partial W(\hat{u}_1,...,\hat{u}_{\mathbb{P}},\hat{v}_1,...,\hat{v}_N)}{\partial u_i} \leq 1$ *wherever* $W(\hat{u}_1, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N)$ *is differentiable.*

Lemma 13.2 *For each j=1,...N, fixing all* $v_{k\neq j}$ *and all* u_i*,*

1. $W(\hat{u}_1, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N)$ *is weakly increasing in* v_j *and*

2. $\frac{\partial W(\hat{u}_1,...,\hat{u}_{\mathbb{P}},\hat{v}_1,...,\hat{v}_N)}{\partial v_j} \geq -1$ *wherever* $W(\hat{u}_1, ..., \hat{u}_{\mathbb{P}}, \hat{v}_1, ..., \hat{v}_N)$ *is differentiable.*

Then, using the reasoning analogous to Proposition 3, the proposition is true. □

References

[1] M. Amador, G. Angeletos, and I. Werning. Commitment vs. flexibility. *Econometrica*, 74:365–396, 2006.

[2] R. F. Baumeister, E.A. Sparks, T.F. Stillman, and K.D. Vohs. Free will in consumer behavior: Self-control, ego depletion, and choice. *Journal of Consumer Psychology*, 18:4–13, 2008.

[3] C. Brown, M. Flinn, and A. Schotter. Real time search in the laboratory and the market. *American Economic Review*, 2011. Forthcoming.

[4] K. Burdett and T. Vishwanath. Declining reservation wages and learning. *Review of Economic Studies*, LV:655–666, 1988.

[5] A. Cooke, T. Meyvis, and A. Schwartz. Avoiding future regret in purchase timing decisions. *The Journal of Consumer Research*, 27:447–459, 2001.

[6] B. De los Santos, A. Hortacsu, and M. R.Wildenbeest. Testing models of consumer search using data on web browsing and purchasing behavior. 2011. Working Paper.

[7] E. Dekel, B. Lipman, and A. Rustichini. Representing preferences with a unique subjective state space. *Econometrica*, 69:891–934, 2001.

[8] E. Dekel, B. Lipman, and A. Rustichini. Temptation-drive preferences. *Review of Economic Studies*, 76:937–971, 2009.

[9] S. DellaVigna and D. Paserman. Job search and impatience. *Journal of Labor Economics*, 23:527–588, 2005.

[10] S. Esteban and E. Miyagawa. Optimal menu of menus with self-control preferences. 2006. Working Paper.

[11] S. Esteban and E. Miyagawa. Temptation, self-control, and competitive nonlinear pricing. *Economics Letters*, 90:348–355, 2006.

[12] S. Esteban, E. Miyagawa, and M. Shum. Nonlinear pricing with self-control preferences. *Journal of Economic Theory*, 135:306–336, 2007.

[13] D. Fudenberg and D. Levine. A dual self model of impulse control. *American Economic Review*, 96:1449–1476, 2006.

[14] F. Gul and W. Pesendorfer. Temptation and self-control. *Econometrica*, 69:1403–1435, 2001.

[15] B. Irons and C. Hepburn. Regret theory and the tyranny of choice. *The Economic Record*, 83:191–203, 2007.

[16] Levav J., H. Heitmann, A. Herrmann, and S. Iyengar. Order in product customization decisions: Evidence from field experiments. *Journal of Political Economy*, 118:274–299, 2010.

[17] D. Kreps. A representation theorem for "preference for flexibility". *Econometrica*, 47:565–577, 1979.

[18] P. Krusell, B. Kuruscu, and Jr. A. Smith. Temptation and taxation. *Econometrica*, 2011. Forthcoming.

[19] G. Loomes and R. Sugden. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92:805–824, 1982.

[20] Y. Masatlioglu, D. Nakajima, and E. Ozdenoren. Revealed willpower. 2011. Working Paper.

[21] J.J. McCall. Economics of information and job search. *Quarterly Journal of Economics*, 84:113–126, 1970.

[22] J. Miao. Option exercise with temptation. *Economic Theory*, 34:473–501, 2008.

[23] J. Noor and N. Takeoka. Uphill self-control. *Theoretical Economics*, 5, 2010.

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

[25] C. Raymond. Revealed search theory. 2010. Working paper.

[26] C. Raymond. Search markets and self-control. 2011. Working paper.

[27] M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82:689–711, 1974.

[28] T. Sarver. Anticipating regret: Why fewer options may be better. *Econometrica*, 76:263–305, 2008.

[29] S. Sharma and S. Bikhchandani. Optimal search with learning. *Journal of Economic Dynamics and Control*, 20:333–359, 1996.

[30] I. Simonson. Anticipating regret and responsibility on purchase decisions. *The Journal of Consumer Research*, 19:105–118, 1992.

[31] N. Stokey, R. Lucas, and E. Prescott. *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge MA., 1989.

[32] J. Stovall. Temptation-driven preferences. *Econometrica*, 78:349–376, 2010.

CHAPTER II

Revealed Search Theory[1]

2.1 Introduction

The class of optimal stopping problems, including search processes, represents an important set of models within economics. Techniques for understanding behavior in these situations, typically using a recursively defined value function, are widespread in economics and finance, in no small part due to the ubiquity of real-world analogs. Search is a key component of many retail, labor, and matching markets. Despite the importance of these situations, and the extensive literature solving optimal stopping and search problems, little has been done to provide behavioral foundations for this class of models. This paper provides necessary and sufficient conditions on observed choice for the class of recursive preferences used in search and optimal stopping models, referred to in this paper as "Generalized Search Representations" (GSRs). The class of GSRs generalizes the classical Bellman equation used in search models.

To preview the results, consider the classical model, where the decision-maker faces a set of currently available outcomes A. She has two actions available to her: stop, or continue searching. If she stops, she can choose any item in A. If she continues searching she will observe an additional outcome drawn from a time-invariant distribution F, and face the same action set. She discounts the future at rate $0 < \delta < 1$. The value of her current choice problem is:

$$V(x) = \max\{\max_{x \in A} u(x), \delta E_F(V(A \cup x'))\}$$

[1] I thank Yan Chen, Josh Cherry, Stephan Lauermann, Dmitry Lubensky, Yusufcan Masatlioglu, Daisuke Nakajima, Dan Silverman and Lones Smith for discussions, as well as seminar participants at the Michigan Theory Lunch and the Midwest Economic Theory Meetings for helpful comments. I also thank Sean Horan for detailed feedback on drafts of this paper, and Efe Ok for suggesting the title. Lastly, I thank the University of Michigan School of Information's Socio-Technical Infrastructure for Electronic Transactions Multidisciplinary Doctoral Fellowship funded by NSF IGERT grant 0654014 for financial support. All remaining errors are my own.

Typically, V is described as the Bellman equation characterizing the value of search.

GSRs include not only the classical model of sequential search, but also many other models, including models where the decision-maker has non-standard preferences. For example, GSRs capture situations where the decision-maker draws from a time-varying distribution, or where the decision-maker cares about multiple outcomes in her history — such as when she purchases multiple items, as in Burdett and Malueg (1981). Moreover, GSRs are flexible enough to accommodate preferences where the decision-maker experiences emotions such as temptation, so that her payoff upon stopping depends not just on the chosen item, but also on unchosen available options (as in Raymond, 2011). For any GSR, the value of a search choice problem will be equal to the maximum of two values: the utility of stopping and choosing an item already observed (which may or may not be equal to the utility of the best item already observed), and the discounted value of waiting a period, receiving an additional option, and facing the same action set.

This paper characterizes GSRs using preferences over sequences of outcomes. Intuitively, the decision-maker is asked whether she would prefer to face sequence x or sequence x'. This is in contrast to the previous literature, which focuses on observing choice from a particular search sequence. For example, suppose that the decision-maker is trying to decide which state to move to receive job offers. In previous papers, the data is the choice of job after the decision-maker chooses to move to a state (for example Michigan). In contrast, I assume that the researcher observes which state the decision-maker chooses to move to — whether the decision-maker would prefer to move to Michigan and choose among currently available jobs there, and have the option of receiving future job offers there; or move to Illinois and choose among currently available jobs there, and have the option of receiving future job offers there. A key characteristic of the observed preferences is that the decision-maker's choice represents a commitment to search within only the chosen stream: if the decision-maker chooses to search using sequence x, she cannot switch at any point in the future to a different sequence x'.

The paper also demonstrates how to distinguish between different sub-classes of GSRs. It also shows under what conditions it is without loss of generality to assume that a decision-maker's discounts using a constant scalar $0 < \delta < 1$. Although actual decisions from within a sequence are not observed (for example, when to stop, and what item to choose), preferences over sequences allow for inference about these behaviors.

The characterization result in this paper has several important benefits. First, it allows for the identification of classes of models that are behaviorally equivalent, and which cannot be distinguished in the data. This can allow for a better understanding of the relationship between alternative search models in the literature.

Second, the characterization result allows for identification of the objective function of the decision-maker, as well as her preference parameters (e.g. the discount rate), solely from observed preferences over sequences. Typically, it is assumed that preferences over final outcomes and the discount rate are known by the researcher, presumably from choice observed on some other domain. Furthermore, it is assumed that the decision-maker will behave rationally and formulate an optimal plan by maximizing a recursively defined value function. In contrast, the approach in this paper allows for direct inference of preferences from behavior in search processes, rather than relying on the assumption that preferences observed in some other domain translate into search settings. Identification of preference parameters can assist in welfare assessments.

Third, the behavioral foundations provided in this paper allow for careful testing of the assumptions underlying models of search. The axiomatic approach can be particularly helpful, because systematic violations of a particular axiom can help identify ways in which the model should be generalized. For example, systematic violations can indicate that the costs and benefits of search may depend not just on the continuation sequence, but also other factors, such as the current history. The axiomatic foundations in this paper allow to test the assumptions in standard search models in three different ways. First, the axioms directly speak to the assumption that the decision-maker has a well defined utility function over final outcomes and continuation sequences which is aggregated using the max operator. This can easily be violated if unchosen, but previously observed alternatives change behavior. Second, the axioms also allow for testing of the assumption that the decision-maker conceives of the value of continuing to search in a recursive fashion — she evaluates the future in the same way she evaluates the present. This may not be true if the decision-maker 'thinks less' about the future, and uses a coarser evaluation. Items which she may exhibit a strict preference over in the present may be indifferent when conceived of in the future.[2] Third, the axioms also can test the assumption that the decision-maker has a preference for receiving payoffs earlier, and furthermore, she has a constant discount rate. This can be violated is she non-time stationary discount rates, or discount rates that depend on either past or future outcomes.

Violations of the assumptions of the standard model of search are empirically relevant. There is direct evidence from psychology of the effects of context on behavior in search: Simonson (1992) and Cooke, Meyvis, and Schwartz (2001) find that either regret about options already observed,

[2] An example of coarsening the future are models of satisficing, as discussed by Simon (1961), Masatlioglu, Nakajima and Raymond (2010), Papi (2010), and Horan (2011). In these papers the decision-maker may have a rational ordering of jobs, but has a threshold that is independent of the future — which implies that the future has only a single equivalence class of outcomes.

or options that will be passed up in the future, can induce changes in the decision-maker's search behavior. Hey (1982, 1987) and Gabaix, Laibson, Moloche, and Weinberg (2006) find evidence for heuristics that are inconsistent with the typical optimal threshold strategies. Using field data collected from consumers who are searching for items among internet retailers De los Santos, Hortacsu, and Wildenbeest (2011) find that subjects' behavior is not adequately described by the standard sequential search model, even when accounting for learning. Brown, Flinn and Schotter (2011) find that even when experimental subjects know that the distribution they are sampling from is stationary, they have a non-stationary threshold policy, ruling out learning as the sole reason for this phenomenon. In a closely related setting, Levav, Heitmann, Herrmann, and Iyengar (2010) find that the order of attribute presentation changes final choices in the product customization of suits and automobiles.

Recognizing the importance of search, several recent papers have begun to explore revealed preferences from behavior in search and stopping settings. Within this literature, including Caplin and Dean (2011) and Horan (2011), as well as the earlier contributions of Rubinstein and Salant (2006), the focus has been on finding conditions that are sufficient for a unique ranking over chosen outcomes, as well as characterizing decision-maker's stopping rule. This literature has not explicitly addressed how to characterize forward-looking beliefs, nor valuation of sequences of outcomes. The approach in this paper not only identifies the ranking over final outcomes (if it exists), but identifies a broader class of preference parameters, including the discount rate, the structure of the aggregator, and the value of continuing to search. This happens because this paper assumes that the observed data are preferences over sequences, rather than choices from within a sequence.

The rest of the paper is organized as follows. First, in order to build intuition, and focus on the role of the key axioms, I consider preferences over sequences of deterministic outcomes. In Section 2, I consider the basic domain of deterministic sequences, define GSRs, relate them to the standard search literature, and provide intuition for identification strategy. I provide the axioms that characterize GSRs and give the main representation theorem (along with uniqueness results). I also compare search preferences across individuals, formalizing a notion of search aversion, when one individual always wants to search less than another, and relates this to the discount rate. In Section 3, I generalize the approach in Section 2 by allowing for future outcomes to be drawn from known distributions, where uncertainty is resolved sequentially, period-by-period. I provide axioms that characterize GSRs in this setting, and also demonstrate how to characterize the discount rate as a constant scalar $0 < \delta < 1$. Section 4 provides a characterization for GSRs that capture standard (rational) search with recall. Section 5 discusses extensions of the model in this paper and concludes.

The appendices provide additional characterization results and proofs of the results.

2.2 Deterministic Sequences

This section provides a characterization results for GSRs when the sequences that the decision-maker (hereafter DM) faces are deterministic. By eliminating uncertainty, this section focuses on providing important intuitions about the axioms that characterize the representation. Because of the dynamic nature of the problem, I consider domains with a history and a future, and a shift operator σ that maps today's history and future into a continuation sequence which represents the future and history the DM faces tomorrow.

Typical models of search assume that DMs are concerned only about the set of items currently available, not the order that they were observed in. Furthermore, they assume that although the future is infinite, the DM has been searching for only a finite history. The model presented here relaxes both of those assumptions. The focus is on sequences where the history and the future are both ordered, infinite sequences. This corresponds to elements that are already available having a date attached to them (the date when they first became available in the past), as well as future elements have a date attached to them (the date that they will become available). This domain accommodates a very general payoff structure — the DM cares not only about what is available to her, but also for how long it has been available. However, the approach in this paper works equally well for other domains, discussed in Appendix A, including situations where the set of currently available items is an unordered set, or where both the past and the future are finite.[3] Furthermore, a key axiom discussed in Section 4 characterizes in terms of choice behavior situations when DM is not concerned about the order of already observed outcomes.

Preferences over sequences have not previously been used to identify search behavior. However, they are observable in a variety real-world settings, can be elicited experimentally, and correspond to existing theoretical constructs. Individuals often choose between locations to search for jobs (e.g. deciding which city to move to after completing university). And although data from lab experiments have focused on choice from sequences, it is relatively easy to design an experiment where DMs first choose among sequences, and then engage in search from within the chosen sequence. Furthermore, preferences over search sequences mirror the assumptions used in the some of the modern macro-

[3]The assumption of infinite sequences is not without loss of generality. Because the domain does not allow for preferences to be observed in situations where the DM has just 'begun' her search (she has no history), as well as preferences over histories of different length, the domain excludes certain phenomenon, such as strategies that are non-time stationary.

labor literature, where DMs direct their search efforts only at a particular sub-market. These models are typically called models of directed search. In that spirit, the sequences of outcomes that the DM's observe are Directed Search Problems (DSPs).

Furthermore, preferences over sequences have a long history of use in understanding dynamic consumption behavior. However, in that literature the sequences are of dated outcomes, with the outcome being consumed at the appropriate date. For example, Koopmans (1960, 1972a, 1972b) considers sequences that begin today and go into the infinite history. In this paper the sequences are still composed of dated outcomes, but the DM does not necessarily consume the outcome x_t at date t. Instead she consumes outcomes only, and decides, endogenously, when to do so. Because of this the DM's preferences are non-time separable.[4]

2.2.1 Domain

I consider a DM who faces an infinite horizon decision problem. At any point in time the DM is in the middle of an infinite sequence of lotteries that extends into the future as well as the past. Outcomes dated in the past, $t \leq 0$, represent outcomes currently available to the DM if she chooses to stop.[5] Outcomes dated in the future, $t \geq 1$, represent outcomes that will become available if the DM continues to search for the appropriate number of periods.

Formally, let Ω be a connected compact metric space, and X be the set of infinite sequences of elements of Ω, which extend from $-\infty$ to ∞. Letting Z represent the set of integers, each $x \in X$ represents a function mapping Z to Ω, where x_t is the entry at time t for sequence x, and x can be denoted as as $(..., x_{-1}, x_0, x_1, ...)$. I will refer to X as the set of directed search problems (DSPs). The DM has preferences over X. The current period when decisions are taking place is $t = 0$. I endow X with the product topology, which is metrizable, and since Ω is compact and connected X is also compact and connected. Since X is a compact metric space, it is separable.

$(y, z) \in X$ is defined as $(y, z)_i = z_i$ if $i \geq 1$ and $(y, z)_i = y_i$ if $i \leq 0$. For every $x \in X$ and $t \in Z$, I denote the subsequence ending at time t as $x^t \equiv (..., x_{t-1}, x_t)$ and the subsequence beginning at time t as ${}^t x \equiv (x_t, x_{t+1}, ...)$. X^t and ${}^t X$ are the sets of those sequences respectively. I will refer to x^0 as the history at x, while ${}^1 x$ is the future. I will denote $\varsigma(x)$ as the unordered set of all elements of x. One extremely important operation is mapping the sequence that the DM faces today into the

[4] In fact, the domain in this paper is very similar to that taken by Rustichini and Siconolfi (2006), who consider sequences that extend into both the infinite past and the infinite future. Relatedly, Rozen (2010) conditions preferences over infinite sequences extending from today into the future on infinite histories.

[5] However, there are situations where an alternative interpretation of the history $(..., x_0)$ is plausible: the history could be the set of possibly available options. Then observed preferences can serve to identify both utility and the extent of the DM's recall. This is discussed in more detail in Appendix B.

sequence that she will face tomorrow — a shift operator. Whenever the DM chooses to continue to search, this operation occurs on the sequence. Let σ map from the domain of DSPs to itself, $\sigma : X \to X$, and be defined as $\sigma(x)_t \equiv x_{t+1}$. Therefore, the outcome that the DM observed today (at $t = 0$), will occupy the $t = -1$ time slot tomorrow (after the σ operation), and so on. I will refer to x as today's DSP (or today's sequence), while $\sigma(x)$ will be the continuation DSP (or tomorrow's DSP). Furthermore, for any subsequence $(x_i, ..., x_j)$, let $\sigma(x)_t = x_{t+1}$.

2.2.2 Representation

Given that the DM faces a sequence of known outcomes $x \in X$, where x_0 is the most recently observed outcome, she has two actions available to her: either she can stop or she can continue to search. If she stops, she receives a payoff $\mu(..., x_0)$ that depends on the set of outcomes already observed $(..., x_0)$. If she continues, she observes outcome x_1 and again faces the same two actions, but discounts the future. Thus, the implicit payoff from searching is the discounted valuation of the continuation sequences $\sigma(x)$. Furthermore, to preserve the recursive structure of the representation, the ranking over $\sigma(X)$ tomorrow must agree with the ranking over X today. Therefore, the value of a sequence x, $V(x)$ is the maximum of two quantities: the value from stopping, $\mu(..., x_0)$, or the discounted value of $V(\sigma(x))$. The type of discounting allowed for is time invariant, but can depend on the value of $\sigma(x)$, and will be denoted by the function ν. [6]

Definition (Generalized Search Representation) *A generalized search representation (GSR) is a triple of functions (V, μ, ν) where $V : X \to \mathbb{R}$, $\mu : X^0 \to \mathbb{R}$, $\nu : \mathbb{R} \to \mathbb{R}$ strictly increasing, such that $\nu(V) < V$, $V(x) = \max\{\mu(..., x_0), \nu(V(\sigma(x)))\}$, and $x \succeq y$ if and only if $V(x) \geq V(y)$*

In order to better understand the representation consider the following graphs. The first graph illustrates how the value of V varies with μ. The second graph illustrates how the value of V varies with $V(\sigma)$.

More formally, to understand Figure 2.1, fix an equivalence class of future outcomes $\{\sigma(y', z') | \sigma(y, z) \sim \sigma(y', z')\}$.[7] Assume there exists a complete and transitive ranking over the set $\{y' \in X^0 |$ there exists a $z' \in {}^1X$ such that $\sigma(y, z) \sim \sigma(y', z')\}$, and represent this ranking with μ. Also assume that it is possible to represent the ranking over all (y', z') with function V. Then Figure 2.1 demonstrates how the value of $V(y', z')$ changes with the value of $\mu(y')$. The horizontal axis represents $\mu(y')$, while the

[6] In fact, this structure of discounting is the strongest that can be identified in the ordinal framework used in this section.

[7] This is analogous, in the classical model of sequential search without recall, to fixing the distribution that is being drawn from in the future.

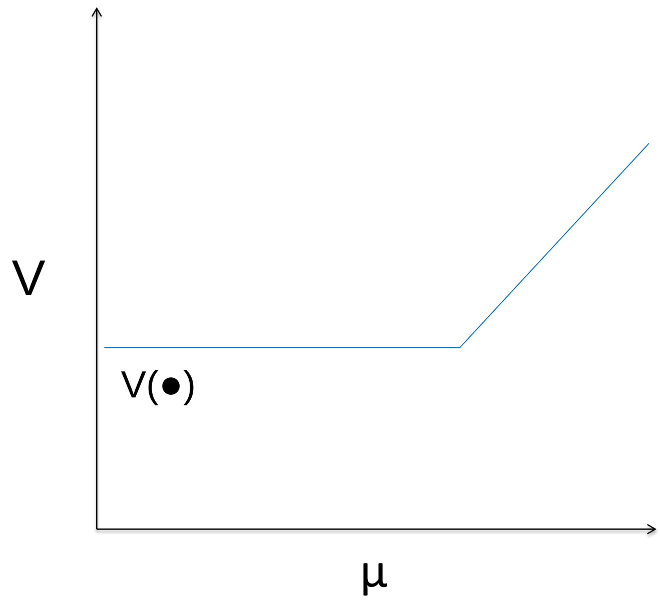

Figure 2.1: Value of sequence x as function of the value of history x^0

vertical axis represents $V(y', z')$. Notice that for low enough values of $\mu(y')$ the value $V(y', z')$ does not depend on $\mu(y')$. At these points the DM will always want to search, and the value of of the sequence today is fixed at the value of the continuation sequence. However, at some point, the value of $V(y', z')$ begins rising monotonically with the value of $\mu(y')$. At these points the DM will want to stop, and so the value of the sequence depends only on $\mu(y')$.

For an alternative view, fix a history $y \in X^0$. Now vary the future $z \in {}^1X$, so that the ranking of $\sigma(y, z)$ varies. Assume it is possible to represent the ranking of $\sigma(y, z)$ and (y, z) with function V. Figure 2.2 shows how the latter varies with the former. The horizontal axis represents $V(\sigma(y, z))$, while the vertical axis represents $V(y, z)$. Notice that for low enough values of $V(\sigma(y, z))$ the value of $V(y, z)$ does not depend on $V(\sigma(y, z))$. It is constant, because the DM will always stop, and receive a payoff that depends only on $\mu(y)$. At some continuation value the value of the current sequence $V(y, z)$ begins to rise monotonically with the value continuation sequence $V(\sigma(y, z))$, and this is where the DM is searching. There can only be one point where she switches from stopping to searching, and at that point the DM is indifferent between stopping and searching.

The 45 degree line (the dotted black line) represents the matching of continuation DSPs and current DSPs that give the same payoff. Because the DM discounts, if she searches, she gets the value of the continuation DSP. Discounting implies that the value of waiting for this DSP is less than the value of the same DSP today. Therefore, the point where the DM is indifferent between

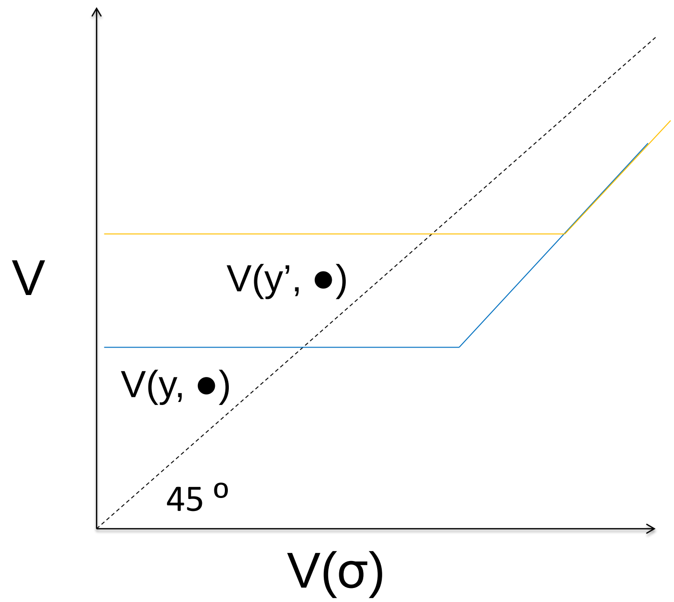

Figure 2.2: Value of sequence x as function of the value of the continuation sequence $\sigma(x)$

searching and stopping must be to the right of the 45 degree line.

Figure 2.2 also shows how the valuations of sequences given different histories related to one another. For any given history where the value of the sequence today does not depend on the value of the continuation sequence the DM chooses to stop. Where the value of the sequence today is increasing in the value of the continuation sequence, the DM chooses to search, and so the value of today's DSP depends only on the value of the continuation DSP and not on the history.

2.2.3 Axioms

The GSR can be characterized by four axioms. The first axiom is the standard requirement for a utility representation on preferences over X.

Continuous Weak Order (CWO) $\succeq$ is is a continuous weak order.

The second axiom, **Limited Stationarity**, states that the value of today's DSP is weakly increasing in the value of the continuation DSP.

Limited Stationarity (LS) For all $y \in X^0$ and $z, z' \in {}^1X$, if $\sigma(y,z) \succeq \sigma(y,z')$ then $(y,z) \succeq (y,z')$

In order to characterize the class of GSRs, it is necessary to identify from preferences, given a particular sequence, whether the DM will choose to stop or continue. If the DM strictly prefers to stop, then it must be that, locally, preferences over the continuation DSPs do not change preferences over today's DSPs. Therefore, it is the case that DM strictly prefers to stop at (y, z), and weakly prefers to stop at (y, w), if $\sigma(y, w) \succ \sigma(y, z)$ but $(y, w) \sim (y, z)$. If the DM strictly prefers to search, then locally preferences over today's DSPs match the preferences over the continuation DSPs. If $(y, w) \succ (y, w')$ then the DM strictly prefers to search at (y, w). She will weakly prefer to search at (y, w) if $(y, w') \succ (y, w)$ and for all z such that $\sigma(y, z) \succ \sigma(y, w)$ it is the case $(y, z) \succ (y, w)$.[8]

The next axiom, **Weak Monotonicity**, applies to situations where the DM is searching. If the DM searches given the current DSP x, then she values x as the discounted value of the continuation stream $\sigma(x)$. Therefore, no other stream which has a continuation value at least as good as $\sigma(x)$ can be worse than x.

Weak Monotonicity (WM) For all $y, y' \in X^0$, $w, w', z' \in {}^1X$ such that (i) $(y, w) \succ (y, w')$, or (ii) $(y, w') \succ (y, w)$ and for all z such that $\sigma(y, z) \succ \sigma(y, w)$ it is the case $(y, z) \succ (y, w)$,

- if $\sigma(y', z') \sim \sigma(y, w)$ then $(y', z') \succeq (y, w)$
- if $\sigma(y', z') \succ \sigma(y, w)$ then $(y', z') \succ (y, w)$

In search settings, it is possible to judge one history y to be better than another y' only if the DM stops when facing history y, and prefers this to stopping at history y'. If the DM has well-defined preferences over histories, it can never be the case that if the DM stops at history y', she prefers this to stopping given a history y. Preferences over histories can be recovered by observing reversals in the preference ranking of a DM between today's DSPs and tomorrow's DSPs. Suppose that tomorrow the DM will prefer to look for a job in Illinois (denoted by sequence (y, w)) rather than in Michigan (denoted by sequence (y', w')). In this case, the value of continuing to search today in Illinois must be better than the value of continuing to search today in Michigan. However, also suppose that today the DM prefers to look for a job in Michigan. In other words the DM prefers $\sigma(y, w)$ to $\sigma(y', w')$, but prefers (y', w') to (y, w). Then it must be that she wants to take a job currently available in Michigan (i.e. she will stop at (y', w')) and second, the value of stopping in Michigan today must be better than the value of stopping in Illinois today. This is because if the DM searches today, her ranking of today's DSPs will agree with her ranking of the continuation DSPs, because if a DM searches, all she cares about is the value of the continuation DSP. A reversal

[8]This is true only if these comparisons are local — for any neighborhood around (y, w) there is a (y, w'). If the comparisons were not local it would not be clear that the DM weakly preferred to search, or strictly preferred to stop at (y, w). Because the set of outcomes is connected, local comparisons are guaranteed to exist.

can only occur if the DM plans on stopping and choosing from the items already available to her. I will restrict the DM's preferences so that the value of stopping today in Michigan, relative to the value of stopping today in Illinois, cannot be changed by job offers the DM has not yet received.

Unfortunately, an ordering over histories induced by only looking at these types of preference reversals may be incomplete. Intuitively, the reason is that preference reversals are observed only when the individual is willing to stop. For some DMs there may be histories where no future would make the DM want to stop today.[9] Because the DM never wants to stop, it is not clear how to rank these sequences. All that is possible is to provide a kind of "upper bound" on the ranking of the history. However, because the DM always wants to search given that history, this upper bound is sufficient for constructing preferences. Therefore, one way to rank histories is by the value of the $\succeq$ minimal element in any particular history.

Definition ($\succeq_0$) *For all $y, y' \in X^0$, $y \succeq_0 y'$ if there exists a $z' \in {}^1X$ such that for all $z \in {}^1X$ $(y, z) \succeq (y', z')$.*

Not surprisingly, preference reversals between today and tomorrow generate strict rankings in $\succeq_0$.

Lemma 1 *If for $y, y' \in X^0$, there exists $z, z' \in {}^1X$, $\sigma(y, z) \preceq \sigma(y', z')$ and $(y, z) \succ (y', z')$, then $y \succ_0 y'$.*

By endowing X^0 with the product topology (which makes it a metrizable separable space) it follows that $\succeq_0$ has all the usual properties of a representable preference relation.

Lemma 2 *If $\succeq$ satisfies **CWO** then $\succeq_0$ is a complete, transitive, and continuous binary relation.*

When a DM searches she believes she will stop at some point. To guarantee that this always happens, the DM must value tomorrow's sequence today less than she values it tomorrow — in other words, she must discount. The last axiom captures the appropriate notion of discounting for preferences over X.[10]

Sooner is Better(SB) For all $y \in X^0$, $z, z' \in {}^1X$, if $(y, z) \succ (y, z')$ then $(y, z) \prec \sigma(y, z)$

These axioms, all of which are necessary, are also sufficient for a GSR.

[9] With classical preferences this would never be the case.

[10] In the dynamic consumption literature, no separate axiom is needed for discounting, because the existence of the additive aggregator between today and tomorrow's payoffs imply a discount factor. In this paper the aggregator is non-additive, and so discounting must be imposed separately.

Theorem 1 *$\succeq$ satisfies axioms **CWO**, **LS**, **WM**, and **SB** if and only if it can be represented using a GSR (V, μ, ν), such that μ represents $\succeq_0$. Furthermore, $V^{'}$ represents $\succeq$ if and only if $V^{'} = \psi \circ V$, for some monotonically increasing function ψ.*

In many applied settings DMs are assumed to have a constant discount rate. Given preferences over deterministic sequences it is only possible to identify a unique discount rate (up to monotonic transformations of the representation) for each equivalence class of continuation DSPs. This type of non-constant discount rate is conceptually different than many other models with non-constant discount rates. Typically, the discount rate varies only with time (as in the case of quasi-hyperbolic discounting). Future outcomes do not change the discount rate.[11] In contrast, the discount rate may vary with future outcomes for GSRs, but, conditional on a particular equivalence class of continuation DSPs, is constant over time.

One way to understand the role of each axiom is to examine situations where it may fail. The following examples provide search preferences that fail to satisfy the axioms. The examples provide intuition for the independence of the axioms.

- **Regret:** Suppose that the DM anticipates that if she stops now, but observes a better outcome in the future, she will experience a loss of utility.

- **Self-Control:** Following Raymond (2011) suppose that the DM experiences self-control problems whenever she observes a tempting item. She pays a cost of self-control every period until she chooses an item.

- **Patient Search:** The DM will search until she comes across the best item in the sequence, no matter how many periods it takes.

The preferences of an agent in the regret example violate **LS** because the DM's stopping payoff can fall if the continuation sequence improves, but satisfies the other axioms. A DM with self-control preferences fails **WM** but satisfies the other axioms. To see why, imagine two sequences x and x' such that DM searches given both sequences, and $\sigma(x) \sim \sigma(x')$. However, the most tempting item observed so far in x imposes a higher cost of self-control than the most tempting item observed so far in x'. In this case the DM prefers x to x', which violates **WM**. An agent who is patient does not discount and so violates **SB**.

[11] There are some models where the discount rate can also vary by past or future consumption, such as Noor (2011) and Epstein (1983).

2.2.4 Uniqueness

Theorem 1 provides a weak uniqueness result regarding V. Stronger uniqueness results are essential in order to make interpersonal comparisons between DMs. Attempts to uniquely identify the functional forms of μ and $\nu \circ V$ (up to the standard monotonic transformations) of the representation are hampered by the structure of the aggregator and possible paucity of the data. Because V aggregates payoffs using the max operator, changes to either μ or $\nu \circ V$ does not change the revealed preference of the DM unless she strictly prefers to stop or search respectively. For example, if the DM strictly prefers stopping, then small changes to the value of the continuation DSP do not change observed preferences, and so it is not clear how the ranking of tomorrow's DSP translates into preferences today. Similarly, if the DM strictly prefers continuing, then small changes to the history (leaving the continuation sequence unchanged in terms of preference) do not change observed preferences. Unique identification of preferences is possible only when changes in μ or $\nu \circ V$ change V, situations that captured by the following definitions. Strict f-relevance of either μ or $\nu \circ V$ (or a history or continuation sequence respectively) imply that they directly affect the value of V, and if a history is j-relevant then it contains futures that make it strictly f-relevant for both μ and $\nu \circ V$.

Definition (Relevance) *Given a GSR representation and a sequence (y, z):*

- μ *is strictly f-relevant at (y, z) if there exists a $z' \in {}^1X$ such that $\sigma(y, z') \succ \sigma(y, z)$ but $\sigma(y, z') \sim \sigma(y, z)$. Furthermore, history y is strictly f-relevant if there exists a future z such that μ is strictly f-relevant at (y, z).*

- $\nu \circ V$ *is strictly f-relevant at $\sigma(y, z)$ if there exists a $w \in {}^1X$ such that $(y, z) \succ (y, z)$, or $(y, w) \succ (y, z)$ and for all z' such that $\sigma(y, z') \succ \sigma(y, z)$ it is the case $(y, z') \succ (y, z)$. $\sigma(y, z)$ is strictly f-relevant if there exists a (y'', z'') such that $\sigma(y'', z'') \sim \sigma(y, z)$ and $\nu \circ V$ is strictly f-relevant at $\sigma(y'', z'')$.*

- *History y is j-relevant if it is strictly f-relevant and there exists a future $z \in^1 X$ such that $\sigma(y, z)$ is strictly f-relevant.*

Theorem 2 shows that the uniqueness of the ranking over histories can be achieved where the μ is strictly f-relevant. It also shows that, within the class of GSRs, $\nu \circ V$ is unique.

Theorem 2 *Assume $\succeq$ satisfies axioms **CWO**, **LC**, **WM**, and **SB**, while GSR (V, μ, ν) and GSR (V', μ', ν') both represent $\succeq$.*

- *If histories y and y' are both strictly f-relevant, then $\mu(y) \geq \mu(y')$ if and only if $\mu'(y) \geq \mu'(y')$.*
- $\nu(V(\sigma(y,z))) \geq \nu(V(\sigma(y',z')))$ *if and only if* $\nu'(V(\sigma(y,z))) \geq \nu'(V(\sigma(y',z')))$.

The uniqueness of $\nu \circ V$ is within the class of GSRs. However, more robust uniqueness results can be obtained. Consider the more general class of 'quasi-GSRs', which take the form $V(y,z) = \max\{\mu(x), \gamma(y,z)\}$. Within this class (which contains GSRs), uniqueness can be achieved by ensuring not only that each history is strictly f-relevant, but also that for each equivalence class of continuation DSPs there is a sequence today, whose continuation sequence is in the equivalence class, and where slightly altering the ranking of the continuation sequence alters the ranking of the current sequence. This is equivalent to the strict f-relevance of $\nu \circ V$. In this case, the rankings that μ and $\nu \circ V$ provide cannot be reversed by μ' and γ for any quasi-GSR wherever strict f-relevance holds.[12]

2.2.5 Comparative Search Aversion

One of the primary goals of identifying unique representation of preferences is that it can assist in making useful statements regarding comparative statics. In particular, it may be useful in applied settings to be able to say when DM A values observing an additional option less than DM B. This section considers such comparative statics. Strong results showing equivalence between comparative search aversion and the discount function can be obtained for situations where A and B have the same preferences over sequences if they stop. Formally, denote $\succeq_A$ as DM A's preferences, and likewise for $\succeq_B$. Do the same for $\succeq_A^0$ and $\succeq_B^0$.

Definition (Search-Aversion) *A is more ordered search-averse than individual B if*

- $\succeq_A^0 = \succeq_B^0$ *and*
- *For any history y, such that y is strictly f-relevant, and there exists a $z \in {}^1X$ such that $\sigma(y,z)$ is strictly f-relevant, for both A and B, if $(y,z) \sim^A (y,z')$ then $(y,z) \sim^B (y,z')$.*

If A is more search averse than B, then fixing a history, if A is indifferent between two DSPs (meaning A wants to stop), then B must also be indifferent (meaning B must also want to stop).

[12] It may be implausible in many settings that all continuation sequences are strictly f-relevant. For example, assume a DM has preferences that can be represented by $V(x) = \max\{u(x_0), \beta V(x_1)\}$ (i.e. that the DM is rational and has no recall). There will be some continuation DSPs (for example, those that feature only the worst outcome for periods $t = 1, 2, ...T$ for large enough T) where the DM will always want to stop, regardless of x_0. Those continuation DSPs will not be strictly f-relevant. Furthermore, it is impossible for all histories and all continuation sequences to be strictly f-relevant at the same time.

If A and B's preferences are sufficiently rich, then this maps directly to a relationship about A and B's relative discounting.

Theorem 3 *Assume $\succeq_i$ for $i = A, B$ satisfies axioms **CWO**, **LC**, **WM**, **SB**, that for all histories y, y is strictly f-relevant for both A and B and for all $(y, z) \in X$, the set of sequences $w \in X$ such that $w \sim (y, z)$, that are strictly f-relevant is dense. Consider representations of A and B's preferences V_A, μ_A, ν_A and V_B, μ_B, ν_B such that $\mu_A = \mu_B$. A is more ordered search-averse than B if and only if $\nu_A \leq \nu_B$ everywhere $\nu_A \circ V_A$ and $\nu_B \circ V_B$ are both strictly f-relevant.*

Theorem 3 focuses on situations where A and B have the same preferences over histories. In this situation, when preferences are sufficiently rich, A being more search-averse than B is equivalent to A discounting the future more than B for any history and any future. Relatedly, Quah and Strulovici (2011) provide conditions on the relationship between DMs' discount rates that are sufficient for search aversion for a broader class of optimal stopping problems than those considered in this paper.

2.3 Resolution of Uncertainty

A common feature of search and optimal stopping processes is that the DM is not facing a known sequence of deterministic outcomes. Instead, she believes that if she continues searching, at each period in the future the offer she will receive is drawn from a distribution. In each period, the uncertainty about the offer that is received resolves independently.

This section extends the results in Section 2 to allow for sequential resolution of uncertainty, as well as the identification of a constant discount rate. Furthermore, the DM's attitude toward uncertainty will be captured in a particular manner — as satisfying expected utility with objectively known probabilities.[13] In order to accommodate these features, the domain of preferences must be extended. First, sequences consist of histories of deterministic outcomes, but futures of known distributions that outcomes will be drawn from independently over time. Second, preferences are not only over sequences, but also over lotteries of sequences, the uncertainty about which will be resolved after the preferences are observed, but before the DM actually begins searching from any given sequence. Key to the approach in this section is that the DM is indifferent, for very particular types of decision problems, over whether uncertainty is resolved before she searches (i.e. the uncertainty

[13]This does not mean that it is not possible to capture other types of attitudes toward uncertainty, such as subjective expected utility or ambiguity aversion.

is contained in lotteries over sequences), or after (i.e. the uncertainty is contained in lotteries over outcomes). The DM cannot *always* be indifferent to the timing of the resolution of information, because payoffs are aggregated using a max operator, which is generally not exchangeable with the expectations operator.[14]

Formally,assume that Ω is a compact connected metric space. Moreover, set of simple lotteries (i.e. with finite support) over the outcomes, $\Delta(\Omega)$, is a compact metric space (when endowed with the weak topology). Furthermore, denote $\hat{X}$ as the set of sequences such that for sequence $x \in \hat{X}$, if $t \leq 0$, then $x_t \in \Omega$, but if $t > 0$, then $x_t \in \Delta(\Omega)$. For the sake of clarity, sequences in $\hat{X}$ will be denoted $(..., x_{-1}, x_0, f_1, ...)$, where x is a deterministic outcome, while f is a lottery over Ω. Furthermore, denote the set of simple lotteries over $\hat{X}$ as $\Delta(\hat{X})$. Last, for any for any subsequence $(x_i, ..., x_j) \in \hat{X}$, it is the case that $\sigma(x)_t = x_{t+1}$.[15]

$\hat{X}$ is the set of sequences such that for every period in the past, and the current period, the outcome is deterministic. However, for every future period, the outcome is as of yet unresolved. In other words, all that is known is the distribution that the outcome will be drawn from. This captures the uncertainty regarding future draws. $\Delta(\hat{X})$ is composed of lotteries of sequences, each of which has a history of deterministic outcomes and a future of independently resolving lotteries. In order to fully characterize standard search behavior in this setting, there will be four axioms that describe preferences over $\Delta(\hat{X})$.[16]

Since there are lotteries operating at multiple levels it is important to be clear about the timing of the resolution of uncertainty. The observable data, preferences over lotteries of sequences, is observed when the DM chooses a member of $\Delta(\hat{X})$, in Epoch 0. Suppose the DM chooses α. Then, all uncertainty about which element of $\hat{X}$ in the support of α the DM is facing is resolved in Epoch 1. Epoch 2 is subdivided into time periods $t \in (0, 1, ...)$. Suppose the DM faces $x \in \hat{X}$. At time $t = 0$, the DM decides whether to stop, facing sequence x or to continue. If she chooses to continue, then at the beginning of $t = 1$ the uncertainty regarding $f_1 \in x$ is resolved, and the DM again chooses whether to stop or to continue, and so on. When the DM decides to stop and choose an item, Epoch 2 ends and Epoch 3 begins. In Epoch 3, the DM's payoff is realized. Note that the DM experiences no elapsed time (i.e. does not discount) between epochs, only between time periods.

[14]The behavioral axioms that characterize attitudes toward the resolution of uncertainty in this section are similar to the Reversal of Order axiom used in subjective expected utility. Seo (2008) has a useful discussion how subjective expected utility requires that the DM be indifference about whether objective uncertainty is resolved before or after subjective uncertainty.

[15]$\sigma({}^ix^j)$ may not necessarily be a subsequence of an element of $\hat{X}$. However, in this section σ will only be applied to subsequences ${}^ix^j$ such that both ${}^ix^j$ and $\sigma({}^ix^j)$ are subsequences of elements of $\hat{X}$.

[16]As rest of the section will make clear, restricting preferences to $\hat{X}$ does not allow interchanging of the expectations and the discount functions.

The resolution of uncertainty regarding $\Delta(\hat{X})$ and the realization of payoffs upon a choice from a sequence are instantaneous. Figure 2.3 illustrates the timeline.

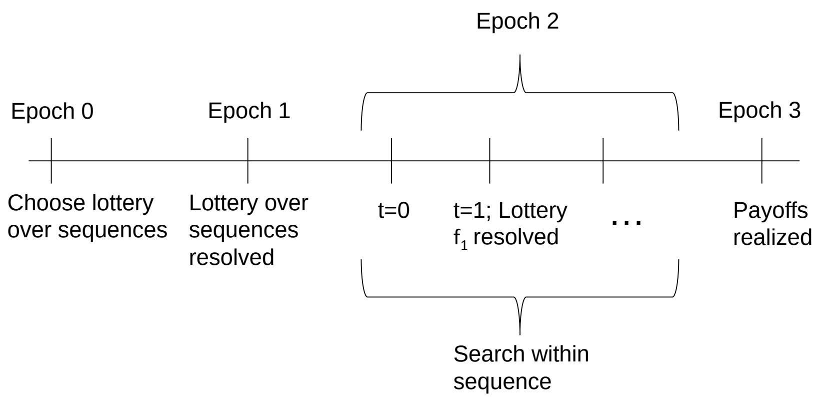

Figure 2.3: Timeline

Expected Utility GSRs extend the definition of GSRs presented in Section 2. They explicitly allow for the resolution of uncertainty regarding the future draws as well as a linear discount rate.

Definition (Expected Utility Generalized Search Representation) *An expected utility generalized search representation (EU-GSR) is a quadruple of functions* (U, V, μ, δ) *where* $U : \Delta(\hat{X}) \to \mathbb{R}$, $V : \hat{X} \to \mathbb{R}$, $\mu : X^0 \to \mathbb{R}$, *and* $\delta : X^0 \to (0, 1)$ *such that*

$$U(\alpha) = \sum_{x \in \alpha} p(x)V(x)$$

$$V(x^0,\ {}^1f) = \max\{\mu(x^0), \delta(x^0)E_{x' \in f_1}(V(\sigma(x^0), x', \sigma({}^2f)))$$

and for all $\alpha, \beta \in \Delta(\hat{X})$, $\alpha \succeq \beta$ *if and only if* $U(\alpha) \geq U(\beta)$

The first axiom ensures that there exists a representation of $\succeq$ that has an expected utility form.

Expected Utility on Lotteries of Sequences (EULS) $\succeq$ is a continuous weak order that satisfies independence.

Consider a DM who is evaluating the continuation value of the sequence $x \in \hat{X}$. She knows she will face resolution of uncertainty about the offer from distribution f_1 at the beginning of next period. In order to characterize the DM's behavior, I will construct a lottery that gives an equivalent value to searching given x, but features no uncertainty at $t = 1$. The first step is to define a lottery over sequences, where all of the sequences in the support of the lottery agree on all elements at all time periods $t \neq 0$. The elements of the sequences at $t \neq 1$ correspond to the elements of x, shifted forward by one period. However distribution over outcomes at time $t = 0$ exactly mimics the distribution f_1. Denote this lottery α_x.

Definition (α_x) *Let $f_1 = (p(z_1), p(z_2), ...)$. For any sequence $x = (..., x_0, f_1, f_2, ...)$, define a lottery over sequences α_x where the support of α_x consists of sequences which:*

- *Agree for entries $t \neq 0$. For $i \neq 0$ let the i^{th} entry of any element in the support of α_x be equal to the $x_(i+1)$. Denote this element $\alpha_{x,i}$.*
- *For entry $t = 0$, the entry is z_i for some z_i in the support of f_1*

Furthermore, the element in the support of α_x that assigns z_i to entry $t = 0$ occur with probability with probability $p(z_i)$.

Next, consider a sequence that agrees with all the sequences in the support of α_x at all time periods $t \neq 0$ and is valued the same as α_x. Because Ω (and so $\hat{X}$) is connected, and since **EULS** holds, there exists a $z^* \in \Omega$ such that $\alpha_x \sim (..., \alpha_{x,-1}, z^*, \alpha_{x,1}, ...)$.[17]

Lemma 3 *If $\succeq$ satisfies **EULS** then for all $x \in \hat{X}$ there exists $x' \in \hat{X}$ such that $\alpha_x \sim x'$, and where $\alpha_{x,i} = x'_i$ for all $i \neq 0$.*

At this point it is possible to construct a new shift operator $\hat{\sigma}$, which maps from $\hat{X}$ to itself. Rather than literally describing how today's sequence will map tomorrow's sequence (as σ did), $\hat{\sigma}(x)$ finds a sequence such that the discounted value of $\hat{\sigma}(x)$ is equal to the continuation value of x. Given this new shift operator, the axioms developed in Section 2 can be adopted to the new domain in this section.

Definition ($\hat{\sigma}$) *Define function $\hat{\sigma}$: $\hat{X} \to \hat{X}$ such that $\hat{\sigma}(..., x_0, f_1, ...)$ has the following properties:*

[17] z^* may not be unique. Any of the potential z^*s can be used arbitrarily.

- $\hat{\sigma}(x)_t = x_{t+1}$ *if* $t \neq 0$.
- $\hat{\sigma}(x)_0 = z^*$, *where* $z^* \in \Omega$ *is an arbitrarily chosen element such that* $\alpha_x \sim (..., \alpha_{x,-1}, z^*, \alpha_{x,1}, ...)$.

Such a function exists and is continuous.

Lemma 4 $\hat{\sigma}$ *exists and is continuous.*

To extend the previous characterization to the setting considered here, simply replace σ with $\hat{\sigma}$ for each axiom **LC**, **WM**, and **SB**. It is worth noting that this construction gives the axioms stronger implications than they had in Section 2. In particular, they imply that the DM evaluates today's sequences using both the value of the history x^0, but also the expected value of the continuation sequence, which is evaluated using certainty equivalents derived from preferences over lotteries of sequences. Because $\hat{\sigma}$ depends on the DM's preferences, axioms that reference it may be more difficult to test.

GSR Existence (GSRE) LC, **WM**, and **SB**, hold but with $\hat{\sigma}$ substituted for σ.

A DM who is an expected utility maximizer will assign the same value to searching given x, which is the discounted value of $\hat{\sigma}(x)$, as she does to the discounted value of α_x. This is because both cases, the DM makes a choice (chooses to search, or chooses a sequence), and experiences resolution of the same uncertainty about the additional information she has before she has to decide whether to stop or search (again). Therefore, the DM will assign the same value to $\hat{\sigma}(x)$ and α_x.

Lottery Indifference (LI) For all $x \in \hat{X}$, $\alpha_x \sim \hat{\sigma}(x)$.

Suppose the DM is facing sequence x. Consider the set z_i in the support of f_1, and assume that for all sequences $(..., x_{-1}, x_0, z_i, f_2, ...)$ the DM prefers to search. In this case, the DM will be indifferent about whether the uncertainty about f_1 is resolved after she decides to search given x, or before she decides to search. x itself captures the situation where the uncertainty about f_1 is resolved after the DM decides to search. The situation where the uncertainty is resolved before the DM decides to search can be captured by construction a lottery over sequences, where all the sequences agree with x at time periods $t \neq 1$, and at $t = 1$ the lottery mimics the distribution induced by f_1. But rather than the uncertainty being contained within a sequence, the uncertainty is contained within a lottery over sequences, and so each sequence in the support of the lottery features no uncertainty at $t = 1$. Call this lottery α_{x,f_1}.

Definition **(α_{x,f_1})** *Let $f_1 = (p(z_1), p(z_2), ...)$. For any sequence $x = (..., x_0, f_1, f_2, ...)$, define a lottery over sequences α_{x,f_1} where the support of α_x is $(..., x_0, z_i, f_2, ...)$, for all z_i in the support of f_1, such that α_x assigns $(..., x_0, z_i, f_2, ...)$ with probability $p(z_i)$.*

A DM who always wants to search regardless of the actual realized outcome at $t = 1$, will be indifferent about when the resolution of uncertainty regarding outcomes at $t = 1$ is resolved. Therefore, she is indifferent between x and α_{x,f_1}. It is important to note that the DM is indifferent about the timing of the resolution of uncertainty only when she will search (at $t = 0$) for all z_i in the support of f_1. If this is not the case, then the DM should prefer earlier resolution of information (i.e. she prefers information to be contained in lotteries over sequences, rather than in f_1).

> **Continuation Linearity (CL)** Let $f_1 = (p(z_1), p(z_2), ...)$. If for all z_i in the support of f_1, it is the case that either there exists an $x' \in \hat{X}$ such that $(..., x_0, z_i, f_2, ...) \succ x'$, or there exists an $x' \in \hat{X}$ such that $x' \succ (..., x_0, z_i, f_2, ...)$ and for all $x'' \in \hat{X}$ such that $\hat{\sigma}(x'') \succ \hat{\sigma}(..., x_0, z_i, f_2, ...)$ it is the case that $(x'') \succ (..., x_0, z_i, f_2, ...)$, then $x = (..., x_0, f_1, f_2, ...) \sim \alpha_{x,f_1}$[18]

Although **LI** and **CL** are very similar, they address different issues. **LI** addresses lotteries over sequences where there is some uncertainty about the current period outcome. **CL** addresses lotteries over sequences where there is uncertainty about the outcomes starting tomorrow. Furthermore, the condition in **LI** applies to any sequence, whereas **CL** only applies to sequences where the DM always wants to search.

It is now possible to characterize EU-GSRs, and furthermore, when it must be the case that the DM's discount rate is not only linear, but also constant.

Theorem 4 *$\succeq$ satisfies **EUSL**, **GSRE**, **LI** and **CL** if and only if there exists an EU-GSR that represents $\succeq$. Furthermore, if there exists a sequence $\hat{\sigma}(y, z) \in X$ such that for all $\hat{\sigma}(y', z') \in X$, if $\hat{\sigma}(y', z') \succ \hat{\sigma}(y, z)$ then there exists a $\hat{\sigma}(y'', z'') \in X$ such that $\hat{\sigma}(y', z') \sim \hat{\sigma}(y'', z'')$ and $\nu \circ V$ is strictly f-relevant for $\hat{\sigma}(y'', z'')$, then $\delta(x^0) = \delta$.*

[18]Recall from **WM** that for all $y, y' \in \hat{X}^0$, $w, w', z' \in {}^1X$ the DM prefers to search at (y, w) if $(y, w) \succ (y, w')$, or if $(y, w') \succ (y, w)$ and for all z such that $\hat{\sigma}(y, z) \succ \hat{\sigma}(y, w)$ it is the case $(y, z) \succ (y, w)$.

This representation captures the standard model of search with sequential resolution of uncertainty and constant discounting. The first part of Theorem 4 describes the existence of a model that allows for sequential resolution of uncertainty, as well as linear (but not necessarily constant) discounting. The second part discusses when the discount factor must be constant. In order to guarantee a constant discount factor the set of continuation payoffs that are strictly f-relevant must be a connected set. The reason is that although **CL** is sufficient for ν to be linear, it does not necessarily impose a constant discount rate. For any given history, the discount rate will be constant. **GSRE** implies that whenever two histories overlap where the DM is searching, the discount rate used for each of those two histories must be the same. However, if there is a 'gap' in the data where histories do not overlap when they are searching, then the discount rate, while always remaining linear, can change.

2.4 Preferences over Histories

The representations studied so far are quite permissive about the payoff that the DM receives upon stopping. These representations only restrict this payoff to be solely a function of the observed history x^0. This section characterizes a model of search with full recall where the DM has a preference relation over the set of outcomes, and is fully rational in evaluating his current choice set. In this case $\mu(x^0) = \max_{x_i \in x^0} u(x_i)$ for a continuous function u. A more detailed discussion of other models of search, including search with no recall, and search with non-standard preferences, is addressed in Appendix B. To make the analysis simpler, I assume that for each history μ is strictly f-relevant.

If all items in the history are available to the DM to choose upon stopping, and their payoff does not depreciate, the order in which items were observed does not matter. Furthermore, if the DM is not learning, but is otherwise rational, the number of times that she has observed a particular option does not change her payoff upon stopping, since she only chooses one item. In this case, if two histories correspond to exactly the same unordered set of outcomes, the DM will be indifferent between them.

Order Independence (OI) For all $y, y' \in X^0$, $z \in {}^1X$, if $\varsigma(y) = \varsigma(y')$ then $(y, z) \sim (y', z)$.

I next define an operator that merges histories.[19]

[19]Strict f-relevance helps clarify how to rank merged histories.

Definition ($\biguplus$) *Denote* $\biguplus : X^0 \times X^0 \to X^0$ *as the operator that merges the two sequences together by alternating the elements from them* $y \biguplus y' = (..., y'_{-1}, y_{-1}, y'_0, y_0)$.

If the DM satisfies **OI** then her history is equivalent to an unordered set of outcomes. Furthermore, if a DM has a preference ordering over outcomes, and only cares about the payoff from her chosen outcomes, then each history's value is the utility value of the best item in that history.

> **Rational Recall (RR)** If for $y, y' \in X^0,\ z, z', w, w', \ddot{z}, \ddot{w} \in X^1,\ (y, z) \sim (y, z')$ and $(y', w) \sim (y', w')$, $(y, z) \succeq (y', w)$, and $(y \biguplus y', \ddot{w}) \sim (y \biguplus y', \ddot{z})$ then $(y, z) \sim (y \biguplus y', \ddot{z})$.

These axioms are necessary and sufficient to characterize the standard representation of search with full recall.

Theorem 5 *Assume that* $\succeq$ *satisfies axioms* ***CWO***, ***LC***, ***WM***, *and* ***SB*** *and* μ *is strictly f-relevant for every history. Then* $\succeq$ *satisfies* ***OI*** *and* ***RR*** *if and only if there exists a continuous function* u *such that* $\mu(x^0) = \max_{x_i \in x^0} u(x_i)$.

2.5 Conclusion

In this paper, I provide behavioral foundations for recursive preferences used in search and optimal stopping problems using preferences over sequences. Furthermore, I demonstrate how to characterize interpersonal comparisons of willingness to search.

The class of GSRs characterized in this paper accommodate a wide range of behavior. That said, they exclude certain interesting models of search. Future work will extend the results of this paper to allow for bounded rationality, interaction effects between the stopping and continuation valuations, and learning. For example, within the class of GSRs, if the DM stops, future draws cannot change her payoff. This can be violated is some situations, such as if the DM anticipates experiencing regret at having stopped and chosen too early. Furthermore, if the DM continues to search, a GSR implies that her payoff can only be influenced by the value of the continuation sequence. This rules out the flow cost of search depending on the current history. Future work will demonstrate how to characterize more general value functions that allow for interaction effects. Future work will also characterize DMs who learn about the distributions they face in the future from past draws. If the DM is learning, the distributions that she faces in the future are not independent of the realizations of the offers in the past.

It is important to understand the behavioral implications of search models. GSRs are a flexible class of representations that allow for testing, not just of the classical model of search of McCall (1970), but also models of search where the DM can purchase multiple items, as in Burdett and Malueg (1981), or where the DM experiences emotions as in Raymond(2011)) or Irons and Hepburn(2007). These papers, and the literature that builds on them, demonstrate that small changes to the DM's utility function can have important implications not only for individual search behavior, but also for market outcomes.

2.6 Appendix A: Alternative Domains

The paper has, in the main text, considers a domain where the past and future are both infinite ordered sequences. However, the representation results in this paper are easily adapted to other domains. Throughout this section I assume that the future is an ordered sequence. This is because the DM 'should' know what distribution, or outcome, she is facing at every time period in the future. If she is uncertain about which distribution she may draw from, because this paper excludes learning she should simply consider the reduced form distribution over outcomes. However, I will vary the way the history the DM faces is structured, as well as the size of the history and future.

If the researcher knows that the DM has limited recall, then this can specifically be incorporated into the domain. The history can be conceived of as a sequence with a finite number of entries —- the full sequence can be conceived of as $(x_{-T}, ..., x_0, ...)$, for $T \in \mathbb{N}$.

If, instead, the researcher knows that the order of outcomes, the length of time an option has been available, and the multiplicity of any particular outcome do not change the DM's evaluation of the history, then it is more appropriate to treat the history as an unordered set, which is the set of outcomes currently available, but with no date attached.[20] In these situations the domain is $(A, x_1, ...)$, where A is the currently available choice set, and x_i are options that will be observed if the decision maker continues to search for an additional i periods.[21]

In the real world, the domain is typically not stationary. If the DM has been searching for n periods, she has only observed n outcomes. The framework developed in this paper can accommodate these situations. Consider Y_1, the set of sequences where only a single outcome has been observed. Next consider Y_2, the set of sequences where only two outcomes have been observed, and so on. To

[20]To relate it to the domain in the main text, the DM will treat any two histories that generate the same unordered set as equivalent. Recall that φ maps ordered histories to unordered histories.

[21]Working with a domain of the type $(A, x_1, ...)$ limits the results. Assume that the DM has limited recall that depends on the length of time that an offer has been available. For example, the DM may only be able to recall options from the current period ($t = 0$) and the last period in the past ($t = -1$). In this case, it matters whether an option was observed this period or last period.

avoid ruling out general forms of context dependence, it is not possible to derive the DM's preferences by looking only at any particular set Y_i, $i < \infty$. Instead, it is necessary to take the set $Y = \cup_{i=1}^{\infty} Y_i$ as the domain. Two additional modifications must be made. First, additional technical conditions are required to ensure that Y possess the necessary properties to ensure a utility representation, since an infinite union of compact sets is not necessarily compact. Second, the axioms must be modified to apply to each Y_i — by defining a modified shift operator $\sigma_i : Y_i \to Y_{i+1}$, and a preference $\succeq_i$ over Y_i and ensuring that the axioms hold for each i and $i+1$. In this case then a representation that directly accounts for finite histories can be achieved.

Furthermore, as in Horan (2011), the future could also be finite. For example, if the DM is searching for a job, there may only be a finite number of employers available to receive offers from. Here, define a set $Y_{i,j}$ which indicates a list with i items in the past ($t \leq 0$) and j elements in the future. I can similarly define the shift operator $\sigma_{i,j}$ and $\succeq_{i,j}$, and imposing the appropriate technical conditions on the domain to guarantee existence of a utility representation, along with the appropriately modified axioms, again characterize the searcher's recursive preferences.

2.7 Appendix B: Preferences Over Histories — Extensions

This section discusses preferences over histories in more detail, and gives additional examples of sub-classes of GRSs. There may be situations where the set of possible past offers are observable (for example, newspaper listings of job offers) to both the DM and the researcher. If the DM has additional information relative to the researcher about exogenous constraints on choice, for example, which options are still available, then the history is not just the set of outcomes currently available, but represents the set of outcomes that could *possibly* be available. For example, if the DM faces some exogenous constraint on recall that is ex-ante unknown to the researcher, such as that offers expire after one period, then by working with the full history ensures that the researcher does not rule out these effects by fiat. A common assumption in papers is that the DM has no recall, which implies μ will depend only on the offer in the current period. More generally, μ will jointly characterize both the extent of recall of the agent (which is typically an exogenous environmental constraint) as well as the type of context-dependence preferences of the DM. [22]

In these situations the history is then not interpretable as the set of items currently available, but rather as the union of all items that were available at some time in the past and present.[23]

[22]However, as previously discussed, if the extent of recall is known, the domain can be appropriately modified to include only a finite number of past observations, and the axioms can be appropriately modified.

[23]Left untouched by this section, although potentially of interest, are models where the value of options varies by how long they have been available. For example, the payoff for any given item could depreciate every period that it

Furthermore, the interpretation of choice behavior is closely tied to whether past options are actually available to the DM, or whether they are simply options that were, but are no longer, available. For example, if all the options in the history are currently available to the DM, then it is be difficult to interpret behavior as influenced by regret at having passed over options.

If the DM cares about options that are not currently available, additional difficulties in interpreting the representation arises, because of the issue of infinite regress. As for all representations, the characterizations in this paper are 'as if': if a consumer obeys the proper axioms, the corresponding representation holds. Despite this, if, given the representation, the DM appears to care about options that are no longer available, it stands to reason that she should also care about options which are not available when she chooses between sequences. Then it should be the case that choices between sequences reflect regret not only about options passed up in the current sequence the DM is searching within, but options from other sequences that could have been chosen. In contrast, context-dependence with full recall is easier to interpret because the DM's utility when she stops depends only on items available to choose.

Preferences over histories are very similar to the menu-choice literature (such as Gul and Pesendorfer, 2001). If the DM has full recall, and cares only about what she has observed, but not the order that she has observed it, then $\succeq_0$ is equivalent to a preference over unordered choice sets.

For the rest of the section assume that the set of outcomes is the set of lotteries over a compact connected metric space Ξ, and so $\Omega = \Delta(\Xi)$.[24] This imposes an additional element of timing with regards to the resolution of uncertainty. Unlike the lotteries over outcomes considered in the previous section, which were only in the future, and resolved sequentially as the DM searched, the lotteries in this section are objects that the DM chooses among when she stops. The uncertainty regarding elements of Ω do not resolve until after the DM stops and chooses among the available options. Formally, their uncertainty resolves at Epoch 3. This assumption is made purely for comparative purposes. Furthermore, in order to simplify the exposition, axioms will be written on preferences over histories $\succeq_0$, as opposed to preferences over sequences $\succeq$. However, the axioms can be rewritten as axioms solely on the observable preference.

First, I will rewrite **OI**.

Order Independence' (OI') For all $y, y' \in X^0$, if $\varsigma(y) = \varsigma(y')$ then $y \sim_0 y'$.

Recall condition **RR**, which says that adding a worse outcome to any history does not change

is available.

[24] I do this to parallel the structure of much of the menu choice literature.

the value of a history. In contrast, **Decision Temptation** says that this is possible, and in a way that mirrors the Set Betweenness condition of Gul and Pesendorfer (2001).

Decision Temptation (DT) If $y \succeq_0 y'$ then $y \succeq_0 y \uplus y' \succeq_0 y$

I also assume, like Gul and Pesendorfer (2001), **Independence**. First, it is important to carefully define what mixing histories means, in an analogous way to mixing menus.

Definition *For any two histories* $y, y' \in X^0$ *and* $\alpha \in [0,1]$ *define* $\alpha y + (1-\alpha)y' = z \in X^0$*, where*

- $z_0 = \alpha y_0 + (1-\alpha)y'_0$
- $z_1 = \alpha y_1 + (1-\alpha)y'_0$
- *and denoting* $z_{i-1} = \alpha y_m + (1-\alpha)y'_n$, $z_i = \alpha y_j + (1-\alpha)y'_k$, *then* $z_{i+1} = \alpha y_q + (1-\alpha)y'_r$ *where*
  - *if* $n = k$ *then* $q = j - 1$ *and* $r = k + 1$
  - *if* $m = j$ *then* $q = j + 1$ *and* $r = k - 1$
  - *if* $n - 1 = k$ *and* $m + 1 = j$ *and* $k > 0$ *then* $q = j + 1$ *and* $r = k - 1$
  - *if* $n - 1 = k$ *and* $m + 1 = j$ *and* $k = 0$ *then* $q = j + 1$ *and* $r = k$
  - *if* $n + 1 = k$ *and* $m - 1 = j$ *and* $j > 0$ *then* $q = j - 1$ *and* $r = k + 1$
  - *if* $n - 1 = k$ *and* $m - 1 = j$ *and* $j = 0$ *then* $q = j$ *and* $r = k + 1$

Then **Independence** can be stated analogously as in Gul and Pesendorfer (2001).

Independence (I) For any histories $y, y', z \in X^0$ if $y \succ_0 y'$ and $\alpha \in (0,1)$ then $\alpha y + (1-\alpha)z \succ_0 \alpha y' + (1-\alpha)z$.

As mentioned, it can possibly be useful to identify exogenous restrictions on choice from behavior. For example, the researcher may observe the set of past offers, but not which ones are currently available. If the DM's ability to recall an offer is directly related to how long ago it was offered, it is easy to identify this from choice behavior. For example, if the DM has no recall, she cares only about the chosen item from her history when she stops, and she can only ever choose the item most recently observed. Therefore, any outcomes observed before $t = 0$ should not change her preferences.

No Recall (NR) For all x^{-1}, y^{-1} and x_0, it is the case that $(x^{-1}, x_0) \sim_0 (y^{-1}, x_0)$

With these axioms in hand, it is possible to characterize two GSRs that specify how payoffs over histories depend on the outcomes available in the history.

Theorem 6 *Assume that the domain is $\Delta(\Omega)$. Furthermore, assume that $\succeq$ satisfies axioms **CWO**, **LC**, **WM**, and **SB** and every history is strictly f-relevant. Then*

- *$\succeq_0$ satisfies **NR** if and only if there exists a function u such that $\mu(x^0) = u(x_0)$*
- *$\succeq_0$ satisfies **OI'**, **I** and **DT** if and only if there exist linear functions u and v from Ω to $\mathbb{R}$ such that $\mu(x^0) = \max_{x \in x^0} u(x) + v(x) - \max_{y \in x^0} v(y)$*

2.8 Appendix C: Choice during Search

This paper takes a very particular approach to deriving the representation — it considers preferences over DSPs. Other search papers that have attempted to characterize preferences in search environments have worked with different data, as well as with different types of identification in mind. Horan (2011) models choices from observable lists — finite sequences of the form $(x_0, x_1 ... x_n)$ and uses choice data. Caplin and Dean (2010) work with both choice and contingent choice data from a potentially infinite sequence.

This raises the question of how to compare the disparate approaches. For example, Horan identifies both preferences over outcomes and a stopping rule when the DM is boundedly rational — she uses a stopping rule that depends only on the history, while Caplin and Dean identify preferences over outcomes. One way to relate the present paper to the previous literature is via the menu-choice literature. In that literature, the actual choices from choice sets are not the object of study. Instead, preferences over choice sets are observed. Given these preferences, a representation is derived. This representation can be interpreted as describing which element will be chosen. Similarly, in this paper, observed preferences over DSPs characterize a representation that can be interpreted as providing information about when DM will stop searching, and which element will be chosen.

Despite the different approaches, this paper shares an interest with the previous literature in a particular sub-class of search policies — policies where the DM uses a reservation policy. However, models characterized in previous papers in the literature and in this paper are not nested. First, this paper can accommodate policies that the previous characterization results cannot. For example, the DMs in this paper have stopping rules that are both forward and backward looking, and the stopping payoff can vary with the history. This is because working only with choice and stopping data, it is impossible to fully identify parameters in the Bellman equation, but identification can be achieved by using preferences over DSPs. The basic intuition for this is that if DM's payoff today can be affected by things they have seen in the past, then it is impossible to compare the payoff of every possible history with the payoff of every possible future. Therefore any data that summarizes when a DM stops or continues searching will not map into uniquely identified preferences. Second, the approach in this paper will not be able to necessarily identify preferences over final outcomes given an arbitrary boundedly rational stopping rules, some of which Horan can accommodate.

First, this section will more fully explore why DSPs are necessary for identification of the preference parameters. As mentioned, the choice data used by Caplin and Dean and Horan are too restrictive if the goal is to identify preferences over histories and continuation sequences. Because of the richer domain, this paper can identify not only rankings over final outcomes, but also payoff-relevant concerns that may not directly change the rankings of final outcomes. In addition, because the domain includes the future, both the recursive nature of preferences and the discount rate can also be identified, and interpersonal comparisons can be me made. Furthermore, so long as the DM's rankings over histories are derived from a preference ordering over outcomes, observed preferences over DSPs can identify the choice data that are the input into Caplin and Dean, as well as the final choice of outcomes that both Horan and Caplin and Dean use as data.

Within Horan and Caplin and Dean's framework, it is not clear how to accommodate sequences that feature resolution of uncertainty, since the DM's expectations are entirely unmodeled. In order to make the comparison as transparent as possible, I will restrict attention to deterministic outcomes.

Using only choice data (rather than additional choice process data) drastically limits the amount of information that can be gained regarding a DM's preferences. Even restricting consideration to sequences of deterministic outcomes, any particular outcome a can occur at different points in the sequence. Even if the list order is known, observing only choice data makes it difficult to ascertain at which point the DM stopped. This creates problems for identifying the stopping rule (i.e. the threshold) of the DM. Horan avoids this issue by assuming that each outcome appears only one time in the list, while Caplin and Dean assume that the data includes a subset of the outcomes considered by the DM during the search process.

Furthermore, even additional choice process data about the DM's behavior while searching does not allow for full identification of preference. Instead of observing choice, suppose that the data is at what point in a sequence the DM stops searching. This data would seem to be much more amenable to ranking the values of histories and continuation sequences, but is still not as powerful as preferences over sequences, and so cannot uniquely identify preferences . In order to clearly demonstrate this, I will consider a simple example where the set of outcomes consists of two alternatives a and b, and so X consists of sequences of these two outcomes. To simplify the example even further, assume that the DM cares about only the two most recent outcomes. Then the domain consists of sequences of the form $(x_{-1}, x_0, x_1, ...)$, while the observable data is whether given a sequence, the DM decides to stop or continue searching. Consider two histories (i, a) and (j, b), where $i, j \in \{a, b\}$. It is impossible to directly compare the DM's value of these two histories directly. Instead, given a future f, the DM compares history (i, a) to continuation sequence (a, f) and history (j, b) to sequence (b, f). However, in order for these rankings to provide information about the relative ranking of (i, a) and (j, b), it is must be the case that there is information about how the DM ranks (a, f) and (b, f). However, it is impossible for the DM to directly compare (a, f) to (b, f). She can only compare (a, f) to the set of histories of the form (i, a), and compare (b, f) to the set of histories of the form (j, b). Then the problem of ranking (i, q) and (b, f) simply reoccurs.

However, if the domain is restricted, choice data becomes tractable. For example, assume that the DM has no recall, and is unaffected by payoffs observed other than the one currently available. Furthermore assume that the observable data is when the DM stops. Given the assumption of no recall, this is equivalent to observing which outcomes the DM chooses.[25] In this case, preferences can be uniquely identified. Intuitively the difference that the restriction makes is that it allows the researcher to know that the DM is comparing every possible history (which are of length 1) to every possible future. In contrast, this cannot occur in the more general model. This limits the ability to uniquely identify the preferences.

However, this paper cannot accommodate all of the patterns of choice that can occur in previous papers in the literature. In particular, the approach in this paper will not be able to necessarily identify preferences over final outcomes given an arbitrary bounded rational stopping rules. To see this, observe that the thresholds that Caplin and Dean and Horan use are not forward looking. Therefore, they implicitly assume that given any particular history, regardless of the future, the continuation value is the same. The way to incorporate this into the model is to have the continuation value $\nu(V(\sigma))$ be constant for all continuation sequences. This requires relaxing **WM**. Given this relaxation, the function μ can accommodate any history dependent stopping rule. However, it is not necessarily the case that the μ function derived from this stopping rule would be derived from an

[25]The assumption of no recall can clearly be falsified given choice data. However, neither choice nor stopping data can falsify the assumption of no context dependence.

interpretable ranking over final outcomes that is time independent. For example, assume there are finite outcomes X, two of which are a and b, and the DM prefers a over b ($u(a) > u(b)$). Furthermore, assume the DM possesses a boundedly rational decision rule regarding stopping such that if the DM observes all b in the past and a in the current period she stops, and then will choose a (since she ranks it higher). However, assume that if she sees all a in the past and then b in the current period, she does not stop. This requires that $\mu(.....b, a) > \mu(.....a, b)$, which cannot be interpreted as incorporating a time independent utility ranking over outcomes (although this could be rationalized by $\mu(....x_{-1}, x_0) = u(x_0) - \beta u(x_{-1})$ with $\beta < 1$). In these situations, the DSP approach needs to be supplemented by actual choice data in order to fully identify preferences over final outcomes.

2.9 Appendix D: Preliminary Results

This Appendix collects some preliminary results that are helpful in proving the main results in the text, and also may be of independent interest.

First, σ is a continuous bijection.

Lemma A1 *σ is a continuous bijection*

I also define several additional conditions on preferences. **Limited Complementarity** ensures that the ranking over histories, derived from observing preference reversals between today and tomorrow, is asymmetric.

> **Limited Complementarity(LC)** For all $y, y' \in X^0$ if for some $w, z \in {}^1X$, $\sigma(y, w) \preceq \sigma(y', z)$ and $(y, w) \succ (y', z)$ then
>
> - for all $z', w' \in {}^1X$, if $\sigma(y, w') \succ \sigma(y', z')$ then $(y, w') \succ (y', z')$
> - for all $z', w' \in {}^1X\, if \sigma(y, w') \sim \sigma(y', z')$ then $(y, w') \succeq (y', z')$

LC is in fact strictly stronger than **LS**.

Lemma A2 *If $\succeq$ satisfies **CWO** and **LC** then $\succeq$ satisfies **LS**.*

Once the DM's value of today's DSP begins to rise in the value of the continuation DSP, it must continue to do so. This is because once a DM begins to search, then for any better continuation DSPs she will continue to search.

> **Future Sensitivity (FS)** For all $y \in X^0$, $z, z', w \in {}^1X$, if $\sigma(y, z) \prec \sigma(y, z')$, and $(y, z) \prec (y, z')$, and $\sigma(y, w) \succ \sigma(y, z')$, then $(y, w) \succ (y, z')$

Because the payoff from stopping does not depend on outcomes not yet observed, fixing a history, a DM is indifferent between sequences if she chooses to stop at both (since the value of the continuation DSP doesn't change the value of today's DSP). Furthermore, if the DM prefers to stop given a history and some continuation sequence, then for the same history, and any worse continuation sequence, she must also prefer to stop.

Future Indifference (FI) For all $y \in X^0$, $z, z', w \in {}^1X$, if $\sigma(y,z) \prec \sigma(y,z')$ and $(y,z) \sim (y,z')$, and $\sigma(y,w) \prec \sigma(y,z')$, then $(y,w) \sim (y,z')$.

FS is a direct implication of the second part of **WM**, while **FI** is implied by the combination of **CWO**, **LS** and **WM**.

Lemma A3 *Assume $\succeq$ satisfies **CWO**.*

- *If $\succeq$ also satisfies **WM** then it satisfies **FS**.*
- *If $\succeq$ also satisfies **LC** then **FS** and **FI** are equivalent.*

In the presence of **WM**, **LS** and **LC** are actually equivalent.

Lemma A4 *Assume $\succeq$ satisfies **CWO**. If $\succeq$ satisfies **LS** and **WM** then it satisfies **LC**.*

Turning to the construction of $\succeq_0$, if for $y, y' \in X^0$, there exists $z, z' \in {}^1X$, $\sigma(y,z) \preceq \sigma(y',z')$ and $(y,z) \succ (y',z')$, then the $\succeq$ minimal element with history y must be valued higher than the $\succeq$ minimal element with history y'.

Lemma A5 *If for $y, y' \in X^0$, there exists $z, z' \in {}^1X$, $\sigma(y,z) \preceq \sigma(y',z')$ and $(y,z) \succ (y',z')$, then there exists a $w' \in {}^1X$ such that for all $w \in {}^1X$ $(y,w) \succ (y',w')$.*

The function that maps elements X to corresponding elements of X^0 is continuous, given the product topology of both spaces.

Lemma A6 *The function $\kappa : X \to X^0$, such that for $y \in X^0$ and $z \in {}^1X$, $\kappa(y,z) = y$ is continuous.*

2.10 Appendix E: Proofs

Throughout the proofs if a sequence is denoted (y,z), it is assumed that $y \in X^0$ and $z \in {}^1X$. Similarly for a sequence denoted $\sigma(y,z)$ it is the case that $y \in X^0$ and $z \in {}^1X$. Elements such as x, y and z represent generic sequences and subsequences.

2.10.1 Proof of Preliminary Results

Proof of Lemma A1 Assume there is a sequence of sequences $(x, y)^n$ that converges to (x, y). Then for each t, x_t^n or y_t^n converges to x_t and y_t respectively. Therefore $\sigma(x, y)^n$ converges to $\sigma(x, y)$.

Assume σ is not one-to-one. Consider some (x'', w) such that $\sigma(x, y) = \sigma(x', y') = (x'', w)$ but $(x, y) \neq (x', y')$. In this case there exists a t such that $(x, y)_t \neq (x', y')_t$. But then $\sigma(x, y)_{t-1} \neq \sigma(x', y')_{t-1}$. This is a contradiction.

To show that σ is onto, consider some (x'', w). Then we can simply construct $\sigma^{-1}(x'', w)$, such that $\sigma^{-1}(x'', w)_{t+1} = (x'', w)_t$.

□

Proof of Lemma A2 Assume **LS** does not hold. Then there exists $y \in X^0$ and $z, z' \in^1 X$, such that $\sigma(y, z) \succeq \sigma(y, z')$ and $(y, z') \succ (y, z)$. In this case the 'if' part of **LC** is satisfied.

Then it is the case that that for all w, w' if $\sigma(y, w) \succ \sigma(y, w')$ then $(y, w) \succ (y, w')$. Furthermore, if $\sigma(y, w) \sim \sigma(y, w')$ then $(y, w) \succeq (y, w')$. Let $z = w$ and $z' = w'$. Then $(y, w) \succeq (y, w')$. This is a contradiction as it violates **WCO**. □

Proof of Lemma A3 Fix a particular history $y \in X^0$. We will prove the claims by contradiction.

- First, assume that **FS** does not hold but **WM**, so there exists $z, z', w \in {}^1X$, such that $\sigma(y, z) \prec \sigma(y, z')$, and $(y, z) \prec (y, z')$, and $\sigma(y, w) \succ \sigma(y, z')$. However, this satisfies the conditions of **WM**. Therefore $(y, w) \succ (y, z')$. So **FS** holds.

- Next, assume that **FI** does not hold but **WM** and **FS** do, so there exists $y, y', w \in {}^1X$, such that $\sigma(y, z) \prec \sigma(y, z')$ and $(y, z) \sim (y, z')$, and $\sigma(y, w) \prec \sigma(y, z')$. By **LS** $(y, w) \preceq (y, z')$. If $(y, w) \prec (y, z')$ then there are three cases. First, $\sigma(y, z) \prec \sigma(y, w)$. But this violates **LS**. Second, $\sigma(y, z) \sim \sigma(y, w)$. This violates **LC**. Third, $\sigma(y, z) \succ \sigma(y, w)$. In this case, by **FS** $(y, z) \nsim (y, z')$. Therefore $(y, w) \sim (y, z')$, so **FI** holds.

- Now assume that **FS** does not hold but **WM** and **FI** do. Then there exist $y, y', w \in {}^1X$ such that $\sigma(y, z) \prec \sigma(y, z')$, and $(y, z) \prec (y, z')$, and $\sigma(y, w) \succ \sigma(y, z')$. In this case, by **LS**, $(y, w) \succeq (y, z')$. If $(y, w) \sim (y, z')$ then by **FI**, $(y, w) \sim (y, z)$, which is a contradiction. Therefore $(y, w) \succ (y, z')$ and so **FS** holds. □

Proof of Lemma A4 Assume **LC** does not hold and so there exists $(y, z) \preceq (y', z') \in X$ such that $\sigma(y, z) \preceq \sigma(y', z')$ and $(y, z) \succ (y', z')$. This implies that there is no z'' such that $\sigma(y, z'') \prec \sigma(y, z)$

and $(y, z'') \prec (y, z)$. If there was, then by **WM** $(y, z) \preceq (y', z')$. Furthermore, by **LS** for all $\sigma(y, z'') \succeq \sigma(y, z)$ and $(y, z'') \succeq (y, z)$. Therefore, for all z'' $(y, z'') \succeq (y, z)$. There are three cases to consider.

- Assume $\sigma(y, x) \succ \sigma(y', x')$. Note that for all x, $(y, x) \succeq (y, z)$. If $(y', x') \prec (y', z')$ then $(y, x) \succ (y', x')$. If $(y', x') \sim (y', z')$, then $(y, x) \succeq (y, z) \succ (y', x')$. If $(y', x') \succ (y', z')$ then by **WM** since $\sigma(y, x) \succ \sigma(y', x')$, $(y, x) \succ (y', x')$. This is a contradiction.

- Assume $\sigma(y, w) \sim \sigma(y', w')$. Note that for all w, $(y, w) \succeq (y, z)$. If $(y', w') \prec (y', z')$ then $(y, w) \succ (y', w')$. If $(y', w') \sim (y', z')$, then $(y, w) \succeq (y, z) \succ (y', w')$. If $(y', x') \succ (y', z')$ then by **WM** since $\sigma(y, w) \sim \sigma(y', w')$ then $(y, w) \succeq (y', w')$. This is a contradiction.

- Assume $\sigma(y', z') \succeq \sigma(y', x')$. Then by **LS** $(y', z') \succeq (y', x')$. Note that for all x, $(y, x) \succeq (y, z)$. If $(y', x') \prec (y', z')$ then $(y, x) \succ (y', x')$. If $(y', x') \sim (y', z')$, then $(y, x) \succeq (y, z) \succ (y', x')$. This is a contradiction.

$\square$

Proof of Lemma A5 If for $y, y' \in X^0$, there exists $z, z' \in {}^1X$, $\sigma(y, z) \preceq \sigma(y', z')$ and $(y, z) \succ (y', z')$, then by **LC** it is the case that for all $w \in {}^1X$, $(y, w) \succ (y', z')$. $\square$

Proof of Lemma A6 Note that if there is a a sequence of DPSs $v^i{}_{i=1}^{\infty}$ which converges to x then it is the case (because of the product topology) that for each j $j = 0, 1, -1, 2, -2, ...$ that the sequence v^i_j converges to x_j. Therefore, again because of the product topology, it is the case that since for $j = 0, -1, -2, ...$ that v^i_j converges to x_j then $v^{0,i}$ must converge to x^0. $\square$

2.10.2 Proof of Main Text Results

Proof of Lemma 1 This is an immediately implication of Lemma A5. $\square$

Proof of Lemma 2 I shall prove each property in turn.

- Completeness: By **WCO** for all $y \in X^0$ there exists a $z \in {}^1X$ such that for all $z' \in {}^1X$ $(y, z) \preceq (y, z')$. Denote this future $z(y)$. For any $y, y' \in X^0$ since CWO is satisfied it is either the case that $(y, z(y)) \succeq (y', z(y'))$ or $(y, z(y)) \preceq (y', z(y'))$.

- Transitivity: Assume that $y' \succeq_0 y$ and $y \succeq_0 y''$. Then there exists a $z' \in {}^1X$ such that for all $z \in {}^1X$ $(y', z) \succeq (y, z')$. Furthermore there exists a $z'' \in {}^1X$ such that for all $z''' \in {}^1X$ $(y, z''') \succeq (y'', z'')$. Therefore $(y, z') \succeq (y'', z'')$. Therefore for all $z \in {}^1X$ $(y', z) \succeq (y'', z'')$ and so $y' \succeq_0 y''$.

- Continuity: I shall show that the upper contour set of $\succeq_0$ is closed. The proof for the lower contour set is analogous. I will prove it by contradiction. Assume that there exists a sequence of histories $\{y^{'n}\}$ that converges to y' such that for each n, $y'n \succeq_0 y$ but $y \succ_0 y'$.

  Then there exists a $z \in {}^1X$ such that for all $z' \in {}^1X$ $(y', z) \prec (y, z')$. Because κ is continuous and $y'n \to y'$, there exists a sequence (y'_n, z_n) that converges to (y', z). By the continuity of $\succeq$ there exists an N' such that for all $N > N'$ $(y'_N, z_N) \prec (y, z')$. Therefore $y'_N \prec_0 y$, a contradiction.

□

Proof of Theorem 1 The proof of the necessity of **CWO** is standard. To show the necessity of the other axioms I will consider each in turn.

- **LS:** Assume that $\sigma(y, z) \succeq \sigma(y, w)$. Then $\nu(V(\sigma(y, z))) \geq \nu(V(\sigma(y, w)))$. If $\nu(V(\sigma(y, w))) \geq \nu(y)$, then $\nu(V(\sigma(y, z))) \geq \nu(y)$ and so $V(y, z) \geq V(y, w)$, which means $(y, z) \succeq (y, w)$. Otherwise $\nu(V(\sigma(y, w))) \leq \nu(y) \leq \max\{\nu(y), \nu(V(\sigma(y, z)))\}$, and so $V(y, z) \geq V(y, w)$, which means $(y, z) \succeq (y, w)$.

- **WM:** There are two parts in the set-up to the axioms. I consider each in turn. Furthermore, for each of those conditions there are the two sub-cases.

  - Assume the first if in WM holds. Then $V(y, w) \geq V(y, w')$. Therefore $V(y, w) = \nu V(\sigma(y, w)$. Then assume that $\nu V(\sigma(y', z') = \nu V(\sigma(y, w)$. There are two possibilities. First is that $\mu(y') \leq \mu(y)$. In this case $\nu V(\sigma(y, w) \geq \mu(y) \geq \mu(y')$, and so $V(y', w') = \nu V(\sigma(y, w)$. In the other case $\mu(y') \geq \mu(y)$ and so by pairwise comparison $\max\{(\mu(y'), \nu(V(\sigma(y', w')))\} \geq \max\{(\mu(y), \nu(V(\sigma(y, w)))\}$.

    Furthermore, if I assume that $\nu V(\sigma(y', w') > \nu V(\sigma(y, w)$ then the reasoning in the previous paragraph still holds. This proves both subcases hold.

  - First it is necessary to show that $V(y, w) = \nu(V(\sigma(y, w))$. Assume not. Then $V(y, w) = \mu(y, w)$. There exists a (y, w') such that $V(y, w') > V(y, w)$, and so there exists a $\sigma(y, w') \succ \sigma(y, w)$. Therefore, by Claim 4 (see below), there exists a sequence of outcomes

$\sigma(y, z^n)$ that converge from above to $\sigma(y, w)$ (because clearly there exiss a sequence that converges from above by the continuity of the preferences, and from Claim 4 there exists a sequence wholly within history y). But since $V(y, w) = \mu(y, w) < \nu V(\sigma(y, w)$ and V is continuous, then there exists some N such that for all $N' > N$ $V(y, z^{N'} < \nu V(\sigma(y, w))$, and so $V(y, z^{N'} = \mu(y)$. But this violates the assumption that everything better tomorrow is better today. Therefore $V(y, w) = \nu(V(\sigma(y, w))$. The rest of the proof proceeds as for the first case.

- **SB:** Assume $V(y, z) > V(y, z')$, then $V(y, z) = \nu(V(\sigma(y, z)))$. By construction $\nu(V(\sigma(y, z))) < V(\sigma(y, z))$

The rest of the proof demonstrates sufficiency of the axioms. Furthermore, note that if the cardinality of Ω is 1, the result is trivial. So I will assume not.

First, there exists a representation of $\succeq$ and $\succeq_0$.

Claim 1 There exists a function V that represents $\succeq$.

Proof Since $\succeq$ satisfies **WCO** and X is a separable metric space, by Debreu (1954) there exists a function V that represents $\succeq$. $\square$

Claim 2 There exists a function $\hat{\mu}$ that represents $\succeq_0$.

Proof Since by Lemma 2, $\succeq_0$ is a complete transitive continuous binary relation, and by endowing X^0 with the product topology, it is a separable metric space, and so by Debreu (1954) there exists a function $\hat{\mu}$ that represents $\succeq_0$. $\square$

Next I present two results that will use repeatedly. First, the histories of two sequences agree, and the DM is indifferent between their respective continuation sequences, then the DM must be indifferent between sequences themselves. Second, given a history, the set of continuation sequences is closed in the preference ranking, and the set of utility values the set maps to is connected.

Claim 3 For all $y \in X^0$, and $z, z' \in {}^1X$ if $\sigma(y, z) \sim \sigma(y, z')$ then $(y, z) \sim (y, z')$.

Proof Assume not. Then $\sigma(y, z) \sim \sigma(y, z')$ and without loss of generality assume $(y, z) \succ (y, z')$. Then by **LC** $(y, z) \preceq (y, z')$. This is a contradiction. $\square$

Claim 4 For all $y \in X^0$ there exists a $\succeq$ maximal and $\succeq$ minimal element of the set $\{\sigma(y,z)|z \in {}^1X\}$ (denoted $\sigma(y,\bar{z})$ and $\sigma(y,\underline{z})$ respectively). Furthermore, if there exists a $(y',z') \in X$ such that $\sigma(y,\bar{z}) \succ \sigma(y',z') \succ \sigma(y,\underline{z})$, then there exists a z such that $\sigma(y',z') \sim \sigma(y,z)$. Similarly, there exists a $\succeq$ maximal and $\succeq$ minimal element of the set $\{(y,z)|z \in {}^1X\}$ (denoted $(y,\bar{z})$ and $(y,\underline{z})$ respectively), and if there exists a $(y',z') \in X$ such that $(y,\bar{z}) \succ (y',z') \succ (y,\underline{z})$, then there exists a z such that $(y',z') \sim (y,z)$.

Proof I shall prove the claims regarding the continuation DSPs. The claim for the current DSPs is analogous. First, I will show that the set $\{\sigma(y,z)|z \in {}^1X\}$ is closed. Assume that the sequence $\sigma(y,z^m)$, for $m = 1....$ converges. In this case the limit must be a sequence $\sigma(y,z) \in X$. To see this assume not — assume it converges to a sequence some other history y': (y',z). In this case, there exists an $\epsilon > 0$ such that the distance between y' and and y is greater than ϵ and so the distance between $\sigma(y,z^m)$ and (y,z) is greater than ϵ for all m.

Therefore $\{\sigma(y,z)|z \in {}^1X\}$ is compact (since any closed subset of a compact space is compact). Next, the image of X under V is compact (since V is continuous), and so it is closed and bounded (by the Heine-Borel Theorem), and the same is true of the image of $\{\sigma(y,z)|z \in {}^1X\}$ under V. Therefore, there exists V maximal and minimal elements of $\{\sigma(y,z)|z \in {}^1X\}$, which are $\succeq$ maximal and minimal elements.

Now assume there exists a $\sigma(y,\bar{z}) \succ \sigma(y',z') \succ \sigma(y,\underline{z})$ but there is no z such that $\sigma(y',z') \sim \sigma(y,z)$. This means it is possible to partition the set $\{\sigma(y,z)|z \in {}^1X\}$ into two pieces — a piece that whose elements are all strictly better than $\sigma(y',z')$ and a piece whose elements are all strictly worse than $\sigma(y',z')$. Denote these two pieces $\bar{Y}$ and $\underline{Y}$ respectively. Therefore for all $\sigma(y,\bar{z}) \in \bar{Y}$ and $\sigma(y,\underline{z}) \in \underline{Y}$, $V(\sigma(y,\bar{z})) > V(\sigma(y',z')) > V(\sigma(y,\underline{z}))$.

$\{\sigma(y,z)|z \in {}^1X\}$ is compact subset of a separable metric space, and since the product of connected spaces is connected, is also connected. Furthermore, $\succeq$ when restricted to $\{\sigma(y,z)|z \in {}^1X\}$ is be continuous. Therefore there exists a continuous utility function ξ that represents $\succeq$ on $\{\sigma(y,z)|z \in {}^1X\}$ that must take on all values between the utility of the $\succeq$ maximal and $\succeq$ minimal element (since ξ is continuous, its image is connected, and so path connected). Therefore for all z, z' there exists a z'' such that $\xi(\sigma(y,z)) > \xi(\sigma(y,z'')) > \xi(\sigma(y,z'))$. Since V also represents $\succeq$ this is a contradiction. $\square$

Next, I will restrict attention to a particular fixed history $x \in X^0$. Note that by the representation theorem there exists a continuous representation of $\succeq$ when it is restricted to the set $\{(x,z)|z \in {}^1X\}$. For the rest of the proof I will only consider, conditional on any particular history $x \in X^0$, the set of outcomes up to unique equivalence classes conditional on a history.

Claim 5 For all $x \in X^0$ there exists a $y \in {}^1X$ that satisfies the following two conditions:

- **Condition 1**: for all $z, z' \in {}^1X$ if $\sigma(x,z') \succeq \sigma(x,y)$ and $\sigma(x,z) \succeq \sigma(x,y)$ then $\sigma(x,z') \succeq \sigma(x,z)$ if and only if $(x,z') \succeq (x,z)$

- **Condition 2**: for all $w, w' \in {}^1X$ if $\sigma(x,w') \preceq \sigma(x,y)$ and $\sigma(x,w) \preceq \sigma(x,y)$ then $(x,w') \sim (x,w)$

Furthermore, if there are two such futures $y, y' \in {}^1X$ that satisfy Conditions 1 and 2, it must be the case that $(x,y) \sim (x,y')$.

Proof First, the claim holds if for all y, y', $(x,y) \sim (x,y')$. So assume that this is not the case. I will show such a $y \in {}^1X$ exists for all $x \in X^0$ by contradiction. To do so, assume that there exists some x where no such y exists.

First I demonstrate that for all $y \in {}^1X$, that if (x,y) satisfies Condition 2 (Condition 1) in the claim then for all $y' \in {}^1X$ such that $\sigma(x,y') \preceq \sigma(x,y)$ $(\sigma(x,y') \succeq \sigma(x,y))$ (x,y') also satisfy Condition 2 (Condition 1) in the claim.

- Assume that (x,y) satisfies Condition 2. I will show that for all $y' \in^1 X$ if $\sigma(x,y') \preceq \sigma(x,y)$ then (x,y') must also satisfy Condition 2. By way of contradiction assume that there exists a y' where the this is not true. Therefore exists $w, w' \in^X$ such that $\sigma(x,w') \preceq \sigma(x,y')$ and $\sigma(x,w) \preceq \sigma(x,y')$, but $(x,w') \nsim (x,w)$. Clearly (x,y) fails the condition as well, since $\sigma(x,w') \preceq \sigma(x,y)$ and $\sigma(x,w) \preceq \sigma(x,y)$.

- Assume that (x,y) satisfies Condition 1 in the claim. I will show that all y' if $\sigma(x,y') \succeq \sigma(x,y)$ then (x,y') must also satisfy Condition 1 in the claim. By way of contradiction assume that there exists a $y' \in^1 X$ where this is not true. Therefore there exists $w, w' \in^1 X$ such that $\sigma(x,w') \succeq \sigma(x,y')$ and $\sigma(x,w) \succeq \sigma(x,y')$, but the ranking of (x,w') and (x,w) disagrees with the ranking of $\sigma(x,w')$ and $\sigma(x,w)$. Clearly (x,y) then fails Condition 1 since $\sigma(x,w') \succeq \sigma(x,y)$ and $\sigma(x,w) \succeq \sigma(x,y)$.

Next I will show that for all $y \in^1 X$, (x,y) must satisfy at least one of the two conditions in the claim.

- First, assume (x,y) fails Condition 2. Then there exists $w, w' \in^1 X$ such that $\sigma(x,w') \preceq \sigma(x,y)$ and $\sigma(x,w) \preceq \sigma(x,y)$ but $(x,w') \nsim (x,w)$. Without loss of generality assume that

$\sigma(x, w') \preceq \sigma(x, w)$. Then by **LS** $(x, w') \preceq (x, w)$, and so $(x, w') \prec (x, w)$. Therefore, by **FS**, for all w'', w''', if $\sigma(x, w'') \succeq \sigma(x, w)$ and $\sigma(x, w''') \succeq \sigma(x, w)$ if $\sigma(x, w'') \succ \sigma(x, w''')$ then $(x, w'') \succ (x, w''')$, and if $\sigma(x, w'') \sim \sigma(x, w''')$ then $(x, w'') \sim (x, w''')$. Therefore, for all $\sigma(x, w'') \succeq \sigma(x, y)$ and $\sigma(x, w''') \succeq \sigma(x, y)$ if $\sigma(x, w'') \succ \sigma(x, w''')$ then $(x, w'') \succ (x, w''')$, and if $\sigma(x, w'') \sim \sigma(x, w''')$ then $(x, w'') \sim (x, w''')$, and so Condition 1 is satisfied.

- Next, assume (x, y) fails Condition 1. Therefore there exists w, w' such that $\sigma(x, w') \succeq \sigma(x, y)$ and $\sigma(x, w) \succeq \sigma(x, y)$ which cause Condition 1 to fail for (x, y). Without loss of generality assume that $\sigma(x, w') \succeq \sigma(x, w)$. For Condition 1 to fail it must be the case that $\sigma(x, w') \succ \sigma(x, w)$ but $(x, w') \preceq (x, w)$. To see this, assume that $\sigma(x, w') \sim \sigma(x, w)$. If $(x, w') \prec (x, w)$ or $(x, w') \succ (x, w)$ then **LS** is violated. Therefore $(x, w') \sim (x, w)$, and Condition 1 does not fail. Next, assume that $\sigma(x, w') \succ \sigma(x, w)$. If $(x, w') \succ (x, w)$ then Condition 1 holds.

  So assume that $\sigma(x, w') \succ \sigma(x, w)$ but $(x, w') \preceq (x, w)$. If $(x, w') \prec (x, w)$ then **LS** is violated. Therefore, let $(x, w') \preceq (x, w)$. Then by **FI**, for all z, if $\sigma(x, z) \preceq \sigma(x, w')$ then $(x, z) \sim (x, w')$. Therefore, Condition 2 must hold for (x, y).

Next I will show that there exists a $y \in {}^1X$ that satisfies both claims simultaneously. There are three cases.

- For all $y \in {}^1X$, (x, y) satisfies Condition 1. In this case, consider w such that $(x, y) \succeq (x, w)$ for all y. Then it must be the case that $\sigma(x, w) \preceq \sigma(x, y)$ (by condition 1 of the claim). Therefore, (x, w) satisfies Condition 2, since for all w' if $\sigma(x, w) \sim \sigma(x, w')$ then $(x, w) \sim (x, w')$, and there is no w'' such that $\sigma(x, w) \succ \sigma(x, w')$.

- For all $y \in {}^1X$, (x, y) satisfies Condition 2. In this case, consider w such that $\sigma(x, w) \succeq \sigma(x, y)$ for all y. Therefore, (x, w) satisfies Condition 1, since for all w' if $\sigma(x, w) \sim \sigma(x, w')$ then $(x, w) \sim (x, w')$, and there is no w'' such that $\sigma(x, w'') \succ \sigma(x, w)$.

- There is a set $C_1 \subset X$ whose members satisfy Condition 1, and a set $C_2 \subset X$ whose members satisfy Condition 2. It was previously shown that fixing an $x \in X^0$, for all $y \in^1 X$, (x, y) must be a member of at least one element of $\{C_1, C_2\}$, and so $C_1 \cup C_2 = X$. I will show that $C_1 \cap C_2 \neq \emptyset$. If $(x, y \in C_1$ but $(x, y) \notin C_2$ then $\sigma(x, y)$ must be preferred to all elements of $\sigma(C_2)$. Furthermore, if $(x, y) \in C_2$ but $(x, y) \notin C_1$ then all elements of $\sigma(C_1)$ must be preferred to $\sigma(x, y)$.

  First I will show that both $\sigma(C_1)$ and $\sigma(C_2)$ have $\succeq$ maximal and minimal elements. $\sigma(C_1)$ clearly has a $\succeq$ maximal element and $\sigma(C_1)$ clearly has a $\succeq$ minimal element since the set

$\{\sigma(x,y)|y \in^1 X\}$ is closed. I shall prove the rest of the statement for C_1 and the analogous proof applies to C_2.

Assume that C_1 does not have $\succeq$ minimal element. Then for all $(x,y) \in C_1$ there exists an $(x,y') \in C_1$, such that $\sigma(x,y') \prec \sigma(x,y)$. Denote $(x,z) = \inf_{\sigma(x,y)\in\sigma(C_1)} \succeq$. Then $(x,z) \notin C_1$. Since $(x,z) \notin C_1$, there exists $\sigma(x,w')$ and $\sigma(x,w)$ that cause Condition 1 to not hold at (x,z). Since it is the case that if $\sigma(x,w'') \succ \sigma(x,y)$, $\sigma(x,w'') \in C_1$, then at least one of $\sigma(x,w')$ and $\sigma(x,w)$ must be indifferent to $\sigma(x,z)$. If both are indifferent then $\sigma(x,w') \sim \sigma(x,w)$, and so $(x,w') \sim (x,w)$. But then Condition 1 holds, a contradiction. Therefore, assume that without loss of generality $\sigma(x,w') \sim \sigma(x,z)$ and $\sigma(x,w) \succ \sigma(x,w')$. Then $(x,w') \preceq (x,w)$ by **LS**. If $(x,w') \prec (x,w)$, then Condition 1 holds, a contradiction. So assume $(x,w') \prec (x,w)$, Since $(x,w) \in C_1$, there exists an (x,w'') such that $(x,w'') \prec (x,w)$ and $(x,w'') \in C_1$, and so $(x,w'') \prec (x,w) \sim (x,w')$, but $\sigma(x,w'') \succ \sigma(x,w')$. This violates **LS**.

Denote (x,c) as the $\succeq$ minimal element of C_1, and (x,C) as the $\succeq$ maximal element of C_2. If these are the same element, then the proof is done. Assume they are not. In this case, there exists no element (x,z') such that $\sigma(x,c) \succ \sigma(x,z') \succ \sigma(x,C)$. Furthermore, since $\sigma(x,c) \succ \sigma(x,C)$ (by construction), by **LS** it is the case that $(x,c) \succeq (x,C)$. Consider both cases separately.

- $(x,c) \succ (x,C)$. Then $(x,C) \in C_1$ and so there is a contradiction.
- $(x,c) \sim (x,C)$. Then $(x,c) \in C_2$ and so there is a contradiction.

Next, I need to show that any $y \in {}^1X$ that satisfies Conditions 1 ans 2 is unique. Assume it is not, and there is $y, y' \in {}^1X$ that satisfy the condition. Without loss of generality let $\sigma(x,y) \prec \sigma(x,y')$, and so by **LS** $(x,y) \preceq (x,y')$.

- If $(x,y) \prec (x,y')$ then (x,y') cannot be an element of C_2
- If $(x,y) \sim (x,y')$ then (x,y) cannot be an element of C_1

This is a contradiction. $\square$

Definition *For any given $x \in X^0$, denote the $y \in {}^1X$ where*

- *for all $z, z' \in X^1$ if $\sigma(x,z') \succeq \sigma(x,y)$ and $\sigma(x,z) \succeq \sigma(x,y)$ then $\sigma(x,z') \succeq \sigma(x,z)$ if and only if $(x,z') \succeq (x,z)$*
- *for all $w, w' \in X^1$ if $\sigma(x,w') \preceq \sigma(x,y)$ and $\sigma(x,w) \preceq \sigma(x,y)$ then $(x,w') \sim (x,w)$*

as $y^*(x)$.

I shall now show that fixing a history $x \in X^0$, for any continuation DSP worse than $\sigma(x, y^*(x))$, the valuation of the DSP today is constant.

Claim 6 For all $x \in X^0$, for all $w \in {}^1X$ if $\sigma(x, w) \preceq \sigma(x, y^*(x))$ then $(x, w) \sim (x, y^*(x))$.

Proof This is directly implied by the previous claim. $\square$

Next, I will show that $(x, y^*(x))$ is weakly the worst sequence in the set of sequences with history $x \in X^0$.

Claim 7 For all $x \in X^0$ for all $w \in {}^1X$, $(x, w) \succeq (x, y^*(x))$.

Proof Note that if $\sigma(x, w) \preceq \sigma(x, y^*(x))$ then $(x, w) \sim (x, y^*(x))$ by the previous claim. Furthermore if $\sigma(x, w) \succ \sigma(x, y^*(x))$ then $(x, w) \succeq (x, y^*(x))$ by **LS**. $\square$

Claim 8 There exists a function η_x such that for all all $x \in X^0$ and $w \in {}^1X$ such that $\sigma(x, w) \preceq \sigma(x, y^*(x))$, $V(x, w) = \eta_x$.

Proof Denote $\eta(x) = V(x, y^*(x))$ Then by the previous claim, for all $w \in {}^1X$ if $\sigma(x, w) \preceq \sigma(x, y^*(x))$ then $(x, w) \sim (x, y^*(x))$, and so $V(x, w) = \eta_x$. $\square$

Next, I will show that for any given x, the valuation of all DSPs of the form $\sigma(x, y)$ in the lower contour set of $\sigma(x, y^*(x))$ defined previously is exactly the ordering induced by $\succeq_0$.

Claim 9 For all $x \in X^0$, $\eta_x \geq \eta_{x'}$ if and only if $\hat{\mu}(x) \geq \hat{\mu}(x')$.

Proof Assume not. There are two cases:

- First assume $\eta_x \geq \eta_{x'}$ but $\mu(x) < \mu(x')$. Note that $V(x, y^*(x)) = \eta_x$. Therefore $V(x, y^*(x)) \geq V(x', y^*(x'))$, and so $(x, y^*(x)) \succeq (x', y^*(x'))$. Therefore, by the previous claim, for all (x, w) and (x', w') such that $\sigma(x, w) \preceq \sigma(x, y^*(x))$ and $\sigma(x', w') \preceq \sigma(x', y^*(x'))$, it is the case that $(x, w) \sim (x, y^*(x)) \succeq \sigma(x', y^*(x')) \sim (x', w')$. Furthermore, if $\sigma(x, z) \succ \sigma(x, y^*(x))$ then by the previous claim $(x, z) \succ (x, y^*(x))$. If $\mu(x) < \mu(x')$ then there exists a $z \in {}^1X$ such that for all $z' \in {}^1X$ $(x, z) \prec (x', z')$. This is a contradiction, since it was just proved that $(x', y^*(x')) \preceq (x, w)$ for all w.

- Now assume $\mu(x) \geq \mu(y)$ but $\eta_x < \eta_y$. Note that $V(x, y^*(x)) = \eta_x$. Therefore $V(x, y^*(x)) < V(x, y^*(x'))$, and so $(x, y^*(x)) \prec (x', y^*(x'))$. Therefore, by the previous claim, for all (x,w) and (x',w') such that $\sigma(x, w) \preceq \sigma(x, y^*(x))$ and $\sigma(x', w') \preceq \sigma(x', y^*(x'))$, it is the case that $(x, w) \sim (x, y^*(x)) \prec \sigma(x', y^*(x'))$. Furthermore, if $\sigma(x', z') \succeq \sigma(x', y^*(x'))$ then by the previous claim $(x', z') \succeq (x', y^*(x'))$. Therefore for all w' it is the case that $(x', w') \succ (x, y^*(x))$. If $\mu(x) \geq \mu(x')$ then there exists a $z' \in {}^1X$ such that for all $z \in {}^1X$ $(x, z) \succeq (x', z')$. This is a contradiction, since was just shown that $(x, y^*(x)) \prec (x, w)$ for all w.

□

Next, I will show that the valuation over histories can be written as the discounted value of the continuation DSP $\sigma(x, y*(x))$.

Claim 10 There exists a monotonic function ϕ such that $\eta_x = \phi(\hat{\mu}(x))$

Proof This is a direct implication of Claim 9. □

Definition *Define the function* $\mu \equiv \phi \circ \hat{\mu}$

Claim 11 For all $x \in X^0$, $\mu(x) \geq \mu(x')$ if and only if $\rho V(\sigma(x, y^*(x))) \geq \rho V(\sigma(x', y^*(x')))$ for all strictly increasing continuous functions $\rho : \mathbb{R} \to \mathbb{R}$.

Proof This is true since $V(\sigma(x, y^*(x))) = \eta_x \geq \eta_{x'} = V(\sigma(x', y^*(x')))$ and $\mu(x) = \eta_x$.□

Next, I will show that fixing a particular history $x \in X^0$ it is possible to represent the value of a DSP as the maximum of either a function that depends on the continuation payoff, or $\mu(x)$. I will first show this for a subset of histories, and then extend it to all histories. Furthermore, fix a strictly increasing continuous function $\rho : \mathbb{R} \to \mathbb{R}$ such that $\rho(V) < V$.

Claim 12 For all $x \in X^0$ and $w \in {}^1X$ such that $\sigma(x, w) \succeq \sigma(x, y^*(x))$, $V(x, w) = \nu'_x(V(\sigma(x, w)))$, where ν'_x is strictly increasing in V, and only maps from the upper contour set of $V(\sigma(x, y^*(x)))$, and $\nu'_x(x, y^*(x)) = \eta_x = \rho(V(x, y^*(x))) = \mu(x)$.

Proof For $x \in X^0$ Consider the set $\{(x, w) | \sigma(x, w) \succeq \sigma(x, y^*(x))\}$. I will show that $V(x, w) \geq V(x, w')$ if and only if $V(\sigma(x, w)) \geq V(\sigma(x, w'))$. Assume not. There are two cases.

- There exists $V(x, w) \geq V(x, w')$ and $V(\sigma(x, w)) < V(\sigma(x, w'))$. In this case, by the representation results, $(x, w) \succ (x, w')$ and $\sigma(x, w) \prec \sigma(x, w')$. Since $\sigma(x, y^*(x)) \preceq \sigma(x, w) \prec \sigma(x, w')$, then $(x, w) \prec (x, w')$. This is a contradiction.

- Now I need to show that $V(x,w) < V(x,w')$ and $V(\sigma(x,w)) \geq V(\sigma(x,w'))$. In this case, $(x,w) \prec (x,w')$ and $\sigma(x,w) \succeq \sigma(x,w')$. Since $\sigma(x,y^*(x)) \preceq \sigma(x,w') \preceq \sigma(x,w)$, then $(x,w) \succeq (x,w')$. This is a contradiction.

The value of $\nu'_x(x,y^*(x)) = \rho(V(x,y^*(x)))$ is simply a normalization. □

Claim 13 For all (x,w), $V(x,w) = \max\{\mu(x), \nu_x(V(\sigma(x,w)))\}$, where $\nu'_x = \nu_x$ for all (x,w) such that $\sigma(x,w) \succeq \sigma(x,y^*(x))$.

Proof Define ν_x as follows.

- $\nu_x(V) = \nu'_x(V)$ for all $V \geq V(\sigma(x,y^*(x)))$
- $\nu_x(V) = \nu'_x(V(\sigma(x,y^*(x))))$ otherwise

Then the proposition is true by construction and the previous claims. □

Claim 14 For all $(x,w) \in X$ if $\max\{\mu(x), \nu_x(V(\sigma(x,w)))\} > \mu(x)$, then $\sigma(x,w) \succ \sigma(x,y^*(x))$.

Proof I will prove the contrapositive. By construction if $\sigma(x,w) \preceq \sigma(x,y^*(x))$ then $V(\sigma(x,w)) \leq V(\sigma(x,y^*(x)))$, and so $\nu_x(V(\sigma(x,w))) = \nu_x V(\sigma(x,y^*(x))) = \mu_x$. □

Now the proof will start to tie together the representation across histories. First, I will identify the set of continuation DSPs that are never 'good enough' for the DM to choose to continue searching.

Claim 15 There exists a $n = (n^0, {}^1 n) \in X$ such that for all $m = (m^0, {}^1 m) \in X$, if $\sigma(m) \preceq \sigma(n)$ then $V(m) = \mu(m^0)$, and if $\sigma(m) \preceq \sigma(n)$, there exists a $(k^0, {}^1 k) = k \in X$ such that $V(k) \neq \mu(k^0)$. Furthermore, for any n, n' that satisfies this definition it must be the case that $\sigma(n) \sim \sigma(n')$.

Proof This is clearly satisfied by definition by considering the worst equivalence classes of continuation payoffs: the set n s.t. for all $m \in X$ $\sigma(m) \succeq \sigma(n)$. n exists because X is a compact set. To show the second part of the claim, assume that without loss of generality $\sigma(n) \succ \sigma(n')$. But then n' does not satisfy the conditions. □

Claim 16 For all $x \in X^0$ the sets $\{(x.y)|(x,y) \preceq (x,y^*(x))\}$ and $\{(x,y)|(x,y) \succeq (x,y^*(x))\}$ are closed.

Proof I will prove the claim for the second set. The proof for the other set is analogous. Recall that $\{(x,y)|y \in^X\}$ is closed. Given history $x \in X^0$, and take a sequence of sequences (x,y^n) that is strictly increasing in $\succeq$. This sequence is bounded (since the set of sequences is compact). Therefore,there exists a subsequence with a limit. Furthermore, this limit must be the limit of the sequence itself, since the sequence is strictly increasing. Next it needs to be shown that this limit has a history $x \in X^0$. Assume not. Consider the limit of a sequence (x, y_n) and assume it converges to (z, y'). In this case x converges to z. Clearly this is a contradiction unless $x = z$. $\square$

Definition *If for $n = (n^0,^1 n) \in X$ is is the case that $m = (m^0,^1 m)$, if $\sigma(m) \preceq \sigma(n)$ then $V(m) = \mu(m^0)$, and if $\sigma(m) \preceq \sigma(n)$, there exists a $(k^0,^1 k) = k \in X$ such that $V(k) \neq \mu(k^0)$, then denote any arbitrary sequence $x \in X$ such that $x \sim n$ as n^*.*

I will first partition the payoffs of the continuation sequences into regions. In each region I will separately construct a function that will map the payoffs of the continuation sequences tomorrow into their payoff today.

Definition *The range of payoffs $\mathbb{V} = \{V | \text{There exists an } x \in X \text{such that } V(x) = V\}$ is a compact connected set. Partition $\mathbb{V}$ into three regions, denoted:*

- $\mathbb{V}_1 = \{V | V \leq V(\sigma(n^*))\}$
- $\mathbb{V}_2 = \{V | V > V(\sigma(n^*))$ *and there exists an* $(x, y \in X)$ *such that* $\sigma(x, y) \succ \sigma(x, y^*(x))$ *and* $V(\sigma(x, y)) = V\}$
- $\mathbb{V}_3 = \{V | V > V(\sigma(n^*))$*such that* $V \notin \mathbb{V}_1 \cup \mathbb{V}_2\}$

Since X is a compact connected set, and V is a continuous function, then the set $\mathbb{V}$ is compact and connected (and in fact path connected). Furthermore, $\mathbb{V}_1$ is a closed set which is path connected. , $\mathbb{V}_2$ does not have to be a path-connected space. However, there exist maximally path-connected subsets of $\mathbb{V}_2$ (i.e. there exist no other path connected subspaces of $\mathbb{V}_2$ that contain them), which partition $\mathbb{V}_2$. Denote an arbitrary one of them $\hat{\mathbb{V}}_2$. Note that $\mathbb{V}_1$ must adjoin, $\mathbb{V}_2$. The $\hat{\mathbb{V}}_2$ that adjoins $\mathbb{V}_1$ is open below, by construction. Furthermore, any $\hat{\mathbb{V}}_2$ is closed above (since for any history, x we have shown there exists a $\succeq$ maximal element of $\sigma(x, u)$). Also, $\mathbb{V}_3$ does not have to be a path-connected space. However, there exist maximally path-connected subsets of $\mathbb{V}_3$ (i.e. there exist no other path connected subspaces of $\mathbb{V}_3$ that contain them), which partition $\mathbb{V}_3$. Denote an arbitrary one of them $\hat{\mathbb{V}}_3$. Consider a $\hat{\mathbb{V}}_3$, then it must be open below (since it adjoins $\hat{\mathbb{V}}_2$). Furthermore, it must be open above (since they adjoin $\hat{\mathbb{V}}_2$), unless if the global $\succeq$ maximal sequence $\mathbb{V}_3$, then the top maximally connected subset of $\mathbb{V}_3$ will be closed above.

Given each of these regions, I will define a function ν over them. For all $V \in \mathbb{V}_1$ ν is the discounted value of V; it maps the value of the continuation payoffs tomorrow into the value of the continuation payoffs today. For all $V \in \mathbb{V}_2$ ν is the minimum value of $\nu_x(V)$ over all histories x. For all $V \in \mathbb{V}_3$ ν is simply an increasing function that connects the values defined in the surrounding regions that are part of $\mathbb{V}_2$. If there doesn't exist any part of $\mathbb{V}_2$ that bounds the particular part of $\mathbb{V}_3$ above, then simply connect it to the minimum valuation that contains a continuation DSP in the highest equivalence class.

First, I will define a region in the set of continuation sequences that will align with the sequences that give payoffs in $\mathbb{V}_2$.

Definition *Define the set* $\mathbb{L} = \{\sigma(x,y) | \textit{there exists a } y' \textit{ such that } (x,y) \succ (x,y')\}$

Claim 17 For all $(x,y) \in X$ if $\sigma(x,y) \in \mathbb{L}$ then

- if $\sigma(w) \sim \sigma(x,y)$ then $w \succeq (x,y)$
- if $\sigma(w) \succ \sigma(x,y)$ then $w \succ (x,y)$

Proof This follows directly from **WM**. □

Claim 18 For all $(x,w) \in X$ if $\sigma(x,w) \succ \sigma(x, y^*(x))$ then $\sigma(x,w) \in \mathbb{L}$.

Proof If $\sigma(x,w) \succ \sigma(x,y^*(x))$ then $(x,w) \succ (x,y^*(x))$, so the definition is satisfied. □

Now I show that $\mathbb{L}$ and $\mathbb{V}_2$ define equivalent regions.

Claim 19 If $V \in \mathbb{V}_2$ then there exists a $\sigma(x,y) \in \mathbb{L}$ such that $V(\sigma(x,y)) = V$. If $\sigma(x,w) \in \mathbb{L}$ then $V(\sigma(x,w)) \in \mathbb{V}_2$

Proof This follows directly from the definitions of $\mathbb{V}_2$ and $\mathbb{L}$.□

Definition *Denote* $\bar{V} = \max_{V \in \mathbb{V}}$

$\bar{V}$ exists since X is compact and V is continuous.

Definition *Consider any* $\hat{\mathbb{V}}_3$. *Given* $\hat{\mathbb{V}}_3$, *denote the maximal connected subset of* $\mathbb{V}_2$ *that adjoins it below as* $\underline{\hat{\mathbb{V}}_2}(\hat{\mathbb{V}}_3)$ *and the maximal connected subset of* $\mathbb{V}_2$ *that adjoins it above as* $\bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)$.

It is now possible to define a discount function. The definition of ν on $\mathbb{V}_1$ and $\mathbb{V}_2$ is straightforward. The difficult part is defining it on $\mathbb{V}_3$ because in $\mathbb{V}_3$ there is no observed ranking of the value of the continuation sequences today. So the discount function ν must be constructed so that they match up with the values assigned in all other regions. There are two cases — whether or not $\hat{\mathbb{V}}_3$ includes the global $\succeq$ maximizing continuation value.

Definition *Define $\nu(V)$ as follows:*

- *If $V \in \mathbb{V}_1$ then $\nu(V) = \rho(V)$.*
- *If $V \in \mathbb{V}_2$ then $\nu(V) = \min_x \nu_x(V)$.*
- *For $V \in \hat{\mathbb{V}}_3$ such that $\bar{V} > V'$ for all $V' \in \hat{\mathbb{V}}_3$ then $\nu(V)$ is a strictly increasing continuous function such that for all $V \in \hat{\mathbb{V}}_3$, $\nu(V) < \nu_x(V)$ and*
  - $\lim_{V \to_+ \sup \hat{\mathbb{V}}_3} \nu(V) = \lim_{V \to_- \inf \bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)} \nu(V)$
  - $\lim_{V \to_+ \sup \underline{\hat{\mathbb{V}}_2}(\hat{\mathbb{V}}_3)} \nu(V) = \lim_{V \to_- \inf \hat{\mathbb{V}}_3} \nu(V)$
- *For $V \in \hat{\mathbb{V}}_3$ such that there exists a $V' \in \mathbb{V}$ where $\bar{V} = V'$ then $\nu(V)$ is an increasing continuous function such that for all $V \in \hat{\mathbb{V}}_3$, $\nu(V) < \nu_x(V)$ and*
  - *$\nu(\bar{V}) = \min_x V(x)$ for all x where $V(\sigma(x)) = \bar{V}$*
  - $\lim_{V \to_+ \sup \underline{\hat{\mathbb{V}}_2}(\hat{\mathbb{V}}_3)} \nu(V) = \lim_{V \to_- \inf \hat{\mathbb{V}}_3} \nu(V)$

It is not clear that such a ν actually exists.

Claim 20 There exists a ν that satisfies the definition

Proof Clearly within $\mathbb{V}_1$ and $\mathbb{V}_2$ a function satisfying the definition exists. However, it remains to be shown that a function satisfying the condition exists in $\mathbb{V}_3$. Consider $\hat{\mathbb{V}}_3$. ν needs to be strictly increasing, continuous, such that $\nu(V) < \nu_x(V)$, and that the limit equalities hold.

First I will show that in $\hat{\mathbb{V}}_3$ it is possible to construct a ν that is strictly rising, continuous, and satisfies the limit requirements. In order to show this it is true, it needs to be shown that either

- $\lim_{V \to_- \inf \bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)} \nu(V) > \lim_{V \to_+ \sup \underline{\hat{\mathbb{V}}_2}(\hat{\mathbb{V}}_3)} \nu(V)$ or
- $\min_x V(x) > \lim_{V \to_+ \sup \underline{\hat{\mathbb{V}}_2}(\hat{\mathbb{V}}_3)} \nu(V)$

is true. Then exists an increasing continuous function that traverses the space between the two limits. Note that there are two cases.

- The first case is when $\hat{\mathbb{V}}_3$ is adjoined both above and below by $\mathbb{V}_2$. A sufficient condition is to show that $\lim_{V \to_{-} \inf \bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)} \nu(V) > \lim_{V \to_{+} \sup \underline{\hat{\mathbb{V}}_2}(\hat{\mathbb{V}}_3)} \nu(V)$. By way of contradiction assume not. By definition there exist $(x, w), (x, w') \in X$ such that $(x, w) \succ (x, w')$ and $\sigma(x, w) \succ \sigma(x, w')$ in the lower adjoining set $\underline{\hat{\mathbb{V}}_2}(\hat{\mathbb{V}}_3)$. There is also a $(x, y), (x, y') \in X$ such that$(x, y) \succ (x, y')$ and $\sigma(x, y) \succ \sigma(x, y')$ in the upper adjoining set $\bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)$. Note that $\sigma(x, y) \succ \sigma(x, y') \succ \sigma(x, w) \succ \sigma(x, w')$ by construction. Therefore, by **WM** $(x, y') \succ (x, w)$.

- The second case is when $\hat{\mathbb{V}}_3$ is adjoined below by $\mathbb{V}_2$ and includes the global $\succeq$ maximal element. In this case, consider the maximal equivalence class of DSPs $\{\sigma(x', y') | \sigma(x', y') \succeq \sigma(r) \forall r \in X\}$, which is in $\hat{\mathbb{V}}_3$). Then the argument in the previous applies in an analogous fashion.

Now I need to show that within the set of potential ν's, there exists one such that $\nu(V) < \nu_x(V)$ in for all $V \in \mathbb{V}_3$. By way of contradiction assume not, and denote an $x \in X^0$ that violates the inequality as $\hat{x}$. Since it has been shown that ν is strictly increasing throughout $\mathbb{V}_3$, this amounts to showing that one of two conditions must hold.

- Denoting the maximal element of $\underline{\hat{\mathbb{V}}_2}(\hat{\mathbb{V}}_3)$ as $\underline{V}_2$ (which exists since each $\hat{\mathbb{V}}_2$ is closed above), then $\nu(\underline{V}_2) \geq \nu_{\hat{x}}(V)$, for some $V \in \hat{\mathbb{V}}_3$. By definition $\underline{V}_2 = V(\sigma(x, y))$ where $\sigma(x, y) \in \mathbb{L}$. Note that for any $V \in \hat{\mathbb{V}}_3$ there exists a w such that $V(\sigma(w)) = V$, and so $\sigma(w) \succ \sigma(x, y)$. Therefore, by **WM** $w \succ (x, y)$. This is a contradiction.

- There exists $x \in X^0$ such that $\lim_{V \to_{+} \sup \hat{\mathbb{V}}_3} \nu_x(V) < \lim_{V \to_{-} \inf \bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)} \nu(V)$. The set $\{\sigma(x, y) | y \in^1\}$ is closed. Denoting $\lim_{V \to_{+} \sup \hat{\mathbb{V}}_3} \nu_x(V)$ as $\nu_x(\hat{V}_3)$, there exists an (x, y) such that $V(\sigma(x, y)) = \nu_x(\hat{V}_3)$.

  First, assume that $\hat{\mathbb{V}}_3$ contains the maximal V value, $\bar{V}$. Then $V(\sigma(x, y))$ is $\bar{V}$. By construction $\nu(\bar{V}) = \min_x V(x)$ for all x where $V(\sigma(x)) = \bar{V}$, and so there is a contradiction.

  Next, assume $\hat{\mathbb{V}}_3$ is bounded above by $\mathbb{V}_2$. If $\bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)$ is closed below, then by continuity, there must exist a (x', y') such that $\sigma(x', y') \sim \sigma(x, y)$. Therefore, by **WM** $(x', y') \preceq (x, y)$. Therefore, $\nu(V(\sigma(x', y'))) \geq \nu_x(V(\sigma(x, y)))$, a contradiction.

  If $\bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)$ is open below, then this means there is some history x' such that there exists a a sequence of futures $y^{',n}$ and a future y such that $\lim_{n\to\infty} \sigma(x', y^{',n}) = \sigma(x', y')$ where $V(\sigma(x', y^{',n})) \in \mathbb{V}_2$ and $V(\sigma(x', y')) \in \mathbb{V}_3$. To see this, note that since $\bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)$ is open below there exists a sequence $\lim_{n\to\infty} \sigma(x', y^{',n})$ that is monotonically falling in $\succeq$, and since it is bounded below, it must have a limit. Clearly this limit is not in $\bar{\hat{\mathbb{V}}}_2(\hat{\mathbb{V}}_3)$ by assumption, and clearly the limit must have history x'. By assumption $V(\sigma(x, y)) = V(\sigma(x', y'))$ and

$V((x,y)) < V((x',y'))$. But, since (x',y') satisfies the if part of **WM**, then it is the case that $V((x,y)) \geq V((x',y'))$, a contradiction.

$\square$

Claim 21 ν is continuous.

Proof Clearly within each region $\mathbb{V}_1$, $\mathbb{V}_2$ and $\mathbb{V}_3$ ν is continuous. The reasoning, respectively, is by the continuity of V and ρ in $\mathbb{V}_1$, because each ν_x is continuous and well defined for $\mathbb{V}$ and so also in $\mathbb{V}_2$ and by construction on $\mathbb{V}_3$.

Therefore, it is sufficient to check that ν is continuous at the boundaries between regions. First, by construction, ν is continuous when it transitions between $\mathbb{V}_2$ and $\mathbb{V}_3$, as well as vice-versa (because the limits agree). Similarly by construction it is continuous transitioning between $\mathbb{V}_1$ and $\mathbb{V}_2$. So the it needs to be shown that ν is continuous when it transitions between $\mathbb{V}_1$ and $\mathbb{V}_2$. Therefore, it needs to be shown that $\rho(V(\sigma(x,w))) = \min_x \nu_x V(\sigma(x,w))$ where $(x,w) \sim n^*$.

Assume $\rho(V(\sigma(x,w))) > \min_x \nu_x V(\sigma(x,w))$ for some history $x \in X^0$. Denote this history as $\underline{x}$. Then it is the case that $V(\underline{x}, y^*(\underline{x})) < \rho(V(\sigma(\underline{x},w))) = \mu(\underline{x}$. This is a contradiction.

Now assume $\rho(V(\sigma(x,w))) < \min_x \nu_x V(\sigma(x,w))$for some history $x \in X^0$. Denote this history as $\underline{x}$. Then $V(\underline{x}, y^*(\underline{x})) < \nu_{\underline{x}} V(\sigma(\underline{x},w)) = V((\underline{x},w)$. But then $(\underline{x}, y^*(\underline{x})) \nsim n*$, a contradiction. $\square$

Claim 22 ν is strictly increasing.

Proof In region $\mathbb{V}_1$ it is strictly increasing by construction. Since ν_x is strictly increasing for each x then in $\mathbb{V}_2$ ν must be increasing. Lastly, it was previously shown that ν is strictly increasing by construction in $\mathbb{V}_3$.$\square$

Claim 23 For all $(x,w) \in X$ if $V(\sigma(x,w)) \in \mathbb{V}_2$ then $\nu_x(V(\sigma(x,w))) = \nu(V(\sigma(x,w)))$.

Proof Assume not. Then it must be the case that there is an $x \in X^0$ such that $\nu_x(V(\sigma(x,w))) > \nu(V(\sigma(x,w)))$ with $V(\sigma(x,w)) \in \mathbb{V}_2$. Therefore, there exists an x' such that $\nu_x(V(\sigma(x,w))) > \nu_{x'}(V(\sigma(x',w')))$. But since $V(\sigma(x,w)) \in \mathbb{V}_2$, then $\sigma(x,w) \in \mathbb{L}$. Therefore, for all $\sigma(x',w')$ such that $\sigma(x',w') \sim \sigma(x,w)$, by **WM** $(x',w') \succeq (x,w)$. Since $\nu_{x'}(V(\sigma(x',w'))) = V((x',w'))$, and $\nu_x(V(\sigma(x,w))) = V((x,w))$, then $\nu_{x'}(V(\sigma(x',w'))) \geq \nu_x(V(\sigma(x,w)))$. This is a contradiction.$\square$

Claim 24 For all $x \in X^0$ $\nu(V) \leq \nu_x(V)$ for all V.

Proof Note that in region $\mathbb{V}_1$, $nu_x(V) = \beta V(\sigma(x, y^*(x)))$ by construction. Since for all $V \in \mathbb{V}_1$, if for some r $V(\sigma(r)) = V$ then $\sigma(r) \preceq \sigma(x, y^*(x))$. Therefore $\nu(V) = \rho(V(\sigma(r))) \leq \sigma(x, y^*(x))$.

It was already shown that in $\mathbb{V}_2$ $\nu \leq \nu_x$. In $\mathbb{V}_3$ by construction $\nu(V) < \nu_x(V)$. $\square$

Claim 25 For all $(x, w) \in X$ $\max\{\mu(x), \nu_x(V(\sigma(x, w)))\} \geq \max\{\mu(x'), \nu_{x'}(V(\sigma(x', w')))\}$ if and only if $\max\{\mu(x), \nu(V(\sigma(x, w)))\} \geq \max\{\mu(x'), \nu(V(\sigma(x', w')))\}$

Proof First, if $\max\{\mu(x), \nu_x(V(\sigma(x, w)))\} = \mu(x)$, then $\mu(x) \geq \nu_x(V(\sigma(x, w)))$. Furthermore, since $\nu_x(V(\sigma(x, w))) \geq \nu(V(\sigma(x, w)))$ then $\max\{\mu(x), \nu(V(\sigma(x, w)))\} = \mu(x)$.

If $\max\{\mu(x), \nu_x(V(\sigma(x, w)))\} > \mu(x)$, then $\sigma(x, w) \succ \sigma(x, y^*(x))$. Therefore, $\sigma(x, w) \in L$ and so, $V(\sigma(x, w)) \in \mathbb{V}_2$, which means $\nu_x(V(\sigma(x, w))) = \nu(V(\sigma(x, w)))$. Therefore $\max\{\mu(x), \nu(V(\sigma(x, w)))\} = \nu(V(\sigma(x, w))) > \mu(x)$.

Thus $\max\{\mu(x), \nu_x(V(\sigma(x, w)))\} = \max\{\mu(x), \nu(V(\sigma(x, w)))\}$. $\square$

Claim 26 There exists a ν such that $\nu(V(\sigma(x, w))) < V(\sigma(x, w))$

Proof Assume not at a point $\hat{V}$. In this case, there exists a point where $\nu(V(\sigma(x, w))) \geq V(\sigma(x, w))$. By construction this cannot happen in $\mathbb{V}_1$. If it occurs in $\mathbb{V}_2$, then for all $x \in X^0$ such that there exists a $z \in^1 X$ where $V(\sigma(x, z)) = \hat{V}$, $\nu_x(V(\sigma(x, z))) \geq V(\sigma(x, z))$. This implies $V((x, z)) \geq V(\sigma(x, z))$. Since $V(\sigma(x, z)) \in \mathbb{V}_2$, then $\sigma(x, z) \in \mathbb{L}$ and so the if part of **SB** holds, which is a contradiction.

In $\mathbb{V}_3$, by construction $\nu_x(V(\sigma(x, w))) < \nu_x(V(\sigma(x, w))) = V(\sigma(x, w))$. $\square$

Therefore a GSR representation exists.

The see that uniqueness holds assume there exists (x, w) and (y, z) such that $V(x, w) \geq V(y, z)$ but $V'(x, w) < V'(y, z)$ (or $V(x, w) < V(y, z)$ but $V'(x, w) \geq V'(y, z)$). But since both V and V' represent $\succeq$ then it is the case that $(x, w) \succeq (y, z)$ and $(x, w) \prec (y, z)$, a contradiction. $\square$

Proof of Theorem 2 The theorem will be claimed in a series of steps.

Claim 1 If μ is strictly f-relevant at some history $x \in X^0$ then μ' is also strictly f-relevant at x.

Proof Without loss of generality assume μ is not strictly f-relevant but μ' is for some history $x \in X^0$. This means that there exists a $y \in {}^1X$ such that $V'(x, y) = \mu'(x) > \nu'(V'(\sigma(x, y)))$ and so there exists y' such that $(x, y) \sim (x, y')$ and $\sigma(x, y) \nsim \sigma(x, y')$.

Since μ is not strictly f-relevant there are two possibilities. It could be that for all $z \in {}^1X$ $V(x, z) > \mu(x)$, but by construction this is impossible. Therefore it must be the case that if $V(x, z) = \mu(x)$ then $\nu(V(\sigma(x, z))) = \mu(x)$.

Since $(x,y) \sim (x,y')$ then it must be the case that $V(x,y) = V(x,y')$. Therefore, $V(x,y) = V(x,y') = \nu(V(\sigma(x,y'))) = \nu(V(\sigma(x,y)))$. But since $\sigma(x,y) \nsim \sigma(x,y')$, then $\nu(V(\sigma(x,y))) \neq \nu(V(\sigma(x,y'))) = \mu(x)$, a contradiction. □

Claim 2 If histories y and y' are both strictly f-relevant, then $\mu(y) \geq \mu(y')$ if and only if $\mu'(y) \geq \mu'(y')$.

Proof This is a direct implication of the uniqueness of V. If (y,z) and (y,z) are strictly f-relevant for μ, then they are for μ'. By way of contradiction, assume $V(y,z) = \mu(y,z) \geq \mu(y,z) = V(y',z')$ but $V'(y,z) = \mu'(y,z) < \mu'(y,z) = V'(y',z')$. This is a contradiction. □

Claim 3 $\nu(V(\sigma(y,z))) \geq \nu(V(\sigma(y',z')))$ if and only if $\nu'(V(\sigma(y,z))) \geq \nu'(V(\sigma(y',z')))$.

Proof To see that this is true, note that ν and ν' are both strictly increasing and V and V' are monotonic transformations of one another. The claim immediately follows. □

This completes the proof. □

Proof of Theorem 3 Necessity is implied by Theorem 1 of Quah and Strulovici (2011). The proof of sufficiency is in a series of steps. Define, for any j-relevant history x, the future $y*_i(x)$, so that for all (x,w), where $\sigma(x,w) \succ_i \sigma(x, y*_i(x))$, it is the case that $(x,w) \succ_i (x, y*_i(x))$.

Recall that if an individual is revealed to weakly prefer to stop at (x,y) then it is the case that there exists a future y' such that $(x,y) \sim (x,y')$ but $\sigma(x,y) \nsim \sigma(x,y')$. Furthermore, if the individual is revealed to strictly prefer to search at (x,y) then there exists a future y' such that $(x,y) \nsim (x,y')$ and the individual is not revealed to weakly prefer to stop at (x,y).

Because continuation sequences are particularly important in this proof, this particular proof will slightly abuse notation in order to make the proof more transparent. Given a sequence (x,y), we will denote the stopping payoff to it's continuation sequence $\sigma(x,y)$ and $\mu(\sigma(x,y))$, rather than $\mu(x,y_1)$.

First, the Theorem will be shown to be true for a subset of the sequences.

Claim 1 Assume that individual B (weakly) prefers to stop at $\sigma(x, y*_B(x))$. Then $\nu_A(V_A(\sigma(x, y*_B(x)))) \leq \nu_B(V_B(\sigma(x, y*_B(x))))$.

Proof Proof by contradiction; assume that $\nu_A(V_A(\sigma(x, y*_B(x)))) > \nu_B(V_B(\sigma(x, y*_B(x))))$. By construction $V_B(\sigma(x, y*_B(x))) = \mu_B(\sigma(x, y*_B(x)_1)) = \mu_A(\sigma(x, y*_B(x))) = V_A(\sigma(x, y*_B(x)))$.

Furthermore, $V_B(x, y*_B(x)) = \mu_A(x) = \mu_B(x)$. Then $V_A(\sigma(x, y*_B(x))) = V_B(x, y*_B(x))$. This is a contradiction. $\square$

Next, the proof will consider all other situations where the assumptions of Claim 1 do not hold. Assume that individual B wants to (strictly) continue searching at $\sigma(x, y*_B(x))$. Note that this implies $V_B(\sigma(x, y*_B(x))) > \mu(x, y*_B(x)_1)$. By way of contradiction assume that $\nu_A(V_A(\sigma(x, y*_B(x)))) > \nu_B(V_B(\sigma(x, y*_B(x))))$. In this case, by construction we can find a sequence w such that $V(\sigma(w))$ is arbitrarily close to $V(\sigma(x, y*_B(x)))$, and which satisfies the assumptions of Claim 1. Note that at this point it must be the case that $\nu_A(V_A(\sigma(w))) > \nu_B(V_B(w)))$, which is a contradiction with Claim 1.

This completes the proof. $\square$

Proof of Lemma 3 Since Ω is connected, then $\Delta(\Omega)$ is connected, and so $\hat{X}$ is connected. Since $\hat{X}$ is compact, then the set of lotteries in the support of α_x is compact. There is a best and worst sequence (considering only $\succeq$ equivalence classes) in the support of α_x which is also the best and worst sequences — denote their utility as $\bar{U}(\alpha_x)$ and $\underline{U}(\alpha_x)$.

Denote the set of sequences who agree on entries for $t \neq 0$, where the t^{th} entry is $\alpha_{x,t}$ for $t \neq 0$, as $\aleph$. Note that $\aleph$ is a compact, connected and closed set. For any real number in $r \in (\bar{U}(\alpha_x), \underline{U}(\alpha_x)$ there exists $x' \in \aleph$ so that $V(x') = r$. Since $V(\alpha_x \in (\bar{U}(\alpha_x), \underline{U}(\alpha_x)$ then for any $\alpha_x \in \Delta(\hat{X})$ there exists an $x' \in \hat{X}$ so that $x' \sim \alpha_x$ and whose entries are $\alpha_{x,i}$ for all $i \neq 0$. $\square$

Proof of Lemma 4 Existence follows directly from Lemma 3. Continuity holds because $\succeq$ is continuous, and therefore z^* is continuous. Given this, the proof is analogous to Lemma A1. $\square$

Proof of Theorem 4 First I shall show the necessity of each axioms:

- **EUSL:** The necessity of **EUSL** is entirely standard.

- **LI:** By construction $U(\alpha_x) = U(..., \alpha_{x,-1}, z^*, \alpha_{x,1}, ...) = V(..., \alpha_{x,-1}, z^*, \alpha_{x,1}, ...) = V(\hat{\sigma}(x))$, which implies **LI**.

- **GSRE:** The necessity of **GSRE** is as in Theorem 1.

- **CL:** If for all z_i in the support of f_1, it is the case that either there exists an $x' \in \hat{X}$ such that $(..., x_0, z_i, f_2, ...) \succ x'$, or there exists an $x' \in \hat{X}$ such that $x' \succ (..., x_0, z_i, f_2, ...)$ and for all $x'' \in \hat{X}$ such that $\hat{\sigma}(x'') \succ \hat{\sigma}(..., x_0, z_i, f_2, ...)$ it is the case that $(x'') \succ (..., x_0, z_i, f_2, ...)$, then the DM prefers to search regardless of the realization z_i. Therefore $U(\alpha_{x,f_1}) = \delta(x^0)E_{x' \in f_1}(V(\sigma(x^0), x', \sigma(^2f))),$

where x' can take on the value of any z_i. By construction $V(x) = \max\{\mu(x^0), \delta(x^0)E_{x'\in f_1}(V(\sigma(x^0), x', \sigma(^2f))),$ and since the DM prefers to search for any $x' \in f_1$ (i.e. each z_i), then $\delta(x^0)V(\sigma(x^0), x', \sigma(^2f)) \geq \mu(x^0)$ for all x' with positive support (i.e. the set of z_i). Therefore $\delta(x^0)E_{x'\in f_1}(V(\sigma(x^0), x', \sigma(^2f))) \geq \mu(x^0)$. Thus $\alpha_{x,f_1} \sim x$.

I will prove sufficiency in a series of steps.

Claim 1 $\succeq$ satisfies **EUSL** if and only if there exists functions $U : \Delta(\hat{X}) \to \mathbb{R}$, $V : \hat{X} \to \mathbb{R}$ such that $\alpha \succeq \beta$ if and only if $U(\alpha) \geq U(\beta)$, where $U(\alpha) = \sum_{x\in\alpha} p(x)V(x)$

Proof This is simply the standard EU representation. □

Claim 2 If $\succeq$ satisfies **EUSL** and **GSRE** then $\succeq$ restricted to $\hat{X}$ can be represented by $V : \hat{X} \to \mathbb{R}$, $\mu : X^0 \to \mathbb{R}$, and an increasing function $\nu : \mathbb{R} \to \mathbb{R}$ such that for any $x \in \hat{X}$

$$V(x^0,{}^1 f) = \max\{\mu(x^0), \nu(V(\hat{\sigma}(x)))$$

Proof This simply is Theorem 1 appropriately modified to the new domain. The proof is exactly analogous, since CWO is also satisfied on $\hat{X}$. □

Claim 3 $\succeq$ restricted to $\hat{X}$ can be represented using $V : \hat{X} \to \mathbb{R}$, $\mu : X^0 \to \mathbb{R}$, and an increasing function $\nu : \mathbb{R} \to \mathbb{R}$ such that for any $(x^0,{}^1 f) \in \hat{X}$

$$V(x^0,{}^1 f) = \max\{\mu(x^0), \nu(E_{x'\in f_1} V(\sigma(x^0), x', \sigma(^2f)))$$

Proof For any $x \in \hat{X}$, let $V(x) = U(x)$. By **LI**, $\alpha_x \sim \hat{\sigma}(x)$, therefore $U(\alpha_x) = E_{x'\in f_1} V(\sigma(x^0), x', \sigma(^2f))) = V(\hat{\sigma}(x))$. □

Claim 4 Fix a history $x^0 \in \hat{X}^0$. $\succeq$ restricted to $\hat{X}$, with history x^0 can be represented by

$$V(x^0,{}^1 f) = \max\{\mu(x^0), \delta_{x^0} E_{x'\in f_1} V(\hat{\sigma}(x^0), x', \hat{\sigma}(^2f))$$

where δ_{x^0} is a scalar strictly less than 1.

Proof First, pick out a potential U, V pair such that for all x, $V(x) > 0$, which exists because we can simply take an affine transformation of any other pair of functions. Since $\mu(x^0) = \delta_{x^0} V(\hat{\sigma}(x^0, y^*(x^0))),$

where δ_{x^0} is a scalar strictly less than 1. Define $\widetilde{V}$ as the set $\{V|$There exists a y such that $\hat{\sigma}(x^0, y) \succeq \hat{\sigma}(x^0, y^*(x^0))$ and $V(\hat{\sigma}(x^0, y) = V\}$. Then **CL** says that

$$\nu(E_{x' \in f_1} V(\sigma(x^0), x', \sigma(^2 f))) = E_{x' \in f_1} \nu(V(\sigma(x^0), x', \sigma(^2 f)))$$

so long as $V(\hat{\sigma}(x^0), x', \hat{\sigma}(^2 f))) \in \widetilde{V}$. Therefore, $\nu(V) = \delta_{x^0} V < V$, so long as $V(\hat{\sigma}(x^0), x', \hat{\sigma}(^2 f))) \in \widetilde{V}$. Furthermore, since $\mu(x^0) = \delta_{x^0} V(\hat{\sigma}(x^0, y^*(x^0)))$, then arbitrarily set $\nu(V(\hat{\sigma}(x^0, y'))) = \delta_{x^0} V(\hat{\sigma}(x^0, y'))$ for all $\hat{\sigma}(x, y') \prec \hat{\sigma}(x, y^*(x))$. $\square$

Claim 5 Assume that there exists a sequence $\hat{\sigma}(y, z) \in X$ such that for all $\hat{\sigma}(y', z') \in X$, if $\hat{\sigma}(y', z') \succ \hat{\sigma}(y, z)$ then there exists a $\hat{\sigma}(y'', z'') \in X$ such that $\hat{\sigma}(y', z') \sim \hat{\sigma}(y'', z'')$ and $\nu \circ V$ is strictly f-relevant for $\hat{\sigma}(y'', z'')$. Then $\delta(x^0) = \delta$.

Proof Consider an arbitrary continuation sequence $\hat{\sigma}(\hat{y}, \hat{z}) \in X$ such that $\hat{\sigma}(\hat{y}, \hat{z}) \succ \hat{\sigma}(y, z)$. Let $\delta = \delta_{\hat{y}}$. Note that for any history $y' \in X^0$ whose continuation sequences all fall in the lower contour set of $\hat{\sigma}(y, z)$ it is possible to arbitrarily set $\delta_{y'} = \delta$.

Consider an arbitrary history y'. Now consider a history $\bar{y'}$ such that there exists no other distinct history $y'' \in {}^1X$ such that there exists a z''' where $\sigma(y'', z''')$ falls strictly between the $\succeq$ minimal element of the set of sequences $\{\hat{\sigma}(\bar{y'}, z'')|\hat{\sigma}(\bar{y'}, z'')$ is strictly f-relevant $\}$ and the $\succeq$ maximal element of $\{\hat{\sigma}(y', z'')|\hat{\sigma}(y', z'')$ is strictly f-relevant $\}$ in the ordering $\succeq$. Also consider history $\underline{y'}$, such that there exists no other distinct history where the $\succeq$ maximal element of the set of sequences $\{\hat{\sigma}(\underline{y'}, z'')|\hat{\sigma}(\underline{y'}, z'')$ is strictly f-relevant $\}$ and the $\succeq$ minimal element of $\{\hat{\sigma}(y', z'')|\hat{\sigma}(y', z'')$ is strictly f-relevant $\}$.

A necessary condition for there to exist a history $\widetilde{y}$ such that $\delta_{\widetilde{y}} \neq \delta$ is that there exists a history y' such that $\delta_{y'} \neq \delta_{\bar{y'}}$ or $\delta_{y'} \neq \delta_{\underline{y'}}$. However, since $\{\hat{\sigma}(\bar{y'}, z'')|\hat{\sigma}(\bar{y'}, z'')$ is strictly f-relevant $\}$ and $\{\hat{\sigma}(y', z'')|\hat{\sigma}(y', z'')$ is strictly f-relevant $\}$ are both closed, the domain is connected, and preferences are continuous, $z^*, \bar{z}^*$ such that $\sigma(\bar{y'}, \bar{z}^*) \sim \sigma(y', z^*)$. Therefore $\delta_{\bar{y}} = \delta_{y'}$.

The analogous argument applies to the y' and $\underline{y'}$. Therefore $\delta(y') = \delta$ for all histories y'. $\square$

This completes the proof. $\square$

Proof of Theorem 5 Necessity of **OI** and **RR** are clear and left to the reader. Sufficiency will first be shown for any sequence with a finite set of unique outcomes. The proof is by induction. Consider sequences with one unique outcome. Denote μ of a history where outcome c appears every period as $u(c) \equiv \mu(....c, c, c)$. Therefore $\mu(.....c) = \max_{x_i \in (....,c)} u(x_i) = u(a)$.

Now assume for all sequence with weakly less than N unique outcomes, that $\mu(.....x_0) = \max_{x_i \in (.....,x_0)} u(x_i)$. Now consider an arbitrary history with $N+1$ unique outcomes, and call it y.

Find an equivalent history $e(y)$, which is composed of the $N+1$ unique elements occupying the last $N+1$ slot (i.e. the slots for t=0, -1,.....-N). Let $a = x_{-N}$ and $b = x_{-N+1}$. For $t = -N-k$ $x_t = a$ for k even (and 0) and b for k odd. This completely describes $e(y)$.

Then there exists y', y'' such that $e(y) = y' \biguplus y''$. Note that by construction both y' and y'' have less than $N+1$ unique outcomes. Without loss of generality assume y' is the better history. By assumption $\max_{x_i \in y'} u(x_i) \geq \max_{x_i \in y''} u(x_i)$ and so $\mu(y) = \mu(e(y)) = \mu(y') = \max_{x_i \in y'} u(x_i) = \max_{x_i \in e(y)} u(x_i)$.

This proves the theorem for all histories with a finite set of unique alternatives. For any history y with a countably infinite set of unique alternatives there exists a sequences of histories y^n each of which contains only a finite set of unique alternatives which converges in the product topology to y. This proves the Theorem for all histories. □

2.10.3 Appendix B Results

Proof of Theorem 6 Each part is proved in turn. Necessity is straightforward and left to the reader.

- $V(x) = V(^0x)$ by **NR**. And so $V(^0x) = \max\{\mu(x^0), \nu(V(\sigma(x)))\}$, therefore $\mu(x^0) = u(x_0)$.
- This claim is directly implied by Gul and Pesendorfer (2001). This is because **OI'** implies that the DM only cares about the unordered sets of histories — i.e. a domain covered by Gul and Pesendorfer. For any particular history $x \in X^0$, if the history has a finite set of unique outcomes then Gul and Pesendorfer's result directly applies. This is because $\succeq_0$ is a continuous preference order, and so satisfies Gul and Pesendorfer's Axioms 1 and 2. Furthermore, I directly impose **I**, which is Gul and Pesendorfer's Axiom 3. Lastly, **DT** is Gul and Pesendorfer's Axiom 4. Thus Gul and Pesendorfer's results directly follow, which imposes the required functional form on μ. If there is a history $x \in X^0$ with a countably infinite set of unique outcomes, then there is a sequence of histories, all with a finite set of unique outcomes that converges to x (simply start with a sequence with one unique outcomes, and then switch each element in reverse chronological order). By the continuity of the preferences, the representation must hold at x as well.

□

References

[1] C. Brown, M. Flinn, and A. Schotter. Real time search in the laboratory and the market. *American Economic Review*, 2011. Forthcoming.

[2] K. Burdett and D. Malueg. The theory of search for several goods. *Journal of Economic Theory*, 24:362–376, 1981.

[3] A. Caplin and M. Dean. Search, choice and revealed preference. *Theoretical Economics*, 6:29–48, 2011.

[4] A. Cooke, T. Meyvis, and A. Schwartz. Avoiding future regret in purchase timing decisions. *The Journal of Consumer Research*, 27:447–459, 2001.

[5] B. De los Santos, A. Hortacsu, and M. R.Wildenbeest. Testing models of consumer search using data on web browsing and purchasing behavior. 2011. Working Paper.

[6] G. Debreu. Representation of a preference ordering by a numerical function. In *Decision Processes*, pages 159–165. John Wiley, 1954.

[7] L. Epstein. Stationary cardinal utilitya and optimal growth under uncertainty. *Journal of Economic Theory*, 31:133–152, 1983.

[8] X. Gabaix, D. Laibson, G. Moloche, and S. Weinberg. Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, 96:1043–1068, 2006.

[9] F. Gul and W. Pesendorfer. Temptation and self-control. *Econometrica*, 69:1403–1435, 2001.

[10] J. Hey. Search for rules for search. *Journal of Economic Behavior and Organization*, 3:65–81, 1982.

[11] J. Hey. Still searching. *Journal of Economic Behavior and Organization*, 8:137–144, 1982.

[12] S. Horan. Sequential search and choice from lists. 2011. Working Paper.

[13] B. Irons and C. Hepburn. Regret theory and the tyranny of choice. *The Economic Record*, 83:191–203, 2007.

[14] Levav J., H. Heitmann, A. Herrmann, and S. Iyengar. Order in product customization decisions: Evidence from field experiments. *Journal of Political Economy*, 118:274–299, 2010.

[15] T. Koopmans. Stationary ordinal utility and impatience. *Econometrica*, 28:287 – 309, 1960.

[16] T. Koopmans. Representations of preference orderings over time. In *Decision and Organization*, pages 57–100. North-Holland, 1972.

[17] T. Koopmans. Representations of preference orderings with independent components of consumption. In *Decision and Organization*, pages 57 – 100. North-Holland, 1972.

[18] Y. Masatlioglu, D. Nakajima, and C. Raymond. Salience and satisficing. 2010. Working Paper.

[19] J.J. McCall. Economics of information and job search. *Quarterly Journal of Economics*, 84:113–126, 1970.

[20] J. Noor. Intertemporal choice and the magnitude effect. *Game and Economic Behavior*, 72:255–270, 2011.

[21] M. Papi. Satisficing choice procedures. 2010. Working Paper.

[22] J. Quah and B. Strulovici. Discounting and patience in optimal stopping and control problems. 2011. Working Paper.

[23] C. Raymond. Search and context dependent preferences. 2011. Working Paper.

[24] K. Rozen. Foundations of intrinsic habit formation. *Econometrica*, 78:1341–1373, 2010.

[25] A. Rubinstein and Y. Salant. A model of choice from lists. *Theoretical Economics*, 1:3–27, 2006.

[26] A. Rustichini and P. Siconolfi. A dynamic theory of preference for variety. 2006. Working Paper.

[27] H. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118., 1955.

[28] I. Simonson. Anticipating regret and responsibility on purchase decisions. *The Journal of Consumer Research*, 19:105–118, 1992.

CHAPTER III

A Model of Non-Belief in the Law of Large Numbers[1]

3.1 Introduction

Most people understand that larger samples are more likely to reflect the proportions of a population than smaller ones. Following Tversky & Kahneman (1971), Rabin (2002) and Rabin & Vayanos (2010) model the notion that people believe in "the Law of Small Numbers," and exaggerate how likely it is that small samples will reflect the underlying population. Yet evidence indicates that people *also* do not believe in the Law of Large Numbers: they believe that even in very large random samples, proportions might depart significantly from the overall population rate. This paper develops a formal model of such "non-belief in the Law of Large Numbers," which we abbreviate by NBLLN.[2] We explore the model's relationship to psychological evidence, its implications for inference, and some of its economic consequences. We identify a few alternative ways that NBLLN may manifest itself, and we confront and explore conceptual challenges with modeling NBLLN that are also likely to arise in modeling other biases in statistical reasoning.

Kahneman & Tversky (1972) find that subjects seem to think sample proportions reflect a "universal sampling distribution," virtually neglecting sample size.[3] In doing so, subjects vastly exaggerate the probability of unbalanced ratios in large samples. For instance, independent of whether a fair coin is flipped 10, 100, or 1,000 times, the median subject thinks that the probability of getting

[1]This chapter is joint work with Matthew Rabin, University of California — Berkeley, and Daniel J. Benjamin, Cornell University. We thank Herman Chernoff, Botond Koszegi, Don Moore, Ted O'Donoghue, Larry Phillips, and Daniel Read and seminar participants at UC Berkeley, Cornell University, Koc University, the LSE-WZB Conference on Behavioral IO, and the Behavioral Economics Annual Meeting for helpful comments. We are grateful to Ahmed Jaber, Greg Muenzen, Desmond Ong, Nathaniel Schorr, Dennis Shiraev, Josh Tasoff, Mike Urbancic, and Xiaoyu Xu for research assistance. Raymond thanks the University of Michigan School of Information'sSocio-Technical Infrastructure for Electronic Transactions Multidisciplinary Doctoral Fellowship funded by NSF IGERT grant #0654014 for financial support.

[2]NBLLN is pronounced letter by letter, said with the same emphasis and rhythm as "Ahmadinejad."

[3]Kahneman & Tversky (1972) associated such sample-size neglect with other biases under the rubric of "representativeness," which are a range of ways that people reason statistically using judgments of similarity between hypotheses and observed evidence. We relate our modeling approach to representativeness and sample-size neglect in Section 7.

between 45% and 55% heads is about .21, and of getting between 75% and 85% heads is about .05.[4] These beliefs are close to right for the sample size of 10, where the correct answers are .25 and .04. But for the sample size of 1,000, there is over a .99 chance of between 45% and 55% heads and a negligible chance of between 75% and 85% heads.

In Section 2, we develop our model of non-belief in the Law of Large Numbers in a simple setting, where a person is trying to predict the distribution of—or make an inference from—a fixed sample size. Throughout, we refer to our modeled non-believer in the Law of Large Numbers as *Barney*, and compare his beliefs and behavior to a purely Bayesian information processor, *Tommy*.[5] Tommy knows the likelihood of different sample distributions of an *i.i.d.* coin biased θ towards heads will be the "θ-binomial distribution." But Barney believes that large-sample proportions will be distributed according to a "β-binomial distribution," for some $\beta \in [0, 1]$ that itself is drawn from a distribution with mean θ. This model directly implies NBLLN: whereas Tommy knows that large samples will have proportions of heads very close to θ, Barney feels that the proportions in any given sample, no matter how large, might not be θ.[6] Even short of the large-sample limit, the model generates some properties largely in line with both the existing empirical evidence and the basic psychology others have proposed for NBLLN. For instance, while the two have identical beliefs about sample sizes of 1, Barney's beliefs about sample proportions will be a mean-preserving spread of Tommy's for samples of two or more signals. Although it embeds *some* sensitivity to sample sizes, the model largely reflects the "universal sample distribution" intuition from Kahneman & Tversky (1972).

Section 2 then draws out the implications of our simple model for *inference*. We show that if Barney applies full Bayesian updating based on his wrong beliefs about the likelihood of different sample realizations, NBLLN implies under-inference: for any priors and any sample size, Barney's expected posterior ratio on different hypotheses will be less extreme on average than Tommy's. Most importantly, for any non-extreme proportions of signals—including the limit proportion generated by the true state—Barney always fails to become fully confident even after infinite data.

In Section 3, we illustrate one economic implication of NBLLN by demonstrating the important role it likely plays in monetary-risk attitudes. In particular, Barney believes that the risk associated

[4]These numbers come from eyeballing histograms contained in the paper.

[5]"Tommy" is the conventional designation in the quasi-Bayesian literature to refer somebody who updates according to the dictums of the Reverend Thomas "Tommy" Bayes.

[6]Although this modeling gimmick is "as if" Barney is unsure that the rate is θ, true parameter uncertainty is not at all our interpretation. Instead, consistent with the underlying NBLLN psychology, we interpret it as Barney's belief that even his certainty that the underlying rate is θ is not a guarantee that the proportion in very large samples will approximate θ. In the dynamic model we develop below, this interpretation will be much more than an aspiration to get the psychology right, but an integral part of the formal model. For instance, in situations where Barney must predict further signals after observing his first 100 signals, we assume his expected proportions are still θ, rather than being influenced by the first signals as the parameter-uncertainty interpretation would suggest.

with a large number of independent gambles is greater than it actually is. This magnifies aversion to repeated risks, whether that risk aversion is due to diminishing marginal utility of wealth or (more relevantly) reference-dependent risk attitudes. Because he does not realize that the chance of aggregate losses becomes negligible, Barney may refuse to accept even infinite repetitions of a small, better-than-fair gamble. This refusal contradicts any plausible preferences combined with belief in the Law of Large Numbers, and yet reflects both the observed choices, most notably in Benartzi & Thaler (1999), who demonstrate clearly the role of what we are calling NBLLN, and the basic psychology of attitudes toward independent risks.

Both Sections 2 and 3 analyze a model of NBLLN when there is a single, "given" sample that Barney will observe. Yet information does not always arrive in a single package of signals. A person may hear a series of individual reports from friends about their car experiences, while also reading large-sample statistics of car performance. If Barney pools each of his friends' tales with the statistics from *Consumer Reports*, his inferences will be very different than if he separately updates his beliefs following each of his friend's stories, and then treats the *Consumer Reports* data as one big sample. Such cases confront us with a conceptual challenge intrinsic to the very nature of NBLLN: because people under-infer worse for larger samples than smaller ones, they will infer differently if they lump observations together versus separately. A model of NBLLN must involve a theory of how Barney groups information as a function of how it is presented to him and other features of his decisionmaking environment. With little empirical research to guide us, in Section 4 we discuss and formalize various combinations of assumptions on how Barney "retrospectively groups" signals—how he interprets evidence once he sees it—and "prospectively groups" signals—how he predicts ahead of time he will interpret evidence he might observe in the future. Different combinations of assumptions may be warranted by different perceptual, framing, and decisionmaking environments. Of special interest is the possibility that Barney retrospectively groups signals differently than he prospectively anticipates he will. He may, for instance, plan to separately ask friends about their experiences, and prospectively focus on each conversation as if it is a separate signal; but then in retrospect, he may pool the conversations together as a large sample.

In Section 5, we explore Barney's behavior in various economic environments involving learning and inference, all of which reflect the central fact that Barney fails to reach appropriately strong confidence based on extensive evidence. When Barney decides as he goes along whether to gather more information, like Tommy, he will plan to quit his costly search once he reaches some threshold of confidence. And, like Tommy, he may stop gathering signals very quickly if information is decisive. But because Barney tends to infer far less from signals than Tommy, when the initial signals are mixed, Barney may continue trying to learn even after many signals. Indeed, if Barney prospectively anticipates separately updating using each arriving signal but actually pools them retrospectively, Barney may become stuck in a "learning trap": he persistently expects to soon be confident enough to stop experimenting, or buying information, but, because he never achieves the confidence he anticipates, continues his costly efforts forever. In such cases, variants of NBLLN predict not only that people will never figure out the truth when they rationally should, but that their efforts to learn may be enormously costly.

In Section 6 we explore possible implications of NBLLN in an environment where Barney learns from observing both his own information directly and also from the actions of others with private information. The most celebrated result in research on this topic is that a sequence of Tommys who each get independent information and observe the full sequence of actions by others who choose before them will surely end up in an "information cascade" in which each ignores his own signal and instead mimics the actions of those he observes. Without thoroughly addressing modeling issues that arise in multi-person settings, we draw out some implications of NBLLN for herding based on a plausible strategic extension of Barneyness. Barney's conservatism in updating means herding is likely to be slower for a society of Barneys than for one of Tommys. But a qualitative difference also arises: although in all cases information cascades occur with positive probability, there are also parameter values where there is a positive probability that a sequence of Barneys will never enter a cascade. If early signals do not generate strong evidence in favor of one action, Barneys may end up in an "information eddy," in which each agent always chooses the action corresponding to his own signal.

In Section 7, we discuss why our model might be more compelling than alternative possible explanations and modeling approaches—both fully rational and not fully rational—that might seem to accommodate the psychology evidence. Our model of course ignores other important departures from Bayesian inference—such as base-rate neglect and belief in the Law of Small Numbers—that seem separable from NBLLN. But it also omits features—including the psychophysics of diminishing sensitivity, as well as unwillingness to hold or express extreme beliefs—that, as alternative sources of under-inference from large samples, are less separable. In Appendix A, in fact, we present a (complicated) formal model embedding some of these other errors along with NBLLN. Guided by this formal model, in Appendix B we attempt to give a fairly exhaustive review of the empirical work on sampling predictions and inference. We believe this review makes clear that our model of NBLLN is capturing a broad empirical reality. Although in many settings there are alternative reasons why a person may rationally disregard a large amount of evidence—e.g., because the evidence is less relevant given the person's preferences than other, smaller-sample sources of information—our review documents NBLLN only in settings where a Bayesian would fully attend to the evidence. But NBLLN is not a phenomenon that has been widely embraced or emphasized by judgment researchers or behavioral economists. We surmise that this is largely because findings of under-inference have been associated with an interpretation called "conservatism" (e.g., Edwards, 1968)—namely, that people tend not to update their beliefs as strongly as Bayesian updating dictates—that does not mesh comfortably with base-rate neglect and other biases that often imply that people infer more strongly than Bayesian. In our view, summarizing people as overly conservative or not conservative enough is manifestly the wrong way to parse human judgment. By focusing on the concrete biases at play and highlighting the co-existence of NBLLN with other biases, we hope to make clear that there is no contradiction. Indeed, as we discuss and illustrate below, NBLLN is often a necessary "enabler" of other psychological biases, including biases where people over-infer from some information. For example, another bias is that people seem to draw excessively strong conclusions from "vivid" signals. But even if observing a traffic accident first-hand has the impact of 100 data points on a person's beliefs about the likelihood of an accident, such a vivid signal would affect beliefs little unless the person *also* drew too weak an inference from the newspaper's report of the past year's accident statistics due to NBLLN. As another example, Barney's failure to learn in an experimentation setting may help explain why people fail to fully learn the extent of their own present bias even after observing their own behavior in many decisionmaking contexts over many years.

Because we think that NBLLN is likely to matter for economic theory, we strive in this paper to develop a model that is both specific in its implications and general in its applicability. We believe the theorems and other results we establish below—in general and in the specific economic applications we study—makes clear its potential to provide an improvement over full rationality models while replicating as much as possible the specificity of the full-rationality predictions. In the spirit of seeing our model as an iterative step towards improving the psychological realism of formal economic models, we highlight throughout the paper cases where we believe our assumptions are restrictive and fail to capture psychological intuitions or existing evidence, and we outline in Appendix A ways that the model can be combined with other psychological biases. Also in this spirit, we conclude Section 7 by briefly discussing an approach to generalizing our model to contexts where Barney observes non-binomial or non-*i.i.d.* data, which would be a crucial step to incorporating NBLLN more broadly into economic models.

3.2 The Single-Sample Model

Throughout the paper, we study a stylized setting where an agent observes a set of binary signals, each of which takes on a value of either a or b. Given a *rate* $\theta \in \Theta \equiv (0,1)$, signals are generated by an *i.i.d.* process where the probability that any given signal is an a is equal to θ. Signals arrive in *clumps* of size N. We denote the set of possible ordered sets of signals of size $N \in \{1,2,...\}$ by $S_N \equiv \{a,b\}^N$, and we denote an arbitrary clump (of size N) by $s \in S_N$.[7] Let A_s denote the total number of a's that occur in the clump $s \in S_N$, so that $\frac{A_s}{N}$ is the proportion of a's that occur in a clump of N signals. For a real number x, we will use the standard notations "$\lceil x \rceil$" to signify the smallest integer that is weakly greater than or equal to x and "$\lfloor x \rfloor$" to signify the largest integer that is weakly less than or equal to than x.[8] For any random variable y that takes as possible values the elements of set Y, let beliefs by Tommy (the perfect Bayesian information processor) be denoted by probability density function $f_Y(\cdot)$, implying cumulative distribution function $F_Y(\cdot)$, expectation $E_Y(\cdot)$, and variance $Var_Y(\cdot)$. Let corresponding beliefs by Barney (the non-believer in the LLN) be denoted by $f_Y^{\psi}(\cdot)$, $F_Y^{\psi}(\cdot)$, $E_Y^{\psi}(\cdot)$, and $Var_Y^{\psi}(\cdot)$, where ψ is a parameter for the degree of NBLLN defined below.

In this section, we develop our model of Barney for the case where he is considering a single clump of N signals. This case corresponds to most of the experimental evidence about NBLLN, which has been collected in settings where subjects were presented with a single, fixed sample of signals or outcomes, where subjects presumably process all the information together. This special case also allows us to lay bare the essential features of how our model captures NBLLN. When generalizing the model, complicating conceptual challenges arise about which there is little existing evidence. Some of our analysis in fact concerns precisely these complications, but we defer discussion of these issues and ways to handle them until Section 4.

According to the Law of Large Numbers, the mean of a random sample equals the rate with probability 1 in the limit as the sample size gets large: For any interval $[\alpha_1, \alpha_2] \subseteq [0,1]$,

$$\lim_{N\to\infty} \sum_{x=\lfloor \alpha_1 N \rfloor}^{\lceil \alpha_2 N \rceil} f_{S_N|\Theta}(A_s = x|\theta) = \begin{cases} 1 & \text{if } \theta \in [\alpha_1, \alpha_2] \\ 0 & \text{otherwise} \end{cases}.$$

How might we capture the possibility that Barney believes (say) that it is reasonably likely that

[7] Note that we forego the conventional strategy of providing notation for a generic signal, indexed by its number. It is less useful here because (within a clump) what matters to Barney is just the number of a signals, not their order. Conserving notation here, in Section 4 we use t to index the clumps of signals.

[8] Although our formal results and proofs fully attend to "integer issues," we ignore them when providing intuition.

at least 600 of 1000 births at a hospital in a given year are boys, even though he knows that boys are born at a rate of 50%? Our key modeling gambit is to assume that Barney believes samples are generated as if a rate of θ, here 50%, means that the rate is θ on average, but might be higher or lower for any given sample. For a given true rate θ, we model Barney as believing that for the sample he is considering: first, a "subjective rate" $\beta \in [0, 1]$ is drawn from a distribution centered at θ. Then the *i.i.d.* sample of 1000 babies is generated using rate β. The key implication is that if a given value of β were the actual rate, it would (by the Law of Large Numbers!) exactly determine the proportion of signals in the limit of a very large sample. Therefore, the probability density that Barney assigns to any proportion β of signals (say, 60% of babies are boys) in a large sample is equal to the probability density that Barney assigns to the possibility that β equals that value.

Formally, we assume that when Barney knows the rate is θ, he believes that signals are generated by an *i.i.d.* process with *subjective rate* $\beta \in [0, 1]$, where β is drawn from a density $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$—which we refer to as Barney's *subjective rate distribution*—that has the following properties.

A1. *For all β and θ, $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$ has full support on $(0, 1)$, is absolutely continuous in β, and is point-wise continuous in θ.*

A2. *For all β and θ, $F^{\psi}_{\text{ß}|\Theta}(\beta|1-\theta) = 1 - F^{\psi}_{\text{ß}|\Theta}(1-\beta|\theta)$.*

A3. *For all θ and $\theta' > \theta$, $\frac{f^{\psi}_{\text{ß}|\Theta}(\beta|\theta')}{f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)}$ is increasing in β.*

A4. *For all β and θ, $E^{\psi}_{\text{ß}|\Theta}(\beta|\theta) = \theta$.*

Assumptions A1 and A2 are mild and consistent with the empirically-observed densities of people's beliefs about the distribution of signals in large samples. Assumptions A3 and A4, however, are substantive assumptions that do not follow easily from the psychology. Assumption A3 is a monotone likelihood ratio property: fixing any two rates, Barney believes that the likelihood of drawing any particular subjective rate given the high rate relative to the low rate is greater the higher is the high rate. This is a substantive assumption,[9] and it is easy to imagine specifications of $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$—especially in the spirit of the type of diminishing-sensitivity evidence discussed in Appendix A—that would violate it. But it is in accord with the most directly relevant evidence, namely Griffin & Tversky's (1992) Study 3, which examines a range of parameters of the sort that seems most likely for violating A3.[10] It holds for our main example of the beta distribution (discussed

[9] For example, Whitt (1979) shows that this monotone likelihood ratio property implies: for all β, θ, and $\theta' > \theta$, $F^{\psi}_{\text{ß}|\Theta}(\beta|\theta')$ first-order stochastically dominates (FOSD) $F^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$. This FOSD relationship is itself useful in proving some of our results.

[10] In particular, Griffin & Tversky asked subjects to infer the likelihood that a coin is biased $\theta_A = .6$ in

below) and more generally is useful for establishing some of our results. Especially because the range of samples for which it is potentially false are inherently very unlikely, we think it is probably not an important caveat to our results. Assumption A4 says that the mean of Barney's subjective rate distribution is the known objective rate. Although we rely on it extensively in the analysis, it is in fact violated in existing data; we discuss below both the violations and the role it is playing in our analysis.

When Barney knows the rate is θ, he believes the likelihood of observing a particular clump of N signals, $s \in S_N$, is

$$f^{\psi}_{S_N|\Theta}(s|\theta) = \int\limits_{\beta\in[0,1]} f_{S_N|\text{ß}}(s|\beta)\, f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)\, d\beta, \tag{3.1}$$

where $f_{S_N|\text{ß}}(s|\beta)$ is the (correct) probability of observing s *if the rate were* β, and this is averaged over the density of subjective rates, $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$. Consequently, Barney's belief that a large sample will have a proportion of a signals in some range $[\alpha_1, \alpha_2]$ is exactly equal to Barney's belief that the subjective rate β is in that range.[11]

Lemma 1. *Barney does not believe in LLN: for any interval* $[\alpha_1, \alpha_2] \subseteq [0,1]$,

$$\lim_{N\to\infty} \sum_{x=\lfloor\alpha_1 N\rfloor}^{\lceil\alpha_2 N\rceil} f^{\psi}_{S_N|\Theta}(A_s = x|\theta) = F^{\psi}_{\text{ß}|\Theta}(\beta = \alpha_2|\theta) - F^{\psi}_{\text{ß}|\Theta}(\beta = \alpha_1|\theta) > 0.$$

Because we assume that Barney's beliefs about the distribution of β puts positive probability density on the entire interval $(0,1)$, the subjective-rate modeling gimmick captures the essence of our interpretation of NBLLN: Barney believes that the proportion of heads from flipping a coin *known to be fair* may not be 50% in any given sample, no matter how large. In keeping with this interpretation, Barney does not believe the realized β is a real feature of the coin, and is certainly not an object he makes inferences about. It is a representation of Barney's subjective uncertainty that the θ will manifest itself in a given sample. In the analysis of Barney's predictions about a single clump of signals, our model is mathematically equivalent to a fully rational model where the agent

favor of heads rather than $\theta_B = .25$ in favor of heads, depending on different possible outcomes from flipping the coin 12 times. There is a psychological (but statistically erroneous) intuition that extreme samples, like 10 heads out of 12, seem so unexpected in the case of either rate that they do not provide strong evidence about which is rate is generating the flips. This diminishing-sensitivity intuition might lead people to think that the evidence in favor of the .6-biased coin is stronger when the observed sample is 7 heads out of 12 than when it is 10 heads out of 12—in violation of the implications of Assumption A3 for inference. However, consistent with A3, Griffin & Tversky find that subjects' posteriors beliefs in favor of the .6-biased coin are monotonically increasing in the number of heads.

[11] All proofs are elevated to Appendix C.

is uncertain about the underlying rate. But both the inference results later in this section and Section 4's dynamic setting demonstrates how different these two interpretations are: if an agent were rationally uncertain about the underlying rate, then in large samples, he would (quickly) become certain about its value. In contrast, NBLLN means that the agent acts as if he is uncertain about the underlying rate even when he knows what the rate is or has seen enough data that he should be sure about it.

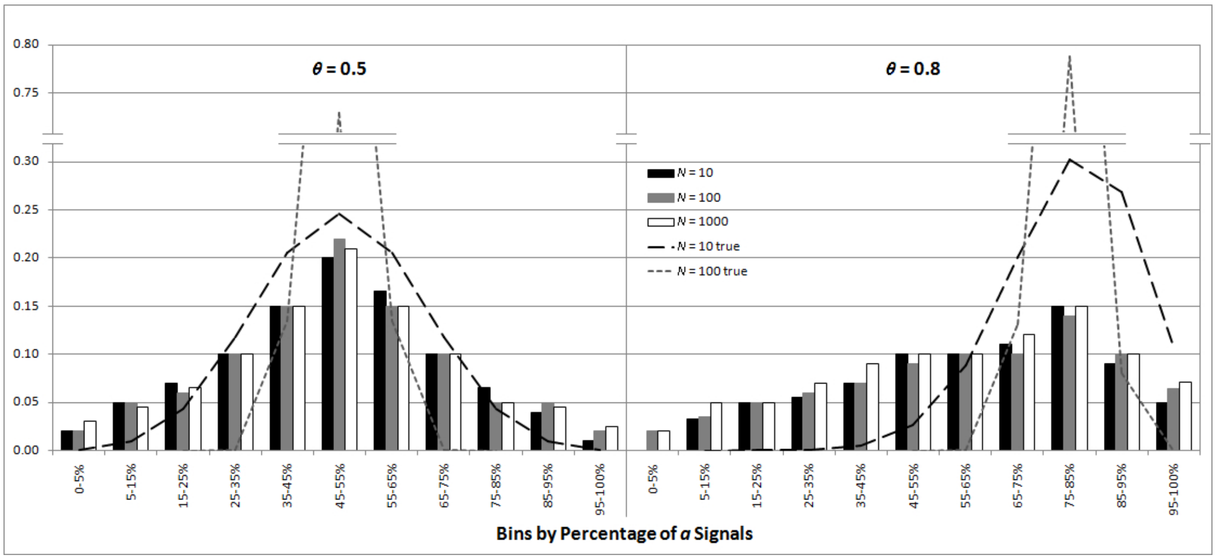

Figure 3.1: Evidence from Kahneman & Tversky (1972)

Since (by the Law of Large Numbers) Barney's belief about the distribution of signals in large samples coincides with his subjective rate distribution, the most appropriate density of $\beta|\theta$ would correspond to the empirical beliefs in studies such as those illustrated in Figure 3.1, drawn from Kahneman & Tversky (1972). The black, gray, and white bars—which correspond to people's reported beliefs regarding samples of size 10, 100, and 1000, respectively—lie on top of each other. This distribution that is common across sample sizes, which presumably corresponds to people's large-sample beliefs about sample proportions, could be directly assumed to be the density of $\beta|\theta$.

Although Assumptions A1 and A2 are consistent with Figure 3.1, A4 is not. Beliefs for $\theta = .5$, depicted in the left panel, naturally have mean approximately equal to .5. However, beliefs for $\theta = .8$, depicted in the right panel, have mean approximately equal to .6. The mean of the distribution of signals is displaced toward .5 apparently because the long tail of the distribution is fat. As we discuss in Appendices A and B, we believe that the fatness of the tail is in turn due to *flatness* of the tail. We attribute the flatness to a psychological bias—"sampling-distribution-tails diminishing sensitivity"— in which people perceive very unlikely outcomes as similar to each other and hence similar in probability. We justify A4 on two grounds: analytical convenience, and our contention that the violation of the assumption is best understood as a distinct psychological bias modeled in Appendix A. Nonetheless, because we omit flat tails from the model, the model will not match some features of the empirical evidence, especially when the agent observes an extreme sample.[12]

A *subjective sampling distribution* specifies an agent's belief about the likelihood of each possible combination of signals when the rate θ is known. Whereas Lemma 1 shows that Barney's subjective sampling distribution (for the number of a-signals) in the large-sample limit equals his "subjective rate distribution," Proposition 1 shows that NBLLN also has implications for Barney's finite-sample subjective sampling distributions.

Proposition 1. *For any $\theta \in (0,1)$ and $N \in \{1, 2, ...\}$:*

1. $E^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) = E_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) = \theta$.
2. $F^{\psi}_{S_N|\Theta}(A_s|\theta)$ *second-order stochastically dominates (SOSD)* $F_{S_N|\Theta}(A_s|\theta)$*, and* $Var^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) \geq Var_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right)$ *with strict inequality for* $N > 1$.
3. $Var^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right)$ *is strictly decreasing in* N*, and* $\lim_{N\to\infty} Var^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) > 0$.
4. $F^{\psi}_{S_N|\Theta}(A_s|\theta')$ *first-order stochastically dominates (FOSD)* $F^{\psi}_{S_N|\Theta}(A_s|\theta)$ *whenever* $\theta' > \theta$.

Part 1 states that Barney, like Tommy, expects the average proportion of a's in the sample to be θ. An immediate and important corollary of Part 1 is that Barney's beliefs coincide with Tommy's when $N = 1$. Part 2 states that Barney has a more-dispersed subjective sampling distribution than Tommy in the sense of second-order stochastic dominance. Combined with the fact that the mean of Barney's subjective sampling distribution is the same as Tommy's, this implies that Barney's subjective sampling distribution is a mean-preserving spread of Tommy's. This naturally implies that the variance of Barney's subjective sampling distribution is larger than Tommy's. Part 3 states that although the variance of Barney's subjective sampling distribution (for the sample proportion) is strictly decreasing in N, it is bounded away from 0. Part 4 states that a higher true rate, besides generating a higher mean proportion of a signals, generates a rightward shift in Barney's entire subjective sampling distribution in the sense of first-order stochastic dominance.

[12]Analogously, we justify omitting the Law of Small Numbers (LSN) from the model for analytical convenience, and our contention that it is a distinct bias. However, our grounds for not incorporating LSN into the model are stronger because LSN will matter most when NBLLN will matter least—when the observed sample is small—whereas sampling-distribution-tails diminishing sensitivity is, like NBLLN, an especially extreme bias when the observed sample is large.

We now turn to *inference* problems, where an agent with prior beliefs must infer from observed signals what the underlying rate is. An example of an inference problem is determining the likelihood that a coin is head-biased rather than tail-biased, after observing a sample of coin flips. Let $\Theta \subseteq (0,1)$ denote the set of rates that have positive prior probability. For simplicity, we assume Θ is a finite set. Without loss of generality, we consider the agent's beliefs about the relative likelihood of two of the rates $\theta_A > \theta_B$, given priors $0 < f_\Theta(\theta_A), f_\Theta(\theta_B) < 1$ with $f_\Theta(\theta_A) + f_\Theta(\theta_B) \leq 1$.

We maintain the standard assumption that an agent draws inferences by applying Bayes' Rule to his subjective sampling distributions.[13] Consequently, Barney's beliefs after observing a particular clump $s \in S_N$ are $f^{\psi}_{\Theta|S_N}(\theta_A|s) = \frac{f^{\psi}_{S_N|\Theta}(s|\theta_A)f_\Theta(\theta_A)}{\sum_{\theta\in\Theta} f^{\psi}_{S_N|\Theta}(s|\theta)f_\Theta(\theta)}$ and $f^{\psi}_{\Theta|S_N}(\theta_B|s) = \frac{f^{\psi}_{S_N|\Theta}(s|\theta_B)f_\Theta(\theta_B)}{\sum_{\theta\in\Theta} f^{\psi}_{S_N|\Theta}(s|\theta)f_\Theta(\theta)}$. Because Barney updates Bayesianly given his model of the data-generating process, his subjective beliefs satisfy the "Law of Iterated Expectations": Barney expects that for any sample size, the mean of his posterior beliefs will equal the mean of his prior beliefs: for any $N \geq 1$, $E^{\psi}_{\Theta|S_N}(\theta|s) = E^{\psi}_{\Theta}(\theta)$.

Due to the LLN, after observing a sufficiently large number of signals, Tommy will be arbitrarily close to knowing the true rate with certainty. In contrast, the central implication for inference of Barney's NBLLN—which plays a large role in many of the applications later in this paper—is that Barney remains uncertain even after observing an *infinite* number of signals. To boot:

Proposition 2. *Let $\theta \in (0,1)$ be the true rate. Then for any $\theta_A, \theta_B \in \Theta$, Barney draws limited inference even from an infinite sample: as $N \to \infty$,*

$$\frac{f^{\psi}_{\Theta|S_N}(\theta_A|s)}{f^{\psi}_{\Theta|S_N}(\theta_B|s)} \to_{a.s} \frac{f^{\psi}_{\mathbb{B}|\Theta}(\beta = \theta|\theta_A)}{f^{\psi}_{\mathbb{B}|\Theta}(\beta = \theta|\theta_B)} \frac{f(\theta_A)}{f(\theta_B)}, \tag{3.2}$$

which is a positive, finite number.

Because Barney's asymptotic sampling distribution coincides with the subjective-rate distribution, his limit inference depends on the relative heights that the pdfs of the subjective-rate distributions for θ_A and θ_B assign to the proportion θ of a's. Since the subjective-rate distributions put positive density on every proportion in $(0,1)$, Barney's likelihood ratio will be finite. An immediate but important corollary of Proposition 2 is that Barney's priors affect his beliefs even in the limit of an infinite sample.

[13] Appendix A discusses and formalizes how to combine NBLLN with an important way people do not draw inferences in accordance with Bayes' Rule: they underweight priors according to the well-established notion of base-rate neglect. In the main model, however, we maintain the standard assumption of Bayesian inference both to highlight the role played *per se* by NBLLN, and because (as we discuss in Appendix B) our reading of the experimental evidence is that *except* for base-rate neglect, people's inferences are actually well-approximated by Bayes' Rule applied to their subjective sampling distributions. Hence in applications where priors are equal, base-rate neglect is neutralized as a factor. In applications where we do not assume equal priors, however, base-rate neglect could modify some of our results.

Proposition 2 shows that, like Tommy's, Barney's limit posterior is a deterministic function of the true rate, but unlike Tommy's, Barney's limit posterior ratio is bounded away from 0 and ∞. Not only will Tommy learn the true rate for sure after observing a sufficiently large number of signals, Tommy also correctly *anticipates* that a sufficiently large number of signals will make him certain of the true rate. In contrast, Barney does not even realize that his limit posterior is a deterministic function of the true rate. Barney mistakenly thinks that his posterior probability of rate θ_A after observing an infinite number of signals is a random function of the true rate. The reason is that, even though Barney knows that his inferences in a large sample will be determined by the proportion of a's, he incorrectly thinks the proportion of a's is determined by the subjective rate, which is a random function of the true rate.

Proposition 3. *Fix any $\theta_A > \theta_B \in \Theta$. Barney believes that, regardless of what the true rate $\theta \in (0,1)$ is, as $N \to \infty$, his posterior ratio $\frac{f^{\psi}_{\Theta|S_N}(\theta_A|s)}{f^{\psi}_{\Theta|S_N}(\theta_B|s)}$ will converge in distribution to a random variable whose support is a simply connected subset of $[0,\infty]$. Moreover, he believes that this subset of $[0,\infty]$ has strictly positive measure.*

The upper bound on Barney's anticipated limit posterior corresponds to what Barney's limit posterior would be if he observed all a-signals; the lower bound corresponds to his limit posterior if he observed all b-signals.

Because Barney's subjective sampling distribution is correct when the sample size is 1, he will draw correct inferences in that case. Barney will also draw correct inferences when exactly half the signals are a-signals. In that case, the sample is uninformative for both Barney and Tommy, and neither updates his beliefs about the rate.

Proposition 4. *Fix rates $\theta_A > \theta_B$. For $N = 1$, Barney and Tommy infer the same. If $\theta_A = 1-\theta_B$, then for any set of $N \in \{1, 2, ...\}$ signal realizations $s \in S_N$, neither Tommy's beliefs nor Barney's beliefs change from the priors when $\frac{A_s}{N} = \frac{1}{2}$.*

For further analysis of the finite-sample case—as well as for theoretical applications and empirical analysis—it is useful to have a parametric model of Barney's subjective rate distribution. For some of our results, we impose a particular functional form for the subjective rate distribution, a beta distribution:

$$f^{\psi}_{\mathbb{B}|\Theta}(\beta|\theta) = \beta^{\theta\psi-1}(1-\beta)^{(1-\theta)\psi-1}\frac{\Gamma(\psi)}{\Gamma(\theta\psi)\Gamma((1-\theta)\psi)}, \tag{3.3}$$

where $0 < \psi < \infty$ is the exogenous parameter of the model, and $\Gamma(x) \equiv \int_{[0,\infty)} y^{x-1}e^{-y}dy$, defined on $x > 0$.[14] The properties A1-A4 are satisfied, and this family of beta densities shares many qualitative features of people's empirically observed large-sample beliefs about the distribution of signals.[15] A major advantage of this formulation is tractability: since the beta distribution is the conjugate prior for the binomial distribution, standard results from probability theory can be used to characterize Barney's beliefs.[16] "Parameterized-Barney" is more biased for smaller ψ—with more dispersed subjective sampling distributions in the sense of SOSD—and Barney coincides with Tommy in the parameter limit $\psi \to \infty$. To interpret the "Barneyness parameter," ψ, consider the problem of inferring the rate from a sample of signals. The beta density (3.3) corresponds to what a Bayesian's posterior about the rate would be if the Bayesian had begun with an improper beta prior on $[0, 1]$ corresponding to the $\psi \to 0$ limit of (3.3), and had observed a total of ψ signals with $\theta\psi$ a's and $(1 - \theta)\,\psi$ b's. For $\psi = 0$, the posterior would remain the improper prior on $[0, 1]$, while for $\psi \to \infty$, the posterior would converge to a point mass at the true rate θ. Although we do not conduct a full and careful structural estimation, using evidence from studies that elicit subjects' subjective sampling distribution as well as those that measure subjects' inference about rates after observing signal realizations, we estimate that ψ falls within a range of 7-15. To illustrate how powerful a bias NBLLN is when $\psi = 10$, it implies that Barney's beliefs about whether a coin is a 40% heads coin or a 60% heads coin are the same after an infinite number of signals of 60% heads as Tommy's are after only 6 heads and 4 tails.

This parameterized model of Barney gives a sense of magnitudes for how Barney's under-inference depends on the rates θ_A and θ_B. Suppose $\psi = 10$, Barney begins with equal priors on the two states, and the true rate is θ_A. If the difference between the rates is relatively large—with $\theta_A = 1-\theta_B = .8$—

[14] The more common way of writing this beta density is $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta) = \beta^{\theta\psi-1}(1-\beta)^{(1-\theta)\psi-1}\frac{(\psi-1)!}{(\theta\psi-1)!((1-\theta)\psi-1)!}$. Our formulation is equivalent, except it allows for non-integer values of ψ. Recall that the Gamma Function, $\Gamma(x)$, is the standard generalization of the factorial function: it has the properties that $\Gamma(x+1) = x\Gamma(x)$ and $\Gamma(1) = 1$, so that for any positive integer x, $\Gamma(x) = (x-1)!$.

[15] We prove that properties A1-A4 are satisfied in Lemma $\beta 5$ in Appendix C. Although we do not assume it in A1-A4, the available empirical evidence is also consistent with the subjective rate distribution being single peaked. Single peakedness, however, is not satisfied for all parameter values of the beta distribution. If $\theta\psi$ and $(1-\theta)\psi$ are both less than 1, the beta distribution has a first derivative with respect to β that is increasing everywhere—the pdf has a minimum in the interior of $(0, 1)$ and rises as it approaches both 0 and 1.

[16] Unfortunately, the functional form (3.3) has a few implications about asymmetric inference (i.e., inference problems where $\theta_A \neq 1 - \theta_B$) that can matter in applications but that do not have general intuitions related to NBLLN: (1) In asymmetric-inference problems, parameterized Barney draws the same inferences as Tommy when the observed sequence is ab or ba; (2) In asymmetric-inference problems, when the sample is all a's or all b's, parameterized Barney always under-infers regardless of the specific values of θ_A and θ_B; and (3) For any true rate $\theta \in (0, 1)$ and for any $\theta_A, \theta_B \in \Theta$, Barney believes that his limit posterior ratio, $\lim_{N\to\infty}\frac{f^{\psi}_{\Theta|S_N}(\theta_A|s)}{f^{\psi}_{\Theta|S_N}(\theta_B|s)}$, will be a random variable that has support on all of $(0, \infty)$. These claims are proved in Lemma $\beta 4$ in Appendix C. Since we believe these properties would not generalize to other models that are equally consistent with existing evidence, we will avoid stating implications that rely on these properties of the functional form.

then the role of NBLLN is relatively small. In an infinite sample, Barney's subjective posterior probability of rate θ_A will converge to .9998. However, if the two rates are closer together—with $\theta_A = 1 - \theta_B = .6$—then in an infinite sample, Barney's subjective posterior probability of rate θ_A will converge to only .69. As a reminder about the role of priors, this in fact means if Barney initially had beliefs more extreme than 2.25:1 in favor of rate θ_B, he will, in an infinite sample, surely end up believing rate θ_B is more likely, even when θ_A is the true rate.

While most dramatic in large samples, NBLLN also has implications for inference in finite samples. Since Barney's subjective sampling distribution is too dispersed when $N > 1$, Barney will "generally" under-infer when the sample size is larger than 1. In order to make that claim precise, we will measure Barney's (and Tommy's) "change in beliefs" by the absolute difference between his posterior probability that θ_A is the true rate and his prior probability: $\left| f^{\psi}_{\Theta|S_N}(\theta_A|s) - f_{\Theta}(\theta_A) \right|$. We say that Barney "under-infers" relative to Tommy if Barney's change in beliefs is smaller than Tommy's.

Unlike in large samples, in small samples it is no longer universally true that Barney under-infers relative to Tommy. For particular realizations Barney can over-infer or under-infer relative to Tommy—or even infer in the opposite direction, so that a sample that causes Tommy to think rate θ_A is more likely, causes Barney to think rate θ_B is more likely![17] Nonetheless, we believe that Barney under-infers *in expectation*, taken with respect to the true sampling distribution. Proposition 5 proves this statement for the case of ψ sufficiently small, but we conjecture that it holds for any $0 < \psi < \infty$.[18]

Proposition 5. *Fix rates $\theta_A > \theta_B$ and a set of $N \in \{1, 2, ...\}$ signal realizations $s \in S_N$. If $f^{\psi}_{ß|\Theta}(\beta|\theta)$ has the functional form (3.3), then regardless of whether the true rate is θ_A or θ_B, for ψ sufficiently small, the expected change in Barney's beliefs is smaller than the expected change in Tommy's beliefs. Furthermore, suppose $\theta_A = 1 - \theta_B$. Then for any sample of $N > 1$ signals such that $\frac{A_s}{N} \neq \frac{1}{2}$ and any ψ, Barney under-infers relative to Tommy. In addition, while Tommy's inference depends solely on the difference in the number of a and b signals, Barney's change in beliefs is smaller from larger samples with the same difference.*

Intuitively, on average Barney under-infers because he partially attributes the information in the

[17]For example, using the parameterized model discussed below, set $\psi = 10$, $\theta_A = .7$, and $\theta_B = .6$, and assume equal priors on the two states. Then if the realizations of 80 signals are 53 a-signals and 27 b-signals, then Tommy believes that state A is more likely, while Barney believes state B is more likely.

[18]We have simulated Barney's and Tommy's expected change in beliefs for a range of parameter values: for each of $\psi \in \{1, 2, ..., 30\}$ and $N \in \{5, 10, 15, 20\}$, we examined each of $\theta_A, \theta_B \in \{.5, .6, .7, .8, .9\}$. We also ran a number of simulations for $\theta_A = .99$ and .999 and for $\psi = 100$. In every case we examined, Barney's expected change in beliefs was smaller than Tommy's.

realized sample to the subjective rate, rather than extracting all of the information about the true rate.

In many of the inference experiments reviewed in Appendix B and in many of our applications involving inference, the two rates are "symmetric" in the sense that $\theta_A = 1 - \theta_B$, e.g., an urn might have either 60% red balls or 40% red balls. In this case, Proposition 5 shows that stronger comparisons can be made between Barney and Tommy: as long as the realized sample is informative, parameterized-Barney will under-infer, not just in expectation. Proposition 5 also notes a key feature of Barney's updating that shows how it leads to a bias towards "proportional thinking" in inference along the lines suggested by researchers such as Griffin & Tversky (1992). Consider samples where the difference between the number of a-signals and the number of b-signals is the same, e.g., aa and $abababaa$. Tommy will draw the same inference from the two samples. But because his asymptotic sampling distributions depend on the proportion of a and b signals rather than their number, Barney infers less from the larger sample.

3.3 Application: Perceived Aggregate Risk

An economics professor at MIT once told his colleague Paul Samuelson that, whereas he would reject a bet for even odds to gain \$200 or lose \$100, he would accept 100 repetitions of that bet. Even though such behavior sounds reasonable to most of us, Samuelson (1963) reports the conversation so as to prove that it violates classical expected-utility theory. That is, a Tommy with classical expected-utility preferences (defined over final wealth) who does not exhibit unrealistically large wealth effects should be willing to take a single bet if and only if he is willing to take $N \geq 1$ independent plays of that bet. Intuitively, in the absence of wealth effects, preferring $K+1$ bets to K bets is the same thing as preferring 1 bet on top of any realization of the K bets. By induction, preferring to take any positive number of the bets is the same as preferring to take one bet.

Yet, it is not just the "switching" that violates classical expected utility preferences, but the aversion to the single bet to begin with. Rabin (2000a, 2000b) and Rabin & Thaler (2001) have followed others in noting that the degree of concavity required for classical expected-utility preferences, defined over wealth, to explain risk-averse behavior over the small stakes involved in a single play of the gamble is calibrationally implausible. Loss aversion—the tendency to feel a loss more intensely than an equal-sized gain—explains why the majority of people who turn down the one-shot gamble do so.[19] A Tommy with a simple, piecewise-linear loss-averse utility function

$$u(w_0, z) = \begin{cases} w_0 + z & \text{if } z \geq 0 \\ w_0 + \lambda z & \text{if } z < 0 \end{cases}, \tag{3.4}$$

where w_0 is initial wealth and z is a monetary gain or loss, will refuse the one-shot bet as long as the coefficient of loss aversion, λ—typically estimated to be around 2.25 (e.g., Tversky and Kahneman (1991))—is greater than 2.[20] However, a Tommy with typical loss-averse preferences would be extremely happy to accept 100 repetitions of the same gamble: while the expected gain is \$5,000, the chance of a net loss is only 1/700, and the chance of losing more than \$1,000 is only 1/26,000.

Despite being loss averse enough to turn down the one bet, nobody with fully rational beliefs would turn down 100 repetitions of this bet. Yet, unlike Samuelson's colleague, many people would do so! In hypothetical questions from one study in Benartzi & Thaler (1999), for instance, 36% of participants said they would turn down a single scaled-down Samuelson type bet (win \$100 or lose \$50); but fully 25% also reject the 100-times repeated gamble.[21] NBLLN helps explain why many people turn down these gambles.[22] Barney exaggerates the probability that the repeated bet will turn out badly. Indeed, Benartzi & Thaler report evidence that NBLLN is implicated in these choices: when asked the probability of losing money after 150 repetitions of a 90%/10% bet to gain \$0.10/lose \$0.50, 81% percent of subjects overestimated the probability—and by an enormous margin. While the correct answer is .003, the average estimate was .24.[23] To show that subjects' mistaken beliefs were driving their choices, Benartzi & Thaler compared subjects' willingness to accept the repeated bet with their willingness to accept a single-play bet that had the histogram

[19] Samuelson himself had speculated that it was the willingness to accept repeated plays of the bet that was the mistake, rather than the refusal to accept a single gamble. Samuelson's conjecture that his colleague's willingness to accept the repeated gamble was the result of a "fallacy of large numbers"—a mistaken belief that the riskiness of the gamble evaporates with a sufficiently large number of repetitions—is the *opposite* of NBLLN, and is contradicted by Benartzi and Thaler's (1999) evidence, reported below, that people exaggerate the probability of a loss in the repeated bet.

[20] Although essentially correct for small gambles, assuming linear consumption utility can become problematic if bets are repeated so many times as to involve large amounts of wealth. However, if in our limit results below, we halve the stakes every time we double the number of repetitions, the linearity assumption is unobjectionable.

[21] In two other subject pools, they find 34% and 23% turn down a simple \$20/\$10 gamble, and *more* people—57% and 50%—turn down the repeated gamble. Keren (1991) finds similar results in incentivized single bets vs five-times-repeated bets; for related hypothetical evidence, see Keren & Wagenaar (1987) and Redelmeier & Tversky (1992).

[22] Our emphasis on how NBLLN helps explain why loss-averse people turn down the repeated bet is because of its calibrational relevance, but it may be worth noting that NBLLN also has implications for how expected-utility-over-wealth agents respond to repetitions of bets. We can extend the "if" part of Samuelson's theorem: if Barney rejects a bet at all initial wealth levels w_0, then he would also reject any $N \geq 1$ independent plays of that bet. The "only if" direction does not extend, and a Barney who is just indifferent between accepting and rejecting a simple bet would, because he exaggerates the risk, strictly prefer to reject repeated versions of the gamble.

[23] Results were essentially identical when subjects were asked the mirror-image problem, the probability gaining money after 150 repetitions of a 10%/90% bet to gain \$0.50/lose \$0.10.

of money outcomes implied by the repeated bet. While only 49% of the college-student subjects accepted 150 repetitions of the bet, 90% accepted the equivalent single-play bet, suggesting that the repeated bet would have been very attractive if subjects had correctly understood the distribution of outcomes.[24,25,26]

Formally, while a loss-averse Tommy will always accept a better-than-fair bet if it is repeated enough times, a loss-averse Barney may—depending on how favorable the bet is and how loss-averse he is—turn down an infinitely-repeated bet.

Proposition 6. *Suppose Barney and Tommy have simple, piecewise-linear loss-averse preferences as specified in (3.4). Fix any gamble* (θ, h, t)*, paying off* $h > 0$ *with probability* θ *and* $-t$ *with probability* $1-\theta$*, that is better than fair:* $\theta h > (1-\theta)\,t$*. For any* $\lambda \geq 1$*, there is some* $N' \geq 1$ *such that if* $N > N'$*, then Tommy will accept* N *repetitions of the gamble. In contrast, for Barney there is some threshold level of loss aversion* $\widehat{\lambda} > 1$ *such that: if* $\lambda < \widehat{\lambda}$*, then there is some* N' *sufficiently large such that Barney will accept* N *repetitions of the gamble for all* $N > N'$*; and if* $\lambda \geq \widehat{\lambda}$*, then there is some* N'' *sufficiently large such that Barney will reject* N *repetitions of the gamble for all* $N > N''$*.*

Moreover, it is possible for Barney to exhibit the *opposite* pattern from Samuelson's colleague, accepting the single bet but rejecting the 100-times repeated bet! Consistent with this possibility, the evidence cited above from Benartzi and Thaler (1999) found behavior in the opposite direction of Samuelson's colleague in two out of their three studies.

Using the one-parameter functional form for Barney, calibrations suggests that NBLLN goes much of the way, but *not* all of the way, in explaining why 25-57% of participants turned down the repeated bets in Benartzi and Thaler's studies. For Tommy, the coefficient of loss aversion required

[24] Note also that something more than the type of "narrow bracketing" stressed by authors such as Tversky & Kahneman (1986), Kahneman & Lovallo (1993), Benartzi & Thaler (1995), Read, Loewenstein & Rabin (1999), Barberis, Huang & Thaler (2006), and Rabin & Weizsäcker (2009) seems to be playing a role. Those papers emphasize that people often react to a combination of risky bets as if they were deciding about each risky bet in isolation from all the others. While such neglect of the effects of aggregating risks may help explain why people reject the repeated gamble, it seems clear that even people who attend to the aggregate effects *misunderstand* these aggregate effects. Benartzi & Thaler (1999) make this especially clear by demonstrating directly that people asked the probability of aggregate loss of independent bets exaggerate along the lines predicted by NBLLN.

[25] Benartzi & Thaler also elicited the effects of showing the histogram in the above hypothetical examples and showed that it reduces rejections from 25%, 57%, and 50% to, respectively, 14%, 10%, and 17%.

[26] Klos, Weber & Weber (2005) replicate and extend Benartzi & Thaler's findings. Klos, Weber & Weber present subjects with four lotteries, each of which may be played singly, repeated 5 times, or repeated 50 times. Subjects generally prefer the repeated gambles and correctly understand that the variance of monetary outcomes increases with number of repetitions, but in the repeated gambles, subjects vastly overestimate the probability of loss as well as the expected loss conditional on losing money. Klos, Weber & Weber also find that, even though the probability of the monetary outcome falling within a given interval around the expected value falls with the number of repetitions, subjects incorrectly believe it increases. This last finding is inconsistent with our model of NBLLN and may reflect a bias from focusing subjects' attention on the expected value, or it may be consistent with "exact representativeness," a bias we discuss in Appendix B.

to explain this data is absurdly high, in excess of 32,000. For Barney with $\psi = 10$, the required loss aversion is approximately 15—many order of magnitudes closer to reality but still much larger than reasonable estimates of λ, such as 2.25.[27]

In addition to predicting that people perceive too much risk in repeated betting, NBLLN also predicts that people perceive too much risk in the closely-related context of long-term investing. Indeed, analogously to the evidence on repeated betting, Benartzi & Thaler (1999) found that university employees vastly overestimated the probability that equities would lose money over a thirty-year horizon. Moreover, the employees stated a far greater willingness to invest in equities when they were explicitly shown the thirty-year returns. While there are many reasons why individuals may invest less in equities for retirement than recommended by standard finance models, we suspect that NBLLN is an important contributing factor. As such, just as in other settings researchers may underestimate risk aversion by ignoring overconfidence when inferring risk preferences from investment behavior, researchers might therefore exaggerate the risk aversion of investors by ignoring NBLLN.

3.4 The Multiple-Sample Model

In Section 2, we described how Barney processes a set of signals that he treats as one sample. In doing so, we deferred two questions that are crucial to answer before reaching conclusions about the implications of NBLLN in most economic settings. First, when and how does Barney separate out data into more than one sample, both as a function of how the data is presented to him, or how he might endogenously group information during his thinking and decision-making? Second, what does Barney believe about how he will process information in the future? In this section, we provide a framework for thinking about about these questions, and we propose possible answers. We review in Appendix B the scant and somewhat contradictory evidence we can find; since there is so little evidence about which assumptions are appropriate, our proposals are tentative. Nonetheless, we formulate plausible approaches in order to provide guidance for future experimental work, and to study the range of possible consequences of NBLLN in some dynamic decision-making environments.[28]

For Tommy, it does not matter whether he treats 20 independent signals as one sample of 20,

[27]Incorporating some of the biases missing from our model of NBLLN—that subjects still have excessively fat tails in their subjective sampling distribution helps reduce the required level of loss aversion even more—although still not to 2.25. Indeed, even if an individual exhibited the most extreme form of NBLLN *and* fat tails, putting equal weight on every possible outcome of the repeated gamble, the required level of loss aversion would be around 4 (because only 1/3 of the outcomes are losses). This exercise implies that some other bias, such as probability weighting, is also implicated in turning down the repeated bet. This may also explain why some subjects turn down the aggregate bet even when it is presented in histogram form.

[28]Moreover, the implications of other non-Bayesian models of judgment biases—such as base-rate neglect—similarly exhibit sensitivity to how data are framed. As such, the range of approaches we outline here may prove useful for studying those other biases.

two samples of 10, or 20 samples of 1. And it is inherent to Tommy that he correctly predicts how he will treat future information. In contrast, and intrinsic to the very meaning of NBLLN, Barney's beliefs about the distribution of and inference from signals depend on how he divides them up into samples. The range of possibilities raises, in turn, the possibility that Barney predicts his future information-processing differently than how he will actually process information.

Several distinctions and definitions are usefully formulated to address these issues. We use the term *clump* to refer to how signals are objectively delivered to Barney by his environment. For example, when Barney asks a sequence of friends about their experience driving a Volvo, each friend's report arrives as a clump, but when Barney reads a summary of 10,000 individuals' experiences in *Consumer Reports*, the 10,000 signals arrive as a single clump. We refer to how Barney subjectively *processes* these clumps for the purposes of making forecasts and inferences as how he *groups* the signals. Barney forms beliefs regarding each group of signals as if a subjective rate β were drawn that applies only to that group, and hence the single-clump model from Section 2 can be applied to each group. In formulating ways that Barney might process information, we will distinguish two facets of how Barney groups data. The first is how he processes clumps into groups *retrospectively*—how he processes clumps he has already received. The second is how he processes clumps into groups *prospectively*. Prospective grouping determines his forecast about what data he will observe given his current beliefs *and* his forecast of what he will infer from that data after he observes it.

In each of the retrospective and prospective directions, we consider three ways that Barney might process clumps of signals. If Barney groups the signals the same way he receives them from the environment, we call him *acceptive*. Acceptive Barney would process *Consumer Reports* as a single sample of 10,000, and then each of his friends' reports as a separate sample. If Barney processes all of the clumps of signals he observes as a single, large group, we call him *pooling*. Pooling Barney would treat *Consumer Reports* data and his friends' stories as a single, larger sample. Finally, if Barney processes each signal individually (as a group of samples of size 1) no matter how they are clumped, we call him *atomizing*. Since Barney treats samples of size 1 identically to Tommy, if Barney is atomizing both prospectively and retrospectively, he is Tommy. If he does so in only one direction but not the other—examples of which we analyze in subsequent sections—he may be very different from Tommy.

To formalize the information-processing assumptions, consider a decision problem in which up to $1 \leq T \leq \infty$ binary signals will be realized in total. To greatly simplify notation and to sharpen the conceptual framework, we assume that how the environment or his own thinking leads him to group signals is independent of the realizations of those signals and of earlier decisions.[29] We define $T+3$ partitions, each represented by a set of time periods, of the set of signals that fully characterize our predictions. We conceive of the first three of these partitions as embodying physical, informational, and perceptual assumptions rather than assumptions about Barney's statistical reasoning. The first partition we define by the set of dates at which the agent knows he must make decisions, $D \subseteq \{0, 1, ..., T\}$. If $\tau \in D$, the agent knows that he made or will make a decision after observing τ signals but before observing $\tau + 1$ signals. The agent's payoff may depend on any or all of the decisions he makes, the signals that are realized, and the underlying state with which the signals are correlated. This opportunity-for-decisions partition, D, is of course specified in every economic model of decisionmaking.

The second partition is characterized by a set $C \subseteq \{0, 1, ..., T-1\}$ of dates where a new clump begins, such that if and only if $\tau \in C$, then signal τ is in a different clump than signal $\tau - 1$. We assume $0 \in C$. For example, if 11,000 signals arrives as a clump of 10,000 signals followed by a clump of 1,000 signals, then $T = 11,000$ and $C = \{0, 10,001\}$. There is one obvious basic restriction on clumping that we must make:

Clumping Assumption 0. $D \subseteq C$.

This assumption states that at any history where the agent makes a decision, subsequent signals arrive as a separate clump than previous signals. This restriction is inherent in the notion of clumps because the signals cannot have "arrived" together (at least in the relevant sense of the agent's knowledge of their realizations) if the agent knows the realizations of only some of the signals. We treat this as a coherence assumption and always impose it.

Another coherence assumption on C is that Barney will not treat signals differently when he does not see the signals distinctly at all. If it is the case, unbeknownst to Barney, that 38 of the people in *Consumer Reports* statistics of 10,000 car owners have the last name Smith, we assume that Barney cannot (even if thusly motivated) treat Smiths as one sample and non-Smiths as another.[30] Let $I \subseteq \{0, 1, ..., T-1\}$ be a set of dates that defines the third partition, a partition of signals

[29] This is a substantive restriction, and we know of examples where it seems unrealistic, but we do not know how relaxing it will improve insights. Importantly, we are not assuming that *whether* Barney gathers more information is independent of what he has learned; our examples in Section 5 revolve around exactly such a decision by Barney. In those examples, Barney faces a decision after each signal whether to pay to observe another signal. To accommodate those examples in the framework of this section, it can be assumed that whenever Barney chooses not to observe further signals, decisions that "occur" at those future signals do not affect his payoff.

[30] Conceivably, Barney could decide to label an indistinguishable group any way he wants, such as ordering them from 1 to 10,000, and then either perceptually or psychologically distinguish the signals based on this labeling. It seems a safe assumption that he will not do so.

into equivalence classes of indistinguishable signals (so that if $\tau \in I$, then the signal at time τ is distinguishable from the signal at time $\tau - 1$). Clearly it must be the case that $D \subseteq I$, but we also impose a second restriction on clumping:

Clumping Assumption 1. $C \subseteq I$.

Although economic models of decisionmaking do not traditionally specify a clumping partition, our aspiration is have the clumping partition be an exogenous assumption that is not *per se* related to NBLLN, and ideally is pinned down by observable characteristics of a situation. In the interest of minimizing the number of assumptions that have to be specified anew for each new economic model (and hence limit degrees of freedom that might reduce the usefulness of the model), it would be especially attractive to pin down some rule for "clumping" that ties it to D or I. The two obvious candidates are:

Clumping Assumption 2(a). $C = D$.

Clumping Assumption 2(b). $C = I$.

While the three partitions characterized by D, C, and I reflect the physical and perceptual environment facing Barney, the remaining T partitions embed Barney's NBLLN psychology of how he separates out data. To capture the important possibility that Barney's grouping might differ from different time perspectives, we assume that the agent may process the signals differently at different dates. At each $t = 0, 1, ..., T$, there is a set $P_t \subseteq \{0, 1, ..., T - 1\}$ of dates where a new group begins. By letting each P_t contain elements both less than t and greater than t, each partition captures both retrospective and prospective grouping of signals at that point in time. For all $\tau \leq t$, if $\tau \in P_t$, then after having observed t signals, the agent processes signal τ as being in a different group than signal $\tau - 1$; while if $\tau \notin P_t$, then after having observed t signals, the agent processes signal τ as being in the same group as signal $\tau - 1$. And for all $\tau > t$, $\tau \in P_t$ means Barney *anticipates* separating out signal τ from signal $\tau - 1$. Whether he actually does so after observing signal τ is determined by $\{P_\tau, P_{\tau+1}, ...\}$. We assume that $0 \in P_t$ for all t.

We impose one restriction on these sets that we believe is necessary for modeling coherence:

Processing Assumption 0. For any $t = 0, 1, ..., T$, if $\tau \in D$ and $\tau + 1 \geq t$, then $\tau + 1 \in P_t$.

This assumption states that at any date where the agent makes a decision, he processes signals before and after that date as being in separate groups—and that before that date, he knows he will do so.

We consider this to be a modeling coherence assumption because it ensures that Barney's NBLLN from the single-clump model in Section 2 generalizes to every decision node in the multiple-clump model. To see this, suppose Barney knows that $\theta = .5$, reads a summary of 10,000 individuals' experiences in *Consumer Reports*, and then must make a prediction about the next 1,000 signals he will observe. If, in violation of Processing Assumption 0, he were to process all 11,000 signals together as a single group after already observing the first 10,000 signals, then he would believe that the same subjective rate β applies to all 11,000 signals. Using the first 10,000 signals, he would update his beliefs about β from $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta = .5)$ to a density that puts almost all the probability mass on the observed proportion of a signals, say 50%. Since the next clump is grouped with the earlier clump, his subjective sampling distribution for the next clump will put negligible weight on a proportion of a signals outside a neighborhood of 50%. In his predictions about future signals, Barney would no longer exhibit NBLLN. In contrast, Processing Assumption 0 requires that Barney forms beliefs as if a new β is drawn from $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$ before the next 1,000 signals, so his subjective sampling distribution is exactly as in the single-clump model for $N = 1000$. Precisely because it rules out learning about the subjective rate, Processing Assumption 0 distinguishes our model of NBLLN from the generalization of the model from Section 2 that one would employ if one interpreted it as a fully-rational model with uncertainty about the parameter β.

We can now formalize various ways that Barney might form beliefs retrospectively and prospectively. Because D is the set of nodes where Barney's beliefs are payoff-relevant, these definitions focus the assumptions on the P_t's where $t \in D$:

- **Retrospective-Pooling**: If $t \in D$ and $0 < \tau \leq t$, then $\tau \notin P_t$.
- **Retrospective-Acceptive**: If $t \in D$ and $0 < \tau \leq t$, then $\tau \in P_t \Leftrightarrow \tau \in C$.
- **Retrospective-Atomizing**: If $t \in D$ and $0 < \tau \leq t$, then $\tau \in P_t$.
- **Prospective-Acceptive**: If $t \in D$ and $\tau > t$, then $\tau \in P_t \Leftrightarrow \tau \in C$.
- **Prospective-Atomizing**: If $t \in D$ and $\tau > t$, then $\tau \in P_t$.

We omit defining "prospective-pooling" Barney because it is ruled out by Processing Assumption 0 in any decision problem with more than one decision node; at an earlier decision node, Barney cannot expect to pool together future signals that come before and after a future decision node. (Barney's predictions regarding a single clump and a single decision node, as in Section 2, are of course covered by the prospective-acceptive case.)

As an implication of all the other ways he is rational, Tommy always processes information the way that he expects to process information. We call this property "processing-consistency": An agent is **processing-consistent** if $P_t = P_{t'}$ for all $t, t' \in D$. Despite his irrationality, Barney shares this property if his retrospective and prospective thought processes coincide. In particular, if Barney is retrospective-acceptive and prospective-acceptive, then he is processing-consistent and accurately forecasts what his own future beliefs will be after he observes a sequence of signals.[31] In contrast, if Barney is retrospective-pooling and prospective-acceptive, then he is not processing-consistent. As a result, he may behave in a time-inconsistent way, e.g., expecting to learn a lot from purchasing a large number of signals, but remaining uncertain after observing the signals and therefore preferring to purchase yet more signals. This time-inconsistency will play a role in some of the applications we study.

Having formalized the processing assumptions we will focus on, we now discuss intuitions not captured by the formalism. When thinking about what processing assumption is psychologically plausible, we try to be guided by the characteristics of the data-generating process. For example, friends' vivid descriptions of their own personal experiences are easily distinguished from each other and therefore are likely to be processed as separate groups. Indeed, here we note questions and clarifications that our framework can add to how some phenomena have been conceptualized in psychological research. A tradition in psychology emphasizes that people may overweight "vivid" evidence in reaching their judgments; a friend's colorful description of the horrors that ensued when her car broke down while trying to pick up her child from school may weigh more heavily in our judgment of the brand of car she drives than summary statistics based on large samples of data like those sometimes published in *Consumer Reports* (e.g., Nisbett & Ross, 1980). Our model predicts a version of this if vivid signals are isolated from less vivid ones. But the *comparative* over-use of vivid evidence may not be over-use of that evidence relative to its proper use, but an indication of under-use of other evidence. The observation that you are using your friend's car experience a lot more than other, similar people's experiences that show up in the data could be, in light of NBLLN, partly or even fully attributable to the under-use of statistical data. [32] Moreover, even when it seems clear that there is genuine over-use of evidence due to vividness, NBLLN is still clearly an "enabling

[31]Note that acceptiveness is the sole processing-consistent case we focus on because we rule out prospective pooling, and a processing-consistent atomizer is named Tommy.

[32]This hypothesis that NBLLN may help explain why some signals may receive attention in ways that seem surprising may shed a favorable light on some important research lines. For instance, the influence on beliefs of friends' information and experiences, learning by one's own experience, and learning in networks and herds, might seem puzzling in light of the amount of other information an agent may have available. But if information from the rest of society is treated as just one sample, the information provided by members of a the agent's more immediate social network may still loom large.

bias." The simple logic that when a person has large data sets she ought to reach a determinative conclusion means that any sort of inference (rational or irrational) from small, additional samples simply should not matter. Counting a friend's car experience a hundredfold still should not influence beliefs by much.

In other cases, it is harder to see whether and when Barney will group signals together, and (importantly in our applications) whether he is more likely to pool signals prospectively or retrospectively. Maybe he pools information prospectively (in violation of Processing Assumption 0) but atomizes it retrospectively: before he talks to the 10 friends (say) he plans to talk to, he may not attend to the time separation of the information and not realize that he will update his beliefs story-by-story as he goes along. But then he may retrospectively atomize, distinguishing colorful details of his friends' stories. Alternatively, maybe he instead processes information more coarsely retrospectively than he anticipates, as we assume in an example in Section 5: if he plans to keep talking to friends one by one until he feels confident, he might think ahead with attention to each separate signal, focusing on how he will update from current beliefs after his next conversation. But then in retrospect, after he has talked to his next friend, he may quite naturally treat that friend's information symmetrically with all the previous conversations and take stock of his current information by thinking together about all the advice he has received.

We hypothesize that many situations are like the last example, and that the psychology of grouping signals leads Barney to believe he will process data more finely in the future than he actually will do retrospectively. For example, while Barney might recognize that he receives a new signal about the average precipitation of San Francisco every day, and atomizes the signals prospectively, when he attempts to infer whether San Francisco is "rainy" or "sunny," he treats all previous outcomes in a coarser fashion, possibly grouping days by weeks, months, or all together.

We have outlined above a few possibilities about how Barney might process clumps of signals, but there are other reasonable possibilities, some of which violate our assumptions. In particular, we emphasize that acceptive Barney is only one plausible type of Barney that is between pooling and atomizing. Barney could process information in a different fashion than presented by his environment, grouping it instead according to the perceived similarity of the information source. After observing 10,000 datapoints from *Consumer Reports* followed by 10 friends' reports obtained sequentially, for instance, Barney may retrospectively process the information as a group of 10,000 followed by a group of 10.

Besides lack of evidence, an additional and major reason the hypotheses in this section are tentative is that we have completely sidestepped the issue of when and how Barney might "think through" his beliefs more fully. Even if Barney is not processing-consistent, given the assumptions above, our *model* of Barney's beliefs is internally consistent. Barney's *beliefs* themselves, however, are not internally consistent, and this raises additional conceptual and practical issues in applying a model of NBLLN.[33] For example, a teacher could elicit Barney's belief about the likelihood that a first signal will be a, the likelihood that a second signal will be a conditional on the first signal being a, and the likelihood that a sequence of two signals will be aa; then the teacher could point out that the product of the first two does not equal the second. In fact, even our coherence assumptions above could fail depending on the questions a teacher asked Barney. For example, suppose Barney expects to observe 10,000 signals from *Consumer Reports*, then make some payoff-relevant decision, and then observe another 1,000 signals. If a teacher asks Barney to forecast all 11,000 signals, then Barney would presumably do so according to the single-clump model with $N = 11,000$—in violation of Processing Assumption 0. Even in the absence of a "teacher," Barney might ask himself such questions.[34] While we flag these issues, and think they are natural subjects of future research, we proceed in subsequent sections with the assumption that Barney does not think through the inconsistencies in his own beliefs.

3.5 Application: Search and Experimentation

This section explores the economic implications of NBLLN inference, applying variants of our model both to classical information acquisition and to experimentation. In the applications here and in Section 6, we use a standard framework, along the lines of many examples above. An agent is uncertain about which of two possible states of the world, $\omega \in \{A, B\}$, is true. State A has prior probability $0 < f_\Theta(\theta_A) < 1$, and state B has prior probability $f_\Theta(\theta_B) = 1 - f_\Theta(\theta_A)$. In state A, the probability of an a-signal is θ_A, and the probability of a b-signal is $1 - \theta_A$. In state B, the

[33]The internal inconsistency we highlight here does not arise in "false-model Bayesian" models of biased beliefs such as Barberis, Shleifer & Vishny (1998), Rabin (2002), and Rabin & Vayanos (2010). In these models, biases are formulated as agents holding the wrong theory as to the statistical structure of the world, but as being fully Bayesian in their interpretation of data within that structure. So long as all events that are possible in the true world are also possible in the agents' imagined world, no internal inconsistency can arise. (Though even in these models, it may be very likely that the agent will observe a sequence of signals that he perceives to be very unlikely.) Processing-consistent variants of Barney likewise reduce to a "false-model Bayesian" theory. The processing-*in*consistent variants of NBLLN, however, assume an intrinsically non-Bayesian thought process. The modeling challenges associated with internally inconsistent beliefs are not specific to NBLLN, and will arise in any model of belief formation that is fundamentally non-Bayesian.

[34]Importantly, a decision Barney faces might itself naturally cause Barney to ask himself such questions. For example, imagine Barney is making a decision whose payoff depends on whether the state is A or B. He could purchase one signal, and then decide whether to purchase a second signal, or he could purchase two signals all at once at a discount. When deciding what to do, it seems natural that Barney would ask himself what he would conclude after observing each of the three possible outcomes: he observes 1 signal, 1 signal followed by 1 signal, and 2 signals together. Having explicitly asked himself about these possibilities, it seems odd that Barney would—as assumed if Barney is prospective-acceptive—expect to conclude less from the 2 signals together than the 2 signals individually.

probability of an a-signal is $\theta_B < \theta_A$, and the probability of a b-signal is $1 - \theta_B$. Depending on the particular application, the agent can take actions, or observe outcomes and signals that can inform him about the state of the world. Here and in Section 6, we assume that Barney's subjective rate distribution has the functional form (3.3) from Section 2, although we believe the intuitions are more general. The applications make use of the varying assumptions defined in Section 4 about how Barney groups signals retrospectively and prospectively. While the analysis of betting and investing focused only on the implications of Barney's subjective sampling distribution, Barney's behavior in the current setting depends both on his subjective sampling distribution—what signals he expects to observe—and his inferences—what he expects to believe and actually believes after observing the signals.

An interesting implication of NBLLN emerges across a range of applications: it makes a bigger difference to his eventual beliefs for Barney than for Tommy whether he happens to observe strong evidence of the true state early or late in his learning process. The basic logic of NBLLN implies that Barney finds observing strong, early evidence, such as the group of signals aa, more persuasive that A is the true state than a group such as $ababaa$, even though these two groups are, objectively, similarly strong signals about the state—indeed, exactly equally strong when $1 - \theta_A = \theta_B$. As a result, if Barney observes strong evidence early on, he may stop trying to learn about the state after only a few signals, while if he observes ambiguous data early on, he may continue trying to learn even after many signals.

Consider a setting of sequential information acquisition. Barney is trying to decide what make of car to buy, a Volvo or a Lada.[35] He could possibly incur costs to acquire signals about the relative quality the cars, e.g., by taking time to ask friends which car is better. Conditional on the responses that he receives, he can decide to ask more friends, or to stop and choose a car to buy.

Each period $t = 1, 2, ...$, the agent can choose to purchase a single signal at cost $c > 0$ or take an action $\mu \in \{\mu_A, \mu_B\}$.[36] If the agent takes an action, he gets payoff $u(\mu, \omega)$, which equals 1 if the action matches the state ω and 0 otherwise, and the agent faces no further decisions. If the agent decides to purchase an additional signal, he sees the realization of the signal, and he proceeds to the next period. The agent will live forever and seeks to maximize the expected action payoff minus expected signal-purchase costs; if the agent purchases κ signals and then takes action μ, his utility

[35] A Lada is a type of car. So is a Volvo.

[36] In all our applications, we restrict Barney to purchasing information a single realization at a time. If we allowed Barney to choose how many signals he could purchase each period, Barney would have to think about what he would learn from purchasing two signals sequentially in order to compare it to purchasing two signals simultaneously. Modeling this thought process raises challenges—the same as those mentioned in footnote 33—that we sidestep in this paper.

is $u(\mu, \omega) - \kappa c$. Note that we assume no discounting, so that the only reason an agent would stop acquiring information before being absolutely certain is the cost c of obtaining an additional signal.

For Tommy, the characterization of optimal behavior is well-known (e.g., Wald, 1947). Each time Tommy purchases a signal, he updates his posterior beliefs. His optimal behavior is characterized by two probabilities, ν_l and ν_h, with $0 < \nu_l < \nu_h < 1$. If and only if the posterior probability of state A exceeds ν_h, he stops and takes action μ_A; if and only if it goes below ν_l, he stops and takes action μ_B. Tommy continues to purchase signals as long as his posterior beliefs remains between ν_l and ν_h. But because his posterior ratio is a martingale process, Tommy will eventually feel strongly enough to take an action almost surely. Importantly, as $c \to 0$, $\nu_l \to 0$ and $\nu_h \to 1$.

Barney's behavior will depend on how he groups signals. Because the signals arrive one at a time, Barney expects to group the signals as samples of size 1, regardless of whether he is prospective-atomizing or prospective-acceptive. Since Barney *expects* to behave exactly like Tommy, his policy is the same as Tommy's, with the *same* thresholds ν_l and ν_h determining when he will stop and take an action. However, Barney's posterior belief may follow a different process than Tommy's, depending on how he groups signals retrospectively. If Barney is retrospective-atomizing or retrospective-acceptive, his beliefs and behavior will be identical to Tommy's.

If he is retrospective-pooling, however, Barney's behavior can differ qualitatively from Tommy's. In this case, the impact of an additional signal on his posterior beliefs is smaller than for Tommy. Hence in expectation, Barney will purchase weakly more signals than Tommy. Moreover, the marginal impact of an additional signal on Barney's posterior beliefs will approach zero as his sample of observed signals grows. This is because Barney's inference becomes more and more driven by the proportion of a signals, which is less affected by an additional signal in a larger sample. However, Barney *believes* that an additional signal will have the same impact, regardless of the sample size he has already observed. As a result, Barney can become stuck in a *learning trap*, in which he purchases signals forever, but they will never change his confidence in the state of the world enough for him to stop. Proposition 7 summarizes this discussion:

Proposition 7. *Fix payoffs $u(\mu, \omega)$, rates $\theta_A > \theta_B$, and priors $f_\Theta(\theta_A)$ and $f_\Theta(\theta_B)$. Suppose Barney is prospective-atomizing or prospective-acceptive.*

1. *If Barney is retrospective-atomizing or retrospective-acceptive, then for any $c > 0$, Barney's behavior exactly coincides with Tommy's, and he will choose an action after observing a finite number of signals almost surely.*

2. *If Barney is retrospective-pooling, then for all $\overline{p} < 1$, there exists $\bar{c} > 0$ such that for all $c \leq \bar{c}$, Barney purchases an infinite number of signals with probability $p > \overline{p}$.*

3. *Suppose Barney is retrospective-pooling, $\theta_A = 1 - \theta_B$, and Barney is willing to purchase an additional signal when his posterior probability of state A equal to q. (a) If Barney's posterior probability of state A is q after observing N signals and again q after observing $N' > N$ signals, then the probability that Barney will purchase an infinite number of signals is weakly higher after he has observed the N' signals than it was after the N signals. (b) Furthermore, there exists a $\hat{N}$ such that for all $N' > \hat{N}$, the probability that Barney will purchase an infinite number of signals is strictly higher after he has observed the N' signals than after the N signals.*

Since the marginal effect of a signal is shrinking toward zero, Barney's posterior probability of state A has a positive probability of always remaining between ν_l and ν_h. If Barney ends up in such a learning trap, then his welfare is unboundedly negative. In fact, because a small c tempts him to wait longer, the probability of this unboundedly negative welfare becomes arbitrarily close to 1 as the signal cost c becomes arbitrarily small. For a given c, the last part of the proposition states that Barney is more likely to get caught in a learning trap the more signals he has already observed, holding constant his posterior belief. For example, if the rates θ_A and θ_B are symmetric, then an equal number of a and b signals does not change Barney's beliefs; so Barney is more likely to end up purchasing an infinite number of signals after having observed *abab* than he was before he observed any signals.[37]

We conclude this section by considering the implications of NBLLN in the related situation where an agent learns through experimentation. Rather than purchasing signals about the quality of the car, Barney could instead take them for test drives, perhaps by renting them. Here, instead of an explicit cost, the cost of information acquisition is the cost of waiting to purchase the correct car. If the Volvo is the better car, then Barney is losing out every day he drives the Lada. For the same reasons that retrospective-pooling Barney could get caught in a learning trap in a

[37]We have found several experiments that set up a dynamic information-purchase setting with a payoff structure similar to the model in the text and that compare subjects' behavior with a Bayesian benchmark, which is calculated assuming expected-value maximization. Tversky & Edwards (1966), Pitz (1968), Wendt (1969), and Hershman & Levine (1970) found that subjects purchased too much information. In contrast, Fried & Peterson (1969) and Pitz & Barrett (1969) found that subjects purchased too little information. Moreover, Pitz & Barrett found that when the already-observed sample size was larger, holding constant the objective strength of evidence, subjects bought *fewer* additional signals. Also contrary to our model's prediction, Sanders & Ter Linden (1967) Studies 1-3 found that, when the already-observed sample size was larger, subjects stopped acquiring information at a point where the objective evidence was weaker. In Sanders & Ter Linden's experiments, however, the signals arrived at a rate of 2, 5, or 10 signals per second, which is so fast that the nature of the inference task is likely quite different than in other studies.

dynamic information-acquisition setting, he could end up taking a suboptimal action forever in an experimentation setting.

Formally, in each of an infinite number of periods, $t = 1, 2, ...$, the agent takes an action $\mu_t \in \{\mu^1, \mu^2, ..., \mu^J\}$ where $J > 1$. After taking an action, the agent receives the (possibly unobserved) payoff $u(\mu_t, \omega)$ and a signal that takes on a value of a or b. An a-signal is generated with probability $0 < \theta_{\mu_t,\omega} < 1$ that is partially revealing of the state (that is, $\theta_{\mu^j,B} < \theta_{\mu^j,A}$ for all μ^j). Here, we assume the agent discounts future utility flows at rate $0 < \delta < 1$. The problem is interesting when the agent faces a tradeoff between taking the action that maximizes his expected payoff subject to his current beliefs—the "myopically best action"—and others which, although giving a low expected payoff, can help reveal the true state of the world, which enables him to choose a better action thereafter.

In this environment, Tommy will eventually stop experimenting, and his behavior will settle down to a single action. To be precise, Aghion, Bolton, Harris & Julien (1991) showed that almost surely at some finite T, Tommy will take the action that gives the highest per-period payoff conditional on his current belief for all periods $t \geq T$.

Just as with information-acquisition, because he observes one signal at a time, Barney will expect to behave like Tommy regardless of whether he is prospective-atomizing or prospective-acceptive. Retrospective-atomizing and retrospective-acceptive Barney will behave identically to Tommy. However, retrospective-pooling Barney may continue to experiment forever with positive probability, with the same possibilities and intuitions as in Proposition 7 above.

There are examples of experimentation environments where there is positive probability that Tommy never learns the true state, e.g., when the myopically best action provides no information about the true state. However, when there are a finite number of possible actions, and each provides some information about the true state, Tommy will eventually discover the true state and take the best action. By contrast, because Barney's learning is limited even when he observes an infinite amount of data, he may never learn enough to take the correct action.

Barney's failure due to NBLLN to learn in an experimentation environment may play an important role in allowing people to maintain misconceptions about themselves despite extensive feedback. For example, naivety about present-biased preferences—a false belief that, even though one weights current utility flows disproportionately relative to future utility flows, one will not do so in the future—is now widely considered to have important implications in a variety of economic settings, differing dramatically from sophisticated present-bias. It is natural to wonder how it can be that people do not eventually learn their true discount function by observing their own behavior.[38] The variant of NBLLN that assumes retrospective-pooling implies that people may not become certain of their own dispositions even after an infinite amount of self-observation. NBLLN may similarly be implicated in the persistence of other biases, such as overoptimism about one's own abilities. NBLLN may also enable the otherwise-puzzling lack of knowledge of one's own preferences that is presupposed in any self-signaling model.

3.6 Application: Observational Learning

Besides learning from his own experience or gathered information, an agent could also choose which car to buy by observing which car his neighbors have bought. Our final application concerns issues in modeling and potential implications of NBLLN in observational learning. Once again, we assume that there are two states, $\omega \in \{A, B\}$, with corresponding rates for a signals, $\theta_A > \theta_B$, and actions $\mu_t \in \{a, b\}$ (here we use the same notation for actions as for signals). The action pays off 1 if it matches the state and 0 otherwise. There is an infinite sequence of agents, $t \in 1, 2, ...$, who each observe a single signal, choose an action in order, and receive a payoff. Agents have common prior beliefs $f_\Theta(\theta_A) = 1 - f_\Theta(\theta_B)$. In addition to their private signal, agents observe the action taken by each previous agent, and the order in which those actions are taken.

Much of the theoretical literature on social observational learning has focused on two closely-related phenomena, "information cascades" and "herds."

Definition 1. *A cascade is said to occur when, following some period, each agent's choice of action does not depend on the agents' private signals.*

Definition 2. *A herd is said to occur when, following some period, all agents choose the same action. It is a good herd if the common action is the true state; a bad herd if the common action is not the true state.*

The characterization of Bayesian Nash equilibria in this simple setting is well-known (e.g., Banerjee, 1992; Bikhchandani, Hirshleifer & Welch, 1992; and Smith & Sørensen, 2000). When the agents are all Tommy, in all equilibria both a cascade and a herd occur with probability 1. To see the

[38] Ali (2011) formalizes this intuition in a planner-doer model and concludes that a Bayesian-rational agent who succumbs to temptations would eventually learn that he has a self-control problem. In analysis of this environment that we do not include here, we find that (a) Barney may forever remain uncertain about his self-control problem, despite succumbing to temptation every period; and (b) the probability of this persistent uncertainty—and hence optimism about his ability to resist temptation—is increasing in Barney's ex-ante optimism.

intuition for why, suppose the common prior puts equal weight on each state, and $1-\theta_A=\theta_B$. In that case, a cascade will occur as soon as the total number of a-actions taken exceeds or falls behind the total number of b-actions by 2. For example, if the actions from the first 8 agents are, in order, *abababaa*, then the 9^{th} agent will believe the state is more likely to be A even if his private signal was b. He will play a regardless of his own signal. The 10^{th} agent understands this logic, so if he observes *abababaaa*; he will likewise ignore his own signal and also understand not to infer anything from the 9^{th} action. Similarly for the 11^{th} agent, and so on. The cascade occurs in period 9, with all agents $t=9,10,...$ "herding" on action a. A good herd is of course more likely than a bad herd because early actions are indicative of the true state, but the herd will be on the wrong action with positive probability.

Although no agent directly observes more than one signal, NBLLN will influence herding under the assumption that agents apply their non-belief to the signals they infer others are getting. Intuitively, if Barney groups together previous agents' actions, then NBLLN implies that he will infer less from their actions than he should. Therefore, he will more often rely on his own signal, making it slower for a herd to form. Moreover, since Barney needs to see more agents following their own signal before ignoring signals than does Tommy, when a herd *does* occur, it is more likely to be on the correct action. Finally, because Barney's learning is limited even in an infinite sample, a qualitatively different kind of behavior is possible. Because Barney may vastly under-infer from a large number of actions, it may be possible that if agents do not herd quickly enough, they will instead form what we call an "eddy":

Definition 3. *An eddy is said to occur when, following some period, every agent chooses the action corresponding to his own signal.*[39]

Because agents do not care about future agents' actions, the way Barney groups signals prospectively is irrelevant for his behavior. For the purposes of this section, we assume that Barney groups all previous agents' actions together and treats his own private signal as separate. This assumption corresponds to being "retrospective-acceptive" if Barney observes the previous actions simultaneously, as well as the order that the actions were taken in, and then observes his own signal. We think this grouping assumption would, in fact, be natural even if agents observed their own signals at the

[39]Smith & Sørensen (2000) show that when agents are all Tommy, but have ordinally different preferences, an eddy can occur. However, in their model, this results from what they describe as "confounded learning": agents' actions, although dependent on their private signal, are uninformative about the state of the world because an action's informational meaning is different depending on the (unobservable) preference of the agent. In contrast, we find that when agents are Barney, eddies occur even when actions are informative. Barney simply never learns enough from this information.

same time as observing actions because other agents' actions may feel like qualitatively different information than the agent's own signal.[40]

Deriving the full implications of NBLLN in this setting requires an extension of the framework introduced in Section 4 to confront an issue that will be important in many economic applications of NBLLN, as it is for other models of cognitive biases. Namely, what is Barney's theory as to how *other* agents draw inferences? While a complete exploration of Barney's beliefs about others' beliefs is beyond the scope of this paper, the herding environment can illustrate some of the assumptions we think may be reasonable. In the analysis here, we compare two different assumptions about Barney's beliefs about other agents' inferences, both of which seem plausible. We say that Barney is *strategically sophisticated* if, when he draws inferences from other agents' behavior, he assumes that they group signals the way they in fact do group signals. That is, even though the 3^{rd} agent groups together the actions of the 1^{st} and 2^{nd} agents, he understands that the 2^{nd} agent does not group his own signal with the 1^{st} agent's action.[41] Alternatively, we say that Barney is *strategically naive* if, when he draws inferences from other agents' behavior, he assumes that they group signals the way he himself does. In that case, the 3^{rd} agent incorrectly believes that the 2^{nd} agent groups his own signal with the 1^{st} agent's action.

Although we believe that strategic naivety is a plausible assumption, in the current framework it generates significant problems. It is possible that strategically naive Barney can observe sequences of actions that he cannot rationalize.[42] For example, consider a set of strategically naive Barneys, where $\theta_A = .51 = 1 - \theta_B$, and $\psi = .10$. The first 100 Barneys receive alternating a and b signals in order (and so take alternating a and b actions). The 101^{st} and 102^{nd} agents both receive a b signal, and so both take a b action. From the perspective of the 104^{th} agent a herd should start on action b with the 103^{rd} agent. However, if the 103^{rd} agent receives an a signal, he will still take an a action, as his own signal outweighs the informativeness of the pooled actions of previous agents. In this case, the 104^{th} agent will not be able to rationalize the observed sequence of outcomes. We sidestep this problem by focusing solely on sophistication.

We begin the formal analysis by characterizing when Barney herds. Although models of ob-

[40] The assumption about grouping is again important. If instead, for example, we assumed Barney pooled his own signal with all other inferred signals, then the standard results in the literature carry through: a cascade will always occur, and therefore a herd will happen, sometimes on the incorrect action.

[41] Although strategic sophistication implies that each agent understands that other agents are using a different partition than themselves, they can support such beliefs by thinking that other agents are wrong.

[42] In a closely related model that we have analyzed—where agents only observe the total number of each type of action, but not the order in which they were taken—this problem never arises. The results presented here for sophisticated agents apply in that setting for both strategically sophisticated and strategically naive Barney. In addition, if all agents are strategically naive Barney, then there are parameter values such that there exist finite histories (reached with positive probability) after which an eddy occurs with probability 1.

servational learning admit multiple equilibira, we will follow the literature in assuming that when an agent is indifferent between actions, which occurs when his own signal balances the information from previous actions (and so the agent is indifferent between actions), the agent follows his own signal.[43] Proposition 8 states that if the prior puts sufficient weight on one of the states, then in equilibrium there is positive probability of a herd on that state. The reason (but not the precise thresholds) is exactly the same for Barney as for Tommy: if early agents receive enough signals in favor of the state, say A, and take action a, then eventually it makes sense for Barney (and all subsequent Barneys) to take action a regardless of his signal. However, if the prior in favor of A is too low, then it is impossible to herd on action a because the game will begin with a herd on b.

Proposition 8. *Suppose Barney is strategically sophisticated. Fix rates $\theta_A > \theta_B$ and Barneyness parameter $0 < \psi < \infty$. There are bounds $0 < \underline{p}_A < \frac{1}{2}$ and $\frac{1}{2} < \overline{p}_A < 1$ such that if the common prior $f_\Theta(\theta_A) \geq \underline{p}_A$, there is positive probability of herding on action a; and if $f_\Theta(\theta_A) \leq \overline{p}_A$, there is positive probability of herding on action b.*

The proposition implies that, like for Tommy, there exists a range of priors ($\overline{p}_A \leq f_\Theta(\theta_A) \leq \underline{p}_A$) such that in equilibrium there is positive probability of a herd on either state.

We next characterize when an eddy might occur. An eddy can occur for Barney but *not* Tommy. If a herd has not occurred after many actions, then the early signals must have provided mixed evidence. For Tommy, as long as $\theta_A = 1 - \theta_B$, a herd starts after one action is played twice in a row, whether this occurs after two actions or eight actions. In contrast, Barney would draw a weaker inference following *abababaa* than *aa*. Hence, if the early evidence is mixed, then the probability of a herd falls because even strong evidence thereafter will not affect Barney's beliefs by as much. In fact, the probability of a herd starting vanishes as the sample size of mixed actions gets large.

For an eddy to occur, it must be the case that beliefs do not become too extreme initially, due to either extreme priors or lopsided initial signals. But it is also necessary that the prior put low enough weight on the true state that even in a very large sample, where the Law of Large Numbers ensures that the proportion of a signals will exactly match the true rate, Barney's inference from other agents' actions is weak enough that he still follows his own signal. For any rates $\theta_A > \theta_B$ and Barneyness parameter $0 < \psi < \infty$, given the true state, there exists a range of prior beliefs such that an eddy occurs with positive probability.[44] Note that this range of prior beliefs is different

[43] Other equilibria, where the agent with some probability takes the action opposite his signal, will reduce the amount of publicly available information. In the extreme case, where an agent always takes an action opposite his signal when indifferent, a herd always starts in the first period.

[44] The intuition that there are parameter values where an eddy occurs with positive probability is robust to alternative models of strategic behavior, such as Eyster & Rabin's (2010) model of inferential naivety.

depending on the true state, and may or may not overlap for the two states.

If the two states are close enough to being equally likely, and if the rates θ_A and θ_B are close enough together, however, then the range of prior beliefs that enables an eddy in state A definitely overlaps with the range in state B, and there is positive probability of an eddy in both states. Under those conditions, Barney's posterior belief will remain weak after observing a large number of previous players' actions, regardless of whether the proportion of a actions is θ_A or θ_B.

Proposition 9. *Suppose Barney is strategically sophisticated. There exist $\delta > 0$ and $\gamma > 0$ such that if $|f_\Theta(\theta_A) - f_\Theta(\theta_B)| < \delta$ and $|\theta_A - \theta_B| < \gamma$, then there is positive probability of an eddy regardless of which state is true.*

In most cases when an eddy occurs, there is positive (but shrinking) probability along the equilibrium path that the potential eddy will end and be superceded by a herd. The next proposition formalizes this: the probability of a herd vanishes as the sample size gets large because the actions of later agents have negligible impact on the proportion of a actions and hence negligible impact on Barney's beliefs. Furthermore, as the sample size gets large, the relative likelihood of a bad herd to a good herd vanishes.

Proposition 10. *Suppose Barney is strategically sophisticated.*

1. *For any $\varepsilon > 0$, there exists $T_\varepsilon > 2$ such that if a herd has not occurred by period T_ε, the probability of herding from then on is less than ε.*

2. *Furthermore, suppose $1 - \theta_A = \theta_B$ and the prior $f_\Theta(\theta_A) \geq .5$ then for any $\nu > 0$ there exists $T_\nu > 2$ such that if a herd has not occurred by period T_ν, the likelihood ratio of a bad herd to a good herd is less than ν.*

Although in an eddy it will never happen that all agents will take the incorrect action (as in a bad herd), it is the case that eddies are on average less efficient than herds. To see this, again consider the case when signals are symmetric. The ratio of correct actions to incorrect actions in a herd is simply the ratio of good to bad herds. This is equivalent to the ratio of the probability of two signals that match the state (which we suppose is A) to the probability of two signals that do not match the state, $\left(\frac{\theta_A}{1-\theta_A}\right)^2$. In an eddy the ratio of correct actions to incorrect actions is simply the ratio of the probability of a correct signal to the probability of an incorrect signal, $\left(\frac{\theta_A}{1-\theta_A}\right)$. Since $\frac{\theta_A}{1-\theta_A} > 1$, the average payoff is higher in a herd than in an eddy.

This is because the situations where Barney's behavior causes eddies—sequences of signals where everyone's action reveals their signals—are those situations where inferential naivety does not affect agents' decisions.

3.7 Conclusion

The logic of NBLLN unambiguously predicts that people will extract far too little information from large samples, but there are strands of literature both within psychology and within economics on "over-confidence" in beliefs. In our reading of the evidence in Appendix B, it is clear that people tend to be over-confident in some environments and under-confident in other environments. Rather than viewing over-confidence and under-confidence as fundamental biases in themselves, we view both as outcomes to be explained as a function of the information a person is confronted with. The model of NBLLN highlights a feature of the decisionmaking environment—namely, sample size—that affects the degree to which an agent will draw too weak an inference from evidence. In Appendix A we combine NBLLN with the Law of Small Numbers (LSN) bias, which generates a bias toward over-confidence in inferences, and the overall pattern we predict is over-inference in small samples and under-inference in large samples.

There are conceptualizations of the tendency to under-infer from large samples that differ from that embedded in our model. One interpretation proposed for many cognitive biases is "ecological mismatch": while a person's thought process leads to biased beliefs for *i.i.d.* processes studied in the laboratory, the same thought process would generate appropriate beliefs for the typical, real-world random processes people encounter. For example, in the case of under-inference, Winkler & Murphy (1973) posit that people treat independent signals as if they were positively correlated because their real-world experience is with positively correlated signals. Such positive correlation would generate excessively-dispersed subjective sampling distributions and under-inference but *not* NBLLN: the Law of Large Numbers applies to positively correlated signals, so an agent who misperceived independent signals as positively correlated would still believe that proportions converge to the population rate. Moreover, while ecological mismatch arguments often have merit, we think the argument is unappealing in this context because the bias we call NBLLN is evident in examples with which subjects have a great deal of real-world experience, such as coin-flipping.[45]

Many have proposed conceptualizing NBLLN as one consequence of the "representativeness heuristic," according to which people draw inferences based on the degree of similarity between features of a sample and features of a population from which the sample might have been drawn. Indeed, Kahneman & Tversky (1972) present evidence for NBLLN in precisely this context. Although NBLLN certainly seems consistent with representativeness, it is not clear how the logic of

[45]We also note that in the case of the Law of Small Numbers, the opposite ecological-mismatch hypothesis is often proposed: that people ordinarily deal with *negatively*-autocorrelated signals. Typical real-world processes would have to have a fairly complicated form involving short-run negative autocorrelation and long-run positive autocorrelation to rationalize both the Law of Small Numbers and NBLLN.

representativeness predicts the prototypical case of NBLLN: e.g., an agent who observes 6 million heads and 4 millions tails continues to put non-trivial probability on the coin being fair. Representativeness could explain this kind of observation if it is interpreted as inferences based on proportions, combined with the additional assumptions of reasonably accurate inferences in small samples and insensitivity to sample size, but that combination of assumptions essentially amounts to our model.

A natural alternative modeling approach would be to build a theory of "sample-size neglect," in which, loosely speaking, an agent forms beliefs about a sample of any size as if it were a "medium-sized" sample of, say, size 7. Such a model would imply under-inference for sample sizes larger than 7 and over-inference for sample sizes smaller than 7. This is a common conceptualization and one which we found compelling enough to consider as our first (and more parsimonious) approach. But there are two reasons why we settled on explaining the data with a combination of NBLLN and LSN instead. First, the approach here makes different predictions regarding inferences an agent draws from a sample of size 1. Because both our model of NBLLN and LSN predict that an agent's beliefs are correct for a sample of size 1, the combined model in Appendix A also makes this prediction. A model of sample-size neglect, by contrast, must predict that over-inference is most extreme for a sample of size 1, a prediction that is not borne out by the evidence we review in Appendix B. Second, a model of sample-size neglect is neither psychologically nor mathematically rich enough to describe people's beliefs about sequences of signals.[46]

Our model of NBLLN is defined only when the signals are *i.i.d.* and binomial. There are some natural approaches to modeling NBLLN for non-*i.i.d.* signal sequences. Consider a binomial random process defined by a mapping from any initial rate, θ_0, and any history of t observed signals, h_t, into a rate that the $(t+1)^{\text{st}}$ signal will be an a-signal, $\theta(\theta_0, h_t)$. When Barney knows the initial rate is θ_0, he forms his beliefs as if the initial rate were β, a random variable drawn from distribution $f^{\psi}_{\mathbb{B}|\Theta}(\beta|\theta_0)$. For the first signal in a group, he believes that the probability of an a-signal is β, and for the $(t+1)^{\text{th}}$ signal within that group, he believes that the probability of an a-signal is $\theta(\beta, h_t)$. This modeling approach can be applied not only when the signals truly are non-*i.i.d.*, but also when an agent falsely believes they are non-*i.i.d.* due to another psychological bias. Indeed, Appendix A applies this approach to combine NBLLN with belief in the gambler's fallacy, which posits that people think that signals are negatively autocorrelated even when they are *i.i.d.*.

There are also natural extensions of our modeling approach to non-binomial cases.[47] Suppose,

[46]The first concern could be addressed by a model of "sample-size misperception," in which an agent correctly perceives a sample of size 1, treats a small sample as if it were larger than it is, but treats all large samples as a sample of size 7. Such a model is similar in spirit to a combination of NBLLN and LSN but does not address the second concern.

[47]While experiments have predominantly focused on the binomial case, in Appendix B we discuss some

for example, that the signals are normally distributed *i.i.d.* with known mean μ and variance σ^2. We can imagine a cousin of Barney believes instead that signals are generated by a two-stage process, where a subjective mean ν is drawn from some distribution centered at μ, and then the signals are drawn from a normal distribution with mean ν and variance σ^2. While Tommy believes that the mean of a large random sample of signals will converge to a point mass at μ, Barney's cousin believes it will converge to the density of ν. We could assume that the density of ν corresponds to the empirical large-sample beliefs, or for analytical tractability, we could assume that ν follows the conjugate prior distribution for the normal distribution, which is itself a normal distribution.

evidence that underinference in large samples applies also to multinomial and normal signals.

3.8 Appendix A: A Combined Model of Non-Bayesian Updating

In this appendix, we briefly outline models of other biases besides NBLLN that also appear to matter for inference, focusing on the same single-clump context of Section 2. These models help organize our review of the evidence in Appendix B. Although these models are far more cursory and preliminary than our model of NBLLN, we hope that formalizing these alternatives can both crisply differentiate them from NBLLN—clarifying in particular that none of them are the "opposite" of or inconsistent with NBLLN—and also be suggestive of how to develop these other biases along lines we have done with NBLLN.

At first glance, NBLLN appears to be directly at odds with another bias in beliefs, the Law of Small Numbers (LSN): Tversky & Kahneman (1971) formulated the term "Law of Small Numbers" to refer to the idea that people exaggerate the likelihood that small samples will reflect the underlying population. While NBLLN generally leads to under-inference, LSN generates over-inference. However, these two biases are neither logically not psychologically inconsistent. Indeed, we believe it is the combination of the two which has led judgment researchers to posit a bias of "sample-size neglect" in which people overestimate the resemblance of small samples and underestimate the resemblance of large samples as if they simply do not see the relevance of sample size. We re-interpret what appears to be sample-size neglect as a combination of these two biases, and we show how the basic under-inference implications of NBLLN goes through after LSN is accounted for.

According to LSN, people exaggerate how much small samples resemble the population. In Rabin's (2002) model, an agent forms beliefs about N *i.i.d.* draws that have known rate θ *as if* signals were drawn without replacement from an "urn" of size M that contains exactly θM a-signals and $(1-\theta)\,M$ b-signals.[48] To make sure that the agent continues to view the draws as random even after many signals have been realized, it can be assumed that $M > 2$, and the urn is "renewed" every odd number of draws. In other words, every odd-numbered draw, the signal is drawn from a refilled urn of size M, and every even-numbered draw, the signal is drawn from the urn of size $M-1$ that is depleted by the previous signal's draw. An agent who believes in LSN is called Freddy. The parameter M governs the strength of LSN, with Freddy becoming Tommy in the parameter limit $M \to \infty$. In Rabin's model, for a given θ, the parameter M, in addition to being an integer larger than 2, must satisfy the constraint that θM is an integer. This constraint becomes problematic when combining the model of LSN with our model of NBLLN. In our model of NBLLN, the agent

[48] Rabin & Vayanos (2010) improves on Rabin's (2002) model that we discuss here, generalizing LSN beyond the binomial case and without assuming the signal-generating process is *i.i.d.* We explore this more contrived simple model of LSN here for ease of combining it with NBLLN.

thinks the outcome of a random sample is determined by the subjective rate β, which is drawn from a distribution with full support on $[0,1]$. Hence for any M, βM will be non-integer-valued with probability 1. For this reason, we propose a variant of Rabin's model that, while still requiring that M is an integer larger than 2, remains well-defined even when θM is not an integer.

Our variant of Rabin's model of LSN is identical to Rabin's model except that instead of believing that the "urn" contains θM a-signals and $(1-\theta) M$ b-signals, Freddy believes that it contains $\widetilde{A}$ a-signals and $\left(M-\widetilde{A}\right)$ b-signals, where $\widetilde{A}$ is an integer-valued random variable that equals $j \in \{0,1,...,M\}$ with probability $\binom{M}{j}\theta^{j}(1-\theta)^{M-j}$. In words, Freddy thinks that signals are drawn without replacement from an urn of size M, but he believes the composition of the urn is random, with the "average" urn containing θM a-signals and $(1-\theta) M$ b-signals. When the rate is known to be θ, Freddy believes that the number of a-signals in the urn is a binomial random variable with parameters (θ, M).

NBLLN is the belief that the mean of a random sample converges to a non-trivial distribution, rather than a precise estimate of the mean of the population, in the limit as the sample size gets large. We refer to an agent who believes in both LSN and NBLLN as Barney-Freddy, with beliefs denoted by $f^{\psi M}$. Like Barney, he predicts that the sample is drawn according to a subjective rate that may not equal the true rate. We assume that the subjective rate $\beta \in [0,1]$ is drawn from a density $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$ determined by θ and Barneyness parameter ψ that satisfies Assumptions A1-A4 from Section 2. In accordance with LSN, Barney-Freddy thinks that signals are drawn without replacement from an "urn" of size M that contains $\widetilde{A}$ a-signals and $\left(M-\widetilde{A}\right)$ b-signals, where $\widetilde{M}$ is a binomial random variable with parameters (β, M). Hence Barney-Freddy thinks that the "average" urn contains βM a-signals and $(1-\beta) M$ b-signals. The integer $M>2$ parameterizes the degree of belief in LSN, and we assume that the urn is "renewed" and $\widetilde{A}$ is re-drawn every odd number of signals. As usual for the model of NBLLN, the same subjective rate β applies to the entire clump.

Barney-Freddy both believes in the gambler's fallacy—that is, he expects recent a-signals to be followed by b's and vice-versa—and believes that the sample mean of a large population will converge toward a full-support limit distribution. He does not observe the subjective rate β, but conditional on any given β, Barney-Freddy expects that if the even-numbered draw t is more likely to be an a signal if the $(t-1)^{st}$ draw was b than if it was a. Hence even without knowing β, Barney-Freddy excessively expects that a and b signals will alternate between an odd draw and an even draw. Consequently, for reasonable calibrations of ψ and M, Barney-Freddy's subjective sampling distribution in a finite sample will be too peaked, putting too little weight in the tails. On the other hand, Lemma A1 states that Barney-Freddy's beliefs about a large sample converge to a full-support limit distribution. Like for Barney, Barney-Freddy's limit density will be equal to $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$, regardless of the degree of belief in LSN. Intuitively, every odd-even pair of draws will have, on average, proportion β of a-signals. Hence by the Law of Large Numbers, the sample as a whole will tend toward having proportion β of a-signals almost surely.

Lemma A1. *Barney-Freddy does not believe in LLN: For any interval* $[\alpha_1, \alpha_2] \subseteq [0,1]$,

$$\lim_{N\to\infty} \sum_{x=\lfloor a_1 N \rfloor}^{\lceil a_2 N \rceil} f^{\psi M}_{S_N|\Theta}(A_s = x|\theta) = F^{\psi}_{\text{ß}|\Theta}(\beta = \alpha_2|\theta) - F^{\psi}_{\text{ß}|\Theta}(\beta = \alpha_1|\theta) > 0.$$

Not only are LSN and NBLLN mutually consistent, but LSN actually *magnifies* NBLLN. For smaller M, Barney-Freddy believes that odd-even signal pairs more frequently alternate, and hence his subjective sampling distribution converges to the limit distribution more quickly.

While it is true that Freddiness generates over-inference while Barneyness tends to generate under-inference, there is a clear pattern to when Barney-Freddy over-infers and when he under-infers. For reasonable calibrated values of ψ and M, Barney-Freddy will over-infer from small samples. For any values of ψ and M, it follows immediately from Lemma A1 that Barney-Freddy will under-infer when the sample size N is sufficiently large.

A well-known bias is base-rate neglect (Kahneman & Tversky, 1973), an underweighting of prior probabilities in drawing inferences. Instead of assuming that an agent updates according to Bayes's Rule applied to his subjective sampling distributions, we can capture base-rate neglect by assuming that the agent draws inferences according to:

$$f^{\psi M b}_{S_N|\Theta}(\theta_A|s) = \frac{f^{\psi M}_{S_N|\Theta}(s|\theta_A)\, f_{\Theta}(\theta_A)^b}{f^{\psi\phi\gamma}_{S_N|\Theta}(s|\theta_A)\, f_{\Theta}(\theta_A)^b + f^{\psi\phi\gamma}_{S_N|\Theta}(s|\theta_B)\, f_{\Theta}(\theta_B)^b}$$

and

$$f^{\psi M b}_{S_N|\Theta}(\theta_B|s) = \frac{f^{\psi M}_{S_N|\Theta}(s|\theta_B)\, f_{\Theta}(\theta_B)^b}{f^{\psi M}_{S_N|\Theta}(s|\theta_A)\, f_{\Theta}(\theta_A)^b + f^{\psi M}_{S_N|\Theta}(s|\theta_B)\, f_{\Theta}(\theta_B)^b},$$

where $f^{\psi M}_{S_N|\Theta}(s|\theta_A)$ and $f^{\psi M}_{S_N|\Theta}(s|\theta_B)$ are the subjective sampling distributions of an agent with Barneyness parameter $0 < \psi < \infty$ and Freddiness parameter $M > 2$, and $0 \leq b \leq 1$ parameterizes the degree of base-rate neglect. If $b = 1$, these formulae specialize to Bayes' Rule, where there is no base-rate neglect. If $b = 0$, the agent ignores base rates altogether, treating any prior probabilities as if they were 50-50. This formulation of base-rate neglect has been previously adopted in empirical work by (e.g., Grether, 1980) and concurrently in theoretical work by Bodoh-Creed (2010). Applying it theoretically in dynamic settings raises many of the same conceptual issues as NBLLN; the way in which the agent processes groups of signals will matter a great deal in how his beliefs evolve. For that reason, we believe the framework we have begun to develop in this paper for analyzing dynamic NBLLN may be of use for analyzing dynamic base-rate neglect as well.

A simple explanation for some of the evidence on people's beliefs is "extreme-belief aversion," an aversion to holding beliefs that are close to certainty. Consider a discrete probability density function, $f_X(\cdot)$, that puts positive probability on a set of possible outcomes $x_1, x_2, ..., x_J$. We capture the idea of extreme-belief aversion by defining a mapping from the true probability density, $f_X(\cdot)$, to a subjective probability density that is less extreme,

$$f_X^{\xi}(x_i) = \frac{.5 + \xi \cdot (f_X(x_i) - .5)}{\sum_{j=1}^{J} (.5 + \xi \cdot (f_X(x_j) - .5))}.$$

The parameter $0 < \xi < 1$ describes the degree of extreme-belief aversion, with smaller values corresponding to greater bias. If $\xi = 1$, the subjective probabilities coincide with the true probabilities, while if $\xi = 0$, all outcomes $x_1, x_2, ..., x_J$ are treated as equally likely.

One interpretation of the transformed probability, $f_X^{\xi}(x)$, is that it represents the agent's truly-held beliefs. Another interpretation is that the agent actually holds beliefs $f_X(x)$ but *reports* beliefs that are transformed by the ξ function. While the latter is certainly plausible when Tommy's beliefs are very extreme—we can easily imagine a person saying she is 99% sure when her true belief is .9999—it does not address the evidence that beliefs inferred from betting behavior also exhibit under-inference.

While extreme-belief aversion does seem to help describe the evidence on people's beliefs, and it may be a confound for other interpretations of biased beliefs, we do not review evidence for extreme-belief aversion, and we know of no evidence that our crude formulation is a close match for people's thinking. Extreme-belief aversion may also lead to internal-inconsistency modeling challenges that we do not address. For example, it seems reasonable to assume that the transformation above could be applied to an agent's sampling distribution or to the agent's inferences, depending on which beliefs are being elicited. In that case, however, subjective sampling distributions and inferences will not in general be linked by Bayes' Rule.

Combining the three biases above with NBLLN gives a complicated model that captures many features that could be applied to predict beliefs and behavior in economic settings. One insight that comes immediately out of the combination is that base-rate neglect—i.e., underweighting priors—is not the opposite of, or contradictory to, the way NBLLN leads people to underweight likelihood information. Indeed, as we have noted in some discussions above about vividness and other biases, NBLLN is, especially in understanding "multi-clump" information processing, likely a contributor to the relevance of other biases. In many information-rich environments where full Bayesians would correctly become very confident independent of their priors, NBLLN is necessary for the question of whether people neglect base rates to be relevant. At the same time, we show above why NBLLN means that unless people completely neglect base rates, people's initial beliefs matter even in the long run. In combination, in fact, NBLLN suggests that it is possible that the real and often important fact that people under-use base rates may be consistent with the possibility that base rates matter more for social and economic phenomena than fully rational models have supposed that they do.

One last bias sits less comfortably with the others, and is harder to integrate. Many experiments make clear that peoples' subjective sampling distributions have flatter tails than our model of NBLLN by itself can explain. We attribute the flatness to "sampling-distribution-tails diminishing sensitivity (SDTDS)," a bias in which people perceive very unlikely outcomes as similar to each other and hence similar in probability. Consider 10 flips of a coin that is biased .8 in favor of heads. We know of no direct evidence but conjecture that most people would judge the likelihood of observing 1 head as very close to the likelihood of observing 0 heads, even though observing 1 head is actually 40 times more likely.

SDTDS can be formalized by assuming that an agent forms beliefs as if the likelihood of sample realizations far from the average are more similar to each other (and to the average) than they truly are. For a sample of size N from a θ-biased coin, let σ denote the standard deviation of the sample proportion $\frac{A_N}{N}$ (which equals $\sqrt{\frac{\theta(1-\theta)}{N}}$ for Tommy but not for Barney-Freddy). Let the perceived distance between the realized $\frac{A_N}{N}$ and θ be $\sigma\gamma\left(\frac{\frac{A_N}{N}-\theta}{\sigma}\right)$, where the twice-differentiable "sample-perception function" $\gamma : (-1,1) \to (-1,1)$ has the properties (1) $\gamma(0) = 0$, (2) $0 < \gamma' < 1$, (3) $\gamma''(x) < 0$ for all $x > 0$, and $\gamma''(x) > 0$ for all $x < 0$, and (4) $\gamma(x) = -\gamma(-x)$ for all x.

Property (1) says that the agent perceives a sample proportion of θ accurately. Property (2) ensures that the agent perceives any other sample proportion as more similar to θ than it actually is. The key concavity and convexity assumption (3) means that neighboring samples are perceived as more similar to each other the further they are from θ. For this reason, the agent makes little distinction between outcomes that fall far in the tails of his subjective sampling distribution. Property (4) specifies that γ is symmetric around 0 so that the misperception is symmetric around θ.

Roughly speaking, a person who exhibits SDTDS with sample-perception function γ judges the probability of a sample $s \in S_N$ with proportion $\frac{A_N}{N}$ of a-signals as if it were the sample $\Lambda(s)$, which has proportion $\theta+\sigma\gamma\left(\frac{\frac{A_N}{N}-\theta}{\sigma}\right)$ a-signals. Formally, $f^{\psi\phi\gamma}_{S_N|\Theta}(s|\theta)=\frac{f^{\psi\phi}_{S_N|\Theta}(\Lambda(s)|\theta)}{\sum_{s'\in S_N} f^{\psi\phi}_{S_N|\Theta}(\Lambda(s')|\theta)}$, where $f^{\psi\phi}_{S_N|\Theta}$ is the subjective sampling distribution for $\psi\phi$-Barney-Freddy, and the denominator is a normalization that ensures that the subjective sampling distribution adds up to 1.[49] While NBLLN by itself leads to subjective sampling distributions that have flat tails, SDTDS is an additional force for flat tails. When the rate θ is not .5, the mean of the agent's subjective sampling distribution no longer equals the mean of the objective sampling distribution because the "long tail" of the distribution is overweighted. Both of these features—flatter tails than can be accommodated with a reasonably-calibrated model of NBLLN and a mean shifted toward the long tail—are present in the subjective sampling distributions measured by Kahneman & Tversky (1972).

Unfortunately, this formulation of SDTDS cannot be so easily integrated with NBLLN—or with a model combining NBLLN with LSN and base-rate neglect—because it does not arm somebody with a theory of the sequence of signals, only the frequency of signals within a sample. While surely a form of it could be specified that embeds a concrete theory of permutations within the sample (allowing for instance a person to believe all different sequences are equally likely), we do not know that the existing evidence provides a guide, nor do we believe the psychology underlying it translates easily into situations where an agent is cognizant of the ordering of signals. Moreover, an improved model or a better formulation than we have found to interpret existing evidence may be more compatible with the models of other biases than we have supposed.

[49] In general, $\Lambda(s)$ could be a sample outside the support of the objective sampling distribution. For example, it may be a sample with 3.2 a's. The expression above is nonetheless well-defined as long as the density $f^{\psi\phi}_{S_N|\Theta}(\cdot|\theta)$ can be evaluated at that sample, as it can for a binomial objective sampling distribution.

Author(s)	Year	Prediction or Inference?	Subjects	Incentives?
Beach, Wise, & Barclay	1970	inference	169 male undergrads	no
Camerer	1987	inference	74 undergrads	financial market
Chinnis & Peterson	1968	inference	40 male undergrads	no
Dave & Wolfe	2003	inference	40 undergrads	BDM for probability
DeSwart	1972	inference	21 male undergrads	no
DeSwart	1972	inference	18 male undergrads	no
Donnell & DuCharme	1975	inference	24 male undergrads	no
Gettys & Manley: Study 1	1968	inference	20 undergrads	no
Gettys & Manley: Study 2	1968	inference	28 undergrads	no
Green, Halbert, and Robinson	1965	inference	32 grad students	paid for guess about state
Grether	1980	inference	341 undergrads	paid for guess about state
Grether: Studies 1 and 2	1992	inference	97 undergrads	paid for guess about state
Grether: Study 3	1992	inference	55 summer students	BDM for probability
Griffin & Tversky: Study 1	1992	inference	35 undergrads	paid for accurate posterior
Griffin & Tversky: Study 2	1992	inference	40 undergrads	paid for accurate posterior
Griffin & Tversky: Study 3	1992	inference	50 undergrads	no
Kahneman & Tversky: prediction	1972	prediction	unclear	no
Kahneman & Tversky: inference	1972	inference	560 high school students	no
Kraemer & Weber	2004	inference	51 students (most grad)	paid for accurate posterior
Marks & Clarkson	1972	inference	68 undergrads	no
Nelson, Bloomfield, Hales, & Libby: Study 1	2001	inference	27 MBA students	financial market
Peterson & Miller	1965	inference	42 undergrads	no
Peterson & Swensson: Study 1	1968	inference	15 male undergrads	no
Peterson & Swensson: Study 2	1968	inference	18 male undergrads	no
Peterson, DuCharme, & Edwards: Study 1	1968	prediction	41 male undergrads	no
Peterson, DuCharme, & Edwards: Study 2	1968	both	24 male undergrads	no
Peterson, Schneider, & Miller	1965	inference	44 undergrads	no
Phillips & Edwards: Study 1	1966	inference	5 male undergrads	no
Phillips & Edwards: Study 2	1966	inference	48 male undergrads	paid for accurate posterior
Phillips & Edwards: Study 3	1966	inference	48 male undergrads	no
Pitz	1967	inference	28 undergrads	no
Sanders	1968	inference	32 undergrads	bets on the state
Sasaki & Kawagoe	2007	inference	1033 employees	no
Strub	1969	inference	12 male undergrads	paid for guess about state
Teigen: Study 1	1974	prediction	22 undergrads	no
Teigen: Study 2	1974	prediction	73 undergrads	no
Wheeler & Beach	1968	both	17 male undergrads	paid for beliefs (prediction), and bets (inference)

Table 3.1: Experimental evidence on NBLLN for binomial signals.

3.9 Appendix B: Experimental Evidence

In this appendix, we report on all papers we could identify with experimental results related to the model's assumptions about subjective sampling distributions and predictions about inference for binomial signals. Table 3.1 lists the papers, which assumptions or predictions they test, their experimental subject population, and their incentive structure. Most of the studies we review did not incentivize subjects' responses; we will also discuss how the evidence from the few incentivized experiments relates to the unincentivized studies.

3.9.1 Evidence on Subjective Sampling Distributions

We begin by assessing the model's assumptions and predictions about subjective sampling distributions. We have found only 6 experiments from 4 papers in which researchers explicitly elicited experimental participants' beliefs about the likelihood of each possible sample.[50] None of these elicitations were incentivized. For 5 of these studies—Kahneman & Tversky (1972), Peterson, DuCharme & Edwards's (1968) Study 1, Wheeler & Beach (1968), and both of Teigen's (1974) studies—the data are displayed in the paper, and we have reproduced the graphs in Figures 3.1, 3.5, 3.6, and 3.2, respectively (we display both of Teigen's studies together), shown below in the order we discuss them.

Among the papers, Kahneman & Tversky (1972) elicited sample-proportion beliefs for the largest sample sizes. As discussed in the Introduction, they find that subjective sampling distributions are "constant in proportions" for $N = 10$, 100, and 1000. There is no noticeable tightening of the distribution even for $N = 1000$; while in fact there is less than a .01 chance of the proportion of heads falling outside the range 45% to 55%, subjects' distributions assign probability .79 to a proportion outside that range.[51]

A straightforward implication of the model is that subjective sampling distributions will be flatter than the objective sampling distributions. In all 7 experiments, with the exception of the $N = 3$ conditions of one experiment, the researchers indeed concluded that subjective sampling distributions are excessively close to uniform.[52]

[50]Cohen and Hansel (1955) also elicited subjective sampling distributions, but we cannot compare their data with our model because they did not tell their subjects the rate that was generating the signals.

[51]At the time of this draft, two of the authors (Benjamin and Rabin, in joint work with Don Moore) have also collected data on people's subjective sampling distributions for $N = 10$ and 1000. We designed the experiment to measure subjective sampling distributions in several different ways to deal with potential confounds such as extreme-belief aversion. When we elicit distributions in the same manner as Kahneman & Tversky, we replicate their results almost exactly. While our preliminary findings support NBLLN for $N = 1000$, we also find that the evidence for NBLLN for $N = 10$—and hence presumably also the evidence about smaller sample sizes reviewed below—is confounded by other explanations.

[52]Unlike the other 6 studies, Teigen (1974) asked subjects about the probability of each possible outcome separately, without requiring that the probabilities sum to 1. He found that the probabilities summed to greater than 1, and subjects often assigned probabilities that were too high to every outcome. When the subjective sampling distributions were normalized to sum to 1, they were too flat.

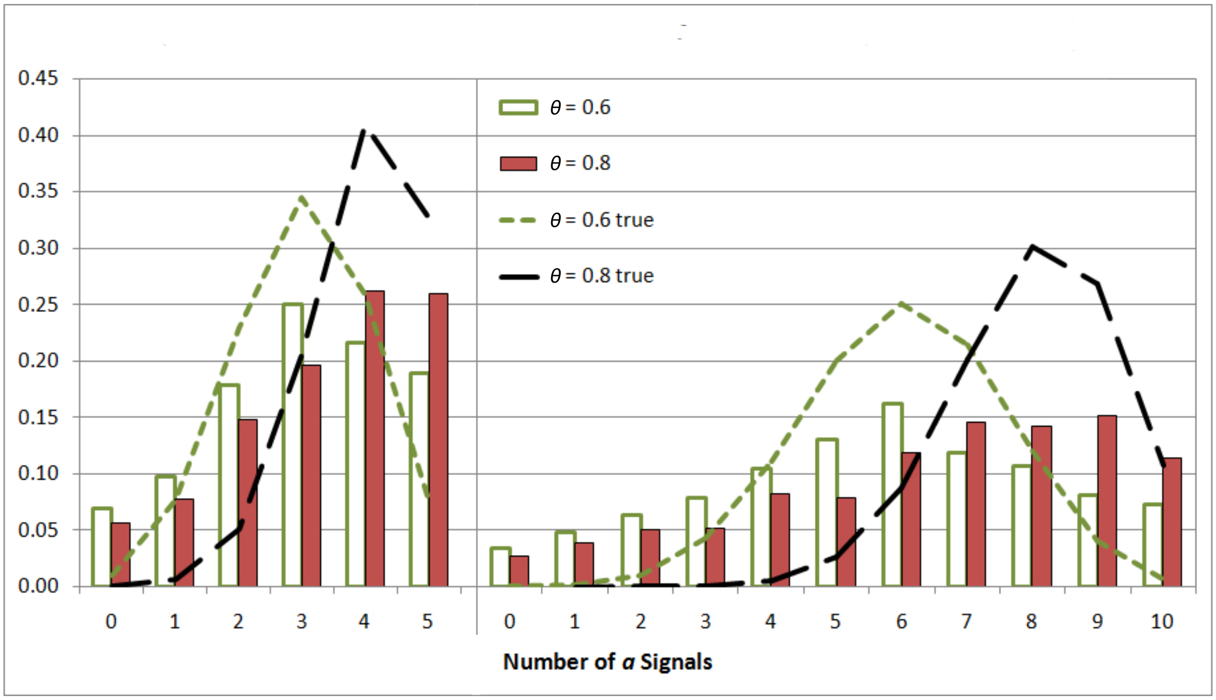

Figure 3.2: Median probability estimates, N=5 and N=10 (Teigen, 1974)

A feature of our model of NBLLN is that people have correct beliefs about samples of size 1. We know of no evidence on this point, but our own introspection suggests it is virtually a tautology that if a person knows the rate of a-signals is θ, the person thinks the probability is θ that a single draw will be a.[53] However, our model of NBLLN implies that for any sample size larger than 1, no matter how small, the subjective sampling distribution will be too close to uniform. Peterson, DuCharme & Edwards's (1968) Study 1 $N = 3$ conditions are the one case where researchers found subjective sampling distributions that are *not* too flat. For all three rates they studied ($\theta = .6, .7, .8$), the $N = 3$ subjective sampling distribution nearly coincides with the objective sampling distribution. We interpret this evidence as consistent with the combined effects of NBLLN and LSN, rather than either of the biases considered alone. The combined model we discuss in Appendix A predicts that subjective sampling distributions will be correct for $N = 1$ and too flat for N large. For plausible parameter values, there will be a non-monotonicity for intermediate N: due to LSN, the subjective sampling distribution may be too peaked when N is larger than 1 but small. The relative strength of NBLLN grows with N, so the subjective sampling distribution may be nearly correct for some

[53] Of course there is a good reason no one has done this experiment: the correct answer about the probability of an a-signal is obvious if the experimenter has just said that it is θ. Nonetheless, we discuss below evidence from experiments on inference from a single signal. The correct answer to an inference problem is not obvious, and indeed subjects often do not get the correct answer, tending somewhat to under-infer on average.

intermediate N.

Assumption A4 of our model is that the subjective sampling distribution has the same mean as the objective sampling distribution, regardless of sample size. In contrast, the evidence indicates that when $\theta \neq .5$, the mean of the subjective sampling distribution is generally between the objective mean, θN, and .5. Kahneman & Tversky (1972) explicitly comment that "the mean is displaced towards the long tail" (p.440), and this pattern is visually evident in all of the figures except the $N = 3$ cases discussed above. From a modeling perspective, we believe it is appropriate for our model to have the counterfactual feature of an accurate mean for the subjective sampling distribution because it allows us to draw out the implications of believing that the limit distribution has full support, without mixing in the implications of an inaccurate mean. Moreover, we speculate that the "displaced mean" is the result of a psychologically distinct bias, "sampling-distribution-tails diminishing sensitivity (SDTSD)," sketched in Appendix A, and our calibrated model of NBLLN generates subjective sampling distributions that come closer to matching the data in Figure 3.1 when we additionally incorporate SDTSD into the model (calculations not shown).

There is mixed evidence about whether training and feedback affects subjective sampling distributions. Wheeler & Beach (1968) elicited subjective sampling distributions for a sample of size $N = 8$ for rates $\theta = .6$ and .8; these are shown in Figure 3.6. Next, their subjects were faced with 100 asymmetric binomial inference problems ($\theta_A = .8$ and $\theta_B = .4$). After each problem, the subject was told the true rate for that problem. Subjective sampling distributions were elicited again, the subjects responded to 100 more symmetric binomial inference problems with feedback, and the subjective sampling distributions were elicited a final time.[54] Comparing the initial subjective sampling distributions with the final ones, the final subjective sampling distributions are less flat. For $\theta = .8$, the final subjective sampling distribution is quite close to the objective sampling distribution. For $\theta = .6$, the final subjective sampling distribution is actually more peaked than the objective sampling distribution. On the other hand, Peterson, DuCharme, and Edwards (1968, Study 2) conducted an experiment with four stages: (1) subjects were faced with symmetric binomial inference problems (with no feedback); (2) subjective sampling distributions were elicited for each combination of $N = 3, 5, 8$ and $\theta = .6, .7, .8$; (3) subjects were shown the objective sampling distributions; and (4) subjects were faced with another series of symmetric binomial inference problems. Subjects' responses in the inference problems were similar in stage 4 as in stage 1, suggesting that showing subjects the objective sampling distributions had little effect on beliefs.

[54]These subsequent elicitations are not shown in Figure 3.6.

3.9.2 Evidence on Inference

We have found 33 studies from 26 papers measuring inferences from samples about which of two equally-likely rates, θ_A and θ_B, generated the samples. Most of these binomial inference problems are symmetric in the sense that $\theta_A = 1 - \theta_B$. Unless otherwise noted, all the studies we mention are symmetric. We focus first on the studies where the prior probabilities, $f_\Theta(\theta_A)$ and $f_\Theta(\theta_B)$, are equal. Equal priors neutralizes the role of base-rate neglect. We study below how inferences are affected by unequal priors.

To compare the degree of under- or over-inference across studies, note that Bayes' Rule can be written as $\frac{f_{\Theta|S_N}(\theta_A|s)}{f_{\Theta|S_N}(\theta_B|s)} = \frac{f_{S_N|\Theta}(s|\theta_A)f_\Theta(\theta_A)}{f_{S_N|\Theta}(s|\theta_B)f_\Theta(\theta_B)}$. Since the signals are binomial, the rates are symmetric, and the priors are equal, Bayes' Rule can be expressed as $\frac{f_{\Theta|S_N}(\theta_A|s)}{f_{\Theta|S_N}(\theta_B|s)} = \left(\frac{\theta_A}{1-\theta_A}\right)^{\frac{2A_s-N}{N}\times N}$. Taking the natural log twice and rearranging,

$$\ln\ln\frac{f_{\Theta|S_N}(\theta_A|s)}{f_{\Theta|S_N}(\theta_B|s)} - \ln\left(\frac{2A_s-N}{N}\right) - \ln\ln\left(\frac{\theta_A}{1-\theta_A}\right) = \ln N. \tag{3.5}$$

It is possible in 9 of the papers to identify the value of θ_A for the inference problem, the actual sample observed by subjects, and subjects' mean or median reported posterior. Using the experimental data, Figure 3.3 plots the left-hand side of equation (3.5) against $\ln N$.[55] If the subjects' beliefs were Bayesian, the points should cluster along the identity line (the dashed line in Figure 3.3). The best-fitting regression line (the solid line in the figure) has a slope smaller than 1, indicating that subjects generally infer less in favor of rate θ_A than a Bayesian would.

To estimate the degree of under-inference and to probe its robustness, we rewrite (3.5) as a regression equation:

$$\ln\ln\frac{f_{\Theta|S_N}(\theta_A|s)}{f_{\Theta|S_N}(\theta_B|s)} = \gamma_0 + \gamma_1\ln N + \gamma_2\ln\left(\frac{2A_s-N}{N}\right) + \gamma_3\ln\ln\left(\frac{\theta_A}{1-\theta_A}\right) + \varepsilon. \tag{3.6}$$

[55] The left-hand side is well-defined only for inference problems such that $\frac{\theta_A}{1-\theta_A} > 1$ (that is, $\theta_A > .5$) and $\frac{A_s}{N} > \frac{1}{2}$ (that is, over half the realized signals are a's). Hence, as written, equation (3.5) only applies to such cases. Although this holds for only 63 of the 99 regression observations in Figure 3.3 and Table 3.2, we can include additional regression observations by relabeling the rates and sample proportions that we plug into formula (3.5). In inference problems such that $\frac{\theta_A}{1-\theta_A} > 1$ and $\frac{A_s}{N} < \frac{1}{2}$, we express Bayes' Rule as $\frac{f_{\Theta|S_N}(\theta_B|s)}{f_{\Theta|S_N}(\theta_A|s)} = \left(\frac{\theta_B}{1-\theta_B}\right)^{\frac{2A_s-N}{N}\times N} = \left(\frac{\theta_A}{1-\theta_A}\right)^{\frac{N-2A_s}{N}\times N}$, so equation (3.5) becomes $\ln\ln\frac{f_{\Theta|S_N}(\theta_A|s)}{f_{\Theta|S_N}(\theta_B|s)} - \ln\left(\frac{N-2A_s}{N}\right) - \ln\ln\left(\frac{\theta_A}{1-\theta_A}\right) = \ln N$. This allows us to use an additional 32 regression observations. Finally, we can use a further 4 regression observations for which $\theta_A < .5$ and $\frac{A_s}{N} < \frac{1}{2}$ by expressing Bayes' Rule as $\frac{f_{\Theta|S_N}(\theta_A|s)}{f_{\Theta|S_N}(\theta_B|s)} = \left(\frac{\theta_A}{1-\theta_A}\right)^{\frac{2A_s-N}{N}\times N} = \left(\frac{1-\theta_A}{\theta_A}\right)^{\frac{N-2A_s}{N}\times N}$, and we can take the log-log of this equation. In the case of 3 inference problems, $\frac{A_s}{N} = \frac{1}{2}$, so the Bayesian posterior ratio is equal to 1, and it is impossible to define what constitutes "over-inference" or "under-inference." Those 3 datapoints are dropped from the Figure 3.3 and Table 3.2.

	(1) Restriction: $\gamma_1 = \gamma_2 = \gamma_3 = 1$	(2) Restriction: $\gamma_2 = \gamma_3 = 1$	(3) Restriction: $\gamma_3 = 1$	(4) Coeffs unrestricted	(5) Coeffs unrestricted Incentivized only
$\ln N$		0.587	0.470	0.485	0.675
		(0.083)	(0.109)	(0.085)	(0.086)
$\ln\left(\frac{2A_s-N}{N}\right)$			0.776	0.914	0.957
			(0.116)	(0.109)	(0.121)
$\ln\ln\left(\frac{\theta_A}{1-\theta_A}\right)$				0.333	0.421
				(0.101)	(0.104)
Constant	-0.753	-0.031	-0.002	-0.145	-0.168
	(0.076)	(0.143)	(0.142)	(0.108)	(0.130)
R^2	0.000	0.426	0.360	0.603	0.763
#obs	99	99	99	99	47
#papers	9	9	9	9	3

Table 3.2: Inference with symmetric rates and equal priors.

The null hypothesis of Bayesian updating is $\gamma_0 = 0$, $\gamma_1 = \gamma_2 = \gamma_3 = 1$. In Table 3.2, we estimate versions of equation (3.6) with several different restrictions on the coefficients and data.

Notes: Results are from OLS regressions, with standard errors in parentheses. The dependent variable is as described in the text. Coefficients for blank entries are restricted to equal 1. The fifth column restricts the data to incentivized experiments.

Column 1 estimates just γ_0, under the restriction that $\gamma_1 = \gamma_2 = \gamma_3 = 1$. The estimate, $\widehat{\gamma}_0 = -.753$, is significantly smaller than zero, indicating that the average pattern is under-inference.

Column 2 estimates both γ_0 and γ_1, while restricting $\gamma_2 = \gamma_3 = 1$. The predicted regression line is plotted as the solid line in Figure 3.3. The regression confirms that the degree of under-inference is related to sample size; $\widehat{\gamma}_1 = .587$ is significantly smaller than 1 (but greater than 0). Moreover, once the degree of under-inference is allowed to depend on sample size, there is no residual under-inference to be picked up by the constant term: $\widehat{\gamma}_0 = -.031$, which is not statistically distinguishable from zero. Breaking down the data by study (not shown in the table), every study that manipulates sample size, while holding constant other features of the experiment—Peterson, Schneider, and Miller (1965); Pitz (1967); Peterson, DuCharme, & Edwards's (1968) Study 2; Kahneman & Tversky (1972); Griffin & Tversky's (1992) Study 1; Nelson, Bloomfield, Hales, & Libby's (2001) Study 1; Kraemer & Weber (2004)—concludes that there is greater under-inference in larger samples.[56]

[56]For Green, Halbert, & Robinson (1965), we can also reach this conclusion by estimating the regression equation (3.6) on the data reported just in that paper. There are a number of other studies that manipulate sample size but do not analyze or display the data in a way that makes it clear how sample size affects the degree of bias in inference: Sanders (1968); Peterson & Swensson's (1968) Study 2; Beach, Wise, & Barclay (1970); Marks & Clarkson (1972); and DeSwart (1972a, 1972b).

Qualitatively, extreme-belief aversion can explain both the excessively-dispersed subjective sampling distributions discussed in section B.1 and the under-inference from column 1. An agent with extreme-belief aversion will have a more dispersed subjective sampling distribution than Tommy in large samples by virtue of compressing beliefs away from 0 and 1. An agent with extreme-belief aversion will also under-infer in situations where Tommy's inference would be extreme, as would almost always occur when the sample is large. Extreme-belief aversion taken alone, however, implies that in any two inference problems where Tommy's posterior is the same, people would hold (or at least report) the same belief. There is evidence contradicting this implication, indicating that extreme-belief aversion is not the *only* deviation from Bayesian belief formation that is going on.

For example, consider the experiment in Griffin & Tversky (1992), where $\theta_A = .6$. Tommy's inference depends only on the difference between the number of a-signals and the number of b-signals. However, when the sample is 4 a's and 1 b, subjects' median belief in favor of θ_A is .80, while when the sample is 10 a's and 7 b's, subjects' median belief in favor of θ_A is .60. Tommy's belief would be .77 in both cases, so people are under-inferring from the sample of size 17 and actually slightly over-inferring from the sample of size 5. Consistent with Proposition 5—but inconsistent with extreme-belief aversion being the only bias in beliefs—people infer less from the same difference in a and b signals when the sample is larger. Kahneman & Tversky (1972) and Kraemer & Weber (2004) also report evidence that beliefs are sensitive to sample size, holding constant the difference in the number of a and b signals.

Column 3 of Table 3.2 relaxes the regression model still further, estimating γ_0, γ_1, and γ_2, with the only remaining restriction being $\gamma_3 = 1$. The coefficient on the proportion of a signals, $\widehat{\gamma}_2 = .776$, is just statistically distinguishable from 1, and the estimates $\widehat{\gamma}_0$ and $\widehat{\gamma}_1$ are not substantially different in column 3 compared with column 2.

Column 4 estimates all four coefficients. Qualitatively, the main effect of relaxing the $\gamma_3 = 1$ restriction on the conclusions from column 3 is that the coefficient on the proportion of a-signals, $\widehat{\gamma}_2 = .914$, is no longer statistically distinguishable from 1. Hence we cannot reject the Bayesian null hypotheses that $\gamma_0 = 0$ and $\gamma_2 = 1$, while $\widehat{\gamma}_1$ is significantly smaller than 1. However, column 4 makes clear that the extremeness of the rates—the extent to which $\theta_A = 1 - \theta_B$ departs from .5—also matters for the degree of biased inference. The coefficient $\widehat{\gamma}_3 = .333$ is much smaller than 1. This means that subjects under-infer by more the further is θ_A from .5. Breaking down the data by study (not shown in the table), every study that manipulates θ_A, while holding constant other features of the experiment either concludes that there is greater under-inference for θ_A further from .5 (Green, Halbert, & Robinson, 1965; Phillips & Edwards's, 1966 Study 1 and 3; Peterson & Miller, 1965; Peterson & Swensson's, 1968 Studies 1 and 2; Peterson, DuCharme, & Edwards's, 1968 Study 2; Sanders, 1968; Donnell & DuCharme, 1975; Kahneman & Tversky, 1972) or finds it without explicitly stating it (Chinnis & Peterson, 1968; Beach, Wise, & Barclay, 1970; Shu & Wu, 2003).[57] Griffin & Tversky (1992) use the term "discriminability" to describe the phenomenon of under-inference becoming more severe when θ_A and θ_B are further apart. In Griffin & Tversky's (1992; Study 3) particularly clear evidence from asymmetric inference problems, subjects were asked to infer the likelihood of rate θ_A where the rates have equal priors, the sample has size $N = 12$, and number of a-signals is $A_s = 7$, 8, 9, or 10. When the rates are close together, $(\theta_A, \theta_B) = (.6, .5)$, the subjects exhibit slight over-inference: a Bayesian's posteriors for these four inference problems would be .54, .64, .72, and .80, respectively, while subjects' median posteriors were .55, .66, .75, and .85. When the rates are further apart, $(\theta_A, \theta_B) = (.6, .25)$, subjects exhibited massive under-inference: whereas a Bayesian's posteriors in these problems would be .95, .98, .998, and .999, respectively, subjects' posteriors were .60, .70, .80, and .90. Griffin & Tversky's evidence cannot be fully explained by extreme-belief aversion (that simply maps an objective posterior into a less extreme subjective one) because, for example, subjects' posterior of .80 is identical whether $\theta_B = .5$ and 10 a-signals were observed or $\theta_B = .25$ and 9 a-signals were observed, but the objective posteriors are quite different in those two cases. Our interpretation of Griffin & Tversky's evidence is that realized samples far in the tails of the subjective sampling generate extreme under-inference because the tails are excessively flat. While NBLLN generates flat tails qualitatively, our one-parameter model of NBLLN cannot fit this evidence quantitatively at plausible calibrated values of ψ. We believe SDTSD captures the right psychology for explaining flat tails and likely explains both Griffin & Tversky's discriminability evidence and the evidence from a number of other studies that find $\gamma_3 < 1$.

Column 5 estimates the same regression as column 4, but with the data restricted to incentivized experiments. There are only three such studies—Green, Halbert, & Robinson (1965), Nelson, Bloomfield, Hales, & Libby (2001), and Kraemer & Weber (2004)—but the results in column 5 are largely similar to column 4, except that both $\widehat{\gamma}_1 = .675$ and $\widehat{\gamma}_3 = .421$ are larger, suggesting greater sensitivity to sample size and to rates when accurate inferences are rewarded. Nonetheless, both coefficients remain far less than 1, indicating substantial biases relative to Bayesian inference.

Training in inference appears to reduce but not eliminate under-inference. When subjects were told after each inference which state actually occurred, they became less biased over time but still under-inferred by the end of the experiment (Phillips & Edwards's, 1966, Study 2; Camerer, 1987).

[57] DeSwart (1972a, 1972b) manipulates how far θ_A is from .5 but does not analyze or display the data in a way that makes it clear how it affects the degree of under-inference.

Strub (1969) found under-inference among subjects who had received 114 hours of lecture sessions, demonstrations, problem-solving sessions, and other training in dealing with probabilities, including prior participation in inference experiments. When subjects were told after each inference what the normatively correct inference is, they very quickly learned to report more extreme beliefs, but they do not seem to have learned to draw better inferences. While reporting more extreme beliefs led to more accurate beliefs in problems where their pre-training beliefs were not extreme enough, it led to less accurate beliefs in problems where their pre-training beliefs were accurate (Donnell & DuCharme, 1975).

A few studies have found over-inference. In all cases, N is relatively small, and θ_A is relatively close to .5.[58] Griffin & Tversky's Study 1 (1992; $\theta_A = .6$) , which compared inference from samples of size 3, 5, 9, 17, and 33, found over-inference for $N = 3$ and 5 and under-inference for the others. Nelson, Bloomfield, Hales, & Libby's Study 1 (2001; $\theta_A = .6$) conducted an experimental asset market where payoffs depended on correct inferences from samples of size 3 and 17. Subjects under-inferred for $N = 17$ and over-inferred for $N = 3$. There is some evidence, though it is weak, that over-inference in favor of a particular state occurs when the realized sample exactly matches the expected sample in that state, a phenomenon that has been called "exact representativeness." In an experimental asset market, Camerer (1987; $N = 3$, $\theta_A = .67$) found that the price of a state-contingent asset that pays off if state A is true was too high—indicating over-inference in favor of state A—when the observed sample contained exactly 2 a-signals and 1 b-signal, and symmetrically over-inference in favor of state B when the sample proportions were reversed. In an experimental asset market with asymmetric rates of $\theta_A = .67$ and $\theta_B = .5$, Grether (1980; $N = 6$) similarly found evidence indicating over-inference in favor of state A when the realized sample was 4 a's and 2 b's and over-inference in favor of state B when the realized sample was 3 a's and 3 b's. In a similar experiment, Grether (1992) found less support for "exact representativeness." Neither NBLLN alone, nor NBLLN combined with SDTSD, can explain over-inference.

Such over-inference can, however, be explained by the Law of Small Numbers (LSN). As shown in Appendix A, when NBLLN and LSN are combined in a single model, small N is a necessary condition for over-inference. Moreover, we have argued that SDTSD will tend to generate under-inference when θ_A and θ_B are far apart.

A distinctive feature of our theory—a feature that differentiates it from alternative theories of under-inference from large samples discussed in Section 7 and Appendix A—is the prediction that

[58] Peterson & Swennson's (1968) Study 1 finds over-inference for $N = 1$ and $\theta_A = .6, .67, .75, .9$ in the first half of their data. In the same inference problems in the second half of their data from Study 1, and in both halves from Study 2, however, they find under-inference.

inferences from a sample of size 1 will be correct. Sample-size neglect predicts over-inference for samples of size 1, while extreme-belief aversion predicts under-inference for samples of any size, including 1. There are 11 experiments that measure inference when $N = 1$. Peterson, Schneider, & Miller (1965; $\theta_A = .6$), Dave & Wolfe (2003; $\theta_A = .7$), Peterson & Swennson's (1968; $\theta_A = .6$, .67, .75, .9) Study 2, and Gettys & Manley's (1968) Studies 1 and 2 (which used a variety of asymmetric inference problems) found substantial under-inference. Chinnis & Peterson (1968; $\theta_A = .67$, .8), Kraemer & Weber (2004; $\theta_A = .6$), and Sasaki & Kawagoe (2007; $\theta_A = .67$) found slight under-inference, very close to Bayesian, and Peterson & Swennson's (1968; $\theta_A = .6$, .67, .75, .9) Study 1 found over-inference in the first half of their data and under-inference in the second half. In a mix of symmetric and asymmetric problems, Peterson & Miller (1965) found under-inference for $(\theta_A, \theta_B) = (.83, .17)$, $(\theta_A, \theta_B) = (.71, .2)$, and $(\theta_A, \theta_B) = (.67, .33)$, and over-inference for $(\theta_A, \theta_B) = (.6, .43)$. Green, Halbert, and Robinson (1965; $\theta_A = .6$, .8) found inferences very close to Bayesian when $N = 1$. The evidence is mixed but with more of the studies leaning toward under-inference.[59]

Notes: The x-axis is depicted on the natural log scale. For all datapoints in the figure, subjects knew that prior probabilities of the two rates were equal. The dotted line represents the null hypothesis of Bayesian updating, and the solid line is the best-fitting regression line from column 2 of Table 3.2. The includes studies are: BWB = Beach, Wise, & Barclay (1970); GHR = Green, Halbert, and Robinson (1965); G3 = Grether's (1992) Study 3; GT = Griffin & Tversky's (1992) Study 1; KT = Kahneman & Tversky (1972); KW = Kraemer & Weber (2004); NBHL = Nelson, Bloomfield, Hales, & Libby's (2001) Study 1; PM = Peterson & Miller (1965); SK = Sasaki & Kawagoe (2007).

As an aside, we can use the same balls-and-urns experiments to study how priors affect inference and thereby measure the prevalence and extent of base-rate neglect in these studies. Taking the log

[59] Presumably, the many papers on base-rate neglect also contain evidence on inferences from samples of size 1. Virtually none of them have 50-50 priors, however, so it is difficult to disentangle biased inference from base-rate neglect. We do not review this literature systematically, but to give a flavor of what it may indicate, we examined Bar-Hillel's (1980) seminal paper. Our impression is that the evidence in Bar-Hillel's paper roughly mirrors the evidence from experiments on single-signal inference reviewed above. The full distribution of subjects' reported posteriors can be eyeballed from histograms reported for Bar-Hillel's Studies 1, 2, 3, 7 and 8, each of which presents an inference problem where a single signal is indicative of the less likely of two states that subjects are given base rates for. We divide the 222 subjects' responses into four categories. Because the signal strength always was in the opposite direction of the base rate, the 33% of subjects whose posteriors equaled the base rate or weaker must have been either under-inferring from the signal or (as is presumably unlikely) "over-using" the base rate. By contrast, 9% of subjects reported posteriors stronger than the signal, almost surely indicating over-use of the signal (since otherwise they must be reversing the base rate). 31% of subjects reported posteriors of exactly the signal strength. Although not logically necessary, we share the presumption of Bar-Hillel and most researchers in this area that these subjects were almost surely simply using the signal strength and ignoring the base rate altogether. The remaining 27% of subjects reported posteriors strictly between the correct Bayesian posteriors and the posteriors that would completely ignore the base rate. It is unclear how many of these subjects were over-using or under-using the signal because anyone under-using the base rate could have been either over-inferring or under-inferring from the signal. From these data taken together, it seems likely that between 9% and 36% of the subjects were over-inferring from the signal, and at least 33% of the subjects were under-using the signal.

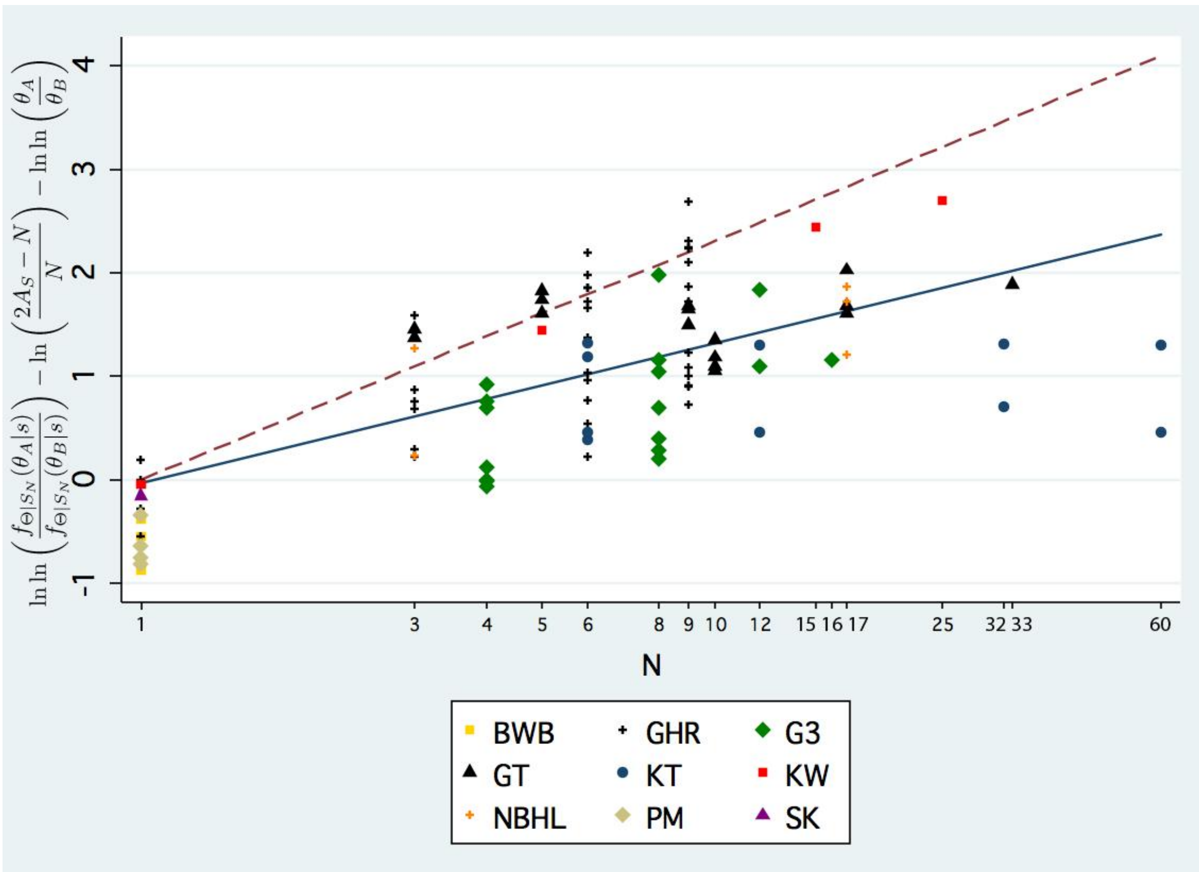

Figure 3.3: Inference with symmetric rates and equal priors

of Bayes' Rule and rearranging:

$$\ln \frac{f_{\Theta|S_N}(\theta_A|s)}{f_{\Theta|S_N}(\theta_B|s)} - \ln \frac{f_{S_N|\Theta}(s|\theta_A)}{f_{S_N|\Theta}(s|\theta_B)} = \ln \frac{f_{\Theta}(\theta_A)}{f_{\Theta}(\theta_B)}. \tag{3.7}$$

As above, from the published experiments, we obtain for each inference problem the values θ_A, θ_B, $f_{\Theta}(\theta_A)$, the actual sample observed by subjects, and subjects' mean or median reported posterior. The right-hand side of equation (3.7) can be readily calculated from $f_{\Theta}(\theta_A)$, as can the first term on the left-hand side, $\ln \frac{f_{\Theta|S_N}(\theta_A|s)}{f_{\Theta|S_N}(\theta_B|s)}$, from subjects' posterior. In order not to confound base-rate neglect with other biases (that affect inference even when priors are equal), we calculate the second term on the left-hand side, $\ln \frac{f_{S_N|\Theta}(s|\theta_A)}{f_{S_N|\Theta}(s|\theta_B)}$, as the predicted value, $\ln \frac{\widehat{f_{S_N|\Theta}(s|\theta_A)}}{f_{S_N|\Theta}(s|\theta_B)}$, from the previously-estimated regression equation (3.6). We use only the symmetric-rate data. Figure 3.4 plots the left-hand side of equation (3.7) against $\ln \frac{f_{\Theta}(\theta_A)}{f_{\Theta}(\theta_B)}$. Regardless of whatever other biases may affect inference, if the subjects correctly incorporate base rates into their inferences, the data should lie along the identity line (the dashed line in Figure 3.4). However, the best-fitting line (the solid line in the figure) has a slope less than 1, indicating that subjects' inferences are too insensitive to the prior probabilities.

To formally investigate the degree of base-rate neglect, we rewrite (3.7) as a regression equation:

$$\ln \frac{f_{\Theta|S_N}(\theta_A|s)}{f_{\Theta|S_N}(\theta_B|s)} - \ln \frac{f_{S_N|\Theta}(s|\theta_A)}{f_{S_N|\Theta}(s|\theta_B)} = \varphi_0 + \varphi_1 \ln \frac{f_{\Theta}(\theta_A)}{f_{\Theta}(\theta_B)} + \varepsilon. \tag{3.8}$$

The null hypothesis of Bayesian updating corresponds to $\varphi_0 = 0$ and $\varphi_1 = 1$.

The first column of Table 3.3 estimates equation (3.8) on the full sample. While we can reject the null hypothesis that $\varphi_0 = 0$, it is clear from Figure 3.6 that the deviation is small. On the other hand, the fitted value $\widehat{\varphi}_1 = .601$ is significantly less than 1 (and larger than 0). The second column estimates the equation using only the data with unequal prior probabilities, but the results are similar. Finally, the third column restricts the data to incentivized experiments. While $\widehat{\varphi}_0$ is now essentially zero, $\widehat{\varphi}_1 = .405$ indicates even stronger base-rate neglect in these data.

Notes: Results are from OLS regressions, with standard errors in parentheses. The dependent variable is as described in the text.

Notes: The dotted line represents the null hypothesis of Bayesian updating, and the solid line is the best-fitting regression line from column 1 of Table 3.3. The includes studies are: BWB = Beach, Wise, & Barclay (1970); C = Camerer (1987); GHR = Green, Halbert, and Robinson (1965); G3 = Grether's (1992) Study 3; GT = Griffin & Tversky's (1992) Study 1; KT = Kahneman & Tversky (1972); KW = Kraemer & Weber (2004); NBHL = Nelson, Bloomfield, Hales, & Libby's (2001)

	(1) All data	(2) Only unequal priors	(5) Only incentivized
$\ln \frac{f_\Theta(\theta_A)}{f_\Theta(\theta_B)}$	0.524	0.510	0.360
	(0.078)	(0.081)	(0.115)
Constant	0.216	0.299	-0.130
	(0.053)	(0.087)	(0.084)
R^2	0.242	0.303	0.176
#obs	209	110	48
#papers	10	5	3

Table 3.3: Base-rate neglect.

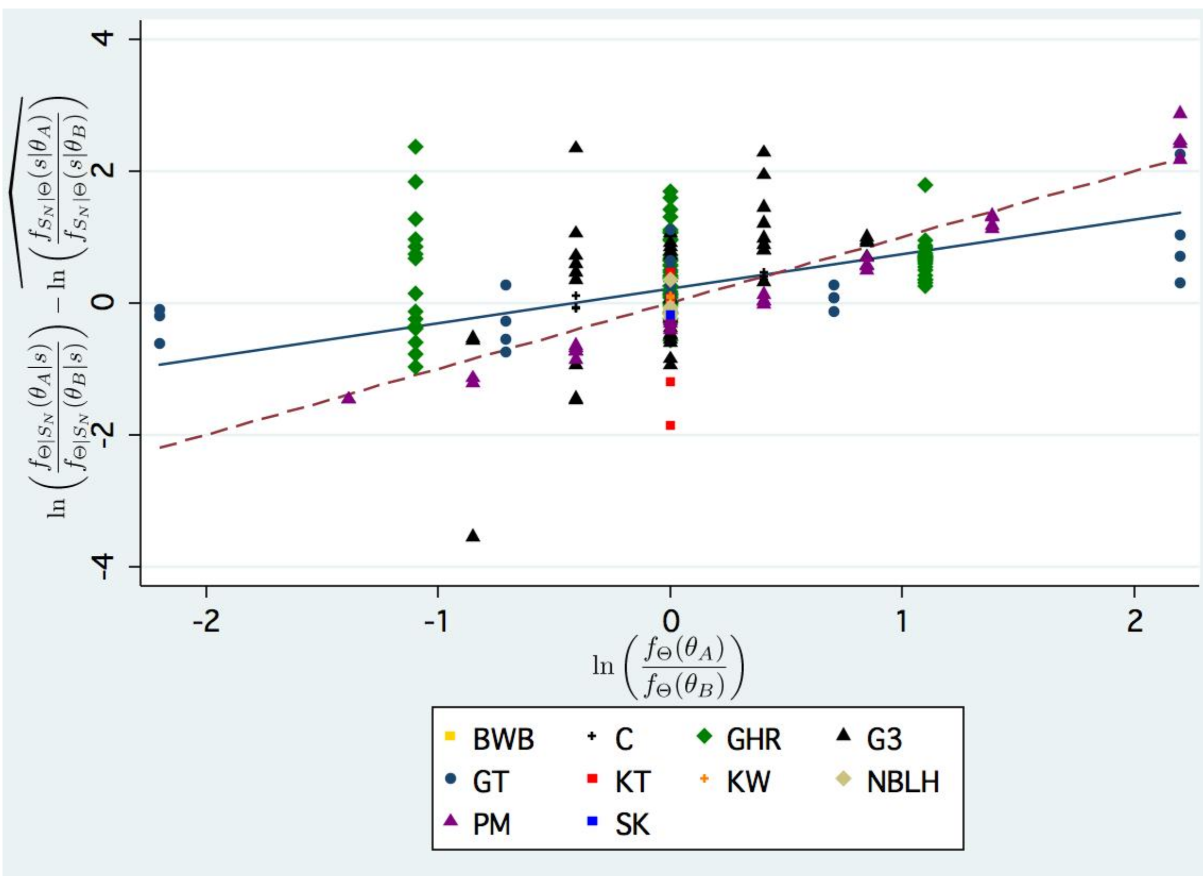

Figure 3.4: Base-rate neglect in symmetric-rate inference problems.

Study 1; PM = Peterson & Miller (1965); SK = Sasaki & Kawagoe (2007).

3.9.3 Consistency Between Subjective Sampling Distributions and Inference

A core feature of the fully rational model that our model retains is that people draw inferences that are consistent with Bayes's Rule applied to their subjective sampling distributions. Although surely not a perfect fit, we believe (like other researchers before us) that this feature is approximately right except insofar as people neglect base rates. There is indirect supportive evidence from the qualitative correspondence between the evidence on subjective sampling distributions (reviewed in section B.1) and the evidence on inferences (reviewed in section B.2). There is direct evidence from two studies that measured subjective sampling distributions and inferences for the same subject.

Peterson, DuCharme, & Edwards (1968, Study 2) conducted symmetric inference experiments with every combination of $N = 3, 5, 8$, and $\theta_A = 1 - \theta_B = .6, .7, .8$. Then subjects drew subjective sampling distributions for the nine binomial distributions (shown in Figure 3.5). Peterson, DuCharme, & Edwards plotted subjects' inferences against what their inferences would be if they applied Bayes's Rule to their subjective sampling distributions. Peterson, DuCharme, & Edwards found that "most points cluster extremely close to the identity line."

Wheeler & Beach (1968) elicited subjects' subjective sampling distributions for a sample of size $N = 8$ for rates $\theta = .6$ and $.8$ (see Figure 3.6) and then asked subjects to make bets in an inference task. Wheeler & Beach inferred subjective posteriors from the bets, under the assumption that subjects sought to maximize expected winnings. The correlation between the median subjective log-posterior (calculated from the first 20 inference problems) and the median subjective log-likelihood (calculated from applying Bayes's Rule to the subjective sampling distribution elicited at the very beginning of the experiment) was .90. For individual subjects' data, the median correlation was .85.

3.9.4 Inference For Sequential Clumps

In Section 4, we lay out various possible dynamic extensions of our model of NBLLN, but there are few experiments aimed at comparing inferences from a sample presented simultaneously with a sample drawn sequentially, and there are no experiments that elicit people's beliefs about what they will infer conditional on observing samples in the future.

Grether (1992, Study 3) confronted subjects with (incentivized) symmetric, binomial inference problems, with rates $\theta_A = 1 - \theta_B = .2, .3, .4, .6, .7, .8$, and priors $f_\Theta(\theta_A) = .3, .4, .5, .6, .7$. The sample size always began as $N = 4$. In some cases, however, after subjects made their inference,

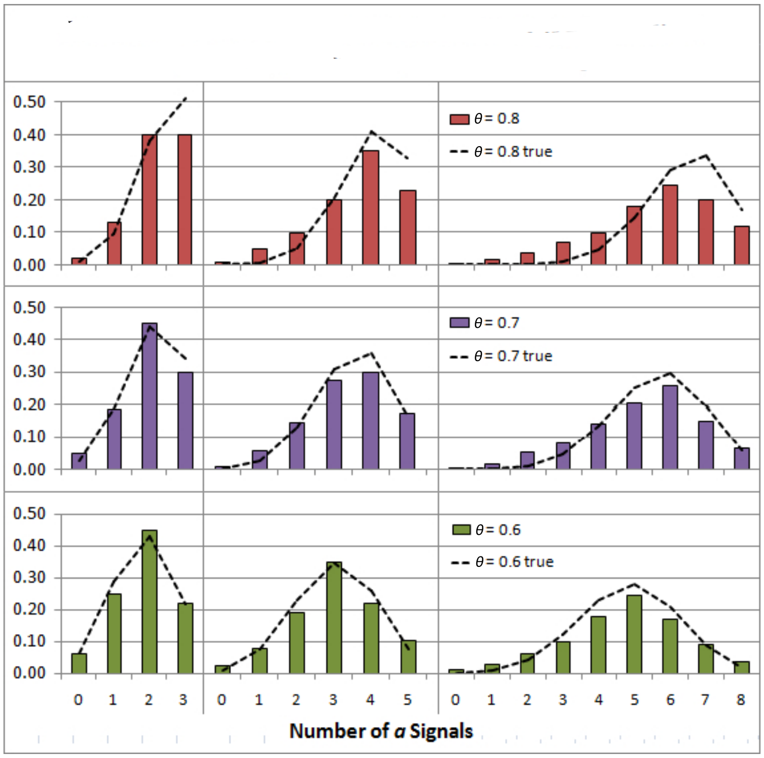

Figure 3.5: Median probability estimates, N=3, N=5, and N=8 (Peterson DuCharme and Edwards 1968)

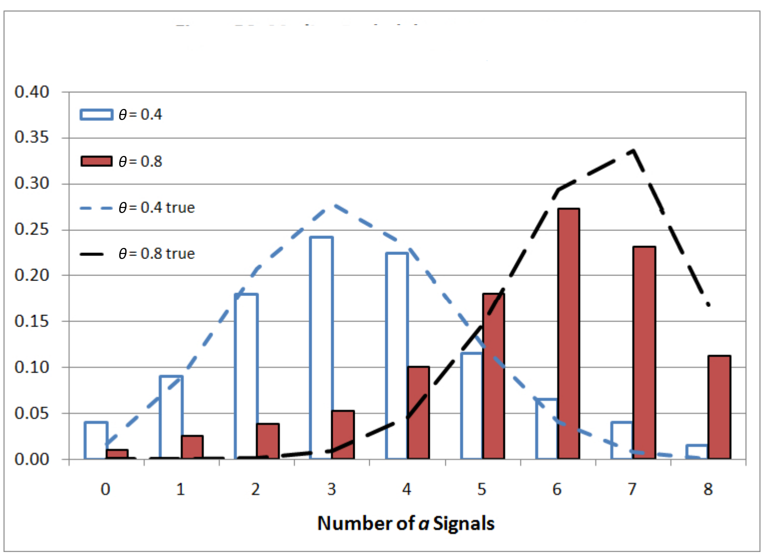

Figure 3.6: Median probability estimates, N=8 (Wheeler and Beach 1968)

they were asked to make an updated inference after an additional 4 signals were drawn, up to a maximum of 12 signals in total. Although only aggregate statistics are reported in the paper, David Grether sent us the subject-level data. In a few cases, we can learn about how subjects process the signals by comparing their inferences before and after they receive a clump. For example, in one situation, the rates were $\theta_A = 1 - \theta_B = .2$, the prior probabilities of the rates were equal, and the first four signals were all b's. The next four signals were 2 a's and 2 b's. The objective posterior probability of rate θ_B is the same after all eight as after the first four: .9961. However, the subjects' subjective posterior (that is, the median across subjects) is .95 after the first four signals and .70 after all eight. This pattern of additional uninformative signals causing the subjective posterior to move toward .5 is consistent with retrospective-pooling Barney, not retrospective-acceptive Barney. The same pattern holds in all three other test cases in Grether's data.[60]

In contrast, Kraemer & Weber (2004) present evidence that supports retrospective-acceptive Barney. For an incentivized, symmetric binomial inference problem with $\theta_A = 1 - \theta_B = .6$ and $N = 5$, subjects presented with a sample of 3 a's and 2 b's gave mean posterior probability for rate θ_A of .585. Other subjects who were instead shown the same signals as two separate samples, one with 3 a's and 0 b's and one with 0 a's and 2 b's, gave mean posterior probability for state A of .56, which is marginally statistically different. Similarly, with $\theta_A = 1 - \theta_B = .6$ and $N = 25$, subjects' mean posterior probability for rate θ_A was .56 when the sample was 13 a's and 12 b's, but .53 (strongly statistically distinguishable from .56) when two samples of 13 a's and 0 b's, and then 0 a's and 12 b's were presented sequentially. The fact that subjects make different inferences in these two cases is inconsistent with being retrospective-pooling, but not inconsistent with being retrospective-acceptive.[61]

[60]In the 2^{nd} test case, $f_\Theta(\theta_A) = .4$, $\theta_A = 1 - \theta_B = .3$, and the first eight signals were 1 a and 7 b's. The next four signals were 2 a's and 2 b's. While the objective posterior probability of rate θ_B is identically .9908, subjects' median subjective posterior fell from .9 to .5. In the 3^{rd} test case, $f_\Theta(\theta_A) = .7$, $\theta_A = 1 - \theta_B = .4$, and the first four signals were 2 a's and 2 b's, and the next four signals were 2 a's and 2 b's. The objective posterior probability of rate θ_A remains .7, but subjects' median subjective posterior fell from .775 to .7. In the 4^{th} test case, $f_\Theta(\theta_A) = .3$, $\theta_A = 1 - \theta_B = .7$, and the first four and next four signals were 2 a's and 2 b's. The objective posterior probability of rate θ_A remains .3, but subjects' median subjective posterior increased from .42 to .51. If we examine subjects' mean posterior probabilities, the pattern is robust in the first three test cases (.87 to .66, .72 to .59, and .73 to .68) and hard to interpret in the fourth (.54 to .58).

[61]Unlike the experimental subjects, however, retrospective-acceptive Barney will under-infer weakly *more* from two clumps than from one. This is because an additional clump provides information about an additional draw of β from the distribution, which is weakly more informative about the true rate than an additional signal about a fixed β. That being said, a combination of retrospective-acceptive Barneyness with SDTSD could explain why the two extreme clumps lead to a weaker overall inference than the single clump.

Kraemer & Weber (2004) also have a third experimental treatment that generates ambiguous evidence. In the $N = 5$ case, when subjects are presented with a sample of 1 a and 1 b, followed by a sample of 2 a's and 1 b, their mean posterior is .575, in between the other two treatments and not statistically distinguishable from either. Similarly, in the $N = 25$ case, the mean posterior is .555 when subjects are presented with a sample of 6 a's and 6 b's, followed by a sample of 7 a's and 6 b's.

Shu & Wu's (2003) Study 3 appears to be inconsistent with any of the dynamic extensions of our model that we consider. They conduct a symmetric binomial inference problem with samples of size $N = 10$ and three different levels of the rates, $\theta_A = 1 - \theta_B = .6$, .75, or .9. In one condition, subjects observed the 10 signals one at a time before stating a posterior belief. In the other conditions, subjects observed the 10 signals in clumps of 2 signals each or 5 signals each. While for some realizations of the 10-signal sets subjects draw less extreme inferences when the signals arrive in larger clumps—as predicted for retrospective-acceptive Barney—the results on average go in the opposite direction.[62] While it may be possible to reconcile Shu & Wu's results with a combination of NBLLN and the dynamics of base-rate neglect, that combined model should be worked out to see if it systematically reverses some of the conclusions we reach in Section 5.

3.9.5 Evidence For Non-Binomial Distributions

While the vast majority of simple inference experiments have been conducted with binomial signals, there are a few studies with other distributions. The results overall are consistent with NBLLN and SDTSD applying beyond binomial subjective sampling distributions.

There are a handful of studies where signals are multinomial. For example, in Beach's (1968) experiment, there were two decks of cards, a Red Deck and a Green Deck. Each card had a letter from A to F written on it. The Red and Green Decks had equal priors, but each deck had different proportions of the lettered cards. The subjects were shown $N = 3$ cards, one card at a time, and reported their subjective probability that the cards were being drawn from the Red Deck, as opposed to the Green Deck, after each draw. The likelihood ratios for the cards ranged from 1:2.5 to 3:1. For example, the likelihood ratio for card F was 1:2. A second group of subjects faced the same inference task with the same likelihood ratios for each card, but with the absolute probabilities scaled down for some cards and scaled up for others. For example, for the first group of subjects, the probability of an F card was .03 for the Red Deck and .06 for the Green Deck; for the second group of subjects, the probability of an F card was .16 for the Red Deck and .32 for the Green Deck. The first main finding is that subjects under-inferred on average. The other finding was that, for a given objective likelihood ratio, Group 1 under-inferred more for cards where Group 1's probabilities for that card were scaled down relative to Group 2's. Our interpretation is that when Group 1's probabilities are scaled down, the observed sample lies further in the tails on the subjective sampling distribution for

[62] Sanders (1968) and Beach, Wise, & Barclay (1970) also compare inferences from simultaneously-presented samples with inferences from sequentially-presented samples, but it is difficult to interpret their results because the results for the simultaneously-presented samples are averaged across different sample sizes.

both decks. SDTSD predicts more extreme under-inference in such cases.

Under-inference is also the general finding in the other multinomial experiments we could find that compared subjects' posteriors with Bayesian posteriors (Phillips, Hays, & Edwards's Study 1, 1966; Dale, 1968; Martin, 1969; Martin & Gettys, 1969; Chapman, 1973). However, there are two exceptions: (1) Phillips, Hays, & Edwards (1966) varied the sample size of signals observed by subjects and, while finding under-inference for $N = 3$, 5, and 9, found essentially Bayesian inference for $N = 1$, and (2) Dale (1968) reported that in one particular trial where the data happened to exactly match one of the multinomial rates, about 1/8 of the subject over-inferred—the most over-inference he observed on any trial. Martin (1969), Martin & Gettys (1969), and Chapman (1973) also reported that subjects' under-inferred by more when an observed sample warranted a more extreme conclusion. We believe this "discriminability" finding is likely due to SDTSD.

DuCharme (1970) conducted a normal-signal inference experiment. He found under-inference when the sample was relatively far in the tails of both distributions, consistent with SDTSD, although he interpreted his results as meaning that people are reluctant to report extreme probabilities. Gustafson, Shukla, Delbecq, & Walster (1973) told subjects the average heights and weights of Midwestern college-age men and women. Subjects were then asked a series of questions such as, "The observed height of a person is 68 inches. Is the person more likely to be a male or female? How much more likely?" Gustafson et al. found that subjects over-inferred when the objective likelihood ratio was relatively small and under-inferred when the objective likelihood ratio was relatively large. Assuming that subjects believed that the sampling distributions for height and weight were normal distributions, this result means that subjects over-inferred when the sample was relatively close to the men's or women's mean height or weight and under-inferred when the sample was relatively far in the tails of both distributions. In two studies, DuCharme & Peterson (1968) familiarized subjects with normal distributions for male and female heights and then elicited subjects' beliefs that a sample was being drawn from the population of men or of women. Subjects' posteriors were nearly Bayesian when $N = 1$, but subjects under-inferred for samples of size $N = 4$.

Peterson & Phillips (1966) conducted an an experiment where the rate generating binary signals was drawn from a uniform distribution on $[0, 1]$. Subjects observed 48 binary signals and after each signal had to specify a 33% confidence interval for the rate. Subjects' confidence intervals were almost always too wide, indicating that subjects under-inferred from the data about the rate.

3.10 Appendix C: Proofs

3.10.1 Preliminary and Appendix A Results

According to Whitt's (1979) Theorem 4, A3 implies that for any N and any $\theta_A, \theta_B \in \Theta$, Barney's likelihood ratio $\frac{f^{\psi}_{S_N|\Theta}(s|\theta_A)}{f^{\psi}_{S_N|\Theta}(s|\theta_B)}$ is strictly increasing in the number of a signals in the sample. We will repeatedly use this fact in the proofs.

The preliminary results assume that $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$ has the functional form of the beta distribution with parameters $\theta\psi$ and $(1-\theta)\psi$. Let $p_0 \equiv \frac{f_\Theta(\theta_A)}{f_\Theta(\theta_B)}$ denote the agent's prior ratio, let $\Pi^{\psi}_{S_N|\Theta\times\Theta}(s|\theta_A,\theta_B) \equiv \frac{f^{\psi}_{S_N|\Theta}(s|\theta_A)}{f^{\psi}_{S_N|\Theta}(s|\theta_B)}$ denote Barney's likelihood ratio after observing a clump of N signals $s \in S_N$, and let $\Pi^{\psi}_{\Theta\times\Theta|S_N}(\theta_A,\theta_B|s) \equiv \frac{f^{\psi}_{\Theta|S_N}(\theta_A|s)}{f^{\psi}_{\Theta|S_N}(\theta_B|s)} = p_0\Pi^{\psi}_{S_N|\Theta\times\Theta}(s|\theta_A,\theta_B)$ denote his posterior ratio.

Lemma β1. *Assume Barney is prospective-acceptive, he believes the rate is θ, and he will observe a clump of N signals $s \in S_N$. Then he believes the probability of observing A_s a-signals in the sequence S_N is:*

$$f^{\psi}_{S_N|\Theta}(s|\theta_A) = \frac{\Gamma(\psi)}{\Gamma(\psi+N)}\frac{\Gamma(\theta\psi+A_s)}{\Gamma(\theta\psi)}\frac{\Gamma((1-\theta)\psi+N-A_s)}{\Gamma((1-\theta)\psi)}\frac{\Gamma(N+1)}{\Gamma(A_s+1)\Gamma(N-A_s+1)}.$$

Proof:

$$\begin{aligned}
& f^{\psi}_{S_N|\Theta}(s|\theta_A) \\
= & \int_0^1 f_{S_N|\text{ß}}(s|\beta)\, f^{\psi}_{\text{ß}|\Theta}(\beta|\theta_A)\, d\beta \\
= & \int_0^1 \frac{\Gamma(N+1)}{\Gamma(A_s+1)\Gamma(A_s-N+1)}\beta^{A_s}(1-\beta)^{N-A_s}\frac{\Gamma(\psi)}{\Gamma(\theta\psi)\Gamma((1-\theta)\psi)}\beta^{\theta\psi-1}(1-\beta)^{(1-\theta)\psi-1}\, d\beta \\
= & \frac{\Gamma(\psi)}{\Gamma(\psi+N)}\frac{\Gamma(\theta\psi+A_s)}{\Gamma(\theta\psi)}\frac{\Gamma((1-\theta)\psi+N-A_s)}{\Gamma((1-\theta)\psi)}\frac{\Gamma(N+1)}{\Gamma(A_s+1)\Gamma(N-A_s+1)}\times \\
& \int_0^1 \frac{\Gamma(\psi+N)}{\Gamma(\theta\psi+A_s)\Gamma((1-\theta)\psi+N-A_s)}\beta^{\theta\psi+A_s-1}(1-\beta)^{(1-\theta)\psi+N-A_s-1}\, d\beta \\
= & \frac{\Gamma(\psi)}{\Gamma(\psi+N)}\frac{\Gamma(\theta\psi+A_s)}{\Gamma(\theta\psi)}\frac{\Gamma((1-\theta)\psi+N-A_s)}{\Gamma((1-\theta)\psi)}\frac{\Gamma(N+1)}{\Gamma(A_s+1)\Gamma(N-A_s+1)},
\end{aligned}$$

where the fourth equality follows because the term being integrated is the pdf of a beta distribution, and the integral of a pdf is equal to 1.

□

Lemma β**2.** *Consider two states of the world, A and B, associated with known rates θ_A and θ_B, and a prior ratio p_0. Assume Barney is retrospective pooling. Barney's likelihood ratio after observing a clump of N signals $s \in S_N$ which has A_s a-signals is*:

$$\Pi^{\psi}_{S_N|\Theta\times\Theta}\left(s|\theta_A,\theta_B\right) = \frac{\Gamma\left(\theta_A\psi + A_s\right)}{\Gamma\left(\theta_B\psi + A_s\right)} \frac{\Gamma\left(\left(1-\theta_A\right)\psi + N - A_s\right)}{\Gamma\left(\left(1-\theta_B\right)\psi + N - A_s\right)} \frac{\Gamma\left(\theta_B\psi\right)}{\Gamma\left(\theta_A\psi\right)} \frac{\Gamma\left(\left(1-\theta_B\right)\psi\right)}{\Gamma\left(\left(1-\theta_A\right)\psi\right)}.$$

Moreover, in the "symmetric inference" case where $\theta_A = 1 - \theta_B$,

$$\Pi^{\psi}_{S_N|\Theta\times\Theta}\left(s|\theta_A,\theta_B\right) = \frac{\Gamma\left(\theta_A\psi + A_s\right)}{\Gamma\left(\theta_B\psi + A_s\right)} \frac{\Gamma\left(\left(1-\theta_A\right)\psi + N - A_s\right)}{\Gamma\left(\left(1-\theta_B\right)\psi + N - A_s\right)}$$

Proof: This is an immediate implication of the previous lemma.

□

Lemma β**3.** *Consider two states of the world, A and B, associated with known rates $\theta_A \geq \theta_B$. Suppose Barney has observed N signals, s_N, and he observes his $N+1^{st}$ signal. Then Barney's likelihood ratio is closer to $\Pi^{\psi}_{S_N|\Theta\times\Theta}\left(s_N|\theta_A,\theta_B\right)$ if he pools the $N+1^{st}$ signal with the others than if he groups it separately: that is, if the $N+1^{st}$ signal is an a-signal, then*

$$\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(a|\theta_A,\theta_B\right)\Pi^{\psi}_{S_N|\Theta\times\Theta}\left(s_N|\theta_A,\theta_B\right) \geq \Pi^{\psi}_{S_{N+1}|\Theta\times\Theta}\left(s_N \cup a|\theta_A,\theta_B\right),$$

and if $N+1^{st}$ signal is a b-signal, then

$$\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(b|\theta_A,\theta_B\right)\Pi^{\psi}_{S_N|\Theta\times\Theta}\left(s_N|\theta_A,\theta_B\right) \leq \Pi^{\psi}_{S_{N+1}|\Theta\times\Theta}\left(s_N \cup a|\theta_A,\theta_B\right).$$

Proof: We will prove the claim in the lemma that pertains if the final signal is an a-signal; the proof for the other case is analogous. Denote by A_N the number of a-signals in s_N. If Barney pools the a-signal with the others, then his likelihood ratio is

$$
\begin{aligned}
& \Pi^{\psi}_{S_{N+1}|\Theta\times\Theta}\left(s_N \cup a|\theta_A,\theta_B\right) \\
= \; & \frac{\Gamma\left(\theta_A\psi + A_N + 1\right)}{\Gamma\left(\theta_B\psi + A_N + 1\right)} \frac{\Gamma\left(\left(1-\theta_A\right)\psi + N + 1 - A_N - 1\right)}{\Gamma\left(\left(1-\theta_B\right)\psi + N + 1 - A_N - 1\right)} \frac{\Gamma\left(\theta_B\psi\right)}{\Gamma\left(\theta_A\psi\right)} \frac{\Gamma\left(\left(1-\theta_B\right)\psi\right)}{\Gamma\left(\left(1-\theta_A\right)\psi\right)} \\
= \; & \frac{\left(\theta_A\psi + A_N\right)\Gamma\left(\theta_A\psi + A_N\right)}{\left(\theta_B\psi + A_N\right)\Gamma\left(\theta_B\psi + A_N\right)} \frac{\Gamma\left(\left(1-\theta_A\right)\psi + N - A_N\right)}{\Gamma\left(\left(1-\theta_B\right)\psi + N - A_N\right)} \frac{\Gamma\left(\theta_B\psi\right)}{\Gamma\left(\theta_A\psi\right)} \frac{\Gamma\left(\left(1-\theta_B\right)\psi\right)}{\Gamma\left(\left(1-\theta_A\right)\psi\right)} \\
= \; & \frac{\theta_A\psi + A_N}{\theta_B\psi + A_N} \Pi^{\psi}_{S_N|\Theta\times\Theta}\left(s_N|\theta_A,\theta_B\right),
\end{aligned}
$$

where the second equality uses Lemma $\beta 2$, and the third equality uses the fact that $\Gamma(n+1) = n\Gamma(n)$. If Barney instead separately groups the final signal, then his likelihood ratio is

$$
\begin{aligned}
\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(a|\theta_A,\theta_B\right)\Pi^{\psi}_{S_N|\Theta\times\Theta}\left(s_N|\theta_A,\theta_B\right) &= \frac{\theta_A\psi}{\theta_B\psi}\Pi^{\psi}_{S_N|\Theta\times\Theta}\left(s_N|\theta_A,\theta_B\right) \\
&= \frac{\theta_A}{\theta_B}\Pi^{\psi}_{S_N|\Theta\times\Theta}\left(s_N|\theta_A,\theta_B\right).
\end{aligned}
$$

By assumption, $\frac{\theta_A}{\theta_B} \geq 1$. Therefore, for all $k > 0$,

$$
\frac{\theta_A}{\theta_B} \geq \frac{\theta_A\psi + k}{\theta_B\psi + k} \geq 1.
$$

The result follows.

$\square$

Lemma $\beta 4$. *Consider two states of the world, A and B, associated with known rates for $\theta_A > \theta_B$, and a prior ratio p_0. Assume Barney is retrospective pooling.*

- *Barney draws the same inferences as Tommy when the observed sequence is ab or ba.*
- *If the sample is all a's or all b's, Barney under-infers relative to Tommy.*
- *For any true state $\theta \in (0,1)$, Barney believes that: as $N \to \infty$, his posterior ratio $\frac{f^{\psi}_{\Theta|S_N}(\theta_A|s)}{f^{\psi}_{\Theta|S_N}(\theta_B|s)}$ will converge in distribution to a random variable that has full support on $(0,\infty)$.*

Proof:

- Barney's posterior after one a and one b signal is

$$
\begin{aligned}
& p_0 \frac{\Gamma(\theta_A\psi+1)}{\Gamma(\theta_B\psi+1)} \frac{\Gamma((1-\theta_A)\psi+1)}{\Gamma((1-\theta_B)\psi+1)} \frac{\Gamma(\theta_B\psi)}{\Gamma(\theta_A\psi)} \frac{\Gamma((1-\theta_B)\psi)}{\Gamma((1-\theta_A)\psi)} \\
= \; & p_0 \frac{\Gamma(\theta_A\psi)\,\theta_A\psi}{\Gamma(\theta_B\psi)\,\theta_B\psi} \frac{\Gamma((1-\theta_A)\psi)(1-\theta_A)\psi}{\Gamma((1-\theta_B)\psi)(1-\theta_B)\psi} \frac{\Gamma(\theta_B\psi)}{\Gamma(\theta_A\psi)} \frac{\Gamma((1-\theta_B)\psi)}{\Gamma((1-\theta_A)\psi)} \\
= \; & p_0 \frac{\theta_A}{\theta_B} \frac{(1-\theta_A)}{(1-\theta_B)},
\end{aligned}
$$

which equals Tommy's posterior.

- Barney's posterior after N a-signals and 0 b-signals is:

$$
\begin{aligned}
& p_0 \frac{\Gamma(\theta_A\psi+N)}{\Gamma(\theta_B\psi+N)} \frac{\Gamma((1-\theta_A)\psi)}{\Gamma((1-\theta_B)\psi)} \frac{\Gamma(\theta_B\psi)}{\Gamma(\theta_A\psi)} \frac{\Gamma((1-\theta_B)\psi)}{\Gamma((1-\theta_A)\psi)} \\
= \; & p_0 \frac{\Gamma(\theta_A\psi+N)}{\Gamma(\theta_B\psi+N)} \frac{\Gamma(\theta_B\psi)}{\Gamma(\theta_A\psi)} \\
= \; & p_0 \frac{(\theta_A\psi+N-1)(\theta_A\psi+N-2)...(\theta_A\psi)\Gamma(\theta_A\psi)}{(\theta_B\psi+N-1)(\theta_B\psi+N-2)...(\theta_B\psi)\Gamma(\theta_B\psi)} \frac{\Gamma(\theta_B\psi)}{\Gamma(\theta_A\psi)} \\
= \; & p_0 \frac{(\theta_A\psi+N-1)(\theta_A\psi+N-2)...(\theta_A\psi)}{(\theta_B\psi+N-1)(\theta_B\psi+N-2)...(\theta_B\psi)}.
\end{aligned}
$$

Note that the numerator and denominator each has N terms in it. Tommy's posterior ratio is

$$p_0 \frac{(\theta_A)^N}{(\theta_B)^N},$$

so also for Tommy, the numerator and denominator each has N terms. Furthermore, since $\frac{\theta_A}{\theta_B} > 1$, for all $k > 0$, $\frac{\theta_A+k}{\theta_B+k} < \frac{\theta_A}{\theta_B}$. Therefore,

$$p_0 \frac{(\theta_A\psi+N-1)(\theta_A\psi+N-2)...(\theta_A\psi)}{(\theta_B\psi+N-1)(\theta_B\psi+N-2)...(\theta_B\psi)} < p_0 \frac{(\theta_A)^N}{(\theta_B)^N}.$$

Moreover, the likelihood ratio clearly favors state A for both Tommy and Barney. Hence, Barney under-infers relative to Tommy. The case of N b-signals and 0 a-signals proceeds analogously.

- Lemma 1 implies that Barney's subjective sampling distribution for an infinite sample has pdf $\lim_{N\to\infty} f^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N} = \alpha|\theta\right) = f^{\psi}_{\text{ß}|\Theta}(\beta = \alpha|\theta)$ a.s. Therefore, after observing proportion α of a-signals in a large sample, Barney anticipates having likelihood ratio $\frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)}$ a.s. By A3, $\frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)}$ is strictly increasing in α.

Furthermore,

$$\begin{aligned}
\frac{f^{\psi}_{ß|\Theta}\left(\beta=0|\theta_A\right)}{f^{\psi}_{ß|\Theta}\left(\beta=0|\theta_B\right)} &= \frac{\frac{\Gamma(\psi)}{\Gamma(\theta_A\psi)\Gamma((1-\theta_A)\psi)}0^{\theta_A\psi-1}\left(1-0\right)^{(1-\theta_A)\psi-1}}{\frac{\Gamma(\psi)}{\Gamma(\theta_B\psi)\Gamma((1-\theta_B)\psi)}0^{\theta_B\psi-1}\left(1-0\right)^{(1-\theta_B)\psi-1}} \\
&= \frac{\Gamma(\theta_B\psi)\Gamma((1-\theta_B)\psi)}{\Gamma(\theta_A\psi)\Gamma((1-\theta_A)\psi)}\frac{0^{(\theta_A-\theta_B)\psi}}{1^{(\theta_A-\theta_B)\psi}} = 0
\end{aligned}$$

and

$$\begin{aligned}
\frac{f^{\psi}_{ß|\Theta}\left(\beta=1|\theta_A\right)}{f^{\psi}_{ß|\Theta}\left(\beta=1|\theta_B\right)} &= \frac{\frac{\Gamma(\psi)}{\Gamma(\theta_A\psi)\Gamma((1-\theta_A)\psi)}1^{\theta_A\psi-1}0^{(1-\theta_A)\psi-1}}{\frac{\Gamma(\psi)}{\Gamma(\theta_B\psi)\Gamma((1-\theta_B)\psi)}1^{\theta_B\psi-1}0^{(1-\theta_B)\psi-1}} \\
&= \frac{\Gamma(\theta_B\psi)\Gamma((1-\theta_B)\psi)}{\Gamma(\theta_A\psi)\Gamma((1-\theta_A)\psi)}\frac{1^{(\theta_A-\theta_B)\psi}}{0^{(\theta_A-\theta_B)\psi}} = \infty.
\end{aligned}$$

According to A1, Barney's beliefs put full support on $\alpha \in (0,1)$. Hence Barney thinks his large-sample likelihood ratio will converge in distribution to a random variable whose support is $\left(\frac{f^{\psi}_{ß|\Theta}(\beta=0|\theta_A)}{f^{\psi}_{ß|\Theta}(\beta=0|\theta_B)}, \frac{f^{\psi}_{ß|\Theta}(\beta=1|\theta_A)}{f^{\psi}_{ß|\Theta}(\beta=1|\theta_B)}\right) = (0,\infty)$. Since Barney's posterior ratio is his likelihood ratio times his prior ratio, the result follows.

□

Lemma β**5.** *Parameterized-Barney, whose subjective-rate distribution has the beta-distribution functional form*

$$f^{\psi}_{ß|\Theta}\left(\beta|\theta\right) = \frac{\Gamma\left(\psi\right)}{\Gamma\left(\theta\psi\right)\Gamma\left(\left(1-\theta\right)\psi\right)}\beta^{\theta\psi-1}\left(1-\beta\right)^{(1-\theta)\psi-1},$$

satisfies A1-A4.

Proof: We will prove each property in turn.

- Clearly the full-support property holds since $f^{\psi}_{ß|\Theta}\left(\beta|\theta\right) = \frac{\Gamma(\psi)}{\Gamma(\theta\psi)\Gamma((1-\theta)\psi)}\beta^{\theta\psi-1}\left(1-\beta\right)^{(1-\theta)\psi-1} > 0$ for all $\beta \in (0,1)$ and $\theta \in (0,1)$. This function is also clearly point-wise continuous in θ since all the components are point-wise continuous. Since $f^{\psi}_{ß|\Theta}\left(\beta|\theta\right)$ is constructed to be a pdf, it is Lebesgue-integrable with respect to β and hence absolutely continuous with respect to β.

- Note that

$$f^{\psi}_{ß|\Theta}\left(\beta|\theta\right) = \frac{\Gamma\left(\psi\right)}{\Gamma\left(\theta\psi\right)\Gamma\left(\left(1-\theta\right)\psi\right)}\beta^{\theta\psi-1}\left(1-\beta\right)^{(1-\theta)\psi-1} = f^{\psi}_{ß|\Theta}\left(1-\beta|1-\theta\right).$$

Therefore,

$$\int_0^{1-x} f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)\,d\beta = \int_0^{1-x} f^{\psi}_{\text{ß}|\Theta}(1-\beta|1-\theta)\,d(1-\beta).$$

The left-hand side is $F^{\psi}_{\text{ß}|\Theta}(1-x|\theta)$. Using the change-of-variables $\widetilde{\beta} = 1-\beta$, the right-hand side becomes

$$\begin{aligned} -\int_1^x f^{\psi}_{\text{ß}|\Theta}\left(\widetilde{\beta}|1-\theta\right)d\widetilde{\beta} &= \int_x^1 f^{\psi}_{\text{ß}|\Theta}\left(\widetilde{\beta}|1-\theta\right)d\widetilde{\beta} \\ &= 1-\int_0^x f^{\psi}_{\text{ß}|\Theta}\left(\widetilde{\beta}|1-\theta\right)d\widetilde{\beta}. \end{aligned}$$

Hence

$$F^{\psi}_{\text{ß}|\Theta}(1-x|\theta) = 1 - F^{\psi}_{\text{ß}|\Theta}(x|1-\theta).$$

- Note that

$$\begin{aligned} \frac{f^{\psi}_{\text{ß}|\Theta}(\beta|\theta')}{f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)} &= \frac{\frac{\Gamma(\psi)}{\Gamma(\theta'\psi)\Gamma((1-\theta')\psi)}\beta^{\theta'\psi-1}(1-\beta)^{(1-\theta')\psi-1}}{\frac{\Gamma(\psi)}{\Gamma(\theta\psi)\Gamma((1-\theta)\psi)}\beta^{\theta\psi-1}(1-\beta)^{(1-\theta)\psi-1}} \\ &= \frac{\Gamma(\theta\psi)\Gamma((1-\theta)\psi)}{\Gamma(\theta'\psi)\Gamma((1-\theta')\psi)}\left(\frac{\beta}{1-\beta}\right)^{(\theta'-\theta)\psi}. \end{aligned}$$

  If $\theta' > \theta$, then this expression is increasing in β.

- The mean of the beta distribution with parameters $\theta\psi$ and $(1-\theta)\psi$ is $\frac{\theta\psi}{\theta\psi+(1-\theta)\psi} = \theta$.

□

Lemma A. *Fix $\theta_A \geq \theta_B$ and some $0 < \underline{\lambda} < \bar{\lambda} < \infty$. If Barney's posterior ratio after one signal is in $(\underline{\lambda}, \bar{\lambda})$,*

$$[p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B), p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)] \subseteq (\underline{\lambda}, \bar{\lambda}),$$

then for any natural number N, we can construct a sequence of N signals, $s \in S_N$, such that for every truncation of s—i.e., for every sequence comprising only the first $n \leq N$ signals of s—the posterior ratio after that truncation, $\Pi^{\psi}_{\Theta\times\Theta|S_n}(\theta_A,\theta_B|s_n)$, is in $(\underline{\lambda}, \bar{\lambda})$.

Proof: We prove this by induction. By hypothesis, the result is true for the truncation $n = 1$. Assume that the likelihood ratio after N signals is in $(\underline{\lambda}, \bar{\lambda})$, and we will show that we can add a signal, and it will still be in $(\underline{\lambda}, \bar{\lambda})$. Suppose that there are A_N a-signals and $N - A_N$ b-signals. Since

$$[p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B), p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)] \subseteq (\underline{\lambda}, \bar{\lambda})$$

and $\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B) \leq 1 \leq \Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)$, it follows that $p_0 \in (\underline{\lambda}, \bar{\lambda})$.

There are two cases. The first is that Barney's posterior ratio after the sequence of N signals is (weakly) less than p_0, and so $\Pi^{\psi}_{S_N|\Theta\times\Theta}(s_N|\theta_A,\theta_B) \leq 1$. In this case add an a-signal. The new posterior ratio is $p_0\Pi^{\psi}_{S_{N+1}|\Theta\times\Theta}(s_N \cup a|\theta_A,\theta_B)$. Now,

$$\begin{aligned}
\bar{\lambda} \geq p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B) &\geq p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)\,\Pi^{\psi}_{S_N|\Theta\times\Theta}(s_N|\theta_A,\theta_B) \\
&\geq p_0\Pi^{\psi}_{S_{N+1}|\Theta\times\Theta}(s_N \cup a|\theta_A,\theta_B) \\
&\geq p_0\Pi^{\psi}_{S_N|\Theta\times\Theta}(s_N|\theta_A,\theta_B) \geq \underline{\lambda}
\end{aligned}$$

The first inequality is by hypothesis, the second is due to the fact that $\Pi^{\psi}_{S_N|\Theta\times\Theta}(s_N|\theta_A,\theta_B) \leq 1$, the third inequality is due to Lemma $\beta 3$, the fourth inequality holds since the posterior ratio increases when an additional signal is an a-signal, and the last inequality is again by hypothesis.

The second case is that the likelihood ratio after N signals is greater than p_0, in which case add a b-signal, and the argument is analogous.

□

Lemma B. *Fix $\theta_A \geq \theta_B$. If $p_0 \in \left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}\right]$, then for any N, there exists a sequence of signals s_i for $i = 0, 1, ..., N$, where s_i is a truncation of s_N, such that for all $i = 0, 1, ..., N$, $p_0\Pi^{\psi}_{S_i|\Theta\times\Theta}(s_i|\theta_A,\theta_B) \in \left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}\right]$.*

Proof: The lemma is a direct implication of the following claim: If

$$p_0 \in \left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}\right],$$

then for any n, there exists a sequence of signals $\{s_i\}_{i=0,1,\ldots,n}$, such that

$$p_0\Pi^{\psi}_{S_{n+1}|\Theta\times\Theta}(s_n \cup a|\theta_A,\theta_B) \geq 1$$

and

$$p_0\Pi^{\psi}_{S_{n+1}|\Theta\times\Theta}(s_n \cup b|\theta_A,\theta_B) \leq 1.$$

In words, we can construct a sequence of signals such that Barney's posterior ratio flips around 1 with each signal. We will prove this claim by induction and then show that it implies the lemma.

Clearly for $n = 0$ the claim is true, since the prior ratio is in $\left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}\right]$. Now we will assume the statement is true up to n and prove it is true for $n + 1$.

There are two cases. In the first case, assume that after n signals, it is the a-signal that flips the posterior ratio around 1: $p_0\Pi^{\psi}_{S_n|\Theta\times\Theta}(s_n|\theta_A,\theta_B) \leq 1$ and $p_0\Pi^{\psi}_{S_{n+1}|\Theta\times\Theta}(s_n \cup a|\theta_A,\theta_B) \geq 1$ (since adding an a-signal must increase the posterior ratio). Assume that the claim is not true for $n+1$: that is, there is no set of $n+1$ signals for which $p_0\Pi^{\psi}_{S_{n+2}|\Theta\times\Theta}(s_{n+1} \cup a|\theta_A,\theta_B) \geq 1$ and $p_0\Pi^{\psi}_{S_{n+2}|\Theta\times\Theta}(s_{n+1} \cup b|\theta_A,\theta_B) \leq 1$. In particular, taking the set of n signals plus an a-signal,

$$p_0\Pi^{\psi}_{S_{n+2}|\Theta\times\Theta}(s_n \cup a \cup a|\theta_A,\theta_B) \geq 1$$

and

$$p_0\Pi^{\psi}_{S_{n+2}|\Theta\times\Theta}(s_n \cup a \cup b|\theta_A,\theta_B) \geq 1.$$

Because the claim is true up to n signals, however, we know that an additional b-signal must flip the posterior ratio below 1,

$$p_0\Pi^{\psi}_{S_{n+1}|\Theta\times\Theta}(s_n \cup b|\theta_A,\theta_B) \leq 1.$$

But then we have identified a set of $n+1$ signals for which the statement is true, namely $s_n \cup b$: we know that $s_n \cup b \cup b$ must also generate a posterior ratio below 1 (since adding another b-signal must decrease the likelihood ratio); and $s_n \cup b \cup a$ generates the same posterior ratio as $s_n \cup a \cup b$, which we know is above 1. So we have a contradiction. Therefore, either $s_n \cup a$ or $s_n \cup b$ must satisfy the statement. The proof for a b-signal proceeds analogously.

To see that the statement implies the lemma, assume WLOG that $p_0\Pi^{\psi}_{S_n|\Theta\times\Theta}(s_n|\theta_A,\theta_B) \leq 1$, and that $p_0\Pi^{\psi}_{S_{n+1}|\Theta\times\Theta}(s_n \cup a|\theta_A,\theta_B) \geq 1$. Since

$$p_0\Pi^{\psi}_{S_{n+1}|\Theta\times\Theta}(s_n \cup a|\theta_A,\theta_B) \leq p_0\Pi^{\psi}_{S_n|\Theta\times\Theta}(s_n|\,\theta_A,\theta_B)\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)$$

(due to Lemma β3), it follows that

$$1 \leq p_0\Pi^{\psi}_{S_n|\Theta\times\Theta}(s_n|\theta_A,\theta_B)\,\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)\,.$$

Hence

$$\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)} \leq p_0\Pi^{\psi}_{S_n|\Theta\times\Theta}(s_n|\theta_A,\theta_B) \leq 1 \leq \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}.$$

The proof for the case that a b-signal flips the posterior ratio below 1 proceeds analogously.

□

Lemma C. *Fix rates $\theta_A \geq \theta_B$, prior ratio p_0, and the true state $i \in \{A, B\}$ (with corre-*

sponding true rate θ_i). Denote Barney's limit likelihood ratio conditional on the true state as $\Pi^{\psi}_{S_\infty|\Theta\times\Theta}(s_\infty|\theta_A,\theta_B,\Omega=i)$. Suppose Barney is retrospective-pooling. For any $0<\underline{\lambda}<\bar{\lambda}<\infty$ satisfying the following two statements:

1. $p_0\Pi^{\psi}_{S_\infty|\Theta\times\Theta}(s_\infty|\theta_A,\theta_B,\Omega=i)\in(\underline{\lambda},\bar{\lambda})$, *and*
2. $[p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B), p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)]\subseteq(\underline{\lambda},\bar{\lambda})$ *or* $p_0\in\left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)},\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}\right]\subseteq(\underline{\lambda},\bar{\lambda})$,

there is a positive probability that for all $N\in\{1,2,...\}$,

$$\Pi^{\psi}_{\Theta\times\Theta|S_N}(\theta_A,\theta_B|s_N)\in(\underline{\lambda},\bar{\lambda}).$$

Proof: The proof proceeds in three steps:

1. Show that the prior ratio is in $(\underline{\lambda},\bar{\lambda})$. This is true under either of the assumptions in condition 2 of the lemma.

   - Since $\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)\le 1$ and $\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)\ge 1$, the assumption

     $$[p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B), p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)]\subseteq(\underline{\lambda},\bar{\lambda})$$

     implies $p_0\in(\underline{\lambda},\bar{\lambda})$.
   - Alternatively, we directly assume that $p_0\in\left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)},\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}\right]\subseteq(\underline{\lambda},\bar{\lambda})$.
   - Note: If $\theta_A=\theta_B$, then Barney's posterior equals his prior for any sequence of signals. Since $p_0\in(\underline{\lambda},\bar{\lambda})$, the conclusion of the lemma follows trivially. Hence, for the remainder of the proof, assume $\theta_A>\theta_B$.

2. Show that for any finite N, we can construct a sequence of signals such that the posterior ratio is always in $(\underline{\lambda},\bar{\lambda})$. There are two ways we can guarantee this:

   - If $[p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B), p_0\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)]\subseteq(\underline{\lambda},\bar{\lambda})$, then it is true by Lemma A.
   - Alternatively, if we assume $p_0\in\left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)},\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}\right]\subseteq(\underline{\lambda},\bar{\lambda})$, then it is true by Lemma B.

3. Show that there exists a large enough number of signals $\hat{N}$ such that if the posterior ratio after $\hat{N}$ signals is in $(\underline{\lambda}, \bar{\lambda})$, then with positive probability the posterior ratio will always be inside $(\underline{\lambda}, \bar{\lambda})$. The rest of the proof demonstrates this. The conclusion of the lemma then follows from combining steps 1-3, taking the sequence in step 2 to have $\hat{N}$ signals, which (because it is a finite sequence) has positive probability.

Let $g(r) \equiv \Pi^{\psi}_{S_\infty|\Theta\times\Theta}(s_\infty|\theta_A, \theta_B, r)$ denote Barney's likelihood ratio in an infinite sample where the proportion of a-signals is r. This function will allow us to map from an observed proportion r to the posterior ratio that it would imply for an infinite sample, $p_0 g(r)$. Even in dealing with finite samples, it will be useful to work with the $g(\cdot)$ function in describing the mapping from the proportion of a-signals to Barney's posterior ratio. In a finite sample, however, $p_0 g(r)$ will only approximate Barney's posterior ratio, but we can bound the approximation error in a large enough sample; formally, using the Law of Large Numbers, for any r, $\varepsilon > 0$, and $\nu > 0$, there exists an $N_{r,\epsilon,\nu}$ such that if $N \geq N_{r,\epsilon,\nu}$, then with probability at least $1-\nu$, the posterior ratio after observing a proportion r a-signals out of N is within ϵ of $p_0 g(r)$. Since all we will need for step 3 is that the posterior ratio remains within $(\underline{\lambda}, \bar{\lambda})$ with positive probability, the ν we pick does not matter. So fix ν.

Since $p_0 g(r)$ is only an approximation of Barney's posterior ratio for a finite sample, we need to impose tighter bounds than $(\underline{\lambda}, \bar{\lambda})$ on where $p_0 g(r)$ can wander in a finite sample in order to ensure that Barney's posterior ratio always remains within $(\underline{\lambda}, \bar{\lambda})$. To that end, fixing ϵ, we will now pick two other numbers, $\underline{\tau} < \bar{\tau}$, such that:

- $\underline{\tau} > \underline{\lambda} + \epsilon$, and $\bar{\tau} < \bar{\lambda} - \epsilon$;
- the limit posterior ratio conditional on the true state, $p_0 \Pi^{\psi}_{S_\infty|\Theta\times\Theta}(s_\infty|\theta_A, \theta_B, \Omega = i)$, is in $(\underline{\tau}, \bar{\tau})$; and
- $g^{-1}(\underline{\tau})$ and $g^{-1}(\bar{\tau})$ are both rational.

Note that $\underline{\tau}$ and $\bar{\tau}$ satisfying the above properties exist because g is continuous and monotonic in r.

We now map these bounds on the posterior ratio onto the implied bounds on the proportion of a-signals. Let $\underline{r}$ denote the proportion of a-signals such that $p_0 g(\underline{r}) = \underline{\tau}$; and $\bar{r}$, the proportion of a-signals such that $p_0 g(\bar{r}) = \bar{\tau}$. By assumption both $\underline{r}$ and $\bar{r}$ are rational. Denote the proportion of realized a-signals after N signals by r_N.

We claim that for any particular $\underline{r}$ and $\bar{r}$ such that the true rate $\theta_i \in (\underline{r}, \bar{r})$, there is a large enough number of signals $\widetilde{N}$ such that if there have already been $\widetilde{N}$ observed with a proportion

a-signals $r_{\widetilde{N}} \in (\underline{r}, \overline{r})$, then with positive probability, $r_n \in (\underline{r}, \overline{r})$ for all numbers of signals $n \geq \widetilde{N}$. This is because by Chebyshev's inequality, for any N and any $\delta > 0$, $P(|r_N - \theta_i| < \delta) \geq 1 - \frac{\theta_i(1-\theta_i)}{N^2\delta^2}$. Thus, the probability that the proportion of a-signals always stays within the bounds $\underline{r}, \overline{r}$ is

$$P\left(|r_n - \theta_i| < \delta \; \forall n \geq \widetilde{N} \text{ s.t. } \left|r_{\widetilde{N}} - \theta_i\right| < \delta\right) \geq \Pi_{n=\widetilde{N}}^{\infty}\left(1 - \frac{\theta_i(1-\theta_i)}{n^2\delta^2}\right).$$

Note that for any $x > 0$ and $y < 1$, it is true that $y(1-x) > y - x$, and so we can bound the right-hand side of the previous equation by iteratively applying this inequality:

$$\Pi_{n=\widetilde{N}}^{\infty}\left(1 - \frac{\theta_i(1-\theta_i)}{n^2\delta^2}\right) > 1 - \sum_{n=\widetilde{N}}^{\infty} \frac{\theta_i(1-\theta_i)}{n^2\delta^2}.$$

Hence there is positive probability that the proportion of a-signals always stays within $(\underline{r}, \overline{r})$. While not relevant for this proof, we point out for reference in proofs below that this infinite series is a known convergent sum; so as $\widetilde{N} \to \infty$, the right-hand side converges to 1.

Therefore, for any ϵ, we can choose a large enough number of signals, $\hat{N} = \max\{\widetilde{N}, N_{\underline{r},\epsilon,\nu}, N_{\overline{r},\epsilon,\nu}\}$, so that if the proportion of a-signals after $\hat{N}$ signals is within $(\underline{r}, \overline{r})$, then with positive probability the proportion always remains within $(\underline{r}, \overline{r})$. Therefore, with positive probability $p_0 g(r)$ always remains within $(\underline{\tau}, \overline{\tau})$. Since $\hat{N} \geq N_{\underline{r},\epsilon,\nu}$ and $\hat{N} \geq N_{\overline{r},\epsilon,\nu}$, whenever $p_0 g(r)$ is within $(\underline{\tau}, \overline{\tau})$, Barney's posterior ratio is within $(\underline{\lambda}, \bar{\lambda})$.

□

Lemma A1. *Barney-Freddy does not believe in LLN: For any interval* $[\alpha_1, \alpha_2] \subseteq [0, 1]$,

$$\lim_{N\to\infty} \sum_{x=\lfloor a_1 N \rfloor}^{\lceil a_2 N \rceil} f^{\psi M}_{S_N|\Theta}(A_s = x|\theta) = F^{\psi}_{ß|\Theta}(\beta = \alpha_2|\theta) - F^{\psi}_{ß|\Theta}(\beta = \alpha_1|\theta) > 0.$$

Proof: First, we claim that conditional on a particular subjective-rate β, the agent believes that the distribution of percentage of a-signals converges to a point mass on β. To see this, fix β, and suppose Barney-Freddy has observed N signals. We begin with the case where N is an even number (i.e., the agent has observed $\frac{N}{2}$ pairs of signals, where recall that the "urn" is renewed after each pair of signals). Considering each pair of signals as a single signal that has 4 possible values and a mean value of 2β, LLN implies that Barney-Freddy believes that the distribution of percentage of a-signals converges to a point mass on β. This proves the claim for the case where N is an even number. Now, notice that for any $[\alpha_1, \alpha_2] \subseteq [0,1]$ and $\varepsilon > 0$, we can find a large enough $\hat{N}$ such that for all $N \geq \hat{N}$, if $\frac{A_{s_N}}{N} \in (\alpha_1, \alpha_2)$, then $\frac{A_{s_N}}{N+1} \in (\alpha_1 - \varepsilon, \alpha_2 + \varepsilon)$ and $\frac{A_{s_N}+1}{N+1} \in (\alpha_1 - \varepsilon, \alpha_2 + \varepsilon)$. This observation, combined with the claim for even number N, proves the claim for odd number N.

The remainder of the proof is identical to the proof of Lemma 1 below.

□

3.10.2 Main Text Results

Lemma 1. *Barney does not believe in LLN: For any interval $[\alpha_1, \alpha_2] \subseteq [0,1]$,*

$$\lim_{N\to\infty} \sum_{x=\lfloor \alpha_1 N \rfloor}^{\lceil \alpha_2 N \rceil} f^{\psi}_{S_N|\Theta}(A_s = x|\theta) = F^{\psi}_{ß|\Theta}(\beta = \alpha_2|\theta) - F^{\psi}_{ß|\Theta}(\beta = \alpha_1|\theta) > 0.$$

Proof: We begin by proving a key fact. Consider some rational $\alpha \in [0,1]$ (which will represent the percentage of a-signals) and an increasing sequence of integers $m_1, m_2, ...$ (which will represent sample sizes) such that αm_j (which will represent numbers of a-signals) is an integer for all $j = 1, 2, ...$ Using the definition of Barneyness,

$$\begin{aligned}
\lim_{j\to\infty} F^{\psi}_{S_{m_j}|\Theta}(\alpha m_j|\theta) &= \lim_{j\to\infty} \int_0^1 F_{S_{m_j}|ß}(\alpha m_j|\beta)\, f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta \\
&= \int_0^1 \lim_{j\to\infty} F_{S_{m_j}|ß}(\alpha m_j|\beta)\, f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta \\
&= \int_0^\alpha \lim_{j\to\infty} F_{S_{m_j}|ß}(\alpha m_j|\beta)\, f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta + \int_\alpha^1 \lim_{j\to\infty} F_{S_{m_j}|ß}(\alpha m_j|\beta)\, f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta \\
&= \int_0^\alpha (1)\, f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta + \int_\alpha^1 (0)\, f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta \\
&= F^{\psi}_{ß|\Theta}(\beta = \alpha|\theta),
\end{aligned}$$

where the second equality follows from the absolute continuity of $f^{\psi}_{ß|\Theta}(\beta|\theta)$, and the fourth follows from the Law of Large Numbers.

Now, by the above fact and the definition of the cdf, for any $[a_1, a_2] \subseteq [0,1]$ with a_1, a_2 rational,

$$\begin{aligned}
\lim_{N\to\infty} \sum_{x=\lfloor a_1 N \rfloor}^{\lceil a_2 N \rceil} f^{\psi}_{S_N|\Theta}(A_s = x|\theta) &= \lim_{N\to\infty} (F^{\psi}_{S_N|\Theta}(\lceil a_2 N \rceil|\theta) - F^{\psi}_{S_N|\Theta}(\lfloor a_1 N \rfloor|\theta)) \\
&= F^{\psi}_{ß|\Theta}(\beta = a_2|\theta) - F^{\psi}_{ß|\Theta}(\beta = a_1|\theta).
\end{aligned}$$

Consider a sequence of pairs of rational numbers, $\left(a_1^{(1)}, a_2^{(1)}\right), \left(a_1^{(2)}, a_2^{(2)}\right), ...$ that converges to the pair of real numbers (α_1, α_2). Taking the limit of the above equality along this sequence gives

$$\lim_{N\to\infty} \sum_{x=\lfloor \alpha_1 N \rfloor}^{\lceil \alpha_2 N \rceil} f^{\psi}_{S_N|\Theta}\left(A_s = x|\theta\right) = F^{\psi}_{\text{ß}|\Theta}\left(\beta = \alpha_2|\theta\right) - F^{\psi}_{\text{ß}|\Theta}\left(\beta = \alpha_1|\theta\right).$$

This is greater than 0 by the full-support assumption in A1.

$\square$

Proposition 1. *For any $\theta \in (0,1)$ and $N \in \{1, 2, ...\}$:*

1. $E^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) = E_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) = \theta$.

2. $F^{\psi}_{S_N|\Theta}\left(A_s|\theta\right)$ *SOSD* $F_{S_N|\Theta}\left(A_s|\theta\right)$*, and* $Var^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) \geq Var_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right)$ *with strict inequality for* $N > 1$.

3. $Var^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right)$ *is strictly decreasing in* N*, and* $\lim_{N\to\infty} Var^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) > 0$.

4. $F^{\psi}_{S_N|\Theta}\left(A_s|\theta'\right)$ *FOSD* $F^{\psi}_{S_N|\Theta}\left(A_s|\theta\right)$ *whenever* $\theta' > \theta$.

Proof:

1. Since A_s is a binomial random variable with rate θ, $E_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) = \theta$. Using the Law of Iterated Expectations, $E^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) = E^{\psi}_{\text{ß}|\Theta}\left[E_{S_N|\text{ß}}\left(\frac{A_s}{N}|\beta\right)|\theta\right] = E^{\psi}_{\text{ß}|\Theta}\left[\beta|\theta\right] = \theta$, where the last equality follows from A4.

2. To keep notation compact, let T denote a random variable whose distribution is binomial distribution with rate θ. Let Y denote the random variable whose distribution is $F^{\psi}_{S_N|\Theta}\left(A_s|\theta\right)$, i.e., the random variable induced by taking a binomial draw using rate β and then integrating over all possible β's using $f^{\psi}_{\text{ß}|\Theta}\left(\beta|\theta\right)$. We will show that Y is a mean preserving spread of the binomial distribution with rate θ, which implies second-order stochastic dominance.

   For Y to be a mean preserving spread of T, it must be the case the $Y = T + Z$, where Z is a random variable with conditional mean $E\left[Z|T\right] = 0$. We will construct Z.

   Recall that for any random variable V, its moment generating function $M_V(t) \equiv E\left[e^{tV}\right]$ (when it exists) completely characterizes the distribution of V and has the following useful properties: $E\left[V\right] = \frac{d}{dt} M_V(t)\big|_{t=0}$; for random variables V, V' and V'', $V = V' + V''$ if and only if $M_V(t) = M_{V'}(t) M_{V''}(t)$; and for random variables V and V', $M_V(t) = E\left[M_{V|V'}(t)|V'\right]$.

Since the random variable $Y|\beta$ has a binomial distribution with parameter β, its moment generating function is $M_{Y|\beta}(t) = 1 - \beta + \beta e^t$. Therefore,

$$M_Y(t) = E\left[M_{Y|\beta}(t)|\beta\right] = \int_0^1 (1 - \beta + \beta e^t) f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta,$$

which clearly exists since $M_{Y|\beta}(t)$ exists for each $\beta \in [0,1]$. If there is some Z, then we know it must have moment generating function

$$M_Z(t) = \frac{M_Y(t)}{M_T(t)} = \frac{\int_0^1 (1 - \beta + \beta e^t) f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta}{1 - \theta + \theta e^t}.$$

(With this moment generating function, the random variable $T + Z$ will have the same distribution as Y.) Hence, we will have proved that $F^{\psi}_{S_N|\Theta}(A_s|\theta)$ SOSD $F_{S_N|\Theta}(A_s|\theta)$ once we verify that $E[Z|T] = 0$. By construction, the random variable Z is independent of T (since Z's moment generating function, and hence its distribution, does not depend on the realization of T). Therefore, $E[Z|T] = E[Z]$. Finally,

$$\begin{aligned}
E[Z] &= \left.\frac{d}{dt} M_Z(t)\right|_{t=0} \\
&= \left.\frac{(1 - \theta + \theta e^t)\int_0^1 (\beta e^t) f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta - \theta e^t \int_0^1 (1 - \beta + \beta e^t) f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta}{(1 - \theta + \theta e^t)^2}\right|_{t=0} \\
&= \frac{\int_0^1 \beta f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta - \theta \int_0^1 f^{\psi}_{ß|\Theta}(\beta|\theta)\, d\beta}{(1 - \theta + \theta)^2} = \theta - \theta = 0.
\end{aligned}$$

This proves the first claim.

Using the Law of Total Variance,

$$\begin{aligned}
Var^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) &= E^{\psi}_{ß|\Theta}\left[Var_{S_N|ß}\left(\frac{A_s}{N}|\beta\right)|\theta\right] + Var^{\psi}_{ß|\Theta}\left[E_{S_N|ß}\left(\frac{A_s}{N}|\beta\right)|\theta\right] \\
&= E^{\psi}_{ß|\Theta}\left[\frac{\beta(1-\beta)}{N}|\theta\right] + Var^{\psi}_{ß|\Theta}[\beta|\theta] \\
&= \frac{\theta}{N} - \frac{Var^{\psi}_{ß|\Theta}[\beta|\theta] + \theta^2}{N} + Var^{\psi}_{ß|\Theta}[\beta|\theta] \\
&= \frac{\theta(1-\theta)}{N} + \frac{N-1}{N} Var^{\psi}_{ß|\Theta}[\beta|\theta],
\end{aligned}$$

where the third equality uses $E^{\psi}_{ß|\Theta}\left(\beta^2|\theta\right) = Var^{\psi}_{ß|\Theta}[\beta|\theta] + \left[E^{\psi}_{ß|\Theta}(\beta|\theta)\right]^2$. Since $Var_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) = \frac{\theta(1-\theta)}{N}$, the result immediately follows.

3. From part 2,

$$Var^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N}|\theta\right) = \frac{\theta\,(1-\theta)}{N} + \frac{N-1}{N} Var^{\psi}_{ß|\Theta}\,[\beta|\theta]\,.$$

Note that the highest variance of any distribution that has support on $[0,1]$ and a mean of θ can have is $\theta(1-\theta)$; this is the variance of a distribution with mass $(1-\theta)$ on 0 and mass θ on 1. Any shifting of weight that preserves the mean must move positive mass closer to the mean. Since $f^{\psi}_{ß|\Theta}$ has full support (by A1) and mean θ (by A4), $Var^{\psi}_{ß|\Theta}\,[\beta|\theta] < \theta\,(1-\theta)$. Also, the full support of $f^{\psi}_{ß|\Theta}$ implies that $Var^{\psi}_{ß|\Theta}\,[\beta|\theta] > 0$. The result immediately follows.

4.

$$\begin{aligned} F^{\psi}_{S_N|\Theta}\,(A_s|\theta) &= \sum_{i=0}^{A_s}\int_0^1 \frac{N!}{(N-i)!i!}\beta^i(1-\beta)^{N-i} f^{\psi}_{ß|\Theta}\,(\beta|\theta)\,d\beta \\ &= \int_0^1 \left[\sum_{i=0}^{A_s} \frac{N!}{(N-i)!i!}\beta^i(1-\beta)^{N-i}\right] f^{\psi}_{ß|\Theta}\,(\beta|\theta)\,d\beta \\ &= \int_0^1 F_{S_N|ß}\,(A_s|\beta)\, f^{\psi}_{ß|\Theta}\,(\beta|\theta)\,d\beta, \end{aligned}$$

where $F_{S_N|ß}\,(A_s|\beta)$ is the cdf of the binomial distribution conditional on a rate β. By A3, $F^{\psi}_{ß|\Theta}\,(\beta|\theta')$ FOSD $F^{\psi}_{ß|\Theta}\,(\beta|\theta)$ whenever $\theta' > \theta$, and therefore $\int_0^1 g\,(\beta)\, f^{\psi}_{ß|\Theta}\,(\beta|\theta')\,d\beta > \int_0^1 g\,(\beta)\, f^{\psi}_{ß|\Theta}\,(\beta|\theta)\,d\beta$ for any function $g\,(\beta)$ that is decreasing in β. Since $F_{S_N|ß}\,(A_S|\beta)$ is decreasing in β, the result follows.

$\square$

Proposition 2. *Let $\theta \in (0,1)$ be the true rate. Then for any $\theta_A, \theta_B \in \Theta$, Barney draws limited inference even from an infinite sample: as $N \to \infty$,*

$$\frac{f^{\psi}_{\Theta|S_N}\,(\theta_A|s)}{f^{\psi}_{\Theta|S_N}\,(\theta_B|s)} \to_{\text{a.s.}} \frac{f^{\psi}_{ß|\Theta}\,(\beta=\theta|\theta_A)}{f^{\psi}_{ß|\Theta}\,(\beta=\theta|\theta_B)}\frac{f\,(\theta_A)}{f\,(\theta_B)},$$

which is a positive, finite number.

Proof: Using the Law of Large Numbers,

$$\lim_{N\to\infty}\frac{f^{\psi}_{\Theta|S_N}\,(\theta_A|s)}{f^{\psi}_{\Theta|S_N}\,(\theta_B|s)} = \left(\lim_{N\to\infty}\frac{f^{\psi}_{S_N|\Theta}\,(s|\theta_A)}{f^{\psi}_{S_N|\Theta}\,(s|\theta_B)}\right)\frac{f\,(\theta_A)}{f\,(\theta_B)} = \frac{f^{\psi}_{ß|\Theta}\,(\beta=\theta|\theta_A)}{f^{\psi}_{ß|\Theta}\,(\beta=\theta|\theta_B)}\frac{f\,(\theta_A)}{f\,(\theta_B)}\text{ a.s.}$$

$\square$

Proposition 3. *Fix any $\theta_A > \theta_B \in \Theta$. Barney believes that, regardless of what the true rate $\theta \in (0,1)$ is, as $N \to \infty$, his posterior ratio $\frac{f^{\psi}_{\Theta|S_N}(\theta_A|s)}{f^{\psi}_{\Theta|S_N}(\theta_B|s)}$ will converge in distribution to a random variable whose support is a simply connected subset of $[0,\infty]$. Moreover, he believes that this subset of $[0,\infty]$ has strictly positive measure.*

Proof: Lemma 1 implies that Barney's subjective sampling distribution for an infinite sample is $\lim_{N\to\infty} f^{\psi}_{S_N|\Theta}\left(\frac{A_s}{N} = \alpha|\theta\right) = f^{\psi}_{\text{ß}|\Theta}(\beta = \alpha|\theta)$ a.s. Therefore, after observing proportion α of a-signals in a large sample, Barney anticipates having posterior ratio $\lim_{N\to\infty} \frac{f^{\psi}_{\Theta|S_N}\left(\theta_A|\frac{A_s}{N}=\alpha\right)}{f^{\psi}_{\Theta|S_N}\left(\theta_B|\frac{A_s}{N}=\alpha\right)} = \frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)f(\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)f(\theta_B)}$ a.s. This large-sample posterior ratio is a random variable because, according to A1, Barney's beliefs put full support on $\alpha \in (0,1)$.

By A3 and Whitt (1979), $\frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)f(\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)f(\theta_B)}$ is strictly increasing in α. It follows that

$$\inf_{\alpha\in(0,1)} \frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)\, f(\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)\, f(\theta_B)} = \lim_{\alpha\to 0} \frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)\, f(\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)\, f(\theta_B)} < 1$$

$$\sup_{\alpha\in(0,1)} \frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)\, f(\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)\, f(\theta_B)} = \lim_{\alpha\to 1} \frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)\, f(\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)\, f(\theta_B)} > 1.$$

Furthermore, by A1 and the Intermediate Value Theorem, the posterior ratio must take on each value in

$$\left(\inf_{\alpha\in(0,1)} \frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)\, f(\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)\, f(\theta_B)}, \sup_{\alpha\in(0,1)} \frac{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_A)\, f(\theta_A)}{f^{\psi}_{\text{ß}|\Theta}(\beta=\alpha|\theta_B)\, f(\theta_B)}\right)$$

for some $\alpha \in (0,1)$. This interval is a simply connected subset of $[0,\infty]$ that Barney believes has strictly positive measure.

□

Proposition 4. *Fix rates $\theta_A > \theta_B$. For $N = 1$, Barney and Tommy infer the same. If $\theta_A = 1-\theta_B$, then for any set of $N \in \{1,2,...\}$ signal realizations $s \in S_N$, neither Tommy's beliefs nor Barney's beliefs change from the priors when $\frac{A_s}{N} = \frac{1}{2}$.*

Proof: The first claim can be seen from the fact that the for a single sample, the subjective sampling distribution of Tommy matches that of Barney.

To see that the second claim is true, suppose that Barney observes exactly half a and b signals,

k of each signal (so that $N-k=k$). Barney's likelihood ratio is

$$\begin{aligned}
\frac{\int_0^1 \frac{N!}{k!k!}\beta^k(1-\beta)^k f^{\psi}_{\text{ß}|\Theta}(\beta|\theta_A)d\beta}{\int_0^1 \frac{N!}{k!k!}\beta^k(1-\beta)^k f^{\psi}_{\text{ß}|\Theta}(\beta|\theta_B)d\beta} &= \frac{\int_0^1 \frac{N!}{k!k!}\beta^k(1-\beta)^k f^{\psi}_{\text{ß}|\Theta}(\beta|\theta_A)d\beta}{\int_0^1 \frac{N!}{k!k!}\beta^k(1-\beta)^k f^{\psi}_{\text{ß}|\Theta}(\beta|1-\theta_A)d\beta} \\
&= \frac{\int_0^1 \frac{N!}{k!k!}\beta^k(1-\beta)^k f^{\psi}_{\text{ß}|\Theta}(\beta|\theta_A)d\beta}{\int_0^1 \frac{N!}{k!k!}\beta^k(1-\beta)^k f^{\psi}_{\text{ß}|\Theta}(1-\beta|\theta_A)d\beta} \\
&= \frac{\int_0^1 \frac{N!}{k!k!}\beta^k(1-\beta)^k f^{\psi}_{\text{ß}|\Theta}(\beta|\theta_A)d\beta}{\int_0^1 \frac{N!}{k!k!}(1-\beta)^k\beta^k f^{\psi}_{\text{ß}|\Theta}(\beta|\theta_A)d\beta} \\
&= 1,
\end{aligned}$$

where the first equality comes from substituting $\theta_B = 1-\theta_A$, and the second equality follows from A2. This directly implies that Barney's likehood ratio is equal to 1 if and only if Tommy's is, since the likelihood ratio is strictly increasing in the number of a-signals for both Barney and Tommy if $\theta_A > \frac{1}{2}$, and equal to 1 for any sample for both Barney and Tommy if $\theta_A = \frac{1}{2}$.

□

Proposition 5. *Fix rates $\theta_A > \theta_B$ and a set of $N \in \{1, 2, ...\}$ signal realizations $s \in S_N$. If $f^{\psi}_{\text{ß}|\Theta}(\beta|\theta)$ has the functional form (3.3), then regardless of whether the true rate is θ_A or θ_B, for ψ sufficiently small, the expected change in Barney's beliefs is smaller than the expected change in Tommy's beliefs. Furthermore, suppose $\theta_A = 1-\theta_B$. Then for any sample of $N > 1$ signals such that $\frac{A_s}{N} \neq \frac{1}{2}$ and any ψ, Barney under-infers relative to Tommy. In addition, while Tommy's inference depends solely on the difference in the number of a and b signals, Barney's change in beliefs is smaller from larger samples with the same difference.*

Proof: We will prove the first claim in two steps. First, we shall show that a sufficient condition is that Tommy's subjective sampling distribution is Blackwell-sufficient for Barney's. Second, we will show that this sufficient condition in fact holds for ψ sufficiently small.

First, suppose $F_{S_N|\Theta}(A_s|\theta_i)$ is Blackwell-sufficient for $F^{\psi}_{S_N|\Theta}(A_s|\theta_i)$ for $i \in \{A, B\}$. Using Blackwell (1951, 1953), note that $F_{S_N|\Theta}(A_s|\theta_i)$ is Blackwell-sufficient for $F^{\psi}_{S_N|\Theta}(A_s|\theta_i)$ if and only if for any continuous, convex function g,

$$\begin{aligned}
&\sum_{k=0}^{N} f_{S_N|\Theta}(A_s=k|\theta_i)\, g\left(\frac{f_{S_N|\Theta}(A_s=k|\theta_i)}{f_{S_N|\Theta}(A_s=k|\theta_{-i})}\right) \\
\geq\ &\sum_{k=0}^{N} f^{\psi}_{S_N|\Theta}(A_s=k|\theta_i)\, g\left(\frac{f^{\psi}_{S_N|\Theta}(A_s=k|\theta_i)}{f^{\psi}_{S_N|\Theta}(A_s=k|\theta_{-i})}\right)
\end{aligned}$$

for $i \in \{A, B\}$ and $-i$ being the other state. Taking $g(x) = |p_0 - p_0 x|$, this implies

$$\begin{aligned} & \sum_{k=0}^{N} f_{S_N|\Theta}(A_s = k|\theta_i) \left| p_0 - p_0 \frac{f_{S_N|\Theta}(A_s = k|\theta_i)}{f_{S_N|\Theta}(A_s = k|\theta_{-i})} \right| \\ \geq \ & \sum_{k=0}^{N} f^{\psi}_{S_N|\Theta}(A_s = k|\theta_{-i}) \left| p_0 - p_0 \frac{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_i)}{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_{-i})} \right| \end{aligned}$$

for $i \in \{A, B\}$.

From Proposition 1, we know that $f^{\psi}_{S_N|\Theta}(A_s|\theta_{-i})$ is a mean-preserving spread of $f_{S_N|\Theta}(A_s|\theta_{-i})$. Since $g(x) = |p_0 - p_0 x|$ is a continuous convex function, it follows that

$$\begin{aligned} & \sum_{k=0}^{N} f^{\psi}_{S_N|\Theta}(A_s = k|\theta_{-i}) \left| p_0 - p_0 \frac{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_i)}{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_{-i})} \right| \\ \geq \ & \sum_{k=0}^{N} f_{S_N|\Theta}(A_s = k|\theta_{-i}) \left| p_0 - p_0 \frac{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_i)}{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_{-i})} \right| \end{aligned}$$

for $i \in \{A, B\}$. Chaining the two inequalities,

$$\begin{aligned} & \sum_{k=0}^{N} f_{S_N|\Theta}(A_s = k|\theta_i) \left| p_0 - p_0 \frac{f_{S_N|\Theta}(A_s = k|\theta_i)}{f_{S_N|\Theta}(A_s = k|\theta_{-i})} \right| \\ \geq \ & \sum_{k=0}^{N} f_{S_N|\Theta}(A_s = k|\theta_{-i}) \left| p_0 - p_0 \frac{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_i)}{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_{-i})} \right|. \end{aligned}$$

Aggregating over states weighted by the priors,

$$\begin{aligned} & f(\theta_A) \left(\sum_{k=0}^{N} f_{S_N|\Theta}(A_s = k|\theta_A) \left| p_0 - p_0 \frac{f_{S_N|\Theta}(A_s = k|\theta_A)}{f_{S_N|\Theta}(A_s = k|\theta_B)} \right| \right) \\ & + f(\theta_B) \left(\sum_{k=0}^{N} f_{S_N|\Theta}(A_s = k|\theta_B) \left| p_0 - p_0 \frac{f_{S_N|\Theta}(A_s = k|\theta_B)}{f_{S_N|\Theta}(A_s = k|\theta_A)} \right| \right) \\ \geq \ & f(\theta_A) \left(\sum_{k=0}^{N} f_{S_N|\Theta}(A_s = k|\theta_A) \left| p_0 - p_0 \frac{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_A)}{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_B)} \right| \right) \\ & + f(\theta_B) \left(\sum_{k=0}^{N} f_{S_N|\Theta}(A_s = k|\theta_B) \left| p_0 - p_0 \frac{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_B)}{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_A)} \right| \right). \end{aligned}$$

This is what we were seeking to prove.

Now we will show that for $i \in \{A, B\}$ and $N > 1$, $F_{S_N|\Theta}(A_s|\theta_i)$ is Blackwell-sufficient for $F^{\psi}_{S_N|\Theta}(A_s|\theta_i)$ if ψ is sufficiently small. Since Barney's subjective sampling distribution is the same as Tommy's for $N = 1$, $F_{S_1|\Theta}(A_s|\theta_i)$ is Blackwell-equivalent to $F^{\psi}_{S_1|\Theta}(A_s|\theta_i)$ (regardless of the value of ψ). Note that as $\psi \to 0$, Barney's subjective sampling distribution converges to point masses at 0 and 1. Hence in the limit $\psi \to 0$, for any N, Barney expects to observe all a's or all b's; therefore, in the limit $\psi \to 0$, $F^{\psi}_{S_N|\Theta}(A_s|\theta_i)$ is Blackwell-equivalent to $F^{\psi}_{S_1|\Theta}(A_s|\theta_i)$. In contrast, for Tommy, $F_{S_N|\Theta}(A_s|\theta_i)$ for $N > 1$ is strictly Blackwell-sufficient for $F_{S_1|\Theta}(A_s|\theta_i)$ (e.g., Shaked and Tong, 1990). Therefore, for any $N > 1$, in the limit $\psi \to 0$, $F_{S_N|\Theta}(A_s|\theta_i)$ is Blackwell-sufficient for $F^{\psi}_{S_N|\Theta}(A_s|\theta_i)$. Using the necessary and sufficient inequality for Blackwell-sufficiency: for any $N > 1$ and any continuous, convex function g, in the limit $\psi \to 0$,

$$
\begin{aligned}
& \sum_{k=0}^{N} f_{S_N|\Theta}(A_s = k|\theta_i)\, g\left(\frac{f_{S_N|\Theta}(A_s = k|\theta_i)}{f_{S_N|\Theta}(A_s = k|\theta_{-i})}\right) \\
> \quad & \sum_{k=0}^{N} f^{\psi}_{S_N|\Theta}(A_s = k|\theta_i)\, g\left(\frac{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_i)}{f^{\psi}_{S_N|\Theta}(A_s = k|\theta_{-i})}\right).
\end{aligned}
$$

Note that the right-hand side of this inequality is continuous in ψ. Therefore, the strict inequality holds for any sufficiently small $\psi > 0$, and so $F_{S_N|\Theta}(A_s|\theta_i)$ is Blackwell-sufficient for $F^{\psi}_{S_N|\Theta}(A_s|\theta_i)$ for any sufficiently small $\psi > 0$.

To see that the second claim in the proposition is true, consider some sample of size N, with A_s realizations of heads, and assume that $\theta_A > \frac{1}{2}$ and $A_s > N - A_s$. (The other cases have analogous proofs.) Tommy's likelihood ratio is

$$
\frac{\frac{N!}{\theta_A!(N-\theta_A)!}\theta_A^{A_s}(1-\theta_A)^{N-A_s}}{\frac{N!}{\theta_A!(N-\theta_A)!}\theta_A^{N-A_s}(1-\theta_A)^{A_s}} = \left(\frac{\theta_A}{1-\theta_A}\right)^{2A_s - N},
$$

which has $2A_s - N$ terms multiplied together. Barney's likelihood ratio is

$$
\frac{\Gamma(\theta_A\psi + A_s)}{\Gamma((1-\theta_A)\psi + A_s)} \frac{\Gamma((1-\theta_A)\psi + N - A_s)}{\Gamma(\theta_A\psi + N - A_s)} = \frac{(\theta_A\psi + A_s - 1)...(\theta_A\psi + N - A_s)}{((1-\theta_A)\psi + A_s - 1)...((1-\theta_A)\psi + N - A_s)}.
$$

This also has $2A_s - N$ terms. Since by assumption $\frac{\theta_A\psi}{(1-\theta_A)\psi} = \frac{\theta_A}{1-\theta_A} > 1$, then for any $k > 0$,

$$
\frac{\theta_A}{1-\theta_A} = \frac{\theta_A\psi}{(1-\theta_A)\psi} > \frac{\theta_A\psi + k}{(1-\theta_A)\psi + k} > 1,
$$

which implies that

$$
\left(\frac{\theta_A}{1-\theta_A}\right)^{2A_s - N} > \frac{(\theta_A\psi + A_s - 1)...(\theta_A\psi + N - A_s)}{((1-\theta_A)\psi + A_s - 1)...((1-\theta_A)\psi + N - A_s)} > 1.
$$

Hence Barney under-infers relative to Tommy.

Finally, consider two samples, one of size N with A_s a-signals, and the other of size $N' > N$ with A_s' a-signals. Suppose both samples have the same difference, d, between the number of a and b signals: $d = A_s - (N - A_s) = 2A_s - N$ and $d = 2A_s' - N'$. Since Tommy's likelihood ratio is $\left(\frac{\theta_A}{1-\theta_A}\right)^d$, it is the same for both samples. To consider Barney's inference, assume that $\theta_A > \frac{1}{2}$ and $d > 0$. (The other cases have analogous proofs.) Barney's likelihood ratio, stated above, has d terms for both samples. Since by assumption $\frac{\theta_A \psi}{(1-\theta_A)\psi} = \frac{\theta_A}{1-\theta_A} > 1$, then for any $k' > k > 0$,

$$\frac{\theta_A \psi + k}{(1-\theta_A)\psi + k} > \frac{\theta_A \psi + k'}{(1-\theta_A)\psi + k'} > 1.$$

Note that $A_s' > A_s$. It follows that each term in Barney's likelihood ratio for the sample of size N is larger than the corresponding term in Barney's likelihood ratio for the sample of size N' (e.g., for the first term, $\frac{\theta_A \psi + A_s - 1}{(1-\theta_A)\psi + A_s - 1} > \frac{\theta_A \psi + A_s' - 1}{(1-\theta_A)\psi + A_s' - 1}$). Hence while Barney infers in favor of state A in both cases (like Tommy does), Barney's change in beliefs is smaller from the larger sample.

$\square$

Proposition 6. *Suppose Barney and Tommy have simple, piecewise-linear loss-averse preferences as specified in (3.4). Fix any gamble (θ, h, t), paying off $h > 0$ with probability θ and $-t$ with probability $1-\theta$, that is better than fair: $\theta h > (1-\theta)\,t$. For any $\lambda \geq 1$, there is some $N' \geq 1$ such that if $N > N'$, then Tommy will accept N repetitions of the gamble. In contrast, for Barney there is some threshold level of loss aversion $\widehat{\lambda} > 1$ such that: if $\lambda < \widehat{\lambda}$, then there is some N' sufficiently large such that Barney will accept N repetitions of the gamble for all $N > N'$; and if $\lambda \geq \widehat{\lambda}$, then there is some N'' sufficiently large such that Barney will reject N repetitions of the gamble for all $N > N''$.*

Proof: WLOG let the reference point be w_0, and fix $\lambda \geq 1$. Denoting $G(z)$ as the distribution of monetary outcomes z, the expected utility of any lottery is:

$$\begin{aligned} & \int_0^\infty (w_0 + z)\, dG(z) + \int_{-\infty}^0 (w_0 + \lambda z)\, dG(z) \\ = \quad & w_0 + \int_0^\infty z dG(z) + \lambda \int_{-\infty}^0 z dG(z). \end{aligned}$$

Clearly, this is better than the option of refusing the lottery if and only if

$$\int_0^\infty z dG(z) > -\lambda \int_{-\infty}^0 z dG(z) \Leftrightarrow \frac{\int_0^\infty z dG(z)}{-\int_{-\infty}^0 z dG(z)} < \lambda.$$

The monetary outcome of an N-times repeated gamble is $z = A_s h - (N - A_s)t = -Nt + A_s(h + t) = N(-t + \frac{A_s}{N}(h+t))$. Hence the agent earns a positive payoff from the gamble if and only if $A_s \geq \left\lceil \frac{t}{h+t} N \right\rceil$. Substituting in, we find that for Tommy,

$$\frac{\int_0^\infty z dG(z)}{-\int_{-\infty}^0 z dG(z)} = \frac{\sum_{A_s = \lceil \frac{t}{h+t} N \rceil}^{N} (-t + \frac{A_s}{N}(h+t)) f_{S_N|\Theta}(A_s|\theta)}{-\sum_{A_s=0}^{\lceil \frac{t}{h+t} N \rceil - 1} (-t + \frac{A_s}{N}(h+t)) f_{S_N|\Theta}(A_s|\theta)},$$

where θ is the known rate of the good outcome (and the N's in the numerator and denominator cancel out).

Since the bet is better-than-fair, the probability of losing money goes to 0 as $N \to \infty$. Therefore, as $N \to \infty$, the denominator goes to 0. As $N \to \infty$, $\frac{A_s}{N}$ converges almost surely to θ, so the numerator goes to $-t + \theta(h+t) \geq 0$. Hence the ratio of the numerator to the denominator goes to infinity. This implies that Tommy always accepts the N-times repeated gamble for N sufficiently large.

Alternatively, for Barney,

$$\begin{aligned} \frac{\int_0^\infty z dG(z)}{-\int_{-\infty}^0 z dG(z)} &= \frac{\sum_{A_s = \lceil \frac{t}{h+t} N \rceil}^{N} (-t + \frac{A_s}{N}(h+t)) f^{\psi}_{S_N|\Theta}(A_s|\theta)}{-\sum_{A_s=0}^{\lceil \frac{t}{h+t} N \rceil - 1} (-t + \frac{A_s}{N}(h+t)) f^{\psi}_{S_N|\Theta}(A_s|\theta)} \\ &\to \frac{\int_{\frac{t}{h+t}}^1 (-t + \beta(h+t)) f^{\psi}_{ß|\Theta}(\beta|\theta)}{-\int_0^{\frac{t}{h+t}} (-t + \beta(h+t)) f^{\psi}_{ß|\Theta}(\beta|\theta)} \end{aligned}$$

almost surely as $N \to \infty$. This is a finite, positive number.

$\square$

Proposition 7. *Fix payoffs $u(\mu, \omega)$, rates $\theta_A > \theta_B$, and priors $f_\Theta(\theta_A)$ and $f_\Theta(\theta_B)$. Suppose Barney is prospective-atomizing or prospective-acceptive.*

1. *If Barney is retrospective-atomizing or retrospective-acceptive, then for any $c > 0$, Barney's behavior exactly coincides with Tommy's, and he will choose an action after observing a finite number of signals almost surely.*

2. *If Barney is retrospective-pooling, then for all $\overline{p} < 1$, there exists $\bar{c} > 0$ such that for all $c \leq \bar{c}$, Barney purchases an infinite number of signals with probability $p > \overline{p}$.*

3. *Suppose Barney is retrospective-pooling, $\theta_A = 1 - \theta_B$, and Barney is willing to purchase an additional signal when his posterior probability of state A equal to q. (a) If Barney's posterior probability of state A is q after observing N signals and again q after observing $N' > N$ signals, then the probability that Barney will purchase an infinite number of signals is weakly higher after he has observed the N' signals than it was after the N signals. (b) Furthermore, there exists a $\hat{N}$ such that for all $N' > \hat{N}$, the probability that Barney will purchase an infinite number of signals is strictly higher after he has observed the N' signals than after the N signals.*

Proof: Since Barney is prospective-atomizing, he conceives of his problem as in the classic sequential information acquisition setting of Wald (1947). Therefore, we can characterize his optimal policy using two thresholds in the posterior ratio. Denote the lower and upper threshold by $\underline{\lambda}$ and $\bar{\lambda}$, respectively. These thresholds are functions of the cost of signals and the payoffs for each action in each state (but not on the current posterior ratio). So long as as Barney's current posterior ratio is in $(\underline{\lambda}, \bar{\lambda})$, he will continue to acquire information. Once the posterior ratio is in the region $(0, \underline{\lambda}] \cup [\bar{\lambda}, 1)$, Barney will stop and take an action.

If Barney is retrospective-atomizing, then he is simply Tommy, who as it is well known, eventually stops in finite time almost surely (Wald, 1947). If Barney is retrospective-acceptive, then because he receives the signals one at a time, he will group them singly—and so his beliefs after any set of signals will be exactly the same as retrospective-atomizing Barney, and so the same results apply.

To prove the second part, now assume Barney is retrospective-pooling, and WLOG that state A is the true state. If Barney is retrospective-pooling, then for any two values m and n, such that $0 < m < n < \infty$, we can find a c' small enough so that if $c \leq c'$, then $\underline{\lambda} < m < n < \bar{\lambda}$. This is because, as in Wald (1947), the upper (lower) threshold is strictly increasing (decreasing) in c without bound.

Since $0 < p_0 < \infty$ and Barney's limit likelihood ratio $\Pi^{\psi}_{S_\infty|\Theta\times\Theta}(s_\infty|\theta_A, \theta_B, \Omega = i)$ is finite, his limit posterior is bounded away from 0 and ∞:

$$0 < p_0 \Pi^{\psi}_{S_\infty|\Theta\times\Theta}(s_\infty|\theta_A, \theta_B, \Omega = i) < \infty.$$

Moreover, since the posterior ratios after just a single a-signal or a single b-signal are bounded away from 0 and ∞, we can find a small enough $\hat{c} > 0$ so that both the following statements are true for all $c < \hat{c}$:

1. $[p_0 \Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B), p_0 \Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)] \subseteq (\underline{\lambda}, \bar{\lambda})$

2. $p_0 \Pi^{\psi}_{S_\infty|\Theta\times\Theta}(s_\infty|\theta_A,\theta_B,\Omega = i)\, p_0 \in (\underline{\lambda}, \bar{\lambda})$.

Therefore, by the proof of Lemma C, if $c < c'' \equiv \min\{c', \hat{c}\}$, then the probability that retrospective-pooling Barney purchases an infinite number of signals, conditional on still purchasing signals after N signals have already been purchased, is increasing in N and converges to 1 as N goes to infinity.

Fix $\overline{p} < 1$. By the above argument, for any $c < c''$, we find an $\overline{N}$ such that the probability that Barney purchases an infinite number of signals, conditional on Barney still purchasing signals after $\overline{N}$ signals have already been purchased, is larger than $\overline{p}$. Now, there exist $0 < m < n < \infty$ sufficiently far apart such that the probability equals 1 that retrospective-pooling Barney's posterior ratio lies the region (m, n) after he observes $\overline{N}$ signals. We can find c''' sufficiently small that for all $c \leq c'''$, $\underline{\lambda} < m < n < \bar{\lambda}$. It follows that if $c < \bar{c} \equiv \min\{c'', c'''\}$, then Barney purchases an infinite number of signals with probability larger than $\overline{p}$.

Turning to *(a)* of the third part of the proposition, denote the situations where there is a current posterior ratio of q, which is a belief where the agent still strictly wants to experiment, after N and $N' > N$ signals as q_N and $q_{N'}$, respectively. Assume that given $q_{N'}$, a sequence of m signals, denoted s_m, causes the agent to stop purchasing signals and (WLOG) take action a, and that no truncation of s_m causes the agent to stop purchasing signals. We will show that given initial situation q_N, the sequence s_m also causes the agent to stop purchasing signals. This proves that the probability of the agent stopping is weakly higher given q_N than $q_{N'}$.

If any truncation of s_m causes the agent to stop given q_N, then we are done, so assume not.

Now we shall show that a sufficient condition for part *(a)* is that the posterior induced by k a-signals in a row is less in favor of state A when starting with $q_{N'}$ than when starting with q_N. To see this, reorder s_m so that it begins with a and b signals paired off in alternating orders. Note that this reordering does not change the posterior after s_m, only the path of posterior ratios. Since by assumption the agent ends up favoring action a, there must be more a-signals than b-signals, so the reordered s_m must end with some set of a-signals in a row. Denote this number as k. Note that after every pair of signals, a, b, the likelihood ratio must still be q, since the signal structure is symmetric. So therefore, starting with either q_N or $q_{N'}$, after $m - k$ signals (the pairs of a and b signals) of the reordered s_m, the likelihood ratio is still q. Hence we are now simply comparing the effect of k a-signals in a row, given $N + m - k$ preceding signals versus $N' + m - k$ preceding signals.

Note that given $q_{N'}$, there must have been more preceding a-signals than given q_N. Assume there were r' and r a-signals, respectively. Hence in the sequence $q_{N'} \cup s_m$, there is a total of $N' + m$

signals, and the number of a-signals is r' (the number of a-signals in the first N') plus $\frac{m-k}{2}$ (the number of a-signals in the alternating set) plus k (the number of a-signals in a row at the end). Therefore, using Lemma $\beta 2$, the agent's posterior ratio after $q_{N'} \cup s_m$ is

$$\begin{aligned}
& \Pi^{\psi}_{\Theta\times\Theta|S_{N'+m}}\left(\theta_A, \theta_B | q_{N'} \cup s_m\right) \\
= \; & p_0 \frac{\Gamma\left(\theta_A \psi + r' + \frac{m-k}{2} + k\right)}{\Gamma\left(\theta_B \psi + r' + \frac{m-k}{2} + k\right)} \frac{\Gamma\left(\left(1-\theta_A\right)\psi + N' - r' + \frac{m-k}{2}\right)}{\Gamma\left(\left(1-\theta_B\right)\psi + N' - r' + \frac{m-k}{2}\right)} \\
= \; & p_0 \frac{\Gamma\left(\theta_A \psi + r' + \frac{m-k}{2}\right)}{\Gamma\left(\theta_B \psi + r' + \frac{m-k}{2}\right)} \frac{\Gamma\left(\left(1-\theta_A\right)\psi + N' - r' + \frac{m-k}{2}\right)}{\Gamma\left(\left(1-\theta_B\right)\psi + N' - r' + \frac{m-k}{2}\right)} \Pi^k_{i=1} \frac{\theta_A \psi + r' + \frac{m-k}{2} + i}{(1-\theta_A)\psi + r' + \frac{m-k}{2} + i} \\
= \; & q \Pi^k_{i=1} \frac{\theta_A \psi + r' + \frac{m-k}{2} + i}{(1-\theta_A)\psi + r' + \frac{m-k}{2} + i}.
\end{aligned}$$

In the sequence $q_N \cup s_m$, there is a total of $N+m$ signals, and the number of a-signals is r (the number of a-signals in the first N) plus $\frac{m-k}{2}$ (the number of a-signals in the alternating set) plus k (the number of a-signals in a row at the end). Therefore, analogously, the agent's posterior ratio after $q_N \cup s_m$ is

$$\Pi^{\psi}_{\Theta\times\Theta|S_{N'+m}}\left(\theta_A, \theta_B | q_N \cup s_m\right) = q \Pi^k_{i=1} \frac{\theta_A \psi + r + \frac{m-k}{2} + i}{(1-\theta_A)\psi + r + \frac{m-k}{2} + i}.$$

Notice that for all i,

$$1 < \frac{\theta_A \psi + r' + \frac{m-k}{2} + i}{(1-\theta_A)\psi + r' + \frac{m-k}{2} + i} < \frac{\theta_A \psi + r + \frac{m-k}{2} + i}{(1-\theta_A)\psi + r + \frac{m-k}{2} + i}$$

since $r' > r$ and both fractions are larger than 1. Therefore, since the agent stops purchasing signals given his posterior in situation $q_{N'}$, he he must stop given his posterior in situation q_N, which favors state A even more strongly.

To show part *(b)*, note that as N and N' grow farther apart, so do r and r'. Fix N, and WLOG, assume that the agent stops purchasing signals after a sequence of all k a-signals and not before. Note that as $N' \to \infty$ with k fixed, $r' \to \infty$, and so the agent's likelihood ratio becomes increasingly dominated by the alternating sequence of a and b signals that precedes the k a-signals:

$$\Pi^k_{i=1} \frac{\theta_A \psi + r' + \frac{m-k}{2} + i}{(1-\theta_A)\psi + r' + \frac{m-k}{2} + i} \to 1.$$

Therefore, the agent's posterior ratio approaches q,

$$\Pi_{i=1}^{k} \frac{\theta_A \psi + r' + \frac{m-k}{2} + i}{(1-\theta_A)\psi + r' + \frac{m-k}{2} + i} \to q,$$

and so the agent will *strictly* want to experiment for a large enough N'.

Proposition 8. *Suppose Barney is strategically sophisticated. Fix rates $\theta_A > \theta_B$ and Barneyness parameter $0 < \psi < \infty$. There are bounds $0 < \underline{p} < \frac{1}{2}$ and $\frac{1}{2} < \overline{p} < 1$ such that if the common prior $f_\Theta(\theta_A) \geq \underline{p}$, then there is positive probability of herding on action a; and if $f_\Theta(\theta_A) \leq \overline{p}$, then there is positive probability of herding on action b.*

Proof: All types of agents, after observing the set of actions taken previously to them, and their own signal, will simply take the action optimal for them given their beliefs. Denoting the private posterior ratio of state A to state B by agent t, conditional on observing the vector of previous actions α_{t-1} and the private signal s as $\Pi^{\psi}_{\Theta\times\Theta|S_t}(\theta_A, \theta_B|\alpha_{t-1}, s)$, an agent will take action a if $\Pi^{\psi}_{\Theta\times\Theta|S_t}(\theta_A, \theta_B|\alpha_{t-1}, s) > 1$ and action b if $\Pi^{\psi}_{\Theta\times\Theta|S_t}(\theta_A, \theta_B|\alpha_{t-1}, s) < 1$. As noted in the text, we focus on the most informative equilibrium: if $\Pi^{\psi}_{\Theta\times\Theta|S_t}(\theta_A, \theta_B|\alpha_{t-1}, s) = 1$, then an agent will take the action that matches his own signal s.

Denote the common prior ratio as p_0. We will show that if $p_0 \geq \underline{p} \equiv \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}$, then there is positive probability that parameterized-Barney will herd on action a, and if $p_0 \leq \overline{p} \equiv \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}$, then there is positive probability that parameterized-Barney will herd on action b. Proving this claim will prove the proposition.

First, consider a common prior ratio in the interval:

$$p_0 \in [0,\infty) \backslash \left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)} \right].$$

If $p_0 \leq \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}$, it will be optimal for the first agent to take action b regardless of his signal, and so a herd on action b begins in period 1. If $p_0 \geq \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}$, it will be optimal for the first agent to take action a regardless of his signal, and so a herd on action a begins in period 1.

Now consider the set of prior ratios for which no herd begins in period 1:

$$p_0 \in \left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)} \right].$$

We will conclude the proof by showing that, for prior ratios in this interval, there is positive probability of a herd occurring on either action. Note that for any number $L \in (1, \infty)$, there exists a long enough sequence of all a-signals such that parameterized-Barney's likelihood ratio is $\geq L$ after observing this sequence (due to the functional form of the beta distribution). Given that p_0 is in the above interval, it follows that for a large enough sample of signals, there is positive probability that beta-binomial-Barney's posterior ratio will be larger than $\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}$, and so a herd will begin on action a. An analogous argument shows that there is positive probability that a herd will begin on action b.

□

Proposition 9. *Suppose Barney is strategically sophisticated. There exist $\delta > 0$ and $\gamma > 0$ such that if $|f_\Theta(\theta_A) - f_\Theta(\theta_B)| < \delta$ and $|\theta_A - \theta_B| < \gamma$, then there is positive probability of an eddy regardless of which state is true.*

Proof: Strategically sophisticated Barney believes that a herd has started if and only if a herd has actually started. If a herd has started, Barney will ignore his own signal and join the herd, since the previous agent faced the same problem and joined the herd, and Barney has accurate beliefs about that agent's beliefs. So we only need to consider situations where Barney believes no herd has started.

For an eddy to exist, we need that agents will never ignore their own signal. As long as the public posterior ratio is in the interval

$$\left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}\right],$$

agents will take actions corresponding to their signal. This is because the agent's private posterior ratio after observing his own signal will be between

$$\frac{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}$$

and

$$\frac{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)},$$

an interval that includes 1.

WLOG, assume that the true state is A. To guarantee that there is positive probability that the public posterior ratio remains in the above interval, it is sufficient that the prior and limit likelihood ratio satisfy two conditions:

- The first agent must not ignore his own signal; this means that

$$p_0 \in \left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(a|\theta_A,\theta_B\right)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(b|\theta_A,\theta_B\right)}\right].$$

- The limit posterior conditional on the true state is:

$$\Pi^{\psi}_{S_\infty|\Theta\times\Theta}\left(s_\infty|\theta_A,\theta_B,\Omega=A\right)p_0 \in \left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(a|\theta_A,\theta_B\right)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(b|\theta_A,\theta_B\right)}\right].$$

Since

$$1 \in \left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(a|\theta_A,\theta_B\right)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(b|\theta_A,\theta_B\right)}\right],$$

we can find some δ such that if $|f_\Theta(\theta_A) - f_\Theta(\theta_B)| < \delta$, then the first condition holds. Furthermore, since the limit likelihood ratio, $\Pi^{\psi}_{S_\infty|\Theta\times\Theta}\left(s_\infty|\theta_A,\theta_B,\Omega=A\right)$, is equal to 1 when $\theta_A = \theta_B$, and is continuous in both rates, it follows that for any $\tau > 0$ there exists a γ such that if $|\theta_A - \theta_B| < \gamma$, then the limit likelihood ratio (conditional on the true state) is within τ of 1. Therefore, conditional on $|f_\Theta(\theta_A) - f_\Theta(\theta_B)| < \delta$, we can find rates θ_A and θ_B such that the second condition is also satisfied.

We can now use Lemma C. Although Lemma C as stated applies to retrospective-pooling Barney, we can use it here to bound the public posterior likelihood ratio. This is because Barney is retrospective-pooling with respect to the public signals he has inferred from previous actions. Therefore, we know that with positive probability, the public posterior ratio stays within $\left[\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)}, \frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B)}\right]$ forever. Therefore, with positive probability, an eddy will result.

□

Proposition 10. *Suppose Barney is strategically sophisticated.*

1. *For any $\varepsilon > 0$, there exists $T_\varepsilon > 2$ such that if a herd has not occurred by period T_ε, the probability of herding from then on is less than ε.*

2. *Furthermore, suppose the true state is A, the prior $f_\Theta(\theta_A) \geq .5$, and $1 - \theta_A = \theta_B$. For any $\nu > 0$ there exists $T_\nu > 2$ such that if a herd has not occurred by period T_ν, the likelihood ratio of a bad herd to a good herd is less than ν.*

Proof:

1. For a herd to begin, the public posterior ratio must leave $\left[\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(b|\theta_A,\theta_B\right),\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(a|\theta_A,\theta_B\right)\right]$. Since an eddy occurs with positive probability, from the proof of Lemma C, we know that if a herd has not yet occurred after $\hat{N}$ signals, the probability of the public posterior ratio leaving $\left[\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(b|\theta_A,\theta_B\right),\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(a|\theta_A,\theta_B\right)\right]$ is less than $1-\sum_{i=\hat{N}}^{\infty}\frac{p(1-p)}{i^2\delta^2}$, where $p = \Pi^{\psi}_{S_1|\Theta\times\Theta}\left(b|\theta_A,\theta_B\right)$ and δ is the distance between the limit proportion of a-signals and the closest proportion that induces a herd. This expression is an upper bound on the probability of a herd ever starting, if it has not already started after $\hat{N}$ signals have occurred. As noted in the proof of Lemma C, this upper bound converges to 0 as $\hat{N}\to\infty$; hence, for any $\epsilon>0$, we can find a T_ε such that if a herd has not occurred by period T_ε, then the probability of a herd occurring is less than ϵ.

2. Denote the prior ratio on state A by $p=\frac{f_\Theta(\theta_A)}{f_\Theta(\theta_B)}$. Denote the probability of a good herd (that is, a herd on state A) given parameters θ_A,θ_B,ψ,p, and with N agents already have taken an action, as $V(A,\theta_A,\theta_B,\psi,N,p)$. Denote the probability of a bad herd (a herd on state B) as $V(B,\theta_A,\theta_B,\psi,N,p)$. For brevity, we will refer to these as $V(A,N,p)$ and $V(B,N,p)$, respectively. We will show that, for any $\nu>0$, there exists $T_\nu>2$ such that if a herd has not occurred by period T_ν, then for all $N\geq T_\nu$, $\frac{V(B,N,p)}{V(A,N,p)}<\nu$. We will do this in several steps:
   (a) First, we will construct $V\left(i,N,p\right)$ in a way that will make it manageable to work with.
   (b) Second, we will fix equal priors $p=1$, and we will show that, for any $\nu>0$, there exists $T_\nu>2$ such that for all $N\geq T_\nu$, $\frac{V(B,N,1)}{V(A,N,1)}<\nu$.
   (c) Third, we will show that for any priors $p>1$ that strictly favor the true state and for all $N\geq T_\nu$, $\frac{V(B,N,p)}{V(A,N,p)}<\frac{V(B,N,1)}{V(A,N,1)}$. This is because relative to the equal priors case, there is a uniform shift of probability from likelihoods that favor an incorrect herd to a correct herd.

Now we shall do each step in order:

(a) Because the signals are symmetric, we can define

$$\lambda\equiv\frac{1}{\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(a|\theta_A,\theta_B\right)}=\Pi^{\psi}_{S_1|\Theta\times\Theta}\left(b|\theta_A,\theta_B\right).$$

The set of public likelihood ratios such that a herd has not yet started after N signals is denoted $Q(N)\subset[\lambda,\frac{1}{\lambda}]$, a set with a finite number of elements. Denote a typical element by q.

We will construct V from a more elementary function: $\phi(A, q, N, M)$, the probability of a herd starting on the a action in the next M signals when the current public likelihood ratio is q, and N signals have already been observed. Given p and N, there is a one-to-one mapping between the public likelihood ratio q and the "number of a signals" previously observed, which may not be a natural number for a particular q. Therefore we will consider the number of a and b signals as fixed at the levels implied by q.

After N agents have taken actions, a herd starts on action a only if the number of a-signals exceeds the number of b-signals by at least some number. Abusing notation slightly, denote by s_M^j an exact sequence of a and b signals from a set of M signals, with ordering j indexed from 0 to $2^M - 1$ (so s_M^j could contain any number of signals from 0 up to M signals). Denote the number of a-signals in s_M^j as $A(s_M^j)$. Define $a(s_M^j, N, q) \in \{0, 1\}$ as an indicator of whether a herd would have started on action a with the exact ordering of signals in s_M^j, given that the public likelihood ratio was q after the first N signals. Now we can construct the probability of a herd starting on action a within the next M signals, given that the public likelihood ratio was q after the first N signals:

$$\phi(A, q, N, M) = \sum_{j=1}^{2^M-1} a(s_M^j, N, q)\theta_A^{A(s_M^j)}(1-\theta_A)^{M-A(s_M^j)},$$

a polynomial in θ_A and $1 - \theta_A$.

To gain some intuition for ϕ, consider $q = 1$. Because the rates are symmetric, if we consider a set of j signals that causes a herd on the a action, and we replace all the θ_A's with $(1 - \theta_A)$'s, and vice-versa, then that set of signals will cause a herd on action b. Similarly, if we have a set of j signals that has not caused a herd, then the same set of signals, but replacing θ_A's with $(1 - \theta_A)$'s and vice-versa, will also not cause a herd. Therefore, $\phi(B, 1, N, M)$ is the same function as $\phi(A, 1, N, M)$ but with all the θ_A's replaced with $(1 - \theta_A)$'s and vice-versa. Now suppose either $q > 1$ or $q < 1$. Now $\phi(B, q, N, M)$ is not simply $\phi(A, q, N, M)$ with the rates switched. Because rates are symmetric, however, if we define $q' = \frac{1}{q}$, then $\phi(B, q', N, M)$ is the same function as $\phi(A, q, N, M)$ but with θ_A and $1 - \theta_A$ switched.

Define $\rho(N)$ as the minimum number of a-signals in a row that would need to occur after an initial set of N signals to put the public likelihood ratio above $\frac{1}{\lambda}$ (the upper bound of the interval $Q(N)$, as defined above), given that the public likelihood ratio was λ (the lower bound of the interval $Q(N)$) after the first N signals. $\rho(N)$ will be a useful function to define because, regardless of the public likelihood ratio after N observations, it is the number of additional signals after which a herd on either action could always possibly start. Note that $\rho(N)$ is increasing in N.

Define $x(q,N,p) < N$ as the number of a-signals out of N such that $x(q,N,p)$ a-signals and $N - x(q,N,p)$ b-signals induces posterior ratio of q when the prior is p.

For each $x(q,N,p)$, we can construct the set of possible sequences that generates x a-signals and $N - x$ b-signals without starting a herd at any point in the sequence. Denote the set of these sequences as $X(q,N,p)$. Each one of these sequences must have the same final number of a and b signals, but it must be the case that the difference between the number of a-signals and b-signals cannot have been too large at any point in the sequence, or else a herd would have started. Let $g(q,N,p) = \gamma(q,N,p)\theta_A^{x(q,N,p)}(1-\theta_A)^{N-x(q,N,p)}$ denote the probability of the set of sequences that leads to x a-signals and $N - x$ b-signals without a herd ever having started, where γ is the number of sequences in the set.

Define $r_N(i)$ for $i = 0, 1, ...$ recursively: $r_N(0) = N$ and $r_N(i) = r_N(i-1) + \rho(r_N(i-1))$ for $i > 0$.

We now construct $V(A,N,p)$ as follows:

$$V(A,N,p) = \sum_{i=1}^{\infty} \sum_{q \in Q(r_N(i-1))} \phi(A, q, r_N(i-1), \rho(r_N(i-1)))g(q, r_N(i-1), p).$$

We define $V(B,N,p)$ analogously.

To understand this definition of V, each term is the probability of a herd starting on the appropriate action action within the next $r_N(i)$ signals, conditional on having observed $r_N(i-1)$ signals without a herd, starting with a prior ratio of p.

(b) Now we will fix the prior ratio at 1 and show that for any $\nu > 0$, there exists $T_\nu > 2$ such that for all $N \geq T_\nu$, $\frac{V(B,N,1)}{V(A,N,1)} < \nu$.

We will begin by proving some key claims:

Claim 10.A *Let $a_i, b_i > 0$ for $i = 1, 2...$, such that $\frac{a_i}{b_i} \leq \frac{a_{i+1}}{b_{i+1}}$. Then $\frac{a_1}{b_1} \leq \frac{\sum_{i=1}^{N} a_i}{\sum_{i=1}^{N} a_i} \leq \frac{a_N}{b_N}$.*

To see this is true, note that for positive a, b, c, d, if $0 < \frac{a}{b} < \frac{c}{d}$, then $\frac{a}{b} < \frac{a+c}{b+d} < \frac{c}{d}$. Induction then shows that $\frac{\sum_{i=1}^{N} a_i}{\sum_{i=1}^{N} b_i}$ is bounded by the maximum and minimum $\frac{a_i}{b_i}$.

Claim 10.B *For any $\nu > 0$, there exists $T_\nu > 2$ such that for all $N \geq T_\nu$,*

$$\frac{\sum_{q \in Q(N)} \phi(B, q, N, \rho(N))g(q, N, p)}{\sum_{q \in Q(N)} \phi(A, q, N, \rho(N))g(q, N, p)} < \nu.$$

Together, Claims 10.A and 10.B imply the result we are seeking. Claim 10.A, combined with the fact that the probability of a herd is shrinking as the number of signals gets large, implies that

$$\begin{aligned}\frac{V(B,N,p)}{V(A,N,p)} &= \frac{\sum_{i=1}^{\infty}\sum_{q\in Q(N)}\phi(B,q,r_N(i-1),r_N(i))g(q,r_N(i-1),p)}{\sum_{i=1}^{\infty}\sum_{q\in Q(N)}\phi(A,q,r_N(i-1),r_N(i))g(q,r_N(i-1),p)} \\ &\le \frac{\sum_{q\in Q(N)}\phi(B,q,N,\rho(N))g(q,N,p)}{\sum_{q\in Q(N)}\phi(A,q,N,\rho(N))g(q,N,p)}.\end{aligned}$$

Claim 10.B then says that given any $\nu > 0$, for large enough N, this ratio is $< \nu$.

A sufficient condition for Claim 10.B is:

Claim 10.C *For any $\nu > 0$, there exists $T_\nu > 2$ such that for all $N \ge T_\nu$ and $q \in Q(N)$,*

$$\frac{\phi(B,q,N,\rho(N))g(q,N,1)+\phi(B,q',N,\rho(N))g(q',N,1)}{\phi(A,q,N,\rho(N))g(q,N,1)+\phi(A,q',N,\rho(N))g(q',N,1)} < \nu.$$

To see that Claim 10.C implies Claim 10.B, notice that $\frac{\sum_{q\in Q(N)}\phi(B,q,N,\rho(N))g(q,N,p)}{\sum_{q\in Q(N)}\phi(A,q,N,\rho(N))g(q,N,p)}$ is

$$\begin{aligned}&\le \max_{q,q'\in Q(N)\text{ s.t. } q'=\frac{1}{q}} \frac{\phi(B,q,N,\rho(N))g(q,N,1)+\phi(B,q',N,\rho(N))g(q',N,1)}{\phi(A,q,N,\rho(N))g(q,N,1)+\phi(A,q',N,\rho(N))g(q',N,1)} \\ &< \nu\end{aligned}$$

where the first inequality follows from Claim 10.A and the second by Claim 10.C.

Therefore, we will have completed part (b) once we prove Claim 10.C. We substitute:

$$\begin{aligned}&\phi(A,q,N,\rho(N))g(q,N,1) \\ =\;& g(q,N,1)\sum_{j=1}^{2^{\rho(N)}-1} a(s^j_{\rho(N)},N,q)\theta_A^{A(s^j_{\rho(N)})}(1-\theta_A)^{\rho(N)-A(s^j_{\rho(N)})} \\ =\;& \gamma(q,N,1)\theta_A^{x(q,N,1)}(1-\theta_A)^{N-x(q,N,1)}\sum_{j=1}^{2^{\rho(N)}-1} a(s^j_{\rho(N)},N,q)\theta_A^{A(s^j_{\rho(N)})}(1-\theta_A)^{\rho(N)-A(s^j_{\rho(N)})}.\end{aligned}$$

Note that some of the sequences of signals starting with a posterior of q will cause a herd before $\rho(N)$ signals. Imagine two sequences begin with the same initial z signals but differ in the signals that occur afterward. If the z^{th} signal begins a herd, then these subsequent signals are irrelevant, and we can aggregate the two sequences together. Let $\zeta(q,\rho(N),N,z,i)$ denote the number of a-signals needed out of an additional $\rho(N)$ signals, starting with a posterior of q after N signals, for a herd on i to begin after exactly z signals, and no herd having started sooner. Let the lowercase letter ς

count the number of permutations of such sequences. Continuing to substitute:

$$\begin{aligned}
& \phi(A,q,N,\rho(N))g(q,N,1) \\
= \quad & \gamma(q,N,1)\theta_A^{x(q,N,1)}(1-\theta_A)^{N-x(q,N,1)} \sum_{z=1}^{\rho(N)} \varsigma(q,\rho(N),N,z,A)\theta_A^{\zeta(q,z,N,A)}(1-\theta_A)^{z-\zeta(q,z,N,A)}.
\end{aligned}$$

Now, define

$$\begin{aligned}
b_z(q',N) &\equiv \gamma(q',N,1)\theta_A^{x(q',N,1)}(1-\theta_A)^{N-x(q',N,1)}\varsigma(q',\rho(N),N,z,B)(1-\theta_A)^{z-\zeta(q',z,N,B)}\theta_A^{\zeta(q',z,N,B)} \\
a_z(q',N) &\equiv \gamma(q,N,1)\theta_A^{x(q,N,1)}(1-\theta_A)^{N-x(q,N,1)}\varsigma(q,\rho(N),N,z,A)(1-\theta_A)^{z-\zeta(q,z,N,A)}\theta_A^{\zeta(q,z,N,A)}.
\end{aligned}$$

We can write

$$\frac{\phi(B,q,N,\rho(N))g(q,N,1)+\phi(B,q',N,\rho(N))g(q',N,1)}{\phi(A,q,N,\rho(N))g(q,N,1)+\phi(A,q',N,\rho(N))g(q',N,1)} = \frac{\sum_{z=1}^{\rho(N)} b_z(q,N)+\sum_{z=1}^{\rho(N)} b_z(q',N)}{\sum_{z=1}^{\rho(N)} a_z(q,N)+\sum_{z=1}^{\rho(N)} a_z(q',N)}$$

Starting from a prior ratio of 1, because the signals are symmetric, if x a-signals and $N-x$ b-signals leaves an agent with a posterior of q, then x b-signals and $N-x$ a-signals leaves an agent with a posterior of q'. Therefore $x(q',N,1) = N - x(q,N,1)$. Similarly, $\gamma(q',N,1) = \gamma(q,N,1)$. Also, starting with a posterior of q, if x additional a-signals and $N-x$ additional b-signals leaves an agent willing to herd on action A, then starting with a posterior of q', x additional b-signals and $N-x$ additional a-signals leaves an agent willing to herd on action B. Therefore $\zeta(q',z,N,B) = z - \zeta(q,z,N,A)$ and $\varsigma(q',\rho(N),N,z,B) = \varsigma(q,\rho(N),N,z,A)$.

Using these relationships,

$$\begin{aligned}
\frac{b_z(q',N)}{a_z(q',N)} &= \frac{\gamma(q',N,1)\theta_A^{x(q',N,1)}(1-\theta_A)^{N-x(q',N,1)}\varsigma(q',\rho(N),N,z,B)(1-\theta_A)^{z-\zeta(q',z,N,B)}\theta_A^{\zeta(q',z,N,B)}}{\gamma(q,N,1)\theta_A^{x(q,N,1)}(1-\theta_A)^{N-x(q,N,1)}\varsigma(q,\rho(N),N,z,A)(1-\theta_A)^{z-\zeta(q,z,N,A)}\theta_A^{\zeta(q,z,N,A)}} \\
&= \frac{\gamma(q,N,1)\theta_A^{N-x(q,N,1)}(1-\theta_A)^{x(q,N,1)}\varsigma(q,\rho(N),N,z,A)(1-\theta_A)^{\zeta(q,z,N,A)}\theta_A^{z-\zeta(q,z,N,A)}}{\gamma(q,N,1)\theta_A^{x(q,N,1)}(1-\theta_A)^{N-x(q,N,1)}\varsigma(q,\rho(N),N,z,A)(1-\theta_A)^{z-\zeta(q,z,N,A)}\theta_A^{\zeta(q,z,N,A)}},
\end{aligned}$$

which is well-defined as long as $\gamma(q,N,1)\varsigma(q,\rho(N),N,z,A) \neq 0$. (If $\gamma(q,N,1)\varsigma(q,\rho(N),N,z,A) = 0$, then b_z drops out from $V(B,N,p)$, and a_z drops out from $V(A,N,p)$. Hence these terms can be ignored.) Canceling and collecting terms:

$$\frac{b_z(q',N)}{a_z(q',N)} = \frac{(1-\theta_A)^{x(q,N,1)+\zeta(q,z,N,A)}\,(\theta_A)^{N+z-x(q,N,1)-\zeta(q,z,N,A)}}{(\theta_A)^{x(q,N,1)+\zeta(q,z,N,A)}\,(1-\theta_A)^{N+z-x(q,N,1)-\zeta(q,z,N,A)}}$$

Using Claim 10.A ,

$$\frac{\sum_z b_z(q',N)}{\sum_z a_z(q',N)} < \max_{z\ s.t.\ \gamma(q,N,1)\varsigma(q',\rho(N),N,z,B)\neq 0} \frac{(1-\theta_A)^{x(q,N,1)+\zeta(q,z,N,A)}(\theta_A)^{N+z-x(q,N,1)-\zeta(q,z,N,A)}}{(\theta_A)^{x(q,N,1)+\zeta(q,z,N,A)}(1-\theta_A)^{N+z-x(q,N,1)-\zeta(q,z,N,A)}}.$$

The numerator is the probability that, starting with a prior ratio of 1, there will be a sequence of a and b signals, $N+z$ in total, that lands the agent just inside the zone where the agent will herd on action B. The denominator is the probability that, starting with a prior ratio of 1, there will be a sequence of a and b signals, $N+z$ in total, that lands the agent just inside the zone where the agent will herd on action A. This expression can be rearranged to yield

$$\left(\frac{1-\theta_A}{\theta_A}\right)^{2x(q,N,1)+2\zeta(q,z,N,A)-N-z},$$

where the exponent, $2x(q,N,1)+2\zeta(q,z,N,A)-N-z \equiv \delta(N+z)$, is the minimum difference between the number of a and b signals out of $N+z$ signals in total at which a herd will start; if there are $\delta(N+z)$ more a than b signals, then there will be a herd on A; and if there are $\delta(N+z)$ more b than a signals, then there will be a herd on B. Notice that $\delta(N+z)$ is not a function of q. Because Barney's likelihood ratio asymptotically depends on the proportion of a signals, not the difference between the number of a and b signals, $\lim_{N+z\to\infty}\delta(N+z)=\infty$. Furthermore, since $\rho(N)$ is growing in N, then it follows that for any $\epsilon>0$ we can find an N' large enough such that for all $q \in \left(\Pi^{\psi}_{S_1|\Theta\times\Theta}(b|\theta_A,\theta_B), \Pi^{\psi}_{S_1|\Theta\times\Theta}(a|\theta_A,\theta_B)\right)$ (and so for any $q\in Q$ for any N) and all $N>N'$,

$$\left(\frac{1-\theta_A}{\theta_A}\right)^{2x(q,N,1)+2\zeta(q,z,N,A)-N-z},$$

is less than ϵ. It then follows directly from the previous claims that the upper bound on $\frac{V(B,N,p)}{V(A,N,p)}$, $\left(\frac{1-\theta_A}{\theta_A}\right)^{\delta(N+z)}$, goes to 0 as $N\to\infty$.

Therefore, although Q varies with N, it must be the case that for any q in Q, the limit result holds. That is, the convergence is uniform over $q\in Q$.

(c) Now assume that the prior rate favors A: $p>1$. Note that nothing in the construction of the ϕ function changes. However, because p is now closer to $q>1$ than it is to $q'<1$, $\frac{g(q,N,p)}{g(q',N,p)} < \frac{g(q,N,1)}{g(q',N,1)}$. Using this fact, and the result proven above for the $p=1$ case: for any $\nu>0$, there exists $T_\nu>2$ such that for all $N\geq T_\nu$, for any $q\in Q(N)$,

$$\frac{\phi(B,q,N,\rho(N))g(q,N,p)}{\phi(A,q',N,\rho(N))g(q',N,p)} < \frac{\phi(B,q,N,\rho(N))g(q,N,1)}{\phi(A,q',N,\rho(N))g(q',N,1)} < \nu,$$

and

$$\frac{\phi(B,q',N,\rho(N))g(q',N,p)}{\phi(A,q,N,\rho(N))g(q,N,p)} < \frac{\phi(B,q',N,\rho(N))g(q',N,1)}{\phi(A,q,N,\rho(N))g(q,N,1)} < \nu.$$

We argued in part (a) that Claim 10.A implies

$$\frac{V(B,N,p)}{V(A,N,p)} \leq \frac{\sum_{q\in Q(N)} \phi(B,q,N,\rho(N))g(q,N,p)}{\sum_{q\in Q(N)} \phi(A,q,N,\rho(N))g(q,N,p)}.$$

Using Claim 10.A again, along with the above inequalities, gives the result:

$$\frac{V(B,N,p)}{V(A,N,p)} < \nu.$$

□

References

[1] P. Aghion, P. Bolton, C. Harris, and B. Julien. Optimal learning by experimentation. *Review of Economic Studies*, 58:621–654, 1991.

[2] S. Nageeb Ali. Learning self-control. *Quarterly Journal of Economics*, 126:857–893, 2011.

[3] Abhijit V. Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, 1992.

[4] Maya Bar-Hillel. The base-rate fallacy in probability judgments. *Acta Psychologica*, 44(3):211–233, 1980.

[5] Nicholas Barberis, Ming Huang, and Richard H. Thaler. Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American Economic Review*, 96(4):1069–1090, 2006.

[6] Nicholas Barberis, Andrei Schleifer, and Robert Vishny. A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–343, 1998.

[7] Lee Roy Beach. Probability magnitudes and conservative revision of subjective probabilities. *Journal of Experimental Psychology*, 77(1):57–63, 1968.

[8] Lee Roy Beach, James A. Wise, and Scott Barclay. Sample proportions and subjective probability revisions. *Organizational Behavior and Human Performance*, 5(2):183–190, March 1970.

[9] Shlomo Benartzi and Richard Thaler. Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics*, 110(1):73–92, 1995.

[10] Shlomo Bernatzi and Richard Thaler. Risk aversion or myopia? choices in repeated gambles and retirement investments. *Management Science*, 45(3):364–381, 1999.

[11] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Econmoy*, 100(5):992–1026, 1992.

[12] D. Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematics, Statistics and Probability*, pages 93–102. University of California Press, 1951.

[13] D. Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24:265–272, 1953.

[14] Aaron L. Bodoh-Creed. The economics of base rate neglect. June 2010. Working paper.

[15] Colin F. Camerer. Do biases in probability judgment matter in markets? experimental evidence. *The American Economic Review*, 77(5):981–997, December 1987.

[16] Richard C. Chapman. Prior probability bias in information seeking and opinion revision. *American Journal of Psychology*, 86(2):269–282, 1973.

[17] James O. Jr. Chinnis and Cameron R. Peterson. Inference about a nonstationary process. *Journal of Experimental Psychology*, 77(4):620–625, 1968.

[18] John Cohen and C. E. M. Hansel. The idea of a distribution. *British Journal of Psychology*, 46(2):111–121, 1955.

[19] H. C. A. Dale. Weighing evidence: An attempt to assess the efficiency of the human operator. *Ergonomics*, 11(3):215–230, May 1968.

[20] Chetan Dave and Katherine W. Wolfe. On confirmation bias and deviations from bayesian updating. March 2003. Working paper.

[21] J.H. De Swart. Effects of diagnosticity and prior odds on conservatism in a bookbag-and-pokerchip situation. *Acta Psychologica*, 36(1):16–31, February 1972.

[22] J.H. DeSwart. A comment on marks and clarkson's explanation of conservatism. *Acta Psychologica*, 36(5):417–419, November 1972.

[23] Michael L. Donnell and Wesley M. Du Charme. The effect of bayesian feedback on learning in an odds estimation task. *Organizational Behavior and Human Performance*, 14(3):305–313, December 1975.

[24] Wesley M. DuCharme. Response bias explanation of conservative human inference. *Journal of Experimental Psychology*, 85(1):66–74, 1970.

[25] Wesley M. DuCharme and Cameron R. Peterson. Intuitive inference about normally distributed populations. *Journal of Experimental Psychology*, 78(2):269–275, 1968.

[26] W. Edwards. Conservatism in human information processing. In B. Kleinmuntz, editor, *Formal representation of human judgment*. Wiley, 1968.

[27] Erik Eyster and Matthew Rabin. Naive herding in rich-information settings. *American Economic Journal: Microeconomics*, 2(4):221–243, 2010.

[28] Lisbeth S. Fried and Cameron R. Peterson. Information seeking: Optimal versus fixed stopping. *Journal of Experimental Psychology*, 80(3):525–529, 1969.

[29] Charles F. Gettys and Charles W. Manley. The probability of an event and estimates of posterior probability based upon its occurence. *Psychonomic Science*, 11(2):47–48, 1968.

[30] Paul E. Green, Michael H. Halbert, and Patrick J. Robinson. An experiment in probability estimation. *Journal of Marketing Research*, 2(3):266–273, August 1965.

[31] David M. Grether. Bayes rule as a descriptive model: The representativeness heuristic. *Quarterly Journal of Economics*, 95(3):537–557, 1980.

[32] David M. Grether. Testing bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior and Organization*, 17(1):31–57, January 1992.

[33] Dale Griffin and Amos Tversky. The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24(3):411–435, July 1992.

[34] David H. Gustafson, Ramesh K. Shukla, Andre Delbecq, and G. William Walster. A comparative study of differences in subjective likelihood estimates made by individuals, interacting groups, delphi groups, and nominal groups. *Organizational Behavior And Human Performance*, 9(2):280–291, April 1973.

[35] Ramon L. Hershman and J.R. Levine. Deviations from optimum information-purchase strategies in human decision-making. *Organizational Behavior And Human Performance*, 5(4):313–329, July 1970.

[36] Daniel Kahneman and Dan Lovallo. Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science*, 39(1):17–31, 1993.

[37] Daniel Kahneman and Amos Tversky. Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3):430–454, July 1972.

[38] Daniel Kahneman and Amos Tversky. On the psychology of prediction. *Psychological Review*, 80(4):237–251, 1973.

[39] Gideon Keren. Calibration and probability judgments: Conceptual and methodological issues. *Acta Psychologica*, 77(3):217–273, 1991.

[40] Gideon Keren and Willem A. Wagenaar. Violation of utility theory in unique and repeated gambles. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13(3):387–391, 1987.

[41] Alexander Klos, Elke U. Weber, and Martin Weber. Investment decisions and time horizon: Risk perception and risk behavior in repeated gambles. March 2005. Working Paper.

[42] Carlo Kraemer and Martin Weber. How do people take into account weight, strength, and quality of segregated vs. aggregated data? experimental evidence. *The Journal of Risk and Uncertainty*, 29(2):113–142, September 2004.

[43] D.F. Marks and J.K. Clarkson. An explanation of conservatism in the bookbag-and-pokerchips situation. *Acta Psychologica*, 36(2):145–160, April 1972.

[44] David W. Martin. Data conflict in a multinomial decision task. *Journal of Experimental Psychology*, 82(1):4–8, 1969.

[45] David W. Martin and Charles F. Gettys. Feedback and response mode in performing a bayesian decision task. *Journal of Applied Psychology*, 53(5):413–418, 1969.

[46] Mark W. Nelson, Robert Bloomfield, Jeffrey W. Hales, and Robert Libby. The effect of information strength and weight on behavior in financial markets. *Organizational Behavior and Human Decision Processes*, 86(2):168–196, November 2001.

[47] Richard E. Nisbett and Lee Ross. *Human inference: Strategies and shortcomings of social judgment*. Prentice Hall, Englewood Cliffs, NJ, 1980.

[48] C. R. Peterson and L.D. Phillips. Revision of continuous subjective probability distributions. *IEEE Transactions on Human Factors in Electronics*, HFE-7(1):19–22, March 1966.

[49] Cameron R. Peterson, Wesley M. DuCharme, and Ward Edwards. Sampling distributions and probability revisions. *Journal of Experimental Psychology*, 76(2):236–243, 1968.

[50] Cameron R. Peterson and Alan J. Miller. Sensitivity of subjective probability revision. *Journal of Experimental Psychology*, 70(1):117–121, 1965.

[51] Cameron R. Peterson, Robert J. Schneider, and Alan J. Miller. Sample size and the revision of subjective probabilities. *Journal of Experimental Psychology*, 69(5):522–527, 1965.

[52] Cameron R. Peterson and Richard G. Swensson. Intuitive statistical inferences about diffuse hypotheses. *Organizational Behavior And Human Performance*, 3(1):1–11, February 1968.

[53] Lawrence D. Phillips and Ward Edwards. Conservatism in a simple probability inference task. *Journal of Experimental Psychology*, 72(3):346–354, 1966.

[54] L.D. Phillips, W.L. Hays, and W. Edwards. Conservatism in complex probabilistic inference. *IEEE Transactions on Human Factors in Electronics*, HFE-7(1):7–18, March 1966.

[55] Gordon F. Pitz. Sample size, likelihood, and confidence in a decision. *Psychonomic Science*, 8(6):257–258, 1967.

[56] Gordon F. Pitz. Information seeking when available information is limited. *Journal of Experimental Psychology*, 76(1):25–34, 1968.

[57] Gordon F. Pitz and Helen Reinhold Barrett. Information purchase in a decision task following the presentation of free information. *Journal of Experimental Psychology*, 82(3):410–414, 1969.

[58] Matthew Rabin. Diminishing marginal utility of wealth cannot explain risk aversion. In Daniel Kahneman and Amos Tversky, editors, *Choices, Values, and Frames*, pages 202–208. Cambridge University Press, 2000.

[59] Matthew Rabin. Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5):1281–1292, 2000.

[60] Matthew Rabin. Inference by believers in the law of small numbers. *The Quarterly Journal of Economics*, 117(3):775–816, August 2002.

[61] Matthew Rabin and Richard H. Thaler. Anomalies: Risk aversion. *The Journal of Economic Perspectives*, 15(1):219–232, 2001.

[62] Matthew Rabin and Dmitri Vayanos. The gambler's and hot-hand fallacies: Theory and applications. *Review of Economic Studies*, 77(2):730–778, 2010.

[63] Matthew Rabin and Georg Weizsacker. Narrow bracketing and dominated choices. *American Economic Review*, 99(4):1508–1543, 2009.

[64] Daniel Read, George Loewenstein, and Matthew Rabin. Choice bracketing. *Journal of Risk and Uncertainty*, 19(1-3):171–197, 1999.

[65] Donald A. Redelmeier and Amos Tversky. On the framing of multiple prospects. *Psychological Science*, 3(3):191–193, 1992.

[66] Paul A. Samuelson. Risk and uncertainty: A fallacy of large numbers. *Scientia*, 98:108–113, 1963.

[67] A.F. Sanders. Choice among bets and revision of opinion. *Acta Psychologica*, 28:76–83, 1968.

[68] A.F. Sanders and W. Ter Linden. Decision making during paced arrival of probabilistic information. *Acta Psychologica*, 27:170–177, 1967.

[69] Shunichiro Sasaki and Toshiji Kawagoe. Belief updating in individual and social learning: A field experiment on the internet. May 2007. Working Paper.

[70] M. Shaked and Y L. Tong. Comparison of experiments for a class of positively dependent random variables. *Canadian Journal of Statistics*, 18:79–86, 1990.

[71] Suzanne Shu and George Wu. Belief bracketing: Can partitioning information change consumer judgments? October 2003. Working paper.

[72] Lones Smith and Peter Sorenson. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.