

DYNAMIC CHOICE AND COMMON LEARNING

RAHUL DEB[§] AND LUDOVIC RENO[†]

SEPTEMBER 5, 2021

ABSTRACT: A researcher observes a finite sequence of choices made by multiple agents in a binary-state environment. Agents maximize expected utilities that depend on their chosen alternative and the unknown underlying state. Agents learn about the time-varying state from the same information and their actions change because of the evolving common belief. The researcher does not observe agents' preferences, the prior, the common information and the stochastic process for the state. We characterize the set of choices that are rationalized by this model and generalize the information environment to allow for private information. We discuss the implications of our results for uncovering discrimination and committee decision making.

1. INTRODUCTION

Different individuals frequently make decisions based on the same information. For instance, in court cases, judges are presented with the same evidence and arguments prior to reaching their decisions. Employers and graduate schools peruse the same curricula vitae and reference letters prior to inviting applicants for interviews. As yet another instance, deliberative committees such as monetary policy committees gather the opinions and views of all members prior to making decisions. Naturally, even though individuals have the same information, they may still make distinct decisions because their preferences differ. But different preferences alone may not be able to rationalize *all* patterns of choices. What are the rationalizable patterns of choices? Can individual choices reveal changing preferences? If so, can this information be used, for instance, to test for discrimination?

[§]DEPARTMENT OF ECONOMICS, UNIVERSITY OF TORONTO, RAHUL.DEB@UTORONTO.CA

[†]SCHOOL OF ECONOMICS AND FINANCE, QUEEN MARY UNIVERSITY OF LONDON, LRENOU.ECON@GMAIL.COM

This paper was conceptualized at the Studienzentrums Gerzensee and we are grateful for their gracious hospitality. We thank Henrique de Oliveira, Laura Doval, Alex Jakobsen, Rohit Lamba, Paola Manzini, Ariel Pakes, Krishna Pendakur and several conference, seminar participants for comments, suggestions and criticisms that greatly improved the paper. Deb is always grateful to the SSHRC for their continued and generous financial support. Renou gratefully acknowledges the support of the Agence Nationale pour la Recherche under grant ANR CIGNE (ANR-15-CE38-0007-01) and through the ORA Project "Ambiguity in Dynamic Environments" (ANR-18-ORAR-0005).

To address these questions as well as others, we consider a simple binary-state environment in which multiple agents make choices repeatedly. A researcher observes the choices over time, but knows neither the preferences of the agents nor their information. The researcher assumes that the agents agree about the evolution of the state and learn about the state from common signals. We fully characterize the testable implications of this model; that is, we derive necessary and sufficient conditions for choice data to be rationalizable.

We begin with a baseline model in which there are two alternatives and all agents get a higher utility from choosing the alternative that matches the state. The characterization ([Theorem 1](#)) of this model is easy to state: choice data can be rationalized if, and only if, there are no two periods and two agents who pick different and opposite alternatives in each period. This shows that the model can be refuted with just two periods of choice data despite the researcher not observing the state or its evolution, the preferences, the prior belief and the common signals that agents use to learn. Moreover, the proof of this result demonstrates that we can extract preference rankings across agents. As we discuss later, this information can be useful for applications.

We then generalize the baseline model along several dimensions. We first demonstrate how to accommodate general preferences: agents may get a higher utility by choosing the alternative that mismatches the state ([Corollary 1](#)). Then, we allow for multiple actions ([Theorem 2](#)). Finally, we relax the assumption of common signals and instead assume that agents' information, while different, is related. Specifically agents observe private signals that differ in their informativeness but favor the same state. In other words, agents agree in each period whether the information is 'good' or 'bad' news but disagree about how conclusive the new evidence is. We characterize the testable implications of this general model when the state is time-invariant ([Theorem 3](#)) and show that all choice data can be rationalized if the state is time varying ([Theorem 4](#)).

An advantage of the simple reduced form nature of our model is that it can capture several seemingly disparate applications. We discuss these in detail in [Section 5](#) but we end the introduction by providing an example of how our framework can be applied to test for discrimination. Consider two employers who receive sequential job applications in the same order from the same six prospective employees. Applications m_1 , m_2 and m_3 , are from male applicants, while applications f_1 , f_2 and f_3 are from female applicants. Employers evaluate applications (the common information) to determine the productivity of each applicant which could either be low or high (the unknown state) and only want to interview high-productivity applicants. The researcher observes which applicants an employer invites for interviews. [Table 1](#) summarizes the researcher's choice data.

TABLE 1. Employer’s decisions: yes = “invite for an interview.”

	m_1	m_2	m_3	f_1	f_2	f_3
Employer 1	yes	yes	no	no	no	yes
Employer 2	yes	no	no	no	yes	yes

Note that the data described above takes the exact form necessary (sequential decision making and common information) to test our model. Data of this form can be easily generated by sending fictitious curricula vitae to job postings; since [Bertrand and Mullainathan \(2004\)](#), such “correspondence studies” are a common method to test for discrimination in economics. We emphasize that our framework is flexible enough to capture either the case where employers view each application to be unrelated to other applications (so the underlying state is independently drawn each time) or where there is learning (correlation of the states). While the former is often assumed in empirical analyses, there are several reasons why an employer might learn from the curricula vitae of previous applicants.¹ For instance, if m_1 and m_2 are successful mid-career applicants who graduated from a certain program, the employer might update that m_3 , who is an early career applicant from the same program, will be more likely to be successful. Alternatively, if f_1 and f_2 are accomplished candidates of a minority group, an employer might update their beliefs about the quality of minority candidates in general which will effect how they evaluate applicant f_3 . Indeed, the whole notion of statistical discrimination is based on learning and belief formation.

While both employers only want to interview high-productivity applicants, they may differ in their disutility of type I and type II errors (because the job details may not be exactly identical). We now explain how to use our characterization to test for taste-based discrimination. The key observation is that we can not only test whether the entire sample is rationalizable, but also whether the two sub-samples of male and female applicants are. It is immediate to check that the entire dataset is not rationalizable as there are two different candidates, m_2 and f_2 , with opposite and distinct interview decisions by both employers. Thus, we reject the hypothesis that interview decisions are made based on common beliefs and employer-specific preferences that are identical over male and female candidates.

Yet, each sub-sample is rationalizable. Thus, we cannot reject the hypothesis that the employers have common beliefs but different preferences over male and female candidates. In fact, we can even make further inferences about the preferences of the employers. Note that whenever Employer 2 invites a male applicant, so does Employer 1. Since with a binary state, an employer

¹If we assume that the employers view the applications as (statistically) independent, then we do not need to assume that the employers process the applications in the same order, since they have the same relevant information when evaluating applications.

invites an applicant only when the common belief is above a threshold, we can infer that Employer 2's threshold for male applicants is strictly above that of Employer 1. (The threshold is the belief at which an employer is indifferent between the two possible choices.) The opposite is true on the sub-sample of female applicants.

A simple application of our tests thus provides suggestive evidence of taste-based discrimination. In Section 5.1, we also discuss how our results can be used to uncover suggestive evidence of statistical discrimination. The fact that these manifest differently is a strength of our framework. Indeed, as [Bertrand and Duflo \(2017\)](#) observe in their recent survey of field experiments on discrimination, the vast majority of the literature only aims to uncover evidence of discrimination but not the source.

Finally, note that there is nothing special about the context of this example; we could replace employers by judges, job applicants by defendants, gender by race and the interview invitations by judicial decisions.

Related Literature

This paper is related to the recent decision theory literature that aims to characterize different models of learning and costly information acquisition. A few recent examples are [Caplin and Dean \(2015\)](#), [Caplin and Martin \(2015\)](#), [Lu \(2016\)](#) and [Frick et al. \(2019\)](#). Perhaps the main difference of our work with this literature is the data on which our respective analyses are conducted. Decision theory papers typically assume that the analyst can observe agents' choices in *every* possible choice scenario. In the case of random choice induced by learning/information acquisition, this amounts to assuming that the analyst can observe the choice *probabilities* of the various alternatives in all choice problems. Needless to say, this is a demanding data requirement.

By contrast, we only assume that finitely many *realized* choices are observed from a single choice problem and that the environment under which these choices are made is changing. While the data requirement for our analysis is modest relative to the above literature, the tradeoff is that it is not possible to identify the information (and identification is often an important goal in a decision theoretic analysis). Our results are driven by the assumption that we observe the choices of multiple individuals whose beliefs are systematically related. In practice, versions of this assumption are routinely made implicitly or explicitly.² Choice probabilities (as opposed to realized choices) are not observed and must be estimated. It is quite unusual to be able to estimate choice

²Outside of decision theory, this assumption is frequently made in the finance and information design literatures in which information often takes the form of public signals.

probabilities from a single individual as this would require many observations (for any meaningful precision) from the identical choice problem (the same set of alternatives in an identical information environment). So instead, choice probabilities can be obtained by averaging across the choices of different (typically observationally equivalent) individuals in the data. This implicitly assumes that (conditioning on covariates) these scenarios are treated as identical or, in other words, that different agents have the same preferences and information (but may get different signal realizations).

The most closely related papers are [Shmaya and Yariv \(2016\)](#) and [de Oliveira and Lamba \(2019\)](#) which also study the testable implications of learning on observed choices. [Shmaya and Yariv \(2016\)](#) study a single agent decision problem in which an agent observes signals and has to report which state she believes to be the most likely. They show that, absent additional structure on information, all choices can be explained by Bayesian updating. [de Oliveira and Lamba \(2019\)](#) consider a related environment but they allow for completely general dynamic preferences that need not be time separable. They assume that the researcher knows the utility function of the agent but does not observe how the agent learns about the underlying state. They show that not all choices are possible and derive a dominance condition that characterizes rationalizability. [Makris and Renou \(2021\)](#) generalize the concept of Bayes correlated equilibrium to multi-stage games. As an application of their results, they reformulate the problem of [de Oliveira and Lamba \(2019\)](#) as characterizing the Bayes' correlated equilibria of a dynamic decision problem, and derive a novel characterization.

The main difference between these papers and ours is that we do not assume that the researcher knows the agents' preferences (although we do assume preferences are known to be time separable). This is an important generalization because preferences are unobserved in all field data. Instead, our main assumption that generates testable predictions is the fact that different agents' information is related. We view this modeling assumption to be one of the main conceptual novelties of our framework.

2. THE BASELINE MODEL

This section presents and discusses the baseline model, which we analyze first. [Section 4](#) generalizes it in several directions.

A set $\mathcal{I} := \{1, \dots, I\}$ of I agents makes choices over a set $\mathcal{T} = \{1, \dots, T\}$ of T periods. In each period t , agent i chooses an *action* $x_{i,t}$ from the binary set $X = \{x, y\}$. A researcher observes *choice*

data $\mathcal{X} := (x_{i,t})_{(i,t) \in \mathcal{I} \times \mathcal{T}}$ that contains the choices of all agents over all time periods. This is all the researcher observes. We explain later how to deal with richer choice data, where the researcher observes additional characteristics of the choice problems.

We now define the model the researcher postulates to rationalize the choice data \mathcal{X} . There is an unknown, possibly time-varying, underlying *state* of the world $\omega_t \in \{\omega^x, \omega^y\}$ at each period t . We deliberately index the state by the alternative to indicate that action x (resp., y) is the best in state ω^x (resp., ω^y). The researcher thus assumes that agent i 's utility function u_i satisfies $u_i(x, \omega^x) \geq u_i(y, \omega^x)$ and $u_i(y, \omega^y) \geq u_i(x, \omega^y)$, with at least one strict inequality. For any utility function u_i , we denote

$$\bar{u}_i := \frac{u_i(y, \omega^y) - u_i(x, \omega^y)}{u_i(x, \omega^x) - u_i(y, \omega^x) + u_i(y, \omega^y) - u_i(x, \omega^y)}.$$

the threshold (belief) at which the agent is indifferent between choosing x and y .

We now turn to the beliefs and the information. The researcher assumes that the agents share a common *prior* p_0 and agree on the evolution of the state. The *transition probability* from state ω_{t-1} to ω_t is $\gamma_t(\omega_t | \omega_{t-1})$. We say the state is *time-invariant* if $\gamma_t(\omega^x | \omega^x) = \gamma_t(\omega^y | \omega^y) = 1$ for all $t \in \mathcal{T}$; otherwise, we say the state is *time-varying*.

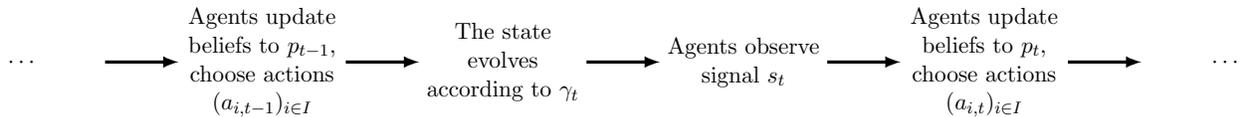
Prior to choosing an action (but after the state has evolved), the agents observe a *common signal* s_t , a realization of the statistical *experiment* $\pi_t : \{\omega^x, \omega^y\} \rightarrow \Delta(S_t)$, with S_t a finite set of signals.³

Finally, the researcher assumes that the agents are expected utility maximizers and update their beliefs by Bayes' rule. We use p_t to denote the *belief* that the state is ω^x at period t (the probability of ω^y is the complementary probability). The common belief at period t is thus:

$$p_t = \frac{[p_{t-1}\gamma_t(\omega^x|\omega^x) + (1-p_{t-1})\gamma_t(\omega^x|\omega^y)] \pi_t(s_t|\omega^x)}{\sum_{\omega \in \{\omega^x, \omega^y\}} [p_{t-1}\gamma_t(\omega|\omega^x) + (1-p_{t-1})\gamma_t(\omega|\omega^y)] \pi_t(s_t|\omega)}. \quad (1)$$

Hence, agent i chooses x at period t if and only if $p_t \geq \bar{u}_i$.

For ease of reference, the timing is summarized below



The aim of this paper is characterize choice data that can be rationalized by this model. In words, we are interested in when a researcher who only observes choice data \mathcal{X} can find some

³Note that this signal is completely general so it can contain the realized payoffs of the agents. In other words, the model permits learning from realized payoffs in addition to other sources of information.

parameters for the model such that, given those parameters, the agents would choose precisely the observed data. Formally, this is defined as follows.

DEFINITION 1. *The choice data \mathcal{X} is **rationalizable** if, for all $i \in \mathcal{I}$ and $t \in \mathcal{T}$, there exist utility functions u_i , transition probabilities γ_t , prior p_0 , beliefs p_t , sets of signals S_t , experiments $\pi_t : \{\omega^x, \omega^y\} \rightarrow \Delta(S_t)$ and realized signals $s_t \in S_t$ such that*

- i. Actions maximize utility: $p_t > \bar{u}_i$ if $x_{i,t} = x$ (resp., $p_t < \bar{u}_i$ if $x_{i,t} = y$),*
- ii. Bayesian Updating: p_t is derived from p_{t-1} by Bayes' rule.*

In what follows, we will sometimes explicitly add a qualifier to signify when the state is time invariant.

Before we proceed to discussing the various assumptions of the model, it is worth making a few comments about our definition of rationalizability. First, observe that this criterion requires choices to be strictly optimal (captured by the strict inequality in point i). This assumption is made to rule out indifference that would make the analysis trivial. If the strict inequality in point i were to be replaced by a weak inequality, then any choice data could be rationalized by a constant belief and common utilities $p_t = \bar{u}_1 = \dots = \bar{u}_I$ for all $t \in \mathcal{T}$. Second, implicit in point ii is the assumption that the realized signals s_t lie in the support of the experiment so that Bayes' rule can be used. In subsequent sections, we will discuss how non-Bayesian belief updating can be accommodated.

2.1. Discussion of the model

The payoffs in the baseline model capture situations where decision-makers are uncertain about which course of action is best and differ in their disutility for type I and II errors. Such preferences are common in economics and political science. Indeed, the model as a whole is fairly canonical in the literature on committee voting. Here, the state captures the correct decision (for instance, maintain the status quo or introduce a reform, convict or acquit a defendant in a court) and the actions are the votes. The repeated choices could either reflect votes on the same issue (deciding whether or not to implement a project that was postponed in a previous vote) or different observationally identical scenarios (such as judges deciding distinct cases with the same basic facts). We now discuss our modeling choices, starting with the assumption of common information.

COMMON SIGNALS: As we discussed in the introduction, there are several multiple-agent decision problems where it is natural to assume that agents learn via common signals. In other contexts such as committee voting, the literature considers both models like ours with common information

and settings where agents receive private signals prior to making choices. Common information is frequently assumed (such as in [Li, 2001](#); [Alonso and Câmara, 2016](#)) when the aim is to analyze the effect of information on committee outcomes. Surveys that describe the foundational work in this strand of the literature are [Gerling et al. \(2005\)](#) and [Li and Suen \(2009\)](#). Note that, even if the agents receive private signals, they often have the opportunity to communicate them to others or, in other words, to deliberate prior to making choices.

Moreover, there is a simple theoretical reason to focus attention on common signals. Without additional constraints, *all* choice data are rationalizable if the model were to allow for private signals (even with a common utility function). In a nutshell, private signals make it possible to decouple the choice problems agent-by-agent. Without additional constraints, we can then generate arbitrary sequences of individual beliefs and, thus, rationalize all choice data. We refer to [Section 4.3](#) for a model with non-trivial testable implications, where agents receive private signals.

REALIZATIONS VERSUS CHOICE PROBABILITIES: As we stressed in the introduction, a strength of our framework is that choice data consists of actual choices as opposed to choice probabilities. This implies that the data requirements to test our model are relatively modest. In particular, this means that we neither need to observe the agents choosing repeatedly in the identical information environment nor do we need to aggregate data across different agents.

The strengths and weaknesses of our framework can be clearly articulated by contrasting with a related decision theory paper. [Lu \(2016\)](#) considers an agent who chooses an alternative after receiving some private information (unobserved by the analyst) and provides a theory in which the information can be fully identified from the random choice rule. For identification, he requires the analyst to observe the choice probabilities from *every* menu of possible acts (mappings from states to lotteries over prizes). By contrast, we only require the analyst to observe the *realized* choices from a *single menu* of two acts; alternative x (and similarly y) can be viewed as an act that generates payoff $u(x, \omega)$ in each state ω . Of course, the downside of an analysis on such parsimonious data is that, unlike [Lu \(2016\)](#), we cannot identify the information.

SEPARATE DECISION PROBLEMS VS. EQUILIBRIUM: We presented our choice data as arising from the separate decisions of I agents. This is the appropriate interpretation for several of the examples from the introduction including that of the two employers. In the context of committee decision making, this corresponds to sincere voting where members vote for their preferred option (and ignore strategic considerations). However, in the case of two alternatives and public information (as in our baseline model), sincere voting is also an equilibrium of the one-shot voting game where alternative x is chosen only when more than a given exogenous threshold of voters choose x .

In the repeated interaction that is implicit in our choice data, sincere voting will also be an equilibrium of the repeated game since it is a stage game equilibrium (although, of course, there may be multiple equilibria). That said, our model also captures certain, more complex dynamic decision making as a special case. We discuss this next.

DYNAMIC OPTIMIZATION: The agents we defined above myopically maximized their per period utility. This assumption is commonly made in the revealed preference literature that builds on [Afriat \(1967\)](#) where the aim is to derive the testable implications of various models of consumption on price quantity data (see [Chambers and Echenique, 2016](#) for a description of the literature). In our setting, behaving myopically is also dynamically optimal if the agents' actions do not affect the evolution of the state and the information they receive. This is true for certain applications like judicial decision making (which we discuss in more detail in Section 5.2) where the decision to convict or acquit does not shed any additional light on the guilt of defendant.

Importantly, certain classes of dynamic models in which actions affect the arrival of information also lead to behavior that is observationally indistinguishable. Loosely speaking, the solution to a Bellman equation can result in cutoff behavior in which agent i 's action choice depends on whether the public belief lies above or below a cutoff \bar{u}_i . In other words, even though agents are actually dynamically optimizing, their behavior is covered by our rationalization criterion ([Definition 1](#)). Recent examples of such dynamic political economy models that fit our framework are [Gul and Pesendorfer \(2012\)](#) and [Chan et al. \(2018\)](#).

RICHER CHOICE DATA: The baseline model assumes that the researcher only observes the agent's choices. We can also consider richer choice data $\mathcal{X} = (x_{i,t}, \theta_t)_{(i,t) \in \mathcal{I} \times \mathcal{T}}$, where the researcher also observes a vector $\theta_t \in \Theta_t$ of characteristics about the choice problem at period t . For instance, in our introductory example, θ_t is the gender of the t -th application. The information agents have is therefore (s_t, θ_t) , with θ_t also observed by the analyst. Our definition of rationalization generalizes to experiments $\pi_t : \{\omega^x, \omega^y\} \rightarrow \Delta(S_t \times \Theta_t)$ and transitions $\gamma_t : \{\omega^x, \omega^y\} \times \Theta_{t-1} \rightarrow \Delta(\{\omega^x, \omega^y\})$, without affecting our main results. Indeed, if the richer choice data $\mathcal{X} = (x_{i,t}, \theta_t)_{(i,t) \in \mathcal{I} \times \mathcal{T}}$ is rationalizable, the simpler choice data $(x_{i,t})_{(i,t) \in \mathcal{I} \times \mathcal{T}}$ is rationalizable with $\hat{\pi}_t(s_t | \omega_t) := \pi_t(s_t, \theta_t | \omega_t) / \pi_t(\theta_t | \omega_t)$ and $\hat{\gamma}_t(\omega_t | \omega_{t-1}) := \gamma_t(\omega_t | \omega_{t-1}, \theta_{t-1})$. Conversely, if the simpler choice data $(x_{i,t})_{(i,t) \in \mathcal{I} \times \mathcal{T}}$ is rationalizable, then any richer choice data is rationalizable: we can simply assume that the agents view the observable characteristics as irrelevant (independent). In Section 5.1, we demonstrate how richer choice data can be used to detect evidence of discrimination.

To summarize, we view our model to be a simple reduced form representation of dynamic behavior that is flexible enough to encompass several distinct dynamic settings as special cases.

3. THE TESTABLE IMPLICATIONS OF THE BASELINE MODEL

We first present a simple example, which shows that not all choice data are rationalized.

EXAMPLE 1. Consider the choice data \mathcal{X} of two individuals i, j and two time periods t_1, t_2 represented in the following table:

	i	j
t_1	x	y
t_2	y	x

In the table, choices in the rows correspond to a given time period and those in the columns correspond to the agents.

We argue that this choice data cannot be rationalized. So suppose, for the purpose of contradiction, that it can. This implies that we can find utilities and beliefs that satisfy

$$p_{t_1} > \bar{u}_i > p_{t_2},$$

$$p_{t_1} < \bar{u}_j < p_{t_2},$$

which is clearly a contradiction. In words, agent i 's actions imply that the belief must be strictly lower in observation t_2 compared to t_1 whereas agent j 's actions imply the exact opposite.

It is worth discussing a few features of the simple argument we employed in the above example. Note that it did not even mention the state transition probabilities. Moreover, this argument did not require the agents to be Bayesian. As long as they share a common belief, they could be learning in *any* way and the data in [Example 1](#) would still not be rationalizable. Finally, the time periods t_1 and t_2 need not be consecutive or ordered chronologically.

The following definition gives a name to the pattern of choices in [Example 1](#).

DEFINITION 2. Choice data \mathcal{X} has a **cycle** if there are distinct time periods $(t_1, t_2) \in \mathcal{T} \times \mathcal{T}$ and agents $(i, j) \in \mathcal{I} \times \mathcal{I}$ such that $x_{i,t_1} = x_{j,t_2} = x$ and $x_{i,t_2} = x_{j,t_1} = y$.

[Example 1](#) shows that a necessary condition for rationalization is the absence of cycles. The following result shows that this condition is also sufficient. In other words, the choice data described in [Example 1](#) is the *only* pattern of choices that cannot be rationalized.

THEOREM 1. Choice data \mathcal{X} is rationalizable if, and only if, it has no cycles.

We present a simple proof of [Theorem 1](#) as it will be useful for the discussion that follows.

PROOF. (\Rightarrow). Already shown.

(\Leftarrow). We define the following binary relation \succsim on the set of agents \mathcal{I} . For any $(i, j) \in \mathcal{I} \times \mathcal{I}$,

$$j \succsim i \quad \text{if, for all } t \in \mathcal{T}, \text{ we have} \quad ([x_{j,t} = x] \implies [x_{i,t} = x]).$$

In words, $j \succsim i$ if whenever agent j chooses x , so does agent i .

We argue that, since \mathcal{X} has no cycles, \succsim is complete. By contradiction, assume that there exists a pair (i, j) such that $j \not\sucsim i$ and $i \not\sucsim j$. Since $j \not\sucsim i$, there exists $t \in \mathcal{T}$ such that $x_{j,t} = x$ and $x_{i,t} = y$. Since $i \not\sucsim j$, there exists $t' \in \mathcal{T}$ such that $x_{i,t'} = x$ and $x_{j,t'} = y$. We thus have a cycle, the required contradiction.

We now argue that \succsim is transitive. This is because for any $(i, j, k) \in \mathcal{I}$, if $([x_{j,t} = x] \implies [x_{i,t} = x])$ for all $t \in \mathcal{T}$ (and so $j \succsim i$) and $([x_{k,t} = x] \implies [x_{j,t} = x])$ for all $t \in \mathcal{T}$ (and so $k \succsim j$), then we must have $([x_{k,t} = x] \implies [x_{i,t} = x])$ for all $t \in \mathcal{T}$ (and so $k \succsim i$).

Since \succsim is complete and transitive, we can find numbers $\bar{u}_i \in (0, 1)$ for all $i \in \mathcal{I}$ (by invoking a textbook utility representation theorem for finite alternatives) such that

$$\bar{u}_j \geq \bar{u}_i \quad \text{if and only if} \quad j \succsim i.$$

We now argue that if $x_{i,t} = x$ and $x_{j,t} = y$, then $\bar{u}_j > \bar{u}_i$. By contradiction, if $\bar{u}_i \geq \bar{u}_j$, then $i \succsim j$ and, therefore, $x_{i,t} = x$ implies that $x_{j,t} = x$, a contradiction. It follows that

$$\min \{ \bar{u}_j \mid x_{j,t} = y \} > \max \{ \bar{u}_i \mid x_{i,t} = x \}$$

for all $t \in \mathcal{T}$.

Finally, let $p_0 = \frac{1}{2}$ and, for each $t \in \mathcal{T}$, choose any $p_t \in (0, 1)$ such that

$$\min \{ \bar{u}_i \mid x_{i,t} = y \} > p_t > \max \{ \bar{u}_j \mid x_{j,t} = x \}.$$

By construction, the data is rationalized with these beliefs and thresholds.

We can easily construct utility functions with the required thresholds, for instance, by setting $u_i(x, \omega^x) = u_i(y, \omega^y) = 1$, $u_i(x, \omega^y) = 1/2$ and $u_i(y, \omega^x) = (3\bar{u}_i - 1)/2\bar{u}_i$.

It remains to construct the transitions probabilities, signals and experiments that generate the given beliefs. So first, we pick transition probabilities $\gamma_t(\omega^x | \omega^x) = \gamma_t(\omega^y | \omega^y) = 1$ for all $t \in \mathcal{T}$ that correspond to the state being time invariant. Then, for all t , pick an experiment with binary

signals $S_t = \{s, s'\}$ and choose any $\pi_t(s|\omega^x) \in (0, 1)$ and $\pi_t(s|\omega^y) \in (0, 1)$ such that

$$\frac{\pi_t(s|\omega^x)}{\pi_t(s|\omega^y)} = \frac{p_t}{1-p_t} \frac{1-p_{t-1}}{p_{t-1}}$$

and assign $\pi_t(s'|\omega^x) = 1 - \pi_t(s|\omega^x)$ and $\pi_t(s'|\omega^y) = 1 - \pi_t(s|\omega^y)$.

It is routine to verify that p_t is obtained from p_{t-1} by Bayesian updating when signal s is realized. This completes the proof. ■

We end this section by discussing some implications of [Theorem 1](#). First, note that the proof shows that, for rationalizable choice data, the actions reveal information about the preference cutoffs via the relation \succsim . Thus, we can determine when an agent i has a higher preference cutoff than j (if $i \succsim j$ and $j \not\sucsim i$). This preference information is potentially useful; for instance in [Section 5.1](#), we argue that it can be used to uncover evidence of discrimination.

Second, the proof also provides an immediate characterization for the case where the preference cutoffs \bar{u}_i are known to the researcher, but the signals and the statistical experiments are not. As we mentioned in the introduction, this is the environment considered by [Shmaya and Yariv \(2016\)](#) and [de Oliveira and Lamba \(2019\)](#). Simply, a necessary and sufficient condition for choice data \mathcal{X} to be rationalizable with known preference cutoffs $(\bar{u}_i)_{i \in \mathcal{I}}$ is that, for all i, j such that $\bar{u}_i \geq \bar{u}_j$, we have $([x_{i,t} = x] \implies [x_{j,t} = x])$ for all t . Thus, with known preferences, it is possible to refute the model with choice data consisting of a single period.

Finally, observe that the proof shows that every rationalizable data set is also rationalizable with a time-invariant state. In other words, while the model is refutable, we cannot distinguish between environments where the state is time varying or invariant.⁴ Similarly, the proof shows that any beliefs that rationalize the choice data can be obtained by Bayesian updating. Therefore, Bayesian updating itself is not testable. So, as long as agents have common beliefs, they could be using *any* method to form their beliefs.⁵ All we need is for the agents to revise their beliefs as they receive informative signals.

While these observations follow almost immediately from our definition of rationalizability, we nonetheless view them to be important conceptual insights of our framework. They demonstrate the limits of the behavioral analysis of learning that can be conducted from finite choice data which

⁴In fact, if a data set is rationalizable, it is also rationalizable with a time-independent process and uninformative signals. To see this, simply observe that if we let $\gamma_t(\omega^x|\omega^x) = \gamma_t(\omega^x|\omega^y) = p_t$ for all t and assume all S_t to be singletons, the posterior at period t is indeed p_t .

⁵So, it is also possible that these changing beliefs are generated by state transitions alone and that there is no learning whatsoever.

is, more often than not, the available field data. In this sense, these insights are akin to Afriat's theorem (Afriat, 1967) that shows any price quantity data that is consistent with utility maximization with a locally nonsatiated utility function is also consistent with utility maximization with a concave utility (so the convexity of preferences is not testable).

4. GENERALIZING THE BASELINE MODEL

This section generalizes the baseline model along several different directions and derives the analogue of Theorem 1 for each of these generalizations. Each subsection generalizes the baseline model along one particular dimension, but it will be clear that we can also simultaneously allow for all, or a subset of, these distinct generalizations.

4.1. General preferences

When we defined rationalization in Section 2, we assumed that utilities were such that agents prefer to match their action to the underlying state. Needless to say, this assumption might not always be appropriate. Note that allowing for arbitrary utilities does not change the cutoff structure of optimal behavior: an agent chooses x if his belief is above or below a cutoff. Hence, with general preferences, we can define rationalization as follows.

DEFINITION 3. *The choice data \mathcal{X} is **rationalizable with general preferences** if, for all $i \in \mathcal{I}$ and $t \in \mathcal{T}$, there exist utility functions u_i , constants $\eta_i \in \{-1, 1\}$, transition probabilities γ_t , prior belief p_0 , beliefs p_t , sets of signals S_t , experiments $\pi_t : \{\omega^x, \omega^y\} \rightarrow \Delta(S_t)$ and realized signals $s_t \in S_t$ such that*

- i. *Actions maximize utility: $\eta_i p_t > \bar{u}_i$ when $x_{i,t} = x$ (resp., $\eta_i p_t < \bar{u}_i$ when $x_{i,t} = y$).*
- ii. *Bayesian Updating: p_t is derived from p_{t-1} by Bayes' rule.*

The difference is captured in point i. The optimal behavior of an expected utility maximizer with general preferences is characterized by two numbers \bar{u}_i and η_i ; the former determines the cutoff belief and the latter whether agents prefer to match or mismatch the state.

The characterization of choice data that is rationalizable with general preferences follows immediately from Theorem 1. To formally state the result, we need one additional piece of notation. Let $\kappa_i : \{x, y\} \rightarrow \{x, y\}$ be a permutation (bijection) on the actions of agent i , and $\kappa = (\kappa_i)_{i \in \mathcal{I}}$ be a profile of permutations, which we refer to simply as a *permutation*. For any permutation κ , let $\mathcal{X}^\kappa := (x_{i,t}^\kappa)_{(i,t) \in \mathcal{I} \times \mathcal{T}}$ be the *permuted choice data* obtained from \mathcal{X} by permuting the choices; that is, $x_{i,t}^\kappa = \kappa_i(x_{i,t})$ for all (i, t) .

COROLLARY 1. *Choice data \mathcal{X} is rationalizable with general preferences if, and only if, we can find a permutation κ such that the permuted choice data \mathcal{X}^κ has no cycles.*

This result is very intuitive. If the data is rationalizable with general preferences, it must mean that we can switch the choices of every agent for whom $\eta_i = -1$ and the permuted choice data must be rationalizable in the sense of [Theorem 1](#). The next example generalizes [Example 1](#) to show that not every choice data set is rationalizable with general preferences.

EXAMPLE 2. Consider the choice data \mathcal{X} of two individuals i, j and four time periods represented in the following table:

	i	j
t_1	x	y
t_2	y	x
t_3	x	x
t_4	y	y

Observe that the choices in the shaded cells are identical to those of [Example 1](#) and constitute a cycle. [Corollary 1](#) states that the data cannot be rationalized with general preferences when there is a cycle for every permutation of the choice data. Clearly, the cycle in periods t_1 and t_2 can be removed by permuting the choices of either agent. However, this would then introduce a cycle into periods t_3 and t_4 . Below, we display the choice data from the three possible profiles (κ_i, κ_j) of permutations other than the identity permutation (which does not change the choice data): from left to right, the permutations are on the choices of agent i only, agent j only and both. Shaded cells highlight the cycles and this demonstrates that the choice data in this example is not rationalizable with general preferences.

	i	j
t_1	y	y
t_2	x	x
t_3	y	x
t_4	x	y

	i	j
t_1	x	x
t_2	y	y
t_3	x	y
t_4	y	x

	i	j
t_1	y	x
t_2	x	y
t_3	y	y
t_4	x	x

4.2. More than two alternatives

There are many situations where there may be more than two alternatives despite the state being binary. For instance, a member of a committee may choose to abstain as opposed to voting for/against a proposal.

We begin by discussing the case where the alternatives are ordered. So suppose there are N alternatives $\{y_1, \dots, y_N\}$ and each agent i picks $x_{i,t} \in \{y_1, \dots, y_N\}$ in period t . Every agent i 's

preferences are captured by cutoffs $0 =: \bar{u}_{i,1} < \dots < \bar{u}_{i,N+1} := 1$ and the agent chooses action $x_{i,t} = y_n$ whenever $\bar{u}_{i,n} < p_t < \bar{u}_{i,n+1}$.

Because agents are expected utility maximizers, for any utility function u_i , their optimal choices will always take the above cutoff form for *some* ordering of the alternatives. Therefore, we are implicitly making two assumptions above. The first is that all agents order alternatives in the same way (although their belief cutoffs might differ). The second is that alternative y_1 is picked at the lowest beliefs and vice versa for y_N . As with [Theorem 1](#), these assumptions can be dispensed with and we describe how to do so immediately after the statement of [Theorem 2](#) below.

Rationalization with multiple alternatives is analogously defined to [Definition 1](#).

DEFINITION 4. *The choice data \mathcal{X} is **rationalizable with multiple alternatives** if, for all $i \in \mathcal{I}$ and $t \in \mathcal{T}$, there exist utility functions u_i , transition probabilities γ_t , prior belief p_0 , beliefs p_t , sets of signals S_t , experiments $\pi_t : \{\omega^x, \omega^y\} \rightarrow \Delta(S_t)$ and realized signals $s_t \in S_t$ such that*

- i. Actions maximize utility: $\bar{u}_{i,n} < p_t < \bar{u}_{i,n+1}$, when $x_{i,t} = y_n$,*
- ii. Bayesian Updating: p_t is derived from p_{t-1} by Bayes' rule.*

Similarly, we can define what it means to have a cycle when there are multiple alternatives. We use the same terminology as this is the obvious generalization of [Definition 2](#).

DEFINITION 5. *Choice data \mathcal{X} with multiple alternatives has a **cycle** if there are distinct time periods $(t_1, t_2) \in \mathcal{T} \times \mathcal{T}$ and agents $(i, j) \in \mathcal{I} \times \mathcal{I}$ such that $x_{i,t_1} = y_{n_{i,t_1}}$, $x_{j,t_1} = y_{n_{j,t_1}}$ and $x_{i,t_2} = y_{n_{i,t_2}}$, $x_{j,t_2} = y_{n_{j,t_2}}$ where $n_{i,t_1} > n_{i,t_2}$ and $n_{j,t_1} < n_{j,t_2}$.*

Note that the above definition implicitly leverages the order of alternatives that we fixed when we defined preferences.

The absence of cycles is a necessary condition for rationalization even with multiple alternatives. To see this, observe that when agent i switches from action $y_{n_{i,t_1}}$ to $y_{n_{i,t_2}}$ where $n_{i,t_1} > n_{i,t_2}$, this implies that we must have $p_{t_1} > \bar{u}_{i,n_{i,t_1}} > p_{t_2}$. However, the choices of agent j imply exactly the opposite inequality for the beliefs $p_{t_1} < \bar{u}_{j,n_{j,t_2}} < p_{t_2}$, which generates the required contradiction. Indeed, like [Theorem 1](#), this condition is also sufficient.

THEOREM 2. *Choice data \mathcal{X} is rationalizable with multiple alternatives if, and only if, it has no cycles.*

The proof of [Theorem 2](#) can be found in [Appendix A](#).

We can relax the particular ordering of the actions that we chose. As we discussed above, any utility function u_i implies cutoff behavior for some order of actions and so, as in [Corollary 1](#), we can define agent specific permutations $\kappa_i : \{y_1, \dots, y_n\} \rightarrow \{y_1, \dots, y_n\}$ and consider the permuted choice data \mathcal{X}^κ . Then, the choice data is rationalizable with multiple alternatives and general preferences if, and only if, we can find a permutation κ such that the permuted choice data \mathcal{X}^κ has no cycles.

4.3. Relaxing common beliefs

Perhaps the strongest assumption in the baseline model is that agents have common beliefs. This section relaxes this assumption. As we mentioned earlier, every choice data set is rationalizable if we allow for arbitrary private signals and, thus, for arbitrary private beliefs. There needs to be some relation between the information of the agents in order for the model to have non-trivial testable implications.

The restriction we impose is that agents receive “co-monotone” private signals. Informally, an experiment is co-monotone if the joint distribution of the signals is such that agents agree on the state that their individual signal is more likely to arise from, but disagree on the strength of the signal. As an example, this assumption captures a deliberative committee where agents all agree that the information is good news for one of the two states, but disagree in terms of how good the news actually is.

Formally, the experiment $\pi_t : \{\omega^x, \omega^y\} \rightarrow \Delta(\times_{i \in \mathcal{I}} S_{i,t})$ is *co-monotone* if for all profiles of signals $(s_{1,t}, \dots, s_{I,t})$ in the support of π_t ,

$$[\pi_{i,t}(s_{i,t}|\omega^x) - \pi_{i,t}(s_{i,t}|\omega^y)][\pi_{j,t}(s_{j,t}|\omega^x) - \pi_{j,t}(s_{j,t}|\omega^y)] \geq 0$$

for all (i, j) , where $\pi_{i,t}(\cdot|\omega)$ is the marginal of $\pi_t(\cdot|\omega)$ over $S_{i,t}$.⁶ Experiments are *strictly* co-monotone if the above inequalities are all strict.

Before defining the rationalization criterion, we present the main implication of assuming co-monotone experiments. To do so, we need some additional notation. We let $p_{i,t}$ be agent i 's belief about the event $[\omega_t = \omega^x]$ at period t after having received the private signal $s_{i,t}$. Similarly, we let

$$q_{i,t+1} := p_{i,t}\gamma_{t+1}(\omega^x|\omega^x) + (1 - p_{i,t})\gamma_{t+1}(\omega^x|\omega^y)$$

be agent i 's belief about the event $[\omega_{t+1} = \omega^x]$ at period $t + 1$ prior to receiving the private signal $s_{i,t+1}$. An immediate implication of assuming co-monotone experiments is that, for any two

⁶The support of the kernel π_t is the union of the supports of each probability $\pi_t(\cdot|\omega)$.

Bayesian agents i and j ,

$$[p_{i,t} > q_{i,t} \implies p_{j,t} \geq q_{j,t}] \text{ and } [p_{i,t} < q_{i,t} \implies p_{j,t} \leq q_{j,t}]. \quad (2)$$

In words, the private beliefs are co-monotone. Conversely, for all posteriors $p_{i,t}$ and $p_{j,t}$ that satisfy (2), we can find a co-monotone experiment and signals $(s_{i,t}, s_{j,t})$, which generate these posteriors. (See the proof of [Theorem 3](#) for details.)

We now define the rationalization criterion. Throughout, we write S_t for the set of signals $\times_{i \in \mathcal{I}} S_{i,t}$.

DEFINITION 6. *The choice data \mathcal{X} is **rationalizable with (strictly) co-monotone experiments** if there exist utility functions u_i , transition probabilities γ_t , prior p_0 , beliefs $p_{i,t}$, sets of signals $S_{i,t}$, (strictly) co-monotone experiments $\pi_t : \{\omega^x, \omega^y\} \rightarrow \Delta(S_t)$ and realized signals $s_t \in S_t$ such that*

- i. *Actions maximize utility: $p_{i,t} > \bar{u}_i$ if $x_{i,t} = x$ (resp., $p_{i,t} < \bar{u}_i$ if $x_{i,t} = y$).*
- ii. *Bayesian Updating: $p_{i,t}$ is derived from $p_{i,t-1}$ by Bayes' rule conditioning on $s_{i,t}$.*

We stress that all the information an agent receives is encoded in the private signal. Since private signals are not correlated with past decisions, states and signals, this implies that agents neither observe the past decisions of others nor their private signals. This assumption is appropriate in some applications, less so in others. We can allow for general experiments, which condition on the current state *and* past signals, actions, and states. As long as we restrict attention to co-monotone experiments, the same testable implications would obtain.

We first argue that the model has testable implications when the state is *time-invariant*. To do so, we revisit [Example 1](#) making one minor alteration.

EXAMPLE 3. *Consider the choice data \mathcal{X} of two individuals i and j and two consecutive time periods t and $t + 1$ represented in the following table:*

	i	j
t	x	y
$t + 1$	y	x

Recall that [Example 1](#) had exactly this pattern of choices, but with arbitrary time periods.

We argue that this choice data cannot be rationalized with co-monotone experiments and a time-invariant state. Suppose to the contrary that we could. This implies that we can find thresholds and beliefs for agent i that satisfy

$$p_{i,t} > \bar{u}_i > p_{i,t+1}.$$

Since agent i 's belief decreases, it must be the case that the information she receives satisfies $\pi_{i,t+1}(s_{i,t+1}|\omega^x) < \pi_{i,t+1}(s_{i,t+1}|\omega^y)$. Since experiments are co-monotone, agent j 's information must satisfy $\pi_{j,t+1}(s_{j,t+1}|\omega^x) \leq \pi_{j,t+1}(s_{j,t+1}|\omega^y)$ for all signals $s_{j,t}$ that agent j can receive with positive probability. This implies that agent j 's belief cannot increase and so we cannot have

$$p_{j,t} < \bar{u}_j < p_{j,t+1},$$

the required contradiction.

The following definition provides a name to such consecutive time period cycles.

DEFINITION 7. Choice data \mathcal{X} has a **consecutive cycle** if there are agents $(i, j) \in \mathcal{I} \times \mathcal{I}$ and a time period $t \in \mathcal{T}$ such that $x_{i,t} = x$, $x_{j,t} = y$ and $x_{i,t+1} = y$, $x_{j,t+1} = x$.

Example 3 shows that a necessary condition for rationalization is the absence of consecutive cycles in the choice data. The next result shows that this condition is also sufficient. The theorem actually shows a stronger result. First, any data that is rationalizable with co-monotone experiments is also rationalizable with strictly co-monotone experiments. Second, if the choice data is rationalizable, then it is rationalizable with an arbitrary profile of utility functions. Put differently, this implies that, if we relax the assumption of common information, then choices no longer provide any ranking of the agents' cutoff beliefs.

THEOREM 3. Fix any arbitrary profile of utility functions $(u_i)_{i \in \mathcal{I}}$. Given the profile $(u_i)_{i \in \mathcal{I}}$, choice data \mathcal{X} is rationalizable with (strictly) co-monotone experiments and a time-invariant state if, and only if, it has no consecutive cycles.

PROOF. Since we have already argued the necessity, we only need to argue the sufficiency. Fix any arbitrary profile of utility functions $(u_i)_{i \in \mathcal{I}}$. The proof is by induction on T .

If $T = 1$, there are no consecutive cycles. We therefore need to argue that we can rationalize all choice data. Given the thresholds $(\bar{u}_i)_{i \in \mathcal{I}}$, choose p_0 and $p_{i,1}$ as follows. For all i such that $x_{i,1} = x$, choose $p_{i,1} > \bar{u}_i$. Similarly, for all i such that $x_{i,1} = y$, choose $\bar{u}_i > p_{i,1}$. Choose $p_0 < \min_{i \in \mathcal{I}} p_{i,1}$. For all i , choose strictly positive numbers $(\pi_{i,1}^x, \pi_{i,1}^y)$ such that $\sum_i \pi_{i,1}^x < 1$, $\sum_i \pi_{i,1}^y < 1$, and

$$\frac{\pi_{i,1}^x}{\pi_{i,1}^y} = \frac{p_{i,1}}{1 - p_{i,1}} \frac{1 - p_0}{p_0} > 1.$$

Such numbers exist. Moreover, since $\pi_{i,1}^x / \pi_{i,1}^y > 1$ for all i , we have that

$$[\pi_{i,1}^x - \pi_{i,1}^y][\pi_{j,1}^x - \pi_{j,1}^y] > 0,$$

for all (i, j) . From [Lemma 1](#) in [Appendix B](#), there exist a strictly co-monotone experiment and a profile of signals $(s_{i,1}, \dots, s_{I,1})$ such that $(\pi_{i,1}^x, \pi_{i,1}^y) = (\pi_{i,1}(s_{i,1}|\omega^x), \pi_{i,1}(s_{i,1}|\omega^y))$. Hence, upon receiving the signal $s_{i,1}$, agent i 's belief is $p_{i,1}$.

As the induction hypothesis, suppose we can rationalize any T period choice data that has no consecutive cycles. For the induction step, we will argue that we can also do this for any $T + 1$ period choice data that has no consecutive cycles.

So consider choice data $\mathcal{X} := (x_{i,t})_{(i,t) \in \mathcal{I} \times \mathcal{T}}$ with no consecutive cycles and $|\mathcal{T}| = T + 1$. By the induction hypothesis, we can rationalize the first T observations $(x_{i,t})_{(i,t) \in \mathcal{I} \times \{1, \dots, T\}}$. Let the period T beliefs be $p_{i,T}$.

If the choices of all agents are the same in periods T and $T + 1$, that is, $x_{i,T} = x_{i,T+1}$ for all $i \in \mathcal{I}$, then the choice data is rationalized by all agents receiving uninformative signals.

It remains to consider the case where at least one agent changes their choice. So, assume that there is an agent \tilde{i} who switches from $x_{\tilde{i},T} = x$ to $x_{\tilde{i},T+1} = y$. (The case where an agent \tilde{i} switches from $x_{\tilde{i},T} = y$ to $x_{\tilde{i},T+1} = x$ is treated analogously and, therefore, omitted.) Since there are no consecutive cycles, there is no agent j switching from $x_{j,T} = y$ to $x_{j,T+1} = x$. So for all $i \in \mathcal{I}$, choose any $p_{i,T+1}$ that satisfy

$$\begin{aligned} p_{i,T} > p_{i,T+1} > \bar{u}_i & \quad \text{if } x_{i,T} = x_{i,T+1} = x, \\ \bar{u}_i > p_{i,T} > p_{i,T+1} & \quad \text{if } x_{i,T} = x_{i,T+1} = y, \\ p_{i,T} > \bar{u}_i > p_{i,T+1} & \quad \text{if } x_{i,T} = x, \quad x_{i,T+1} = y. \end{aligned}$$

Such period- $(T + 1)$ beliefs can always be chosen because, by the induction hypothesis the period- T beliefs satisfy $p_{i,T} > \bar{u}_i$ when $x_{i,T} = x$ (respectively, $p_{i,T} < \bar{u}_i$ when $x_{i,T} = y$).

Observe that the constructed period $T + 1$ beliefs are strictly lower than the period T beliefs and hence, we can construct co-monotone experiments that generate these beliefs. For all i , choose strictly positive numbers $(\pi_{i,T+1}^x, \pi_{i,T+1}^y)$ such that $\sum_i \pi_{i,T+1}^x < 1$, $\sum_i \pi_{i,T+1}^y < 1$, and

$$\frac{\pi_{i,T+1}^x}{\pi_{i,T+1}^y} = \frac{p_{i,T+1}}{1 - p_{i,T+1}} \frac{1 - p_{i,T}}{p_{i,T}} < 1.$$

To complete the proof, we invoke [Lemma 1](#) in [Appendix B](#). ■

Unlike [Theorem 1](#), [Theorem 3](#) requires the state to be time-invariant. Why should the state evolution matter in this case? In the case of common beliefs, by definition, beliefs must move in the same direction for all agents across observations. This is also true for co-monotone experiments as

long as the state is time-invariant. When the state is time-varying however, this need not be true. To see this, suppose that the evolution of the state causes the belief $q_{i,T+1}$ to be lower than $p_{i,T}$ for all agents. Now, consider a co-monotone experiment where the belief of each agent increases, that is, $p_{i,T+1} > q_{i,T+1}$ for all i . The different strengths of the signals may, however, cause the belief $p_{i,T+1}$ of some agents to increase above their belief $p_{i,T}$, while for others it remains below. [Theorem 4](#) states that this is in fact enough to rationalize all choice data.

THEOREM 4. *Every choice data \mathcal{X} is rationalizable with strictly co-monotone experiments (and a time-varying state).*

PROOF. The proof is by induction on T . Throughout, we fix arbitrary thresholds $(\bar{u}_i)_{i \in \mathcal{I}}$. If $T = 1$, all choice data are clearly rationalizable (as shown in the proof of [Theorem 3](#)).

As the induction hypothesis, suppose that we can rationalize all choice data with T periods and, for the induction step, consider any choice data $(x_{i,t})_{(i,t) \in \mathcal{I} \times \mathcal{T}}$ with $|\mathcal{T}| = T + 1$ periods. Since the choice data $(x_{i,t})_{(i,t) \in \mathcal{I} \times \{1, \dots, T\}}$ is rationalizable, there are period- T beliefs $p_{i,T}$ such that $p_{i,T} > \bar{u}_i$ when $x_{i,T} = x$, and $p_{i,T} < \bar{u}_i$ when $x_{i,T} = y$.

Choose any $0 < \tilde{q} < \min_{i \in \mathcal{I}} \bar{u}_i$ and consider the following state transitions

$$\gamma_{T+1}(\omega^x | \omega^x) = 1 - \gamma_{T+1}(\omega^y | \omega^x) = \gamma_{T+1}(\omega^x | \omega^y) = 1 - \gamma_{T+1}(\omega^y | \omega^y) = \tilde{q}.$$

By construction of the transition probabilities, $q_{i,T+1} = \tilde{q}$ for all i , regardless of the beliefs $p_{i,T}$. The evolution of the state at period $T + 1$ is independent of its realization period T .

For all $i \in \mathcal{I}$, choose any $p_{i,T+1} > \tilde{q}$ that satisfies

$$\begin{aligned} p_{i,T+1} &> \bar{u}_i && \text{if } x_{i,T+1} = x, \\ p_{i,T+1} &< \bar{u}_i && \text{if } x_{i,T+1} = y. \end{aligned}$$

These numbers exist because $\tilde{q} < \min_{i \in \mathcal{I}} \bar{u}_i$.

Finally, as in the proof of [Theorem 3](#), we can construct a strictly co-monotone experiment $\pi_{T+1} : \{\omega^x, \omega^y\} \rightarrow \Delta(\times_i \mathcal{S}_{i,t})$, which generates these beliefs. This completes the proof. \blacksquare

The observant reader may have wondered why we assume that agents have the same prior belief and agree about state transitions. It should be clear from the proof of [Theorem 3](#) that nothing changes if we allow the prior belief to vary by agent (since the induction step does not involve the prior belief). [Theorem 4](#) shows that, if we allow for time varying states, anything goes. Thus,

any further relaxation that allows for heterogeneous prior beliefs or private state transitions (even those that are systematically related) would also result in there being no testable implications.

Finally, we can generalize [Theorem 3](#) to accommodate multiple alternatives and general preferences in exactly the same way we did for the baseline model. With multiple alternatives, a consecutive cycle would then just be a cycle in terms of [Definition 5](#) over consecutive time periods. General preferences would once again be the equivalent of permuting the choice data. Obviously [Theorem 4](#) also generalizes and all choice data is rationalizable with strictly co-monotone experiments, a time-varying state and, either or both of, multiple alternatives and general preferences.

We conclude this section with two final remarks. First, an alternative to assuming that agents receive private signals is to assume that all agents observe the same signal, but have different (co-monotone) beliefs about its likelihood. That is, agents observe the signal s_t at period t , but disagree about the likelihood of s_t . Agent i believes that the likelihood is $\pi_{i,t}(s_t|\omega)$ when the state is ω . In turn, this translates into agent i having the belief $p_{i,t}$. Differences in beliefs are thus coming from differences in personal views about the data generating process. This alternative interpretation does not alter our results.

Second, we assume that the agents are Bayesian. Again, this is not needed. All we need is that the agents revise their beliefs in a way that satisfies (2). There are several models in the behavioral economics literature that have this feature. We now describe a couple of them.

To start with, consider the following updating behavior first motivated by [Grether \(1980\)](#). Posteriors are derived by the formula

$$p_{i,t} = \frac{q_{i,t} (\pi_{i,t}(s_{i,t}|\omega^x))^{\beta_{i,t}}}{q_{i,t} (\pi_{i,t}(s_{i,t}|\omega^x))^{\beta_{i,t}} + (1 - q_{i,t}) (\pi_{i,t}(s_{i,t}|\omega^y))^{\beta_{i,t}}}$$

where $\beta_{i,t} \geq 0$. Here, the interpretation is that agents distort (that is, either over or under react to) the signal via their individual, time specific parameter $\beta_{i,t}$. Observe that when $\beta_{i,t} = 1$, this corresponds to standard Bayesian updating but the flexibility accommodates two common biases in belief updating: “under-inference” when $\beta_{i,t} < 1$ and “over-inference” when $\beta_{i,t} > 1$.

Observe that $p_{i,t} > q_{i,t}$ implies that $\pi_{i,t}(s_{i,t}|\omega^x) > \pi_{i,t}(s_{i,t}|\omega^y)$ (resp., $p_{i,t} < q_{i,t}$ implies that $\pi_{i,t}(s_{i,t}|\omega^x) < \pi_{i,t}(s_{i,t}|\omega^y)$). Therefore, if agent j follows the same updating rule (but with parameter $\beta_{j,t}$), his belief $p_{j,t}$ moves co-monotonically with agent i 's belief if the experiment is co-monotone. Our characterization thus continues to hold. The significant theoretical and experimental literature that builds on this framework is described in several surveys; two recent examples are [Benjamin \(2019\)](#) and [Gabaix \(2019\)](#).

Another framework is that of “coarse” updating. Here, in every period t , each agent i has belief cutoffs $0 =: \beta_{i,t}^0 < \beta_{i,t}^1 < \dots < \beta_{i,t}^{N_{i,t}} < \beta_{i,t}^{N_{i,t}+1} := 1$ and $N_{i,t}$ ordered beliefs $0 \leq \hat{q}_{i,t}^1 < \dots < \hat{q}_{i,t}^{N_{i,t}} \leq 1$. Roughly speaking, belief updating works as follows. The updated belief is $p_{i,t} = \hat{q}_{i,t}^{n_i}$ if the belief obtained by Bayesian updating (given by the right side of equation 1) lies between the cutoffs $\beta_{i,t}^{n_i-1}$ and $\beta_{i,t}^{n_i}$. In words, agents update coarsely by assigning intervals of beliefs obtained by Bayesian updating to single beliefs. There are several models that have this feature. Two recent examples are [Wilson \(2014\)](#) and [Jakobsen \(2021\)](#); the latter paper contains a detailed description of the literature. Note that a co-monotone experiment combined with such updating generates posteriors that are ordered in the sense of (2) and so this behavior will also be captured under [Definition 6](#). Additionally, observe that we can allow all agents to use different updating rules with distinct behavioral biases as long as their beliefs satisfy the property (2).

5. APPLICATIONS

5.1. Testing for discrimination

Disentangling channels of discrimination is known to be empirically difficult: in addition to the discussion in the aforementioned [Bertrand and Duflo \(2017\)](#), this is also highlighted in the survey of [Guryan and Charles \(2013\)](#). Our results can provide distinct suggestive evidence for different types of discrimination. To see this, consider richer choice data $(x_{i,t}, \theta_t)_{(i,t) \in \mathcal{I} \times \mathcal{T}}$, where $x_{i,t}$ is agent i 's decision to either approve or reject application t , and θ_t is a vector of observable characteristics (which are observable to both the agents and the analyst). As an example, job seekers routinely apply to several jobs and only receive interview for a subset of them. Here, the dataset consists of the interview decisions along with some observable characteristics of the applicants; for example, sex, age, or ethnicity. Other examples include applications for degree programs, mortgages, and credit cards.

As should be clear from the previous sections, only a specific type of data set is appropriate to test our model on. Specifically, we not only require all the agents to learn from common information, we also require the common signals to arrive in the same order.⁷ It is easy to construct examples (where the state is correlated across time) where cycles may arise because agent i sees the common signal s_{t_1} , chooses x , then sees s_{t_2} and chooses y , while agent j receives these signals and makes choices in the opposite order. In field data, a common set of applications that arrive at the same time (such as those for graduate school) may be viewed and decided upon in different

⁷For the special case where the state is drawn independently in each period, the order does not matter. It is worth mentioning that this assumption is commonly made in empirical work on discrimination. See, e.g., [Kline and Walters \(2021\)](#) and references therein.

orders by the prospective employers and these orders are typically unobserved by the researcher. Therefore, appropriately collected experimental data is perhaps more apt to test our model. Correspondence studies can easily be designed in which fictitious applications are sent to employers sequentially such that each fresh round of applications is only sent after decisions on the previous round are made.⁸

Now, for simplicity, suppose that the choice data $(x_{i,t}, \theta_t)_{(i,t) \in \mathcal{I} \times \mathcal{T}}$ is such that each θ_t corresponds to a single characteristic which can only take one of two values θ and θ' that, for example, correspond to being white and black or male and female. Let θ be the characteristic corresponding to the “favored” group. It is worth reiterating that, uncovering evidence of discrimination with our model does not require the particulars of the applications in different periods to be the same or even comparable.

There are two main theories of discrimination: statistical and taste-based. We consider the former first. Consider the sub-sample where we condition on $[\theta_t = \theta]$, that is, we restrict attention to the observations having the favored characteristic. Assume that it can be rationalized. This means that we cannot reject the hypothesis that the agents share the same belief in this sub-sample. Now, consider the sub-sample where we condition on $[\theta_t = \theta']$. Suppose, we cannot rationalize this sub-sample. Then, this implies that, in contrast to the favored group, there are at least two agents (i, j) who have different beliefs in the sub-sample of the “disadvantaged” group. This is suggestive evidence of *statistical* discrimination.

The evidence is, however, not fully conclusive. The two agents (i, j) may have different beliefs because they have different priors or because they interpret the common signal differently. We cannot reject that hypothesis with our simplest test. However, we can if we assume that the agents receive private but co-monotone signals (but this additionally requires us to assume the state is time-invariant and to therefore design the experiment appropriately). Two additional remarks are worth making. First, we cannot definitively state whether it is agent i , agent j or both, who are discriminating. We can only conclude that at least one is discriminating. Second, even if the complete sample $(x_{i,t})_{(i,t) \in \mathcal{I} \times \mathcal{T}}$ is rationalizable, we cannot reject the null hypothesis of statistical discrimination. This is because it is possible that agents have common beliefs because they are *all* statistically discriminating against the disadvantaged group.

⁸In a correspondence study, the researcher also knows the details of each application. However, the researcher does not know how employers interpret applications or, in other words, how a given application maps onto a signal. This can, of course, be formally modeled but we abstract away from these details in the discussion that follows.

We now consider the case of *taste-based* discrimination. Once again, consider the two sub-samples, where we condition on $[\theta_t = \theta]$ and $[\theta_t = \theta']$. If the two sub-samples are rationalizable, we can infer the ordering over thresholds in the two sub-samples and compare them. If the two orderings differ, there exists (at least) a pair of agents (i, j) such that $\bar{u}_i^\theta \geq \bar{u}_j^\theta$ and $\bar{u}_j^{\theta'} > \bar{u}_i^{\theta'}$. This implies that if we assume that the two agents have common beliefs in each sub-sample—an assumption we cannot reject since the two sub-samples are rationalizable—then agent j more often approves applications with characteristic θ than agent i , while the converse is true when the characteristic is θ' . In other words, this is evidence of *taste-based* discrimination. Finally, note that, if the two orderings coincide, we cannot reject the null hypothesis of no taste-based discrimination because this discrimination, while present, may not be reflected in the choice data.

5.2. Committees

While there is a large theoretical literature on committee voting, there is substantially less structural empirical work. [Iaryczower and Shum \(2012\)](#) estimate a model of committee voting by using the judicial decisions in the US Supreme Court. Here, both states (whether the meaning of the law favors the plaintiff or defendant) and alternatives (the ruling) are binary. [Hansen and McMahon \(2016\)](#) study individual voting behavior in Bank of England’s Monetary Policy Committee. They model the decision making environment using a binary state (more/less inflationary state), binary alternative (vote for either high or low interest rates) and assume that the state is drawn independently from an identical distribution in each period.

Our model is quite similar to these papers albeit with two key distinctions. They allow for private information but do not allow for learning ([Hansen and McMahon, 2016](#) instead focus on the role of signaling) and they effectively assume that (conditioning on observable covariates) each decision problem that a committee faces is ex ante identical (prior to the observation of signals). Moreover, the aim in these papers is to estimate the underlying model parameters and for this they require additional functional form assumptions. By contrast, although our model allows us to determine some ordinal information (such as the ranking of voters’ cutoffs), we cannot estimate any underlying parameters because they are not identified.⁹

One feature that is common to all these papers is that they assume committee members have the same prior beliefs in each decision problem before receiving their private signals. In other words, the heterogeneity across agents is in terms of their preferences or “ideology” instead of their prior

⁹In this sense, our paper is closer to [Iaryczower et al. \(2018\)](#) who develop a model of communication within the committee and estimate the effects of deliberation on the decisions of US appellate courts. Their model has multiple equilibria and is partially identified.

beliefs or “bias.”¹⁰ A natural question is whether these two different sources of (single parameter) heterogeneity lead to distinct testable implications? This would be informative because we might think the sources of bias and ideology differ and so different interventions might be required should we aim to correct them.

Theorem 3 showed that, if we allow for co-monotone signals, then there is no additional generality in assuming different preferences for each agent. So instead, suppose we assume that agents have the same preferences, information in each period arrives via common signals (as in **Definition 1**) but agents have different prior beliefs. This effectively amounts to assuming that the committee is deliberative but that members differ in terms of bias as opposed to ideology. We can show that heterogeneity in prior beliefs or preferences have the same testable implications when the state is time-invariant but the former has strictly weaker (but nontrivial) restrictions when the state is time varying.¹¹ In other words, differences between these sources of heterogeneity are testable.

5.3. Counterfactual analysis

A natural question is whether we can use the observed choice data \mathcal{X} to make counterfactual predictions about actions in a hypothetical period $T + 1$. To fix ideas, consider the baseline model. **Theorem 1** implies that we can rule out the subset of period $T + 1$ actions that would generate a cycle. However, a consequence is that we can never rule out the outcome where all agents pick the same alternative, that is, $x_{i,T+1} = x_{j,T+1}$ for all (i, j) . In the context of committee voting, this implies that a unanimous vote for either alternative is always possible in the counterfactual scenario.

That said, our model allows for considerably sharper *conditional* predictions. To be more precise, if we knew the alternative that one agent plans to choose in period $T + 1$, this can impose a lot more structure on the actions of the other agents. For instance, if a member i of a committee announces that she plans to vote x in period $T + 1$, we might be able to rule out the event where the alternative y gets a majority of votes and therefore is enacted. We end the paper by demonstrating how these observations can be employed in a simple example.

EXAMPLE 4. Consider the choice data \mathcal{X} of three agents i, j, k and two time periods represented in the following table:

¹⁰Iaryczower and Shum (2012) discuss in footnotes 11 and 18 why allowing for heterogeneity in the prior beliefs lead to poorly behaved estimates.

¹¹Formal statements and proofs available upon request.

	i	j	k
t_1	x	y	y
t_2	x	x	y

Since there are no cycles, this data is rationalizable in the sense of [Definition 1](#). Now suppose we want to predict the possible action profiles in a third period t_3 . The following table describe all possible counterfactual action profiles (that do not generate cycles):

	i	j	k
t_1	x	y	y
t_2	x	x	y
t_3	x	x	x

	i	j	k
t_1	x	y	y
t_2	x	x	y
t_3	x	x	y

	i	j	k
t_1	x	y	y
t_2	x	x	y
t_3	x	y	y

	i	j	k
t_1	x	y	y
t_2	x	x	y
t_3	y	y	y

There are eight possible action profiles that are possible in period t_3 and the observed choice data eliminates half of those. Note that the example is such that the choices in the first two time periods completely order the cutoffs of all three agents ($\bar{u}_i < \bar{u}_j < \bar{u}_k$). Therefore, observing the choices from any additional time periods would not change the counterfactual predictions and we would still only be able to eliminate half the possible action profiles.

Now suppose that we knew that either agent i or j will choose y in period t_3 . We can immediately conclude that only the right two action profiles are possible and so x cannot be chosen by a majority. In other words, knowing the choices of a single agent in a counterfactual time period can allow us to considerably sharpen our predictions. Finally, the insights from this example can easily be generalized to demonstrate how counterfactual analysis can be conducted when there are general preferences and/or multiple actions and/or co-monotone experiments.

6. CONCLUDING REMARKS

We develop a simple yet rich framework via which we can test different assumptions on preferences and the underlying learning environment with very parsimonious data. The theoretical results can be used to uncover evidence for different types of discrimination and to detect differences in bias versus ideology in committee decision making. The model allows for counterfactual analysis that can be considerably sharpened if the actions of a subset of agents is known. Finally, we can test the assumption of common beliefs against other models of private learning (because

they have distinct testable implications). This shows that it is possible to distinguish between different learning behaviors. This is a potentially useful result for the large behavioral economics literature that studies non-Bayesian learning.

One assumption of the baseline model that we maintained throughout was that the underlying state is binary. We end the paper with a short discussion on generalizing beyond binary states. As we have argued, testable restrictions in our model arise from the fact that we can map the choices of the agents to the underlying information: recall that in the baseline model, if agent i switches from x to y over two periods this implies that the (common) belief must have gone down. With more than two states, the set of preferences $u_i(\omega_t, x_{i,t})$ expands greatly and so the relation between choices and beliefs can be considerably more complex.

That said, the analysis can be extended provided some structure is imposed on the set of possible preferences. As an example, suppose the unknown underlying state ω can now take (a continuum of possible) values in the set $[0, 1]$. Agents' preferences are as follows. Each agent i has a threshold \bar{u}_i and chooses x whenever $\mathbb{E}[\omega_t] > \bar{u}_i$ (and chooses y whenever $\mathbb{E}[\omega_t] < \bar{u}_i$). As before, the state can be stochastic and learning occurs via signals realized from a statistical experiment.

Note that in this simple setting, if an agent i chooses x in period t_1 and y in period t_2 , it must imply that $\mathbb{E}[\omega_{t_1}] > \bar{u}_i > \mathbb{E}[\omega_{t_2}]$. Therefore, once again rationalizable choice data \mathcal{X} cannot have a cycle. It is simple to show that the absence of cycles is also a sufficient condition for rationalizability; loosely speaking, the expected value of the state $\mathbb{E}[\omega_t]$ plays the identical role as the belief p_t in the baseline model.¹²

Of course, the above is one particular formulation of a multiple-state extension of our model; in future work, we hope to further develop these results. Learning occurs in almost all dynamic choice environments and there is considerable scope to study what can be revealed about preferences and information from observed choices.

¹²The sufficiency can be argued by restricting the distribution of the state ω_t to be supported on only two points $\{0, 1\}$. Note that this implies $\mathbb{E}[\omega_t] = p_t$ where p_t is the belief that the agents assign to the state being $\omega_t = 1$. The argument then follows identically to the proof of [Theorem 1](#).

APPENDIX A. PROOF OF THEOREM 2

PROOF. (\Rightarrow .) In the text.

(\Leftarrow .) We first define an incomplete binary relation \succ on pairs (i, n) of agents $i \in \mathcal{I}$ and indices of actions $n \in \{1, \dots, N\}$ as follows. For all $i \in \mathcal{I}$, for all $n \in \{2, \dots, N\}$,

$$(i, n) \succ (i, n-1).$$

For all $i \neq j$, for all $n \in \{2, \dots, N\}$,

$$(i, n) \succ (j, n'),$$

if there exists $t \in \mathcal{T}$ such that $x_{i,t} = y_{n-1}$ and $x_{j,t} = y_{n'}$. All other pairs are not comparable.

This relation reflects the order of the preference cutoffs. The first line captures the order $\bar{u}_{i,n} > \bar{u}_{i,n-1}$ of actions. As for the second line, observe that, when $x_{i,t} = y_{n-1}$, $x_{j,t} = y_{n'}$, if the data is rationalizable, then $\bar{u}_{i,n} > p_t > \bar{u}_{i,n-1}$ and $\bar{u}_{j,n'+1} > p_t > \bar{u}_{j,n'}$ which imply $\bar{u}_{i,n} > \bar{u}_{j,n'}$.

We now argue that the relation \succ does not have a cycle of any length, that is, there do not exist sequences $((i_1, n_1), \dots, (i_k, n_k))$ for some $k \geq 2$ such that

$$(i_1, n_1) \succ (i_2, n_2) \succ \dots \succ (i_k, n_k) \succ (i_1, n_1).$$

By contradiction, suppose that such a sequence exists. By definition of \succ , $(i_1, n_1) \succ (i_2, n_2)$ implies that there is a time period t_2 at which $x_{i_1, t_2} = y_{m_2}$ where $m_2 \leq n_1 - 1$ and $x_{i_2, t_2} = y_{n_2}$ (if $i_2 = i_1$ then $x_{i_1, t_2} = x_{i_2, t_2} = y_{m_2} = y_{n_2}$ with $m_2 = n_2 = n_1 - 1$).

Similarly, $(i_2, n_2) \succ (i_3, n_3)$ implies that there is a time period t_3 at which $x_{i_2, t_3} = y_\ell$ where $\ell \leq n_2 - 1$ and $x_{i_3, t_3} = y_{n_3}$. Now let us compare periods t_2 and t_3 . Observe that agent i_2 chooses a strictly lower action in t_3 (y_ℓ as opposed to y_{n_2}) and therefore the no cycles condition implies that agent i must also choose a weakly lower action $x_{i_1, t_3} = y_{m_3}$ where $m_3 \leq m_2 < n_1$.

Iterating the argument, there exists a period t_k where agent i_k chooses $x_{i_k, t_k} = y_{n_k}$ and agent i chooses action $x_{i_1, t_k} = y_{m_k}$ where $m_k < n_1$.

Now observe that if $i_k = i_1$, then $x_{i_k, t_k} = y_{n_k} = x_{i_1, t_k} = y_{m_k}$ which is a contradiction because $(i_k, n_k) \succ (i_1, n_1)$ means that $n_k = m_k = n_1 + 1$.

Therefore, it must be the case that $i_k \neq i_1$. Consequently, there exists period t_1 in which $x_{i_k, t_1} = y_{n_k-1}$ and $x_{i_1, t_1} = y_{n_1}$. Compare periods t_k and t_1 . Agent i_k chooses a strictly lower action (y_{n_k-1} versus y_{n_k}) in t_1 whereas agent i_1 chooses a strictly higher action (y_{n_1} versus y_{m_k}). These choices constitute a cycle, a contradiction.

In addition, we define the symmetric part \sim as $(i, n) \sim (i, n)$ for all (i, n) . Thus, the relation $\succ \cup \sim$ is reflexive and consistent. By Theorem 3 (that generalizes Szpilrajn's Theorem) in [Suzumura \(1976\)](#), there exists a complete and transitive binary relation \succ^* on the set of pairs (i, n) that extends $\succ \cup \sim$. In turn, this implies that we can find numbers $\bar{u}_{i,n}$ for all (i, n) that satisfy $\bar{u}_{i,n} > \bar{u}_{j,n'}$ when $(i, n) \succ^* (j, n')$. In particular, note that

- i. $\bar{u}_{i,n} > \bar{u}_{i,n'}$ for $n > n'$ and
- ii. $\bar{u}_{i,n} > \bar{u}_{j,n'}$ if there is a time period $t \in \mathcal{T}$ such that $x_{i,t} = y_{n-1}$ and $x_{j,t} = y_{n'}$.

Item i gives us the correctly ordered cutoffs for each agent, so these cutoffs generate valid preferences. Item ii implies that, for every $t \in \mathcal{T}$, we can find a p_t such that $\bar{u}_{i,n+1} > p_t > \bar{u}_{i,n}$ where $x_{i,t} = y_n$ for all $i \in \mathcal{I}$. To see this, note that, by point ii above, $x_{i,t} = y_n$ and $x_{j,t} = y_{n'}$ implies $\bar{u}_{i,n+1} > \bar{u}_{j,n'}$. Therefore,

$$\min_{i,n} \{\bar{u}_{i,n+1} \mid x_{i,t} = y_n\} > \max_{j,n'} \{\bar{u}_{j,n'} \mid x_{j,t} = y_{n'}\},$$

and any p_t between the left and right sides will suffice.

To complete the proof, we can construct the signals and experiments by following the same steps as in the proof of [Theorem 1](#) from the main text. ■

APPENDIX B. FROM MARGINALS TO JOINT DISTRIBUTIONS

LEMMA 1. *Suppose that there exist strictly positive numbers $(\pi_i^x, \pi_i^y)_{i \in \mathcal{I}}$ such that $\sum_i \pi_i^x < 1$, $\sum_i \pi_i^y < 1$, and*

$$[\pi_i^x - \pi_i^y][\pi_j^x - \pi_j^y] > 0,$$

for all (i, j) . Then, there exists a strictly co-monotone experiment $\pi : \{\omega^x, \omega^y\} \rightarrow \Delta(\times_{i \in \mathcal{I}} S_i)$ and a profile of signals $s = (s_1, \dots, s_I)$ such that

$$(\pi_i(s_i | \omega^x), \pi_i(s_i | \omega^y)) = (\pi_i^x, \pi_i^y),$$

for all i . Moreover, the profile of signal s is in the support of the experiment.

PROOF. We construct $\pi(\cdot | \omega^x)$. For all $\varepsilon \geq 0$, let $\pi_i^x(\varepsilon)$ be a solution to

$$(1 - \varepsilon)\pi_i^x(\varepsilon) + \varepsilon \frac{1}{I} = \pi_i^x.$$

Choose ε small enough such that $\pi_i^x(\varepsilon) \in (0, 1)$ and $\sum_i \pi_i^x(\varepsilon) < 1$. Since $\pi_i^x(0) = \pi_i^x$, this is possible by continuity.

Let $S_i = \{s^0, s^1, \dots, s^I\}$ for all $i \in \mathcal{I}$ and $S^* = \times_{i \in \mathcal{I}} (S_i \setminus \{s^0\})$. We define $\pi(\cdot | \omega^x)$ as follows:

$$\pi(s_1, \dots, s_n | \omega^x) = \begin{cases} (1 - \varepsilon)\pi_i^x(\varepsilon) + \varepsilon \frac{1}{I} & \text{if } s_1 = \dots = s_n = s^i, i \in \mathcal{I} \\ \varepsilon \frac{1}{I} & \text{if } s \in S^* \setminus \bigcup_{i \in \mathcal{I}} \{(s^i, \dots, s^i)\} \\ (1 - \varepsilon)(1 - \sum_{i \in \mathcal{I}} \pi_i^x(\varepsilon)) & \text{if } s_1 = \dots = s_n = s^0 \\ 0 & \text{otherwise.} \end{cases}$$

It is routine to check that $\pi(\cdot | \omega^x)$ is indeed a probability over $\times_{i \in \mathcal{I}} S_i$. (Note that the cardinality of $S^* \setminus \bigcup_{i \in \mathcal{I}} \{(s^i, \dots, s^i)\}$ is $I^I - I$.) The support of $\pi(\cdot | \omega^x)$ is $S^* \cup \{(s^0, \dots, s^0)\}$.

For all $s \in S^*$, agent i 's marginal $\pi_i(s_i | \omega^x)$ is

$$(1 - \varepsilon)\pi_\ell^x(\varepsilon) + \varepsilon \frac{1}{I} = \pi_\ell^x$$

whenever $s_i = s^\ell$, $\ell \in \{1, \dots, I\}$.

The probability $\pi(\cdot | \omega^y)$ is constructed analogously. Note that the epsilons we use to construct $\pi(\cdot | \omega^x)$ and $\pi(\cdot | \omega^y)$ may differ. However, for any epsilon smaller than the minimum of the two of them, we continue to have two well-defined probabilities with the same epsilon. Thus, we can choose the same epsilon for both constructions.

It remains to check that the experiment is co-monotone. Consider any profile of signals $s \in S^*$. We have that

$$[\pi_i(s_i | \omega^x) - \pi_i(s_i | \omega^y)][\pi_j(s_j | \omega^x) - \pi_j(s_j | \omega^y)] = [\pi_\ell^x - \pi_\ell^y][[\pi_\ell^x - \pi_\ell^y]] > 0,$$

with $(s_i, s_j) = (s^\ell, s^{\ell'})$, $(\ell, \ell') \in \mathcal{I} \times \mathcal{I}$.

Finally, if the signal profile is (s^0, \dots, s^0) , we have that

$$[\pi_i(s^0 | \omega^x) - \pi_i(s^0 | \omega^y)][\pi_j(s^0 | \omega^x) - \pi_j(s^0 | \omega^y)] = [\pi_i(s^0 | \omega^x) - \pi_i(s^0 | \omega^y)]^2 > 0,$$

for all pairs (i, j) . This follows immediately from the fact that all agents have the same marginal over s^0 , regardless of the state.

The experiment is therefore co-monotone. This completes the proof. ■

We can illustrate the construction when there are two agents. When the state is ω^x , the joint distribution over signals is given by

	s^1	s^2	s^0
s^1	$(1 - \varepsilon)\pi_1^x(\varepsilon) + \varepsilon\frac{1}{4}$	$\varepsilon\frac{1}{4}$	0
s^2	$\varepsilon\frac{1}{4}$	$(1 - \varepsilon)\pi_2^x(\varepsilon) + \varepsilon\frac{1}{4}$	0
s^0	0	0	$(1 - \varepsilon)(1 - \pi_1^x(\varepsilon) - \pi_2^x(\varepsilon))$

Thus, conditional on any signal, the two agents have the same marginals. Moreover, we construct the numbers ε , $\pi_1^x(\varepsilon)$ and $\pi_2^x(\varepsilon)$ such that

$$(1 - \varepsilon)\pi_i^x(\varepsilon) + \varepsilon\frac{1}{4} = \pi_i^x$$

for all i , and the probability is well-defined.

We do a similar construction when the state is ω^y . It is then immediate to verify that the experiment is co-monotone. Finally, if agent i receives the signal s^i , the ratio of marginals is

$$\frac{\pi_i^x}{\pi_i^y},$$

as required in the main text.

REFERENCES

- AFRIAT, S. N. (1967): "The Construction of Utility Functions from Expenditure Data," *International Economic Review*, 8, 67–77.
- ALONSO, R. AND O. CÂMARA (2016): "Persuading Voters," *American Economic Review*, 106, 3590–3605.
- BENJAMIN, D. J. (2019): "Errors in Probabilistic Reasoning and Judgment Biases," *Handbook of Behavioral Economics: Applications and Foundations 1, 2*, 69–186.
- BERTRAND, M. AND E. DUFLO (2017): "Field Experiments on Discrimination," *Handbook of Economic Field Experiments*, 1, 309–393.
- BERTRAND, M. AND S. MULLAINATHAN (2004): "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economic Review*, 94, 991–1013.
- CAPLIN, A. AND M. DEAN (2015): "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105, 2183–2203.
- CAPLIN, A. AND D. MARTIN (2015): "A Testable Theory of Imperfect Perception," *Economic Journal*, 125, 184–202.
- CHAMBERS, C. P. AND F. ECHENIQUE (2016): *Revealed Preference Theory*, vol. 56, Cambridge University Press.
- CHAN, J., A. LIZZERI, W. SUEN, AND L. YARIV (2018): "Deliberating Collective Decisions," *Review of Economic Studies*, 85, 929–963.
- DE OLIVEIRA, H. AND R. LAMBA (2019): "Rationalizing Dynamic Choices," *Working Paper*.
- FRICK, M., R. IIJIMA, AND T. STRZALECKI (2019): "Dynamic Random Utility," *Econometrica*, 87, 1941–2002.
- GABAIX, X. (2019): "Behavioral Inattention," in *Handbook of Behavioral Economics: Applications and Foundations 1*, vol. 2, 261–343.
- GERLING, K., H. P. GRÜNER, A. KIEL, AND E. SCHULTE (2005): "Information Acquisition and Decision Making in Committees: A Survey," *European Journal of Political Economy*, 21, 563–597.
- GREYER, D. M. (1980): "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, 95, 537–557.
- GUL, F. AND W. PESENDORFER (2012): "The War of Information," *Review of Economic Studies*, 79, 707–734.
- GURYAN, J. AND K. K. CHARLES (2013): "Taste-Based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots," *Economic Journal*, 123, F417–F432.

- HANSEN, S. AND M. MCMAHON (2016): "First Impressions Matter: Signalling as a Source of Policy Dynamics," *Review of Economic Studies*, 83, 1645–1672.
- IARYCZOWER, M., X. SHI, AND M. SHUM (2018): "Can Words Get in the Way? The Effect of Deliberation in Collective Decision Making," *Journal of Political Economy*, 126, 688–734.
- IARYCZOWER, M. AND M. SHUM (2012): "The Value of Information in the Court: Get it Right, Keep it Tight," *American Economic Review*, 102, 202–37.
- JAKOBSEN, A. M. (2021): "Coarse Bayesian Updating," Tech. rep., Working paper.
- KLINE, P. AND C. WALTERS (2021): "Reasonable Doubt: Experimental Detection of Job-Level Employment Discrimination," *Econometrica*, 89, 765–792.
- LI, H. (2001): "A Theory of Conservatism," *Journal of Political Economy*, 109, 617–636.
- LI, H. AND W. SUEN (2009): "Decision-making in committees," *Canadian Journal of Economics*, 42, 359–392.
- LU, J. (2016): "Random Choice and Private Information," *Econometrica*, 84, 1983–2027.
- MAKRIS, M. AND L. RENO (2021): "Information Design in Multi-Stage Games," Tech. rep., Working Paper.
- SHMAYA, E. AND L. YARIV (2016): "Experiments on Decisions under Uncertainty: A Theoretical Framework," *American Economic Review*, 106, 1775–1801.
- SUZUMURA, K. (1976): "Remarks on the Theory of Collective Choice," *Economica*, 43, 381–390.
- WILSON, A. (2014): "Bounded Memory and Biases in Information Processing," *Econometrica*, 82, 2257–2294.