

Temptation and Persuasion

Alexander M. Jakobsen*

December 4, 2024

Abstract

I study the identification and comparative-static properties of Bayesian Persuasion in terms of the receiver’s choices, preferences and welfare. By interpreting sender and receiver as different selves, the model delivers new insights into the interaction between temptation, self-signaling, demand for commitment, and the value of information. All model parameters—and, thereby, endogenous signals/beliefs—are identified from receiver’s behavior, enabling simple comparative statics regarding the degree of conflict between agents and measures of optimism or pessimism. My results differ substantially from those of related decision models, highlighting the importance of “tempting” beliefs in the analysis of intrapersonal conflict.

1 Introduction

1.1 Motivation

The Bayesian Persuasion model (Kamenica and Gentzkow, 2011) has received considerable attention and become a central framework in the economics of information design and disclosure. In the standard setup, one agent (Sender) selects an information structure and commits to revealing its signal realization to another (Receiver). The ensuing literature has examined numerous extensions and variations of the baseline model, with results typically focusing on Sender’s choice of information and whether he benefits from this opportunity. This paper examines the interaction from Receiver’s perspective. In particular, I develop a decision-theoretic analogue of Bayesian Persuasion that takes as primitive Receiver’s preferences or choices; Sender’s behavior is not observed and must be inferred from Receiver’s.

*Kellogg MEDS, Northwestern University. I am grateful to Nemanja Antic, Ian Ball, Peter Klibanoff, Jay Lu, Philipp Sadowski, Ron Siegel, Vasiliki Skreta, Max Stinchcombe, Ina Taneva, and Ben Young, as well as seminar and conference participants at Arizona, Columbia, Duke/UNC, the Hitotsubashi Summer Institute 2022, Northwestern, NYU, Penn State, SAET 2023, Toronto, UCLA, UPenn, and UT Austin for helpful conversations and feedback. This paper was previously entitled “Revealed Persuasion.”

My motivation is twofold. First, the decision-theoretic approach allows several basic questions about the identification and comparative statics of persuasion to be resolved at a high level of generality. For example, when might Receiver benefit from commitment (shrinking the action set) or from additional public information? Each avenue serves to level the playing field against Sender’s commitment power but distorts incentives in nontrivial ways. Nonetheless, these, and other, questions resolve quite naturally in my framework. Identification is also of central importance: can one identify parameter values (utility functions and prior beliefs), and thereby information provision, merely by observing Receiver’s behavior, or is richer data needed? It turns out that various types of Receiver data are sufficient to identify all parameters, and my results provide methods to perform such elicitation.

Second, the Receiver-based model of persuasion, reinterpreted as a “dual-selves” interaction, provides a theory of behavior delivering new insights into widely-studied phenomena such as temptation, selective attention and motivated cognition. Importantly, my approach does not involve non-Bayesian updating, belief-based utility, or other modeling devices that, while intuitively appealing, often lead to conceptual or technical problems. The analysis yields new implications for, and links between, the value of commitment and of information for agents facing temptation or other competing motives when making decisions.

To illustrate, consider an individual deciding between two options at a restaurant: a burger or a salad.¹ The burger is always the same but the salad varies with exogenous factors like the availability of ingredients, how busy the kitchen is, and so on; in effect, there are two types of salad and an exogenous state determines which type will arrive if the individual orders one. Suppose the individual prefers the first type of salad to the burger and the burger to the second type, and that prior beliefs are such that the salad is optimal *ex ante*. These preferences are represented by a utility index u and prior beliefs μ . “Temptation” preferences involve the same prior but a different utility index v ; in this case, v ranks the burger strictly preferred to either type of salad. What will the individual do?

To analyze this situation, I treat u as Receiver’s utility function in Bayesian Persuasion and v as Sender’s. This means information is chosen to maximize v in anticipation of signal-contingent action choices that maximize u . Thus, the effect of temptation is to garble beliefs at the time of choice in a way that increases the probability of the burger being chosen; u and μ are such that the salad is preferred *ex-ante* but selective attention can give the individual (Receiver) an “excuse” to choose the burger at least some of the time.

The interpretation of this behavior is important and differs from that of related models. In my model, perfect information is freely available at the time of choice but, due to temptation, the individual selectively pays attention and/or mentally distorts signals. For example, the

¹This is a slightly-modified version of the leading example in Gul and Pesendorfer (2001).

individual may ask the server what kind of salad is available (perfect information) or opt for a noisier signal (peripheral glances at other tables or the kitchen line, consulting recent online reviews, etc). Even if an objectively informative signal arrives, the individual may mentally distort it by pretending not to have seen it, combining it with biased recollections of previous experiences at the restaurant, or engaging in other methods of self-deception. The key is that belief distortion is driven by v while consumption choices maximize u conditional on posterior beliefs; thus, the individual succumbs to temptation only when doing so is “justified” as a u -optimal choice. The ability to engage in such self-deception is constrained by Bayesianism: the distribution of posteriors must average out to the prior, so the individual cannot systematically hold misspecified beliefs.

In contrast, Strotzian models, or what Gul and Pesendorfer (2001) refer to as settings with *overwhelming temptation*, result in consumption choices that maximize v ; this means information plays no role and the burger is chosen with probability 1. In the Gul-Pesendorfer model, the individual may resist temptation but incur a cost from doing so; again, information plays no role. Here the individual engages in belief distortion or selective attention to justify tempting choices. Consequently, choices are stochastic—sometimes the individual succumbs to temptation, sometimes not—and the agent incurs an indirect “cost of temptation” by garbling freely-available information when making choices. This channel of “tempting” belief formation has rather different implications for behavior and welfare. I expand on these ideas, and the connection to other related models, below.

1.2 Results

I consider a single decision maker, Receiver, who chooses among *acts* (Anscombe and Aumann, 1963): profiles $f = (f_\omega)_{\omega \in \Omega}$ of lotteries $f_\omega \in \Delta X$ where Ω is a finite set of states, X a finite set of outcomes and ΔX the set of objective lotteries over X . If Receiver chooses f and state ω realizes, an outcome is generated by the lottery f_ω . In the example above, for instance, there are two states and ordering a burger corresponds to the act $(Burger, Burger)$ (the same deterministic outcome in each state) and ordering a salad to $(Salad1, Salad2)$ (different outcomes in different states).

A *menu* is a finite set A of acts available to Receiver. The analysis involves different types of choice or preference primitives based on menus:

1. *Menu preferences* $\succsim$ indicating Receiver’s ex-ante ranking of menus; $A \succsim B$ means Receiver would rather commit to menu A than to menu B .
2. *Random choices* ρ indicating Receiver’s unconditional choice distributions from menus; $\rho^A(f) \in [0, 1]$ is the frequency with which Receiver chooses f from A .

3. *State-contingent random choices* λ_ω indicating Receiver's choice distributions in different states; $\lambda_\omega^A(f) \in [0, 1]$ is the frequency with which Receiver chooses f from A when the true state is ω .
4. *Choice correspondence* data c indicating which actions are chosen with positive probability; $c(A)$ is the support of ρ^A (or the union over all ω of the support of λ_ω^A).

I study representations of these primitives—*Persuasion Representations*—involving three parameters (μ, u, v) : prior beliefs, Receiver's utility function and Sender's utility function, respectively. Such representations reflect Receiver's preferences or choices given that an (unobserved) Sender controls the information available to Receiver. At signal realizations, Receiver updates beliefs via Bayes' rule and chooses actions to maximize (the expectation of) u ; anticipating this, Sender chooses information to maximize expected v . For Persuasion Representations of $\succsim$, $A \succsim B$ means Receiver's expected utility from the interaction is greater at menu A than at menu B . In a Persuasion Representation for ρ , distributions ρ_A match the distribution of Receiver's choices generated by Sender's optimal information structure at menu A ; similarly, Persuasion Representations of λ require λ_ω^A to match the distribution generated by the interaction in state ω .

The first results establish whether the parameters (μ, u, v) can be identified from different types of primitives. Theorem 2 states that the parameters are revealed up to standard uniqueness notions by $\succsim$; in particular, μ is unique and u, v are unique up to positive affine transformation (denoted $\approx$). Similarly, Theorem 3 establishes unique identification from ρ . Theorem 4 states that u and v can be identified from λ or c and that μ is identified by either of these primitives only if $u \not\approx v$. Thus, outside of this particular case, each type of primitive provides enough information to identify all parameters. Since c suffices in almost all cases and is derived from ρ or λ , most subsequent results in the paper are established for $\succsim$ and c .

Section 4 analyzes comparative statics in terms of these primitives. Sections 4.1 and 4.2 revolve around two questions: (i) when does Receiver benefit from hard commitment, and (ii) when does Receiver benefit from additional public information? As in much of the decision-theoretic literature on temptation, my approach to (i) involves Receiver's preferences over (or choices from) nested menus $A \supseteq B$. To study question (ii), I introduce a new operation on menus that simulates public signals. For any A and Blackwell experiment σ , the menu σA mixes acts in A in such a way that it is *as if* Receiver (i) chooses from the original set A , and (ii) before doing so, observes a signal from σ in addition to, and independently of, the signal generated by Sender. Thus, σ acts as a type of lower bound on the experiments available to Sender to choose from; in the behavioral interpretation, this is a constraint on the agent's ability to engage in self-deception or selective attention. By construction, $\sigma A \supseteq A$; thus,

such bounds effectively provide flexibility to Receiver.²

Section 4.1 characterizes, and establishes links between, preferences for flexibility and for information. Menu preferences $\succsim$ satisfy *Preference for Flexibility* (Kreps, 1979) if $A \supseteq B$ implies $A \succsim B$. Since $\sigma A \supseteq A$, Preference for Flexibility implies *Preference for Information*: $\sigma A \succsim A$ for all A and σ . Both properties hold for standard, rational decision makers but are typically violated by Persuasion Representations. Proposition 1 states that, for Persuasion Representations, (i) Preference for Flexibility and Preference for Information are equivalent, and (ii) either property holds if and only if either $u \approx v$ or $u \approx -v$. This finding highlights a substantial difference between existing models of temptation and the belief-based approach considered here—in Persuasion Representations, Preference for Flexibility need not indicate the absence of conflict ($u \approx v$) but may in fact indicate the most extreme conflict ($u \approx -v$).³ Subsequent results characterize the $u \approx v$ and $u \approx -v$ cases separately and in terms of both flexibility and information.

I also study how flexibility and information impact the choice correspondence c (hence ρ and λ). For example, Proposition 1 states that $u \approx v$ or $u \approx -v$ if and only if c satisfies *Sen's Condition α* (Sen, 1971): $A \subseteq B$ implies $c(B) \cap A \subseteq c(A)$. Thus, for Persuasion Representations, Sen's α is analogous to Preference for Flexibility. Just as Preference for Information is weaker than Preference for Flexibility but equivalent in Persuasion Representations, a weaker property—*Informational Sen's α* ($c(\sigma A) \cap A \subseteq c(A)$)—turns out to be equivalent to Sen's α . In general, my results show how key properties regarding flexibility have equivalent information-based analogues, and also how a key property for one primitive (say $\succsim$) translates into an equivalent statement for another (c).

Section 4.2 develops finer measures of conflict. In particular, I study how *more-aligned* preferences—formalized by replacing v with a mixture $\alpha u + (1 - \alpha)v$ —impact Receiver's preferences and choices. Once again, I provide full characterizations for $\succsim$ and c . For example, Proposition 6 finds that preferences are less aligned if and only if there are fewer instances where Receiver benefits from committing to a single act; for c , the analogous condition is for there to be more menus where only a single act is chosen.

Finally, section 4.3 considers the possibility that Receiver holds incorrect beliefs about Sender's utility function. The idea is that the index v revealed by $\succsim$ reflects Receiver's *beliefs* about Sender's preferences while that revealed by (say) λ reflects Sender's *actual* preferences. I adopt (and slightly modify) the approach of Ahn et al. (2019) by comparing Receiver's ex-

²As explained in section 4, however, acts in σA offer a kind of commitment power by giving Receiver the ability to delegate choices conditional on σ -signals to another party (eg, to the server at the restaurant).

³In contrast, related models like that of Gul and Pesendorfer (2001) or (modern reformulations of) Strotz (1955) reduce to standard rationality when Preference for Flexibility is imposed—there is no conflict between selves if commitment offers no benefit.

ante preferences $\succsim$ to ex-post choices λ . This yields simple tests of sophistication (correct beliefs about v) and comparative static characterizations for two natural forms of naivete—optimism and pessimism—utilizing the notion of more-aligned preferences described above.

Throughout, my results assume the primitives have Persuasion Representations but that the analyst does not know the values of the parameters (μ, u, v) ; these must be inferred from the primitives. It is natural to wonder what conditions (axioms) on the primitives ensure existence of Persuasion Representations. To some extent, I leave this as an open problem. Theorem 5 in Appendix D provides an axiomatic characterization of Persuasion Representations for menu preferences $\succsim$; this, in turn, can be adapted to a characterization for state-contingent choice data λ (Theorem 6). However, I do not axiomatically characterize the model in terms of ρ or c , and the characterizations for $\succsim$ and λ take the prior μ as given.

1.3 Related Literature

Bayesian Persuasion and Comparative Statics

The framework developed in this paper is a decision-theoretic analogue of the Bayesian Persuasion model of Kamenica and Gentzkow (2011). A key distinction between my approach and that of the persuasion literature is that my focus is on Receiver’s (not Sender’s) choices and welfare and that model parameters (utility functions, prior beliefs) are not imposed from the outset but revealed by Receiver’s behavior. My analysis involves comparisons between different persuasion games (menus or action sets), yielding new insights into how Receiver might benefit from commitment power of his own or from additional public information.

Curello and Sinander (2022) establish rich and general results on the comparative statics of Bayesian Persuasion—specifically, they characterize conditions on model parameters (eg, the convexity of Sender’s interim payoff function) under which Sender chooses a more informative (Blackwell, 1951) structure. This paper also studies comparative statics, but the questions considered are different, as is the methodology. In particular, I focus on the relationship between Receiver’s welfare and the degree of conflict between the agents. My characterizations express patterns in choice or preference data that link demand for commitment with the value of information and, in turn, map to relationships between model parameters (u and v). Changes to Receiver’s welfare need not be due to Blackwell-comparable changes to information, so neither study nests the other.

Self-signaling and Strategic Ignorance

The behavioral interpretation developed in this paper is inspired by models of self-signaling. Bénabou and Tirole (2002) model a time-inconsistent agent who holds beliefs

about the returns to costly effort and may engage in selective information acquisition and/or costly memory retrieval to motivate such effort. Their model is more structured and involves parameters for the cost of effort, the benefit of success, β - δ discounting, costly memory retrieval and prior beliefs. My approach involves three parameters, (μ, u, v) , that apply to arbitrary decision problems (menus) and, thereby, to the full range of stakes, trade offs, and degrees of conflict faced by the agent. On a technical level, the commitment power of Sender makes the persuasion framework somewhat simpler; this enables identification of model parameters as well as the unobserved, endogenous cognition and/or information acquisition underlying such models of intrapersonal behavior.

Carrillo and Mariotti (2000) analyze a model of intrapersonal conflict where a present-biased individual, in every period, makes a binary consumption choice. Consumption exerts a negative externality on future selves with some probability and the agent can choose to learn about that probability. Lacking the ability to commit to future consumption plans, the agent sometimes chooses not to acquire freely-available information—strategic ignorance acts as a commitment device. Jakobsen (2021) examines this channel in arbitrary decision problems by studying Sender’s preferences for information in Bayesian persuasion; the decision-theoretic framework is similar to this paper but involves direct observation of Sender’s informational preferences. Here, Sender’s behavior is unobserved and must be inferred from Receiver’s, so the model, results, and interpretation are quite different.

Temptation, Multiple Selves, Demand for Commitment

There is a rich literature involving multiple-selves as a driver of temptation and related behavior. Such models typically involve utility functions that either vary with time or otherwise directly determine how intrapersonal conflict is resolved. A Strotzian (Strotz, 1955) agent, for example, initially has utility function u but anticipates their future-self having a different function v . The Gul and Pesendorfer (2001) model, by contrast, considers individuals who maximize utility u net of self-control costs determined by v . Crucially, consumption choices in such models are made using either a time-sensitive (as in Strotz) or modified (as in Gul-Pesendorfer) objective function and this generates demand for commitment. In my model, temptation influences *beliefs* but not the utility index u used to make consumption choices. This fundamentally-different mechanism need not generate demand for commitment and will, in some cases, actually generate demand for flexibility—an important implication for empirical or experimental researchers interested in temptation-driven behavior.

Random Choice and Inattention

An interesting feature of Persuasion Representations is that they involve a new mech-

anism generating random choice. Random utility models (Falmagne, 1978; Gul and Persendorfer, 2006) cannot rationalize behavior generated by Persuasion Representations since the information structure chosen by Sender (hence, the distribution of expected utility functions governing Receiver’s choices) varies with the menu of alternatives. For the same reason, random choices generated by private information as in Lu (2016) cannot rationalize choices generated by Persuasion Representations, nor can the random-Strotz framework of Dekel and Lipman (2012). Models of costly contemplation (Ergin and Sarver, 2010) or rational inattention (Ellis, 2018; Caplin and Dean, 2015) can rationalize Persuasion Representations only if one allows the cost of information to vary freely with the menu under consideration; if the cost of information is menu-independent, increasing in the Blackwell order, and non-constant, the resulting model cannot rationalize Persuasion Representations.⁴

Identification and Characterization of Bayesian Persuasion

Several studies derive conditions under which plausible choice data are consistent with rational choice under some information structure; see, for example, Lu (2016), Azrieli and Rehbeck (2022), Rehbeck (2023), or Doval et al. (2024). Such studies, as well as the those related to inattention discussed above, do not fully characterize sender-receiver communication models or analyze how the parameters of such models might be identified from choice data. In that sense, the closest work to this paper is Jakobsen (2021); as noted above, the contribution of that paper is quite different because it focuses on environments where Sender’s informational preferences are directly observed.

2 Persuasion Representations

2.1 Framework

States, Outcomes, Acts

The outcome generated by Receiver’s action depends on an exogenous state of the world. To capture this, I model actions as Anscombe-Aumann **acts** $f : \Omega \rightarrow \Delta X$ where:

- Ω is a finite set of states (with generic members ω),
- X is a finite set of outcomes (generic members x, y), and

⁴Informally, one can construct menus where (i) perfect information is chosen by Sender in a persuasion representation, but (ii) the stakes are so low that the utility difference between perfect information and no-information does not outweigh the (menu-independent) cost of acquiring perfect information.

- ΔX is the set of lotteries over X (generic members p, q); lottery p delivers outcome x with probability $p(x)$.

An act f may be written as a profile $f = (f_\omega)_{\omega \in \Omega}$ where $f_\omega := f(\omega)$. In state ω , act f returns lottery f_ω which in turn generates outcome x with probability $f_\omega(x)$. Acts are sometimes referred to as state-contingent lotteries. Constant acts $(p, \dots, p)$ are typically denoted p . Let F denote the set of all acts and $\mathcal{A}$ the set of all finite, nonempty subsets of F . A set $A \in \mathcal{A}$ serves as an action set, or **menu**, of acts available for Receiver to choose from. Singleton menus $\{f\}$ are typically denoted f .

Lotteries and acts are equipped with standard mixing operations. In particular, $\alpha p + (1 - \alpha)q$, where $\alpha \in [0, 1]$, denotes a lottery r such that $r(x) = \alpha p(x) + (1 - \alpha)q(x)$ for all x . This operation extends to acts by defining $\alpha f + (1 - \alpha)g$ as the act h such that, for all $\omega \in \Omega$, $h_\omega = \alpha f_\omega + (1 - \alpha)g_\omega$. These operations generalize to finite mixtures $\alpha_1 p^1 + \dots + \alpha_n p^n$ or $\alpha_1 f^1 + \dots + \alpha_n f^n$, where $\alpha_i \geq 0$ and $\alpha_1 + \dots + \alpha_n = 1$, in the natural way.

Experiments and Signals

A Blackwell **experiment** is a matrix with $|\Omega| = N$ rows, finitely many columns, and entries in $[0, 1]$ such that each row constitutes a probability distribution and no column consists entirely of zeros. Let $\mathcal{E}$ denote the set of all experiments, with generic members σ . Each column of an experiment represents a message that may be generated and each row a state-contingent distribution over messages. For example, the $N \times N$ identity matrix, denoted σ^* , associates a distinct message to each state and therefore represents **perfect information**.

An experiment can be expressed in terms of its columns. To do so, let S denote the set of all profiles $s = (s_\omega)_{\omega \in \Omega}$ of numbers $s_\omega \in [0, 1]$ such that $s_\omega \neq 0$ for at least one $\omega \in \Omega$. Elements of S , **signals**, represent columns that may be present in an experiment. Abusing notation slightly, ' $s \in \sigma$ ' indicates that s is a column of σ . As is easily verified, a matrix $[s^1, \dots, s^n]$ of signals qualifies as an experiment if and only if $s^1 + \dots + s^n = e$, where $e = (1, \dots, 1) \in S$ denotes an **uninformative signal** (or **uninformative experiment** since e itself qualifies as an experiment).

Signals and experiments yield additional mixture operations on acts. If $s \in S$, let $sf + (1 - s)g$ denote the act h such that $h_\omega = s_\omega f_\omega + (1 - s_\omega)g_\omega$; this operation is similar to the α -mixture of f and g defined above but allows potentially different weights s_ω to be applied in different states ω . More generally, if $\sigma = [s^1, \dots, s^n]$ is an experiment, $s^1 f^1 + \dots + s^n f^n$ denotes the act h such that $h_\omega = s_\omega^1 f_\omega^1 + \dots + s_\omega^n f_\omega^n$; h_ω is a well-defined lottery because σ is a Blackwell experiment and, thus, $s_\omega^1 + \dots + s_\omega^n = 1$.

Choice Primitives

My analysis involves different kinds of preference or choice data:

1. **Menu preferences** $\succsim$ over $\mathcal{A}$ where $A \succsim B$ means Receiver prefers committing to menu A over committing to menu B .
2. **Random choice** data $\rho = (\rho^A)_{A \in \mathcal{A}}$ where ρ^A is a probability distribution over A and $\rho^A(f) \in [0, 1]$ is the probability Receiver chooses f from A .
3. **State-contingent random choice** data $\lambda = (\lambda_\omega^A)_{\omega \in \Omega, A \in \mathcal{A}}$ where λ_ω^A is a probability distribution over A and $\lambda_\omega^A(f) \in [0, 1]$ is the probability Receiver f chooses from A in state ω .
4. **Choice correspondence** data $c : \mathcal{A} \rightarrow \mathcal{A}$ where $c(A) \subseteq A$ is the set of acts Receiver chooses from A with positive probability.

Note that c can be derived from ρ or λ but that no other pairs of primitives are nested in such a way. Some primitives might be more appealing or realistic depending on the application. For example, $\succsim$ may be better-suited to the behavioral interpretation of the model, as analysis of demand for commitment is a popular approach for studying dual-self interactions. In contrast, ρ or λ may be more appealing in settings where Sender and Receiver are distinct individuals and Sender's choice of information is not observable. Combining preferences $\succsim$ between menus with data ρ or λ on choices from menus can facilitate comparisons between models and provide insight into the agent's sophistication or naivete.

2.2 Persuasion Representations

This section defines Persuasion Representations for each type of primitive described above. Let $\Delta\Omega$ denote the set of probability distributions over Ω . The representations involve three parameters, (μ, u, v) , where $\mu \in \Delta\Omega$ is a prior and $u, v : X \rightarrow \mathbb{R}$ are utility indices for Receiver and an (unobserved) Sender, respectively.

For any $\hat{\mu} \in \Delta\Omega$ and state ω , let $\hat{\mu}_\omega$ denote the probability of state ω . Throughout, prior beliefs μ have full support and the utility indices u, v are non-constant. The indices u, v apply to lotteries as follows: if $p \in \Delta X$, then $u(p) := \sum_{x \in X} u(x)p(x)$ is Receiver's expected utility of p ; similarly, $v(p)$ denotes Sender's expected utility of p . For acts, Receiver's expected utility is given by $U : F \rightarrow \mathbb{R}$ where $U(f) := \sum_{\omega \in \Omega} u(f_\omega)\mu_\omega$. More generally, for any signal $s \in S$, let $U^s(f) := \sum_{\omega \in \Omega} u(f_\omega)s_\omega\mu_\omega$; this represents Receiver's expected utility conditional

on s .⁵ Note that $U^e = U$. Replacing u with v leads to functions $V, V^s : F \rightarrow \mathbb{R}$ representing Sender's expected utility conditional on signal realizations.

In a persuasion game at A , Receiver chooses an act from A after observing a signal provided by Sender. Sender provides information by freely selecting an experiment σ and committing to revealing its signal realization. Sender correctly forecasts Receiver's signal-contingent choices and constructs σ to maximize his own expected payoff. The agents share a common prior but typically differ in their preferences over outcomes generated by the chosen act(s); this is the source of conflict and it varies with the menu—some menus consist of acts for which Sender and Receiver largely agree on the best course of action, others less so.

To begin, it is useful to express the range of state-contingent lotteries that can be achieved by varying the information available to Receiver. At menu A , an experiment σ transforms into an act as follows. Fix a state ω . In this state, σ generates a distribution over signals ($s \in \sigma$ is generated with probability s_ω) and at every $s \in \sigma$ Receiver chooses a U^s -optimal act $f^s \in A$. In state ω , this act delivers a lottery f_ω^s . Thus, the state-contingent distribution over signals becomes a distribution over lotteries, which reduces to a single lottery in the natural way. Repeating this procedure for each state yields an **induced act**: a state-contingent lottery over outcomes generated by Receiver's choices under information σ at menu A .

The above procedure associates a unique induced act to an experiment σ if there is a unique U^s -optimal act $f^s \in A$ for each $s \in \sigma$. If there are multiple U^s -optimal acts for some s , different tie-breaking selections produce different induced acts. To capture the full range of possibilities, the set of induced acts at a menu A is defined as

$$F(A) := \left\{ \sum_{s \in \sigma} s f^s : \sigma \in \mathcal{E}, f^s \in \text{co}(A), U^s(f^s) \geq U^s(g) \forall g \in A \right\}, \quad (1)$$

where $\text{co}(A)$ is the convex hull of A . This set contains all induced acts generated by varying both σ and Receiver's tie-breaking behavior: if Receiver finds two or more acts optimal at s , then f^s is permitted to be any convex combination of those acts. Note that only parameters (μ, u) are needed to construct $F(A)$.

Since Sender anticipates Receiver's choices, Sender's choice of information at A is effectively a choice from $F(A)$. This leads to the following definition for menu preferences:

Definition 1. Parameters (μ, u, v) constitute a **Persuasion Representation** for $\succsim$ if u, v

⁵This holds because the Bayesian posterior of μ at s assigns probability $\frac{s_\omega \mu_\omega}{s \cdot \mu}$ to state ω , where $s \cdot \mu = \sum_{\omega' \in \Omega} s_{\omega'} \mu_{\omega'}$. Thus, expected utility conditional on s is $\sum_{\omega \in \Omega} \frac{u(f_\omega) s_\omega \mu_\omega}{s \cdot \mu}$. The function U^s multiplies this value by the constant $s \cdot \mu$ and therefore provides the same ordinal ranking of acts.

are non-constant, μ has full support, and the function $U : \mathcal{A} \rightarrow \mathbb{R}$ given by

$$U(A) := \max_{f \in \arg\max_{g \in F(A)} V(g)} U(f) \quad \text{subject to } f \in \arg\max_{g \in F(A)} V(g) \quad (2)$$

represents $\succsim$, where $U(f) = \sum_{\omega} u(f_{\omega})\mu_{\omega}$ and $V(f) = \sum_{\omega} v(f_{\omega})\mu_{\omega}$.

In a Persuasion Representation for $\succsim$, Sender correctly forecasts Receiver's signal-contingent choices from A and selects an information structure, hence an induced act $f \in F(A)$, that maximizes his own expected utility. When evaluating a menu A , Receiver correctly forecasts Sender's choice and assigns the value $U(f)$ to A , where f is the induced act associated with the chosen information structure. Thus, $U(A)$ is Receiver's ex ante expected utility from the interaction when the action set is A . This is well-defined by the following lemma⁶:

Lemma 1. *For all menus A and parameters (μ, u) , the set $F(A)$ is compact and convex.*

Implicitly, formula (2) makes two assumptions about tie-breaking behavior. First, the requirement that $f \in \arg\max_{g \in F(A)} V(g)$ means that if multiple acts maximize U^s at some s , Sender expects Receiver to select a V^s -maximal act among the U^s -maximizers. This is the standard ‘‘Sender-preferred’’ tie-breaking rule in the Bayesian Persuasion literature and it ensures existence of a Sender-optimal information structure at every A . On a technical level, it emerges from maximization over $F(A)$ because that set includes all possible induced acts that come about by varying both information and Receiver's tie-breaking selections.

Second, the ‘‘max’’ in (2) means that if Sender finds multiple information structures optimal at A , he selects from such structures a Receiver-optimal one. Formally, let

$$V^{\sigma}(A) := \max_{s \in \sigma} \sum_{f \in A} V^s(f^s) \quad \text{subject to } f^s \in \arg\max_{f \in A} U^s(f)$$

and

$$U^{\sigma}(A) := \sum_{s \in \sigma} U^s(f^s) \quad \text{where } f^s \in \arg\max_{f \in A} U^s(f).$$

These functions capture Sender's and Receiver's value, respectively, of information σ at menu A . That is, if Sender chooses σ at A , he expects payoff $V^{\sigma}(A)$ and Receiver expects $U^{\sigma}(A)$. Clearly, $V^{\sigma}(A)$ incorporates the Sender-preferred tie-breaking rule described above. The ‘‘Receiver-preferred’’ rule implied by (2) means that if multiple experiments σ maximize $V^{\sigma}(A)$, Sender breaks the tie by maximizing $U^{\sigma}(A)$. The following theorem makes this

⁶More precisely, Lemma 1 implies $\arg\max_{g \in F(A)} V(g)$ is nonempty; it is also compact because it consists of the maximizers of the continuous function V over the compact set $F(A)$. Thus, there exists a U -maximal act in $\arg\max_{g \in F(A)} V(g)$.

explicit by re-expressing Persuasion Representations for $\succsim$ in a more familiar form; I omit the straightforward proof.

Theorem 1. *Parameters (μ, u, v) , where u, v are non-constant and μ has full support, constitute a Persuasion Representation for $\succsim$ if and only if, for all $A \in \mathcal{A}$, the function U defined by (2) above satisfies*

$$U(A) = \max_{\sigma} U^{\sigma}(A) \text{ subject to } \sigma \in \operatorname{argmax}_{\sigma' \in \mathcal{E}} V^{\sigma'}(A). \quad (3)$$

By Theorem 1, formulas (2) and (3) coincide; (3) more directly expresses Receiver's payoff as the result of Sender's information design problem but, as we shall see, (2) facilitates comparisons to related models and is useful for deriving various results.

Persuasion Representations for ρ , λ , and c do not involve Receiver's ex ante value $U(A)$ but rather the actual choices from A stemming from Sender's information structure. Let $\mathcal{E}^*(A) \subseteq \mathcal{E}$ denote the set of solutions to the maximization problem (3). An experiment $\sigma \in \mathcal{E}^*(A)$ is **A -minimal** if there is no $\sigma' \in \mathcal{E}^*(A)$ such that σ' is a garbling of σ and $\sigma' \neq \sigma$. Finally, given A , a **behavioral strategy** is a profile $\beta^A = (\beta^{A,s})_{s \in S}$ such that $\beta^{A,s} \in \Delta A$ for all $s \in S$; that is, $\beta^{A,s}$ is a distribution of choices from A at signal s .

Definition 2. Parameters (μ, u, v) , where u, v are non-constant and μ has full support, constitute a **Persuasion Representation** for ρ if for every A there is a behavioral strategy β^A and A -minimal experiment $\sigma \in \mathcal{E}^*(A)$ such that

(i) for all $s \in S$,

$$\operatorname{supp}(\beta^{A,s}) = \operatorname{argmax}_f V^s(f) \text{ subject to } f \in \operatorname{argmax}_{g \in A} U^s(g), \text{ and}$$

(ii) for all $f \in A$, $\rho^A(f) = \sum_{s \in \sigma} (s \cdot \mu) \beta^{A,s}(f)$.

Informally, parameters (μ, u, v) constitute a Persuasion Representation for ρ if, for every A , choices frequencies ρ^A coincide with those generated by Sender's chosen experiment and Receiver's signal-contingent choices for that experiment; these choices must be optimal given parameters (μ, u, v) . Part (i) of Definition 2 requires that, for every $s \in \sigma$, Receiver chooses the act(s) that are consistent with his own optimization and the Sender-preferred tie-breaking criterion. Part (ii) requires that the observed probability of choosing f from A , $\rho^A(f)$, coincides with the total probability of choosing f given σ and β^A ; in particular, $s \cdot \mu$ (the dot product) is the total probability of generating signal $s \in \sigma$ under prior μ , and $\beta^{A,s}(f)$ is

the probability of choosing f conditional on signal s .

Since the chosen experiment σ is a member of $\mathcal{E}^*(A)$, it satisfies the Receiver-preferred tie-breaking criterion described above. It is also required to be A -minimal. While neither agent's payoffs or incentives are affected by the A -minimality requirement, it implies that if e (no information) is both Sender- and Receiver-optimal, then e is chosen by Sender. This simplifies the statements and proofs of several results.

Definition 3. Parameters (μ, u, v) , where u, v are non-constant and μ has full support, constitute a **Persuasion Representation** for λ if for every A there is a behavioral strategy β^A and A -minimal experiment $\sigma \in \mathcal{E}^*(A)$ such that

(i) for all $s \in S$,

$$\text{supp}(\beta^{A,s}) = \underset{f}{\operatorname{argmax}} V^s(f) \text{ subject to } f \in \underset{g \in A}{\operatorname{argmax}} U^s(g), \text{ and}$$

(ii) for all $f \in A$ and $\omega \in \Omega$, $\lambda_\omega^A(f) = \sum_{s \in \sigma} s_\omega \beta^{A,s}(f)$.

The definition of a Persuasion Representation for λ is nearly identical to that of ρ . The only difference is that *state-contingent*, as opposed to *total*, choice frequencies must agree with those generated by the persuasion game with parameters (μ, u, v) ; condition (ii) reflects this.

Definition 4. Parameters (μ, u, v) constitute a **Persuasion Representation** for c if there exists a random choice rule $\rho = (\rho^A)_{A \in \mathcal{A}}$ such that (μ, u, v) constitute a Persuasion Representation for ρ and, for all $A \in \mathcal{A}$, $c(A) = \text{supp}(\rho^A)$.

Intuitively, parameters (μ, u, v) constitute a Persuasion Representation for c if, for every A , $c(A)$ coincides with the support of ρ^A where ρ has a Persuasion Representation with parameters (μ, u, v) . Alternatively, one could modify Definition 2 by replacing ρ with c and condition (ii) with $c(A) = \bigcup_{s \in \sigma} \text{supp}(\beta^{A,s})$. Similarly, one could modify Definition 3 by replacing λ with c and condition (ii) with $c(A) = \bigcup_{s \in \sigma, \omega \in \Omega} \text{supp}(\lambda_\omega^A)$. Either way, $c(A)$ contains all acts in A that are chosen with positive probability in at least one state; any additional frequency information is discarded.

Before concluding this section with some comments on Persuasion Representations, a brief discussion of condition (i) in Definitions 2–4 is in order. Condition (i) requires that, at each signal realization, every act in A that survives the Sender-preferred tie-breaking rule is chosen with positive probability. There are other ways of refining choices, but the advantage of (i), together with the A -minimality condition, is that special cases of the model reduce

to familiar representations in decision theory. For example, if $u \approx -v$, Sender chooses no-information and Receiver simply chooses based on his prior. By (i), then, $c(A)$ consists of *all* prior-optimal acts in A , as in standard choice models, rather than some arbitrary subset of optimal acts. For the purposes of this paper, such well-behaved special cases are not strictly necessary but allow a cleaner exposition and simpler proofs; this is the main reason for imposing tie-breaking conventions beyond the standard Sender-preferred criterion.

Persuasion Representations: Comments

Strotz vs Persuasion. The representation of Definition 1 is similar to a Strotzian representation (Strotz, 1955). In the notation of this paper, such representations take the form

$$W_{Strotz}(A) := \max U(f) \text{ subject to } f \in \operatorname{argmax}_{g \in A} V(g). \quad (4)$$

A standard interpretation is that one self, with utility U , anticipates their future-self choosing from A via maximization of V . The Persuasion Representation of Definition 1 has a similar structure, with one key difference: the other self chooses an act not from A , but from $F(A)$. This captures the idea that the other self does not exert full control over future choice; instead, he influences choice through Bayesian belief distortion.

The Role of Sophistication. Most decision-theoretic models of temptation are defined and characterized in terms of ex-ante menu preferences $\succsim$. Crucially, such representations involve forecasting of future behavior. For example, Strotzian representations of the form (4) require the ex-ante agent, with utility U , to anticipate future choices made with V ; in this sense, the function V in the representation reflects the initial self's *beliefs* about future tastes. Similarly, Persuasion Representations of $\succsim$ require Receiver to forecast Sender's choice of information; thus, the index v reflects Receiver's beliefs about Sender's tastes.

In contrast, Persuasion Representations of ρ , λ , and c do not involve such forecasting or beliefs—they merely express Receiver's choices resulting from the interaction. Consequently, the index v in such representations reflects the true preferences of Sender. Comparisons between $\succsim$ and (say) λ , therefore, can detect whether the ex-ante agent is sophisticated (holds correct beliefs about v) or naive (optimistic or pessimistic about v); see section 4.3.

Timing and Interpretation of Information Choice. As explained in the introduction, information choice in the behavioral interpretation amounts to selective attention and/or belief distortion at the time of consumption (act) choice. Since action choices maximize u and information is chosen to maximize v in anticipation of action choices, it is as if two compet-

ing forces exist at the time of choice: the agent seeks to maximize u but allows temptation (preferences v) to influence his beliefs. When ranking menus (preferences $\succsim$), the agent does not experience temptation—the value of a menu is the expected u -utility from the interaction. Thus, in the “cold” (ex ante) state, there is no belief distortion, whereas in the “hot” (action choice) state, choices maximize u subject to beliefs that are influenced by v . This captures the idea that temptation is transient—it occurs when making final consumption decisions—and that it may compel individuals to engage in selective attention or motivated reasoning to “justify” tempting choices.

The Cost of Temptation. In the Gul and Pesendorfer (2001) model, the decision maker may resist temptation but suffer a cost of self-control from doing so. In the notation of this paper, their representation may be written as

$$W_{GP}(A) = \max_{f \in A} U(f) + V(f) - \max_{g \in A} V(g).$$

This means the decision maker chooses f to maximize U minus $\max_{g \in A} V(g) - V(f)$, with the latter component—the gap in V -utility between f and any V -optimal act in A —representing the degree of self-control exerted by the agent. Depending on A , the agent may choose a U -maximal option, a V -maximal option or, more typically, a “compromise” option.

In my model, the agent behaves stochastically: at a typical menu A , Receiver sometimes succumbs to temptation and sometimes not—the distribution of choices is, in effect, the “compromise” option. There is no cost of self control, but Receiver effectively suffers a cost of *temptation* given by $U^{\sigma^*}(A) - U(A)$ (equivalently, $U^{\sigma^*}(A) - U^{\sigma}(A)$, where $\sigma \in \mathcal{E}^*(A)$). Intuitively, perfect information σ^* is freely available but temptation makes the agent garble or distort it to some $\sigma \in \mathcal{E}^*(A)$, resulting in lower expected U -utility.

Cheap Talk vs Persuasion. Other communication models, like the cheap-talk framework of Crawford and Sobel (1982), are potentially amenable to decision-theoretic analysis and the self-signaling interpretation studied in this paper. Moreover, cheap talk does not involve the oft-criticized assumption of Sender commitment power that lies at the core of Bayesian persuasion models. Cheap talk interactions, however, have a more severe equilibrium-selection problem: equilibria are sustained by off-path beliefs about Sender’s communication strategy, giving rise to qualitatively-distinct equilibria. The commitment assumption of Bayesian persuasion eliminates the role of off-path beliefs, leaving only particular tie-breaking issues that are fairly commonplace in decision-theoretic models. The commitment assumption is also more appealing in the behavioral interpretation of the model, where signals are the result of

deliberate (in)attention or belief distortion: the agent sees what he chooses to see.

3 Identification

This section establishes identification results for each type of choice primitive; section 3.1 provides a proof sketch and illustrates methods to elicit parameter values from the primitives.

Theorem 2. *If (μ, u, v) and (μ', u', v') are Persuasion Representations of a preference $\succsim$ on $\mathcal{A}$, then $u \approx u'$, $v \approx v'$, and $\mu = \mu'$.*

Theorem 2 states that Persuasion Representations of menu preferences are unique: if $\succsim$ has a Persuasion Representation, there is a unique prior μ and unique (up to positive affine transformation) utility indices u, v for which (μ, u, v) constitute a Persuasion Representation of $\succsim$. Thus, menu preferences $\succsim$, alone, are sufficient to identify all parameters. An analogous result holds for representations of random choice rules:

Theorem 3. *If (μ, u, v) and (μ', u', v') are Persuasion Representations of a random choice rule ρ , then $u \approx u'$, $v \approx v'$, and $\mu = \mu'$.*

A slightly weaker result holds for λ and c :

Theorem 4. *If (μ, u, v) and (μ', u', v') are Persuasion Representations of either a state-contingent random choice rule λ or a choice correspondence c , then: (i) $u \approx u'$ and $v \approx v'$, and (ii) if $u \not\approx v$, then $\mu = \mu'$.*

Theorem 4 states that u and v are identified from λ or c but that μ is identified by either primitive only if $u \not\approx v$; that is, if there is some conflict between Sender and Receiver. Since c is nested by ρ and λ , this means that outside the hairline case $u \approx v$, both ρ and λ contain more information than is actually required to uniquely identify all three parameters. For this reason, most subsequent results in the paper are established only for $\succsim$ and c .

To see what goes wrong when $u \approx v$, observe that such preferences make Sender choose perfect information at every menu A .⁷ Consequently, in state ω , Receiver chooses precisely those acts $f \in A$ for which $u(f_\omega) \geq u(g_\omega)$ for all $g \in A$. These choices depend on ω but not the probability μ_ω with which ω realizes. Thus, $c(A)$ does not depend on μ . Similarly, when $u \approx v$, state-contingent choices λ_ω do not reveal anything about μ unless one makes

⁷More precisely, perfect information is Sender-optimal at every A ; Sender may choose coarser information in some menus due to the A -minimality requirement, but this does not affect the argument.



Figure 1: Identifying u and v .

arbitrary assumptions about how tie-breaking varies with μ .⁸

The significance of Theorems 2–4 is that full identification can be achieved using any of the four types of preference or choice data, subject to the restriction $u \not\approx v$ for λ and c . Since (μ, u, v) are identified, so is Sender’s choice of information—at any A , the optimal information structure(s) can be derived from (μ, u, v) . Applying the behavioral interpretation, this means unobserved attention or belief distortion can be uniquely identified from standard choice data; to my knowledge, this is the first general result on such identification in dual-selves signaling models. Interestingly, ρ and λ provide more information than is actually required: simply knowing which acts are chosen with positive probability (as recorded by c) is enough to deduce (μ, u, v) , making measurement error potentially less problematic.

3.1 Eliciting Parameters

Given lotteries p and q , a menu A is a pq -**menu** if, for every $f \in A$, there exists $s^f \in S$ such that $f = s^f p + (1 - s^f)q$. A menu A is a pq -**bet** if there is a nonempty $E \subsetneq \Omega$ such that $A = \{pEq, qEp\}$, where $h = p'Eq'$ satisfies $h_\omega = p'$ for $\omega \in E$ and $h_\omega = q'$ for $\omega \notin E$.

Menu Preferences

Receiver’s parameters (μ, u) are easily identified by considering the restriction of $\succsim$ to singleton menus. In particular, $\{f\} \succsim \{g\}$ if and only if $U(f) \geq U(g)$, so the Anscombe-Aumann theorem applies and μ and u are identified.

To elicit v , the key is to determine which pairs (p, q) of lotteries are ranked the same way by u and v . Figure 1b illustrates the idea: for any p , the index u gives a linear indifference curve through p . To identify v , it is enough to determine the slope of v ’s indifference curve through p and the direction of increasing utility. If $A_{pq} = \{pEq, qEp\}$ is a pq -bet where

⁸For example, one could specify a menu A where, in state ω_1 , there is a tie between two acts and Receiver chooses one of them with probability μ_{ω_1} . If Receiver employs such a tie-breaking criterion and the analyst knows which particular act f is chosen with probability μ_{ω_1} , then $\lambda_{\omega_1}^A(f)$ coincides with μ_{ω_1} . However, there is no reason to suspect that Receiver would correlate tie-breaking selections with μ or, if he did, that the analyst would know the correlation structure.




Figure 2: Identifying μ from c .

$u(p) > u(q)$, Sender either agrees with the ranking ($v(p) \geq v(q)$) or disagrees ($v(p) < v(q)$). If Sender agrees, he chooses perfect information and so $A_{pq} \sim p$; otherwise, Sender disagrees and chooses no information, yielding $p \succ A_{pq}$. Thus, fixing p and eliciting all q such that $A_{pq} \sim p$ reveals the indifference curve for v through p and the direction of increasing utility.

Choice Distributions & Correspondences

For ρ , λ and c , elicitation is slightly more involved. The first step is to identify u by analyzing choices from menus $\{p, q\}$ of constant acts. Fixing p , we have $c(\{p, q\}) = p$ if $u(p) > u(q)$. If $u(p) \geq u(q)$, the representation breaks the tie in favor of v . Thus, $p \in c(\{p, q\})$ if and only if $u(p) > u(q)$ or both $u(p) = u(q)$ and $v(p) \geq v(q)$. As illustrated in Figure 1a, fixing p and eliciting all q such that $p \in c(\{p, q\})$ reveals the lower contour set of p for u ; in particular, the closure of the set of all such q is the weak lower contour of u through p . Given u , one can identify v using similar techniques developed for menu preferences.⁹

To see how μ may be identified from c , consider the two-state case and suppose $u \not\approx v$ and $u \not\approx -v$ (the proof provided in the appendix is different and only requires $u \not\approx v$). With such preferences, there are lotteries p, q, r such that the menu $A = \{(p, q), (r, r)\}$ gives rise to the value functions depicted by Figure 2; the horizontal axis' contain all possible posterior beliefs (ordered by the weight assigned to state 1) and the values are those induced by Receiver's choices from A at those beliefs.¹⁰ At μ^* , Receiver is indifferent between the two acts, creating a discontinuity in Sender's value function. If prior beliefs satisfy $\mu_1 \leq \mu_1^*$,

⁹To circumvent tie-breaking issues for the case of two states and a uniform prior, the proof uses the more general class of pq -menus rather than pq -bets; see the appendix for details.

¹⁰The desired lotteries exist because when $u \not\approx v$ and $u \not\approx -v$, one may fix indifference curves (utility levels) for one agent and move along them to set utility levels for the other; see Lemma 4 in the appendix.

Sender chooses no-information; consequently, $c(A) = \{(r, r)\}$. If instead $\mu_1 > \mu_1^*$, Sender maximizes his payoff by choosing an experiment yielding two posteriors: one at μ^* , the other at $\hat{\mu}_1 = 1$. Consequently, $c(A) = A$ since Receiver chooses (r, r) at μ^* and (p, q) at $\hat{\mu}_1 = 1$. Thus, $c(A)$ indicates whether $\mu_1 \leq \mu_1^*$ or $\mu_1 > \mu_1^*$, providing objective information about μ since μ^* is pinned down by u . It follows that μ can be identified by examining $c(A)$ for all binary menus A —for example, by moving r along its v -indifference curve and thereby perturbing only $u(r)$ and, thus, μ_1^* .

4 Comparative Statics from Receiver's Perspective

This section develops comparative static results regarding the degree of conflict between Sender and Receiver. In contrast to most of the persuasion literature, the focus is on how the interaction affects Receiver's (not Sender's) choices and welfare.

The characterizations take two forms: (i) how Receiver's ex ante value and subsequent choices vary with increased flexibility, and (ii) how such values and choices vary with additional public information. The results thereby establish tight links between Receiver's value of flexibility and of information in persuasion models. While analysis of (i) involves standard comparisons between $\subseteq$ -comparable menus, analysis of (ii) involves a new operator on menus that simulates public information. For any A and σ , let

$$\sigma A := \left\{ \sum_{s \in \sigma} s f^s : f^s \in A \right\}.$$

The menu σA simulates an environment where Receiver chooses from A but is able to condition this choice on the realization from σ in addition to the signal generated by Sender. Sender recognizes this but cannot correlate realizations from his chosen experiment with those of σ . The example below provides concrete illustrations.

Example 1. Suppose there are two states and let $A = \{f, g\}$. If $\sigma = \sigma^*$ (the identity matrix), then σ perfectly reveals the state and $\sigma A = \{(f_1, f_2), (g_1, f_2), (f_1, g_2), (g_1, g_2)\}$. Suppose $u(f_1) > u(g_1)$ and $u(g_2) > u(f_2)$; this means neither act dominates the other and, in particular, that f is preferred in state 1 and g is preferred in state 2. At σA , then, Receiver chooses (f_1, g_2) regardless of the information Sender provides.

If instead σ is a noisy structure $\sigma = [s, t]$, we have $\sigma A = \{sf + tf, sf + tg, sg + tf, sg + tg\} = \{f, sf + tg, sg + tf, g\}$. If Receiver prefers f at s and g at t , his prior-optimal act in σA is $sf + tg$; this represents the average state-contingent lottery for Receiver if Sender provides no additional information, thereby forming a lower bound on Receiver's welfare in the persuasion

game with public information σ . ◆

When Sender and Receiver are different people, σA represents an environment where Receiver chooses from A but has access to additional public information σ ; alternatively, σ represents a lower bound on Sender's choice of information at A . In the behavioral interpretation, however, other interpretations of σA are more appropriate. For example, σ might represent information that is too salient to misinterpret or ignore; features of the choice environment like the framing of information can affect such salience. More generally, σA represents a situation where Receiver can commit to conditioning choices on realizations from σ in addition to whatever signal their “other self” provides. In addition to psychological mechanisms like willpower, such commitment might be achieved through delegation. In the restaurant example of the introduction, for instance, Receiver might ask the server to acquire a signal about the salad and to enter an order conditional on the signal; acts of the form $sf + tg$ in Example 1 above provide just that kind of delegation. Note that Sender chooses information knowing such delegation opportunities are available—the menu σA is known to both agents. Thus, σA might represent a situation where the server presents Receiver with the menu A *and* an offer to delegate the choice; Sender's choice takes that offer into account.

4.1 The Value of Flexibility and Information

This section characterizes special cases of the model involving standard rationality postulates. Under what circumstances does Receiver benefit from increased flexibility or from public information, and how is this reflected in choice data c ?

Proposition 1. *Suppose (μ, u, v) represents $\succsim$ and c . Then $u \approx v$ or $u \approx -v$ if and only if any of the following conditions hold:*

- | | |
|---|--|
| <p>(i) $\succsim$ satisfies <i>Preference for Flexibility</i>:
 $A \supseteq B$ implies $A \succsim B$.</p> | <p>(ii) c satisfies <i>Sen's condition α</i>: $A \subseteq B$ implies $c(B) \cap A \subseteq c(A)$.</p> |
| <p>(iii) $\succsim$ satisfies <i>Preference for Information</i>: for all σ and A, $\sigma A \succsim A$.</p> | <p>(iv) c satisfies <i>Informational Sen's α</i>:
 $c(\sigma A) \cap A \subseteq c(A)$.</p> |

Proposition 1 states that in Persuasion Representations, Receiver's preferences and choices satisfy standard rationality postulates if and only if one of two extreme cases holds: the conflict between Sender and Receiver is either non-existent ($u \approx v$) or total ($u \approx -v$). To aid the discussion, note that if $u \approx v$, Sender chooses perfect information at all menus,

reducing the representation of $\succsim$ to $\bar{U}(A) := \sum_{\omega \in \Omega} \max_{f \in A} u(f_\omega) \mu_\omega$ and that of c to $\bar{c}(A) = \bigcup_{\omega \in \Omega} \operatorname{argmax}_{f \in A} u(f_\omega)$. If $u \approx -v$, Sender chooses e (no information) at all menus, reducing the representations to $\underline{U}(A) := \max_{f \in A} U(f)$ and $\underline{c}(A) = \operatorname{argmax}_{f \in A} U(f)$.

Preference for Flexibility is the key axiom of Kreps (1979). In dual-self models such as Gul and Pesendorfer (2001) or Strotz (1955), this axiom typically characterizes the case of no conflict (the two selves have the same utility index). Here, Preference for Flexibility permits the opposite case of total conflict. Intuitively, this is so because the “other self” (Sender) influences choice only via belief distortion. When $u \approx -v$, Sender chooses no-information at all menus, reducing Receiver’s behavior to that of standard expected utility maximization, which satisfies Preference for Flexibility.

Preference for Information requires that public information never harms Receiver. Note that since $A \subseteq \sigma A$, this condition is implied by Preference for Flexibility; Proposition 1 establishes that for Persuasion Representations, it is in fact equivalent to Preference for Flexibility. The condition is satisfied by the $u \approx -v$ case since this reduces the representation to $\underline{U}$ (standard expected utility maximization) and, by Blackwell (1951,1953), the value of every decision problem increases with the availability of information. The $u \approx v$ case also satisfies Preference for Information since the representation reduces to $\bar{U}$ (Receiver’s value under perfect information) making $\sigma A \sim A$ for all A . It is less obvious that menus A and experiments σ satisfying $A \succ \sigma A$ exist when $u \not\approx v$ and $u \not\approx -v$. As illustrated in Figure 3, the idea is that Sender may choose to provide nontrivial information $\hat{\sigma}$ at A but—for some public σ less valuable to Receiver than $\hat{\sigma}$ —be unwilling to provide additional information.

Conditions (ii) and (iv) provide choice-correspondence analogues of Preference for Flexibility and Preference for Information, respectively. Condition (ii), Sen’s α , is a basic property of rational choice; since the menus referenced by the axiom are $\subseteq$ -comparable, the axiom constrains choice patterns that can arise under increased flexibility. Condition (iv) weakens Sen’s α by requiring the original axiom to hold only when the increased flexibility is induced by the availability of public information; nonetheless, the Proposition establishes that it is equivalent to Sen’s α in Persuasion Representations.

Simple axioms characterize the $u \approx v$ and $u \approx -v$ cases. For any menu A and state ω , let $A_\omega := \{f_\omega : f \in A\}$. An experiment σ is **interior** if $0 < s_\omega < 1$ for all $s \in \sigma$ and $\omega \in \Omega$.

Proposition 2. *Suppose (μ, u, v) represents $\succsim$ and c . Then $u \approx v$ if and only if any of the following conditions hold:*




Figure 3: Illustration of Proposition 1(iii) when $u \not\approx v$ and $u \not\approx -v$. When no public information is provided, Sender chooses information yielding posteriors at $\hat{\mu}'$ and $\hat{\mu}_1 = 1$; for Receiver, the corresponding value is given by the upper red dot. If public information generating posteriors at $\hat{\mu}'$ and $\hat{\mu}''$ is provided, Sender chooses not to provide any additional information: doing so can only decrease payoffs (strictly so at posterior $\hat{\mu}'$). Consequently, Receiver's payoff decreases.

- | | |
|---|---|
| <p>(i) $\succsim$ satisfies <i>Preference for Statewise Flexibility</i>: $A_\omega \supseteq B_\omega$ for all ω implies $A \succsim B$.</p> | <p>(ii) c is <i>monotone in statewise flexibility</i>: $A_\omega \supseteq B_\omega$ for all ω implies $c(A) \subseteq c(A \cup B)$.</p> |
| <p>(iii) $\succsim$ is <i>indifferent to information</i>: for all A and σ, $A \sim \sigma A$.</p> | <p>(iv) c is <i>invariant to imperfect information</i>: for all A and interior σ, $c(\sigma A) = c(A)$.</p> |

Proposition 2 characterizes the $u \approx v$ case in terms Receiver's preferences for, and choices under, increased flexibility or information. Condition (i) requires Receiver to prefer menus that offer more flexibility (more potential lotteries) in each state of the world; this condition is implied by, but not equivalent to, Preference for Flexibility. The analogous requirement for c , condition (ii), is that any act chosen at A is also chosen after expanding A in a way that does not expand the set of possible lotteries in any state. Conditions (iii) and (iv) capture the idea that Receiver's welfare and choices are not impacted by public information if $u \approx v$; intuitively, these properties are consistent with $u \approx v$ because Sender chooses perfect information when there is no conflict with Receiver.

Proposition 3. Suppose (μ, u, v) represents $\succsim$ and c . Then $u \approx -v$ if and only if any of the following conditions hold:

- (i) $\succsim$ is Independent of Irrelevant Alternatives: $A \succsim B$ implies $A \sim A \cup B$.
- (ii) c satisfies WARP: $c(A) \cap B \neq \emptyset$ implies $c(B) \cap A \subseteq c(A)$.
- (iii) $\succsim$ satisfies Preference for Information and $\sigma A \succ A$ for some A and σ .
- (iv) c satisfies Informational Sen's α and $c(\sigma A) \neq c(A)$ for some A and interior σ .

Parts (i) and (ii) of Proposition 3 establish that familiar axioms—IIA for $\succsim$, WARP for c —characterize the $u \approx -v$ case. Traditionally, these axioms ensure $\succsim$ and c can be represented by maximization of a utility function. In Persuasion Representations, this is consistent with $u \approx -v$ because Sender's resulting choice of no-information at all menus reduces Receiver's behavior to standard expected utility maximization. Parts (iii) and (iv) characterize $u \approx -v$ in terms of behavior under public information. In particular, satisfying Preference for Information or Informational Sen's α non-trivially is necessary and sufficient for $u \approx -v$ in Persuasion Representations.

For condition (iv) in each of Propositions 1–3, some care is needed in the interpretation of $c(\sigma A)$. Recall that σA merely *simulates* an environment where the choice set is A and Receiver observes a signal generated by σ in addition to that generated by Sender. To understand the difference, consider Example 1 above with public information $\sigma = \sigma^*$ (perfect information). In that example, Receiver only chooses the act (f_1, g_2) from σA . The interpretation is that perfect information would lead Receiver to choose f in state 1 and g in state 2. Thus, both f and g would be chosen from A with perfect public information, but only (f_1, g_2) is chosen from the menu σA .

To translate $c(\sigma A)$ into a statement about choices from A under public information σ , two additional definitions are needed. For any A and σ , let $c_A(\sigma A) := \{f \in A : \exists \sum_{s \in \sigma} s f^s \in c(\sigma A) \text{ and } s \in \sigma \text{ such that } f^s = f\}$ denote the **projection** of $c(\sigma A)$ to A . For an experiment σ , let $\text{supp}(\sigma) := \{\frac{s}{\|s\|} : s \in \sigma\}$ denote its **support**. With this notation in place, the above characterizations involving choice correspondence data under public information can be reformulated as follows:

Proposition 4. *Suppose (μ, u, v) represents c . Then:*

- (i) $u \approx v$ or $u \approx -v$ if and only if c is expansive in signals: $\text{supp}(\sigma) \subseteq \text{supp}(\sigma')$ implies $c_A(\sigma A) \subseteq c_A(\sigma' A)$ for all A .
- (ii) $u \approx v$ if and only if, for all A and σ , $c(A) = c_A(\sigma A)$.
- (iii) $u \approx -v$ if and only if c is expansive in signals and $c(A) \neq c_A(\sigma A)$ for some A and σ .

Proposition 4 characterizes the extreme cases of Sender-Receiver conflict in terms of choices from A (not σA) under public information σ . Part (i) states that there is either no conflict or total conflict if and only if expanding the set of posteriors induced by the public structure results in a larger set of acts being chosen from A . Parts (ii) and (iii) differ in whether this expansiveness property holds trivially or non-trivially: there is no conflict if and only if choices are unresponsive to public information, and there is total conflict if choice data is non-trivially expansive in public information.

4.2 Measures of Conflict

The characterizations in the previous section establish key relationships between Receiver's ex-ante value of (and ex-post choice under) flexibility and public information but are limited to extreme cases ($u \approx v$ or $u \approx -v$) regarding the conflict between the agents. This section develops finer comparisons between Sender and Receiver.

Definition 5. Let $u, v, \dot{v}$ be utility indices. Then $\dot{v}$ is **more u -aligned than** v (and v is **less u -aligned than** $\dot{v}$) if either $u \approx -v$ or $\dot{v} \approx \alpha u + (1 - \alpha)v$ for some $\alpha \in [0, 1]$.

Definition 5 is the key definition of Ahn et al. (2019). The idea is that if, starting from v , one forms a mixture $\dot{v} \approx \alpha u + (1 - \alpha)v$, then $\dot{v}$ more closely resembles u than v does. Consequently, holding u fixed, a Sender with index $\dot{v}$ “disagrees less” with Receiver than a Sender with index v , softening the conflict between the agents.

The aim of this section is to characterize more- or less-aligned utilities in terms of Receiver's value of, and choice under, increased flexibility or public information. To begin, consider the following comparative notions for menu preferences:

Definition 6. Let $\succsim$ and $\dot{\succsim}$ denote preferences on $\mathcal{A}$.

- (i) $\dot{\succsim}$ **values flexibility more than** $\succsim$ if, for all $A \supseteq B$, $A \succ B$ implies $A \dot{\succ} B$.
- (ii) $\dot{\succsim}$ **values information more than** $\succsim$ if, for all A and σ , $\sigma A \succ A$ implies $\sigma A \dot{\succ} A$.

The idea of Definition 6 is that one agent (represented by $\dot{\succsim}$) values flexibility more than another (represented by $\succsim$) if there are more instances where he strictly prefers an expanded option set. Similarly, he values information more than the other agent if there are more instances where he strictly benefits from public information. Since $\sigma A \supseteq A$, an agent who values flexibility more than another necessarily values information more than the other.

Proposition 5. Let (μ, u, v) and $(\mu, u, \dot{v})$ represent $\succsim$ and $\dot{\succsim}$, respectively.




Figure 4: Illustration of Proposition 5(i). When no public information is provided, both Sender preferences (green and blue curves) result in no additional information provision; consequently, Receiver's payoff is the lower of the two red dots. Consider a public structure generating posteriors at $\hat{\mu}'$ and $\hat{\mu}_1 = 1$. The more-aligned Sender utility (green curve) compels Sender to provide additional information, yielding posteriors at $\hat{\mu}_1 = 0$, $\hat{\mu}''$, and $\hat{\mu}_1 = 1$, increasing Receiver's payoff to the higher red dot. However, the less-aligned Sender utility (blue curve) results in no additional information provision. Thus, less-aligned preferences do not increase Receiver's value of public information.

- (i) If $\dot{\succsim}$ values information more than $\succsim$, then $\dot{v}$ is less u -aligned than v is. However, the converse does not hold.
- (ii) If $\dot{\succsim}$ values flexibility more than $\succsim$, then $\dot{v}$ is more u -aligned than v is and $\dot{\succsim}$ values information more than $\succsim$ does; consequently, $\dot{v} \approx v$.

Part (i) of Proposition 5 states that if Receiver's value of public information increases, then there is greater conflict between the agents. Intuitively, the change in value reflects Receiver's expectation that Sender provides less information when there is greater conflict. However, the converse does not hold: less-aligned preferences do not guarantee that Receiver values information more in the sense of Definition 6. Figure 4 provides an example where $v = \frac{1}{2}u + \frac{1}{2}\dot{v}$. Proposition 6 below provides a finer condition that fully characterizes the degree of preference alignment in terms of Receiver's value of public information.

Part (ii) of Proposition 5 establishes that if Receiver's value of flexibility increases, then utilities become more aligned. However, as noted above, increased value of flexibility implies increased value of information; by part (i), then, utilities also become *less* aligned. Since $\dot{v}$ is both more and less u -aligned than v , it follows that $\dot{v} \approx v$. Thus, Definition 6(i) is too strong to characterize finer changes to the degree of conflict. The next result provides a remedy.

Proposition 6. *Let (μ, u, v) represent $\succsim$ and $(\mu, u, \dot{v})$ represent $\dot{\succsim}$. Then $\dot{v}$ is less u -aligned than v is if and only if any of the following conditions hold:*

- (i) *If $f \succsim A$, then $f \dot{\succsim} A$.*
- (ii) *If $f \sim \sigma A \succ A$ and $f \in \sigma A$, then $\sigma A \dot{\succ} A$.*

Proposition 6 modifies Definition 6 to provide full characterizations of utility alignment in Persuasion Representations. In particular, part (i) establishes that utilities are less aligned if there are more instances where full commitment (to a specific act) is preferred to a given menu A .¹¹ This captures the intuition that increased conflict lowers Receiver’s expected payoff at A , enlarging the set of commitment options f that are preferred to playing the game with the full set A . Put differently, the result states that the concept of more-aligned utility fully characterizes the comparative statics of Receiver’s welfare in persuasion games: his expected payoff increases at all menus if and only if utilities become more-aligned.

Part (ii) refines Definition 6(ii) to require that $\sigma A \dot{\succ} A$ if $\sigma A \succ A$ and $f \sim \sigma A$ for some $f \in \sigma A$. In a Persuasion Representation, the latter requirement means f is prior-optimal at σA , indicating Sender does not provide additional information beyond the public structure σ . In this sense, the public structure is “binding” and it follows that $f \dot{\sim} \sigma A$ if $\sigma A \dot{\succ} A$. Thus, part (ii) states there is greater conflict if and only if, for any A , there is a larger set of binding information structures that Receiver strictly benefits from.

To conclude this section, the next result characterizes the comparative statics in terms of choice data c .

Proposition 7. *Let (μ, u, v) represent c and $(\mu, u, \dot{v})$ represent $\dot{c}$. The following are equivalent:*

- (i) *$\dot{v}$ is less u -aligned than v is.*
- (ii) *If $c(A) = f$, then $\dot{c}(A) = f$.*

To understand Proposition 7, observe that $c(A) = f$ implies f is prior-optimal for Receiver at A . This implies Sender does not disclose any information at A , indicating substantial disagreement regarding the value of outcomes generated by acts in A . In line with this intuition, the proposition states there is greater conflict between the agents if there are more menus where Receiver chooses only the prior-optimal act.

¹¹Although f is not required to be an element of A , the characterization holds even if one imposes this restriction.

4.3 Naivete and Sophistication

So far, Persuasion Representations of menu preferences have been interpreted as capturing ex-ante values for a Receiver who correctly forecasts Sender's choice of information. The index v inferred from such rankings thus reveals Receiver's *belief* about Sender's utility function, which in turn affects Receiver's forecast of Sender's choice of information. In contrast, ex-post choice data such as ρ or λ is independent of Receiver's beliefs about v , and instead depends only on μ , u , and the true information structure chosen by Sender. This section examines how comparisons between ex-ante preferences and ex-post choices reveal whether Receiver has correctly-specified beliefs about v and, if not, whether beliefs indicate optimism or pessimism.

State-contingent random choice data λ is ideal for such comparisons. Given λ and A , let $f_\lambda^A := (\sum_{g \in A} \lambda_\omega^A(g) g_\omega)_{\omega \in \Omega}$; this is the induced act generated by the true information structure and Receiver's signal-contingent choices at A .

Definition 7. Given λ , preferences $\succsim$ are **sophisticated** if $A \sim f_\lambda^A$ for all A .

The idea of Definition 7 is that λ , and therefore f_λ^A , reflects what actually happens at A while $\succsim$ captures Receiver's ex-ante expectations about behavior at A . If $A \sim f_\lambda^A$, there is no disconnect between expectations and reality: Receiver's ex-ante value of A coincides with the value of the induced act generated by subsequent behavior at A . Thus, sophistication (or lack thereof) is a property of $\succsim$ given some λ .¹²

Proposition 8. Let (μ, u, v') represent $\succsim$ and (μ, u, v) represent λ . Then $\succsim$ is sophisticated if and only if $v' \approx v$.

Proposition 8 verifies the intuition that a sophisticated Receiver holds correct beliefs about Sender's utility function: when $\succsim$ and λ have Persuasion Representations with common Receiver parameters (μ, u) but potentially different Sender utility functions v' and v , respectively, Receiver is sophisticated in the sense of Definition 7 if and only if $v' \approx v$. The remainder of this section studies the following two natural departures from sophistication.

Definition 8. Given λ , preferences $\succsim$ are **optimistic** if $A \succsim f_\lambda^A$ for all A . If instead $A \precsim f_\lambda^A$ for all A , preferences $\succsim$ are **pessimistic**.

¹²This definition is in the spirit of Ahn et al. (2019), who define sophistication as $A \sim c(A)$; in Persuasion Representations, this property turns out to be too weak, and the notion of Definition 7 involving state-contingent random choices is needed to derive meaningful results.

Optimistic Receivers err only in an optimistic direction: if their ex-ante belief regarding the value of A differs from that of f_λ^A , it is because they expect a better outcome than f_λ^A . Similarly, pessimistic Receivers err only in the opposite direction. The next result provides a simple parametric representation of optimism and pessimism in persuasion models.

Proposition 9. *Let (μ, u, v) and (μ, u, v') represent λ and $\succsim$, respectively. Then:*

- (i) *$\succsim$ is optimistic if and only if v' is more u -aligned than v is.*
- (ii) *$\succsim$ is pessimistic if and only if v' is less u -aligned than v is.*

Proposition 9 formalizes the intuition that optimism and pessimism correspond to Receiver's beliefs v' about Sender's utility being more or less u -aligned, respectively, than the true utility v . Thus, the requirements of Definition 8 characterize tight conditions on parameters in Persuasion Representations.

Proposition 10. *Let (μ, u, v) represent λ while (μ, u, v') and $(\mu, u, \dot{v})$ represent $\succsim'$ and $\dot{\succsim}$, respectively. Suppose $\succsim'$ and $\dot{\succsim}$ are optimistic. Then $\dot{v}$ is more u -aligned than v is (that is, $\dot{\succsim}$ is more optimistic than $\succsim'$) if and only if $A \succ' f_\lambda^A$ implies $A \dot{\succ} f_\lambda^A$.*

Proposition 10 states that, conditional on being optimistic, increased optimism manifests as more instances of A being strictly preferred to f_λ^A . A symmetric result holds for pessimistic agents. This can be combined with the results of section 4.2 to establish, for example, that increased pessimism means greater value of public information (more instances where σA is preferred to A) and greater value of full commitment.

5 Discussion

This paper has developed a decision-theoretic analogue of the Bayesian Persuasion model in terms of Receiver's choices, preferences and welfare. The results establish that persuasion interactions can be understood entirely from Receiver's perspective: his behavior reveals all parameter values and, thereby, Sender's (unobserved) choice of information. While Sender has commitment power in choosing information, the framework is sufficiently rich to compare and contrast two ways of leveling the playing field for Receiver: hard commitment and public information provision. This leads to comparative static characterizations that, alongside the identification results, resolve new and basic questions about the persuasion framework.

Viewed as an intrapersonal interaction, the model provides a general framework of self-signaling that, relative to familiar decision-theoretic models, highlights the role of unob-

served information acquisition and/or belief distortion when making choices. This channel of “tempting” belief formation, where the agent seeks to justify or excuse tempting choices, generates rather different implications for the agent’s welfare, demand for commitment, and value of information—important considerations for the design of experiments or the analysis of real-world data. The model also features full identification via choices from or choices between menus—to my knowledge, the first such results for communication models.

A Proofs for Sections 2 and 3

A.1 Proof of Lemma 1

Throughout this section, fix parameters (μ, u) and a menu A .

To establish convexity of $F(A)$, let $f, g \in F(A)$. This means there are experiments σ^f, σ^g and selections $f^s, g^t \in \text{co}(A)$ such that $f = \sum_{s \in \sigma^f} s f^s$, $g = \sum_{t \in \sigma^g} t g^t$ and, for all $s \in \sigma^f$ and $t \in \sigma^g$, $U^s(f^s) \geq U^s(f')$ and $U^t(g^t) \geq U^t(g')$ for all $f', g' \in A$. Let $\alpha \in (0, 1)$. Consider the experiment $\hat{\sigma} = \alpha \sigma^f + (1 - \alpha) \sigma^g$ defined to be the concatenation of matrices $\alpha \sigma^f$ and $(1 - \alpha) \sigma^g$. To columns αs ($s \in \sigma^f$) and $(1 - \alpha)t$ ($t \in \sigma^g$), associate selections $f^{\alpha s} = f^s$ and $g^{(1 - \alpha)t} = g^t$ from the original induced acts (these selections remain valid because scalar multiplication of signals generate the same Bayesian posteriors). Thus, the associated induced act for $\hat{\sigma}$ is $\sum_{s \in \sigma^f} \alpha s f^s + \sum_{t \in \sigma^g} (1 - \alpha) t g^t = \alpha f + (1 - \alpha) g$, so that $F(A)$ is convex.

The remainder of this section establishes compactness of $F(A)$. For every $B \subseteq A$, let

$$C_B = \left\{ \hat{\mu} \in \Delta\Omega : \operatorname{argmax}_{f \in A} \sum_{\omega \in \Omega} u(f_\omega) \hat{\mu}_\omega = B \right\}.$$

That is, C_B consists of all beliefs $\hat{\mu}$ such that the set of optimal acts in A (given utility index u) is precisely B . Observe that C_B is convex but potentially empty. Let $\mathcal{C}$ denote the collection of nonempty sets C_B ; clearly, $\mathcal{C}$ is a partition of $\Delta\Omega$.

Since A is finite, the closure of each $C_B \in \mathcal{C}$ has finitely many extreme points; let $\text{ext}(C_B)$ denote the set of extreme points of the closure of C_B . Taking the union of sets $\text{ext}(C_B)$ over all $C_B \in \mathcal{C}$ yields a finite set of points, which we enumerate as $(\hat{\mu}^1, \dots, \hat{\mu}^K)$. For every $k = 1, \dots, K$, let $B_K := \operatorname{argmax}_{f \in A} \sum_{\omega \in \Omega} u(f_\omega) \hat{\mu}^k$ denote the acts in A that are optimal at beliefs $\hat{\mu}^k$.

Given full-support prior beliefs μ , each point $\hat{\mu}^k$ is associated with a signal $s^k \in S$ such that $\hat{\mu}^k$ is the Bayesian posterior of μ at s^k . Scalar multiples of s^k result in the same Bayesian posterior, so without loss we assume $s^k \in S$ but $\lambda s^k \notin S$ for all $\lambda > 1$ (this means $s_\omega^k = 1$ for at least one ω). Consequently, $\lambda s^k \in S$ for all $0 < \lambda < 1$.

Lemma 2. For every $f \in F(A)$, there is a vector $\lambda = (\lambda_1, \dots, \lambda_K) \in [0, 1]^K$ and a profile $(f^1, \dots, f^K)$ of tie-breaking selections ($f^k \in \text{co}(B_k)$ for all k) such that $f = \sum_{k=1}^K (\lambda_k s^k) f^k$ and $\sum_{k=1}^K \lambda_k s^k = e$.

Proof. Let $f \in F(A)$. This means there is an experiment $\sigma = [t^1, \dots, t^n]$ and selections $f^i \in \text{co}(\arg\max_{g \in A} U^{t^i}(g))$, $i = 1, \dots, n$, such that $f = \sum_{i=1}^n t^i f^i$. Observe that if t^i and t^j are scalar multiples of each other (that is, they induce the same Bayesian posterior), replacing them with a single signal $t^i + t^j$ and a selection $f^{t^i+t^j} \in \text{co}(\arg\max_{g \in A} U^{t^i+t^j}(g))$ given by $f_{\omega}^{t^i+t^j} = \frac{t_{\omega}^i f_{\omega}^{t^i} + t_{\omega}^j f_{\omega}^{t^j}}{t_{\omega}^i + t_{\omega}^j}$ yields the same induced act f . Thus, we assume without loss that distinct signals in σ are not scalar multiples of each other; this means $\mu \cdot t^i$ is the total probability of generating posterior μ^{t^i} .

If every t^i is of the form $\lambda_k s^k$ for some k , there is nothing to prove. So, suppose there exists $t^i \in \sigma$ such that no signal s^k is a scalar multiple of t^i . This implies the Bayesian posterior μ^{t^i} does not coincide with any of the points $\hat{\mu}^k$. Let C_B denote the cell of $\mathcal{C}$ containing μ^{t^i} ; this implies $f^{t^i} \in \text{co}(B)$. By definition, μ^{t^i} can be expressed as a weighted sum of the points $\text{ext}(C_B)$; without loss of generality, let $\mu^{t^i} = \alpha_1 \hat{\mu}^1 + \dots + \alpha_L \hat{\mu}^L$ denote such a sum (there may be multiple such mixtures; the argument does not require uniqueness). Without loss of generality, each $\alpha_{\ell} > 0$, which means $\text{supp}(\hat{\mu}^{\ell}) \subseteq \text{supp}(\mu^{t^i})$. We modify σ by replacing t^i with a matrix $[r^1, \dots, r^L]$ and f^{t^i} with a profile $(f^{r^1}, \dots, f^{r^L})$ of selections $f^{r^{\ell}} \in \text{co}(\arg\max_{g \in A} U^{r^{\ell}}(g))$ as follows. First, let $\hat{\sigma} = [\hat{r}^1, \dots, \hat{r}^L]$ be an experiment such that, for prior μ^{t^i} , posterior $\hat{\mu}^{\ell}$ is generated with probability α_{ℓ} ; as is easily verified, this means $\hat{r}_{\omega}^{\ell} = \frac{\alpha_{\ell} \hat{\mu}_{\omega}^{\ell}}{\mu_{\omega}^{t^i}}$ if $\omega \in \text{supp}(\mu^{t^i})$; for $\omega \notin \text{supp}(\mu^{t^i})$, let $\hat{r}_{\omega}^{\ell} = 1/L$. Now let $r^{\ell} := (t_{\omega}^i \hat{r}_{\omega}^{\ell})_{\omega \in \Omega}$; then the Bayesian posterior of μ at r^{ℓ} is $\hat{\mu}^{\ell}$, so $r^{\ell} = \lambda_{\ell} s^{\ell}$ for some $\lambda_{\ell} > 0$. Since $\hat{\sigma} \in \mathcal{E}$, we have $\sum_{\ell=1}^L r_{\omega}^{\ell} = t_{\omega}^i \sum_{\ell=1}^L \hat{r}_{\omega}^{\ell} = t_{\omega}^i$ for all ω . Next, observe that since the selection f^{t^i} is in $\text{co}(B)$, it is a valid selection at each $\hat{\mu}^{\ell}$ (an element of $\text{ext}(C_B)$) because each $g \in B$ is utility maximizing at beliefs $\hat{\mu}^{\ell}$ and utility index u . Thus, we may set $f^{r^{\ell}} = f^{t^i}$ for every ℓ . Replacing $t^i \in \sigma$ with the matrix $[r^1, \dots, r^L]$ yields the same induced act because the component of the induced act affected by this substitution coincides with the $t^i f^{t^i}$ component of the original induced act: $r^1 f^{r^1} + \dots + r^L f^{r^L} = r^1 f^{t^i} + \dots + r^L f^{t^i} = t^i f^{t^i}$. $\square$

Lemma 2 implies that for every induced act $f \in F(A)$, there is an experiment consisting only of (scalar multiples of) signals from $(s^1, \dots, s^K)$ that, with appropriate tie-breaking selections, yields the induced act f . Thus, $F(A)$ is a subset of the set of acts that can be induced using scaled signals from $(s^1, \dots, s^K)$; obviously, the converse implication holds.

Consequently, $F(A)$ is the image of the correspondence $\varphi : D \rightrightarrows F$ defined by

$$\varphi(\lambda_1, \dots, \lambda_K) := \left\{ \sum_{k=1}^K (\lambda_k s^k) f^k : f^k \in \text{co}(B_k) \right\},$$

where $D := \{(\lambda_1, \dots, \lambda_K) \in [0, 1]^K : \sum_{k=1}^K \lambda_k s^k = e\}$; that is, the domain D consists of all vectors $(\lambda_1, \dots, \lambda_K)$ such that the collection $[\lambda_k s^k : \lambda_k \neq 0]$ qualifies as an experiment. Observe that D is compact: it is clearly bounded, and the set of vectors satisfying the constraint $\sum_{k=1}^K \lambda_k s^k = e$ is closed.

Lemma 3. *The correspondence φ is upper hemicontinuous.*

Proof. Observe that, for every $\lambda = (\lambda_1, \dots, \lambda_K) \in D$,

$$\varphi(\lambda) = \sum_{k=1}^K \varphi_k(\lambda) := \left\{ \sum_{k=1}^K g^k : g^k \in \varphi_k(\lambda) \right\}, \quad (5)$$

where $\varphi_k(\lambda) := \{(\lambda_k s^k) f^k = \lambda_k (s_\omega^k f_\omega^k)_{\omega \in \Omega} : f^k \in \text{co}(B_k)\}$. Clearly, φ_k is compact-valued since $\text{co}(B_k)$ is compact and s^k is fixed. Moreover, letting $X_k := \{(s_\omega^k f_\omega^k)_{\omega \in \Omega} : f^k \in \text{co}(B_k)\}$, we have $\varphi_k(\lambda) = \lambda_k X_k = \{\lambda_k x^k : x^k \in X_k\}$; thus, φ_k simply scales a fixed compact set by the factor λ_k and therefore is an upper hemicontinuous correspondence. By equation (5), then, φ is a sum of compact-valued, upper hemicontinuous correspondences and therefore is itself upper hemicontinuous (see Aliprantis and Border, 2006, Lem 17.8). $\square$

We have shown that $F(A)$ is the image of a compact-valued, upper hemicontinuous correspondence defined on a compact set; thus, $F(A)$ is compact (Aliprantis and Border, 2006, Thm 17.32).

A.2 Proof of Theorem 2

Suppose $\succsim$ has a Persuasion Representation. It follows that the restriction of $\succsim$ to singleton menus is represented by subjective expected utility for some parameters (μ, u) . By standard uniqueness arguments for the Anscombe-Aumann expected utility model, μ is unique and u is unique up to positive affine transformation.

To establish uniqueness of v , consider pq -bets $A_{pq} = \{pEq, qEp\}$ where p is interior and $u(p) > u(q)$. If $v(p) \geq v(q)$, Sender chooses perfect information and therefore $p \sim A_{pq}$. If $v(p) \leq v(q)$, Sender chooses e (no information) and $p \succ A_{pq}$. Since p is interior, the set $\{q : p \sim A_{pq}\}$ coincides with $\{q : u(p) > u(q) \text{ and } v(p) \geq v(q)\}$. This set is a region in ΔX bounded by two planes: the indifference planes through p for u and v . Since u has been

identified in the first step above, this reveals the indifference plane for v through p . The direction of increasing utility for v is also revealed by definition of the latter set. Thus, v is identified up to positive affine transformation.

A.3 Proof of Theorems 3 & 4

First, suppose a correspondence c has a Persuasion Representation. To identify u , consider menus $A = \{p, q\}$ of constant acts; for such menus, we have $p \in c(\{p, q\}) \Leftrightarrow u(p) > u(q)$ or $[u(p) = u(q) \text{ and } v(p) \geq v(q)]$. Consequently, $u(p) > u(q)$ if and only if there is a neighborhood around q such that $p \in c(\{p, q'\})$ for all q' in the neighborhood. Given p , then, the set of such q reveals the strict lower contour set of u through p , thereby revealing u up to positive affine transformation.

To identify v , consider pq -menus A where $u(p) > u(q)$; this means that for every $f \in A$, there is a signal s^f such that $f_\omega = s_\omega^f p + (1 - s_\omega^f)q$. Observe that if $v(p) \geq v(q)$, Sender chooses perfect information at A and so $c(A) = \bar{c}(A) := \bigcup_{\omega \in \Omega} \arg\max_{f \in A} u(f_\omega)$. If $v(p) < v(q)$, Sender chooses e (no information) and so $c(A) = \underline{c}(A) := \arg\max_{f \in A} U(f)$. Given $u(p) > u(q)$, it is straightforward to construct pq -menus A where $\bar{c}(A) \neq \underline{c}(A)$ and $c(A) = \bar{c}(A) \Leftrightarrow v(p) \geq v(q)$; for example, there is a pq -menu A where there is a unique prior-optimal act f^e (that is, $U(f^e) \geq U(g)$ for all $g \in A$) and, for each state ω , a unique act $f^\omega \in A$ such that $f_\omega^\omega = p$, making this act the unique optimal choice in state ω . Thus, $\bar{c}(A) = \{f^\omega : \omega \in \Omega\} \neq \{f^e\} = \underline{c}(A)$. Having identified u , then, such menus reveal the set $\{q : u(p) > u(q) \text{ and } v(p) \geq v(q)\}$; by the argument in the proof of Theorem 2 above, this reveals v up to positive affine transformation.

To identify μ , consider once again pq -menus A where $u(p) > u(q)$. Suppose $u \not\approx v$. Then there exists q such that $v(q) > v(p)$; consequently, there exist pq -menus A such that $c(A) = \underline{c}(A)$. Normalize $u(p) = 1$, $u(q) = 0$ and consider a pair of states $E = \{\omega, \omega'\}$ where $\omega \neq \omega'$. To pin down the ratio $\frac{\mu_\omega}{\mu_{\omega'}}$, consider a pq -menu A where all $f, f' \in A$ satisfy $f_{\hat{\omega}} = f'_{\hat{\omega}}$ for all $\hat{\omega} \notin E$ (that is, there exists an act $h = s^h p + (1 - s^h)q$ such that, for all $f \in A$, $f = f E h$). Since $c(A) = \underline{c}(A)$, it follows that $f \in c(A) \Leftrightarrow f \in \arg\max_{g \in A} s_\omega^g \mu_\omega + (1 - s_{\omega'}^g) \mu_{\omega'}$. Thus, $f, g \in c(A) \Leftrightarrow (s_\omega^f - s_\omega^g) \mu_\omega = (s_{\omega'}^g - s_{\omega'}^f) \mu_{\omega'}$. Appropriate choices of A thereby pin down $\frac{\mu_\omega}{\mu_{\omega'}}$. For example, letting $g = h = \frac{1}{2}p + \frac{1}{2}q$, one can elicit $s_\omega^f, s_{\omega'}^f$ such that $c(\{f, g\}) = \underline{c}(\{f, g\}) = \{f, g\}$, revealing $(s_\omega^f - \frac{1}{2}) \mu_\omega = (\frac{1}{2} - s_{\omega'}^f) \mu_{\omega'}$. Repeating this procedure for all pairs of states pins down all likelihood ratios and, therefore, pins down μ . This completes the proof of Theorem 4.

For Theorem 3, suppose ρ has a Persuasion Representation. Uniqueness of u and v follows from the argument for c , as does identification of μ if $u \not\approx v$. So, suppose $u \approx v$ and let

$u(p) > u(q)$. Then $v(p) > v(q)$ and Sender chooses perfect information at all pq -menus A . In particular, for each state ω , consider a pq -bet $A = \{pEq, qEp\}$ where $E = \{\omega\}$. Since Sender chooses perfect information, it follows that $\rho^A(pEq) = \mu_\omega$, pinning down μ .

B Proofs for Section 4.1

Lemma 4. *Let u, v be non-constant utility indices such that $u \not\approx v$ and $u \not\approx -v$. For any pair of vectors $(u_1, \dots, u_K), (v_1, \dots, v_K) \in \mathbb{R}^K$, there is a set $\{p_1, \dots, p_K\} \subseteq \Delta X$ and constants $A > 0, B, C \in \mathbb{R}$ such that, for all $k = 1, \dots, K$, $u(p_k) = Au_k + B$ and $v(p_k) = Av_k + C$.*

Proof. Observe that ΔX can be identified with a subset of $\mathbb{R}^{N-1}$ (namely, the unit simplex in $\mathbb{R}^N$). Since u, v are non-constant linear functions on $\Delta X \subseteq \mathbb{R}^{N-1}$, their domains can be extended to all of $\mathbb{R}^{N-1}$ via linearity. For every k , the values u_k and v_k correspond to unique level sets (planes) of u and v in $\mathbb{R}^{N-1}$, respectively; since $u \not\approx v$ and $u \not\approx -v$ the normal vectors of these planes are linearly independent. Thus, for every k , there is a point $z^k \in \mathbb{R}^{N-1}$ such that $u(z^k) = u_k$ and $v(z^k) = v_k$. Pick a lottery p in the interior of ΔX . There is a scalar $\alpha \in (0, 1)$ sufficiently close to 1 such that, for all k , $\alpha p + (1 - \alpha)z^k \in \Delta X$; letting $p_k := \alpha p + (1 - \alpha)z^k$, $A := (1 - \alpha)$, $B := \alpha u(p)$ and $C := \alpha v(p)$ completes the proof. $\square$

The significance of Lemma 4 is that it allows acts and menus to be constructed by selecting utility values for Receiver independently of the values for Sender when $u \not\approx v$ and $u \not\approx -v$. In particular, acts can be defined by specifying arbitrary profiles of utilities $(u_\omega)_{\omega \in \Omega}$ and $(v_\omega)_{\omega \in \Omega}$ (not necessarily in the range of u or v) and applying the lemma to obtain $f = (f_\omega)_{\omega \in \Omega}$ such that, for all ω , $u(f_\omega) = Au_\omega + B$ and $v(f_\omega) = Av_\omega + C$. More generally, menus are obtained by first specifying (for each act in the menu) utility profiles for Sender and Receiver and then applying the lemma to the full set of profiles—each act requires $|\Omega| = W$ lotteries, so a menu of M acts requires $K = MW$ lotteries. This way, the same constants $A > 0, B, C \in \mathbb{R}$ apply to every act, so in the resulting menu it is as if agents compare acts with the desired utility profiles. This simplifies the construction of many examples since value-function (or concavification) arguments only depend on the utility profiles, not the underlying lotteries.

B.1 Proof of Proposition 1

Proof of $u \approx v$ or $u \approx -v \Leftrightarrow (i)$. If $u \approx v$, then at every B Sender chooses perfect information; consequently, in every ω , both agents receive their most-preferred lottery in B_ω . If $B \subseteq A$, then the best outcome in each state can only improve under menu A . Thus, $A \succeq B$.

If instead $u \approx -v$, then e (no information) is Sender-optimal at every menu A ; consequently, Receiver chooses their prior-optimal act(s) from A . Thus, $\succsim$ reduces to a standard expected utility preference and therefore satisfies Preference for Flexibility.

For the converse, suppose $u \not\approx v$ and $u \not\approx -v$. Then, as is easily verified, there exist lotteries p, q, r such that $u(p) > u(r) > u(q)$ and $v(r) > v(p) > v(q)$. Let $E \neq \Omega$ be a nonempty subset of Ω and $B = \{pEq, qEp\}$. Since Sender and Receiver agree on the ranking of p and q , it follows that Sender chooses perfect information at menu B , yielding lottery p in every state. Thus, $U(B) = u(p)$.

Now let $A = \{pEq, qEp, r\}$. Clearly, $A \supseteq B$. We may choose r so that r is prior-optimal for Receiver in menu A ; in particular, $U^e(pEq)$ and $U^e(qEp)$ belong to the open interval $(u(q), u(p))$ because $\mu(E) \in (0, 1)$. Thus, we may choose r near p (without reversing any inequalities above) so that $U^e(r) > \max\{U^e(pEq), U^e(qEp)\}$, making r prior-optimal for Receiver in menu A . Since $v(r) > v(p) > v(q)$, it follows that e (no information) is Sender-optimal at A (more generally, any Sender-optimal σ must yield Bayesian posteriors making r Receiver-optimal). Thus, outcome r is realized with probability 1, so that $U(A) = u(r) < u(p) = U(B)$, violating Preference for Flexibility. $\square$

Proof of $u \approx v$ or $u \approx -v \Leftrightarrow (ii)$. First, suppose $u \approx v$ or $u \approx -v$. If $u \approx -v$, then (by the argument in part (i) above) c is rationalized by expected utility maximization with parameters (μ, u) and therefore satisfies Sen's α . If instead $u \approx v$, then (also by part (i) above) for all $\hat{A}$, we have $c(\hat{A}) = \{f \in \hat{A} : \exists \omega \text{ such that } u(f_\omega) \geq u(g_\omega) \forall g \in \hat{A}\}$. Let $f \in c(B) \cap A$ where $B \supseteq A$. Since $f \in c(B)$, there exists a state, say ω^* , such that $u(f_{\omega^*}) \geq u(g_{\omega^*})$ for all $g \in B \supseteq A$. Thus, $u(f_{\omega^*}) \geq u(g_{\omega^*})$ for all $g \in A$. Since $f \in A$, this implies $f \in c(A)$.

The converse is established by way of contradiction. So, suppose $u \not\approx v$ and $u \not\approx -v$. Choose an event E such that $0 < \mu(E) < 1$ and lotteries r, p, q, p', q' such that $u(p') > u(p) > u(r) > u(q') > u(q)$, $v(p) = v(q) > v(r) > v(p') = v(q')$, and $U^e(pEq) > u(r)$. Let $A = \{pEq, r\}$. Then e (no information) is Sender-optimal at A because pEq is prior-optimal for Receiver and $v(p) = v(q) > v(r)$; thus, $c(A) = \{pEq\}$. Now let $B = \{pEq, p'Eq', r\}$. Observe that, for Receiver, $p'Eq'$ dominates pEq . Thus, $p'Eq'$ is prior-optimal for Receiver; pEq is not chosen by Receiver at any signal; and r is chosen by Receiver at some signals because $0 < \mu(E) < 1$ and $u(p') > u(r) > u(q')$. Since $v(r) > v(p') = v(q')$, Sender selects an information structure where both $p'Eq'$ and r are chosen with positive probability; hence, $c(B) = \{p'Eq', r\}$. Thus, $r \in c(B) \cap A$ but $r \notin c(A)$, violating Sen's α . $\square$

Proof of $u \approx v$ or $u \approx -v \Leftrightarrow (iii)$. First, suppose $u \approx v$ or $u \approx -v$. By part (i), $\succsim$ satisfies Preference for Flexibility. Since $\sigma A \supseteq A$, it follows that $\sigma A \succsim A$.

For the converse, suppose $u \not\approx v$ and $u \not\approx -v$. Consider first the case $|\Omega| = 2$. Since μ has full support, we may construct (by Lemma 4) a menu $A = \{f, g, h\}$ such that the value functions are of the form depicted by Figure 3 in the main text. In particular, Receiver finds f optimal on $[0, \hat{\mu}']$, g optimal on $[\hat{\mu}', \hat{\mu}'']$, and h optimal on $[\hat{\mu}'', 1]$, where the prior μ satisfies $\hat{\mu}' < \mu < \hat{\mu}''$. Sender is indifferent between g and h at all beliefs, prefers g (and h) to f on $[0, \hat{\mu}']$, and f to g and h on $[\hat{\mu}'', 1]$, with indifference between all three acts at $\hat{\mu}'$. Let $\hat{\sigma}$ be an experiment such that, given μ , either posterior $\hat{\mu}'$ or $\hat{\mu} = 1$ is generated (such an experiment exists because $\hat{\mu}' < \mu < 1$). Observe that, at menu A , Sender finds both e and $\hat{\sigma}$ optimal because the concavification at $\hat{\mu}'$, μ , and $\hat{\mu} = 1$ coincides with Sender's value function. By Receiver-preferred tie breaking, Sender chooses $\hat{\sigma}$ as this yields the highest Receiver payoff among all Sender-optimal structures at A . Now consider a public structure σ that generates posteriors at $\hat{\mu}'$ and $\hat{\mu}''$. Again, Sender's values at these posteriors coincide with the concavification. However, Sender now prefers e because any non-trivial information structure creates a mean-preserving spread around both $\hat{\mu}'$ and $\hat{\mu}''$; in particular, any such spread around $\hat{\mu}'$ strictly lowers Sender's payoff. Thus, Sender chooses e and Receiver's value at σA is the prior-value of g and, thus, lower than the value at A .

For the general case $|\Omega| \geq 3$, let p be an arbitrary lottery, $E = \{\omega_1, \omega_2\}$ and consider the menu $A' := \{\hat{f}Ep : \hat{f} \in A\}$, where A is a menu of the form constructed above for the case $|\Omega| = 2$; this can be done because μ has full support. It is easy to see that $F(A') = \{\hat{f}Ep : \hat{f} \in F(A)\}$. Thus, at A , Sender selects an induced act f^*Ep , where $f^* \in F(A)$ is the induced act selected at A . Let σ' be an information structure generating posteriors corresponding to $\hat{\mu}'$ and $\hat{\mu}''$ in the 2-state construction above (conditional on E , the posteriors coincide with $\hat{\mu}'$ and $\hat{\mu}''$) and that leaves $\hat{\mu}'_\omega = \hat{\mu}''_\omega = \mu_\omega$ for all $\omega \notin E$. Then $F(\sigma'A') = \{\hat{f}Ep : \hat{f} \in \sigma A\}$, so Sender selects an induced act $g^*Ep \in F(\sigma'A')$ where $g^* \in F(\sigma A)$ is the induced act selected at σA . Thus, $A' \succ \sigma'A'$. $\square$

Proof of $u \approx v$ or $u \approx -v \Leftrightarrow (iv)$. First, suppose $u \approx v$ or $u \approx -v$. By part (ii), c satisfies Sen's α . Since $\sigma A \supseteq A$, it follows that $c(\sigma A) \cap A \subseteq c(A)$, as desired.

For the converse, suppose $u \not\approx v$ and $u \not\approx -v$. Consider first the case $|\Omega| = 2$. Let $A = \{f, g, h\}$ be a menu such that Sender's value function is of the form depicted in Figure 4 (this is possible via Lemma 4 and the fact that μ has full support). Clearly, e is Sender-optimal at A and so $c(A) = \{h\}$. Consider the public structure σ generating posteriors at $\hat{\mu}'$ and $\hat{\mu} = 1$. Any additional information chosen by Sender generates a mean-preserving spread around $\hat{\mu}'$ but does not affect the point $\hat{\mu} = 1$. Thus, Sender chooses information $\hat{\sigma}$ to achieve the value of the concavification at $\hat{\mu}'$; this requires $\hat{\sigma}$ to generate posteriors at $\hat{\mu}''$ and at $\hat{\mu} = 0$. The latter posterior results in f being chosen. Thus, $f \in c(\sigma A) \cap A$ but




Figure 5: Illustration of Proposition 1(iv).

$f \notin c(A)$. For the general case $|\Omega| \geq 3$, apply the same technique as the proof of part (ii) above to construct A' and $\sigma'A'$ such that $F(A')$ and $F(\sigma'A')$ are isomorphic to the sets $F(A)$ and $F(\sigma A)$ from the case $|\Omega| = 2$. $\square$

B.2 Proof of Proposition 2

Proof of $u \approx v \Leftrightarrow (i)$. Suppose $u \approx v$. Then, for every menu B , perfect information is Sender-optimal because it yields u -maximal (hence v -maximal) lotteries from B_ω in every ω . If $B_\omega \subseteq A_\omega$ for all ω , then a maximal lottery in A_ω is maximal in $A_\omega \cup B_\omega$. Thus, $A \sim A \cup B$.

Conversely, suppose $u \not\approx v$. Then there are lotteries p, q such that $u(p) > u(q)$ and $v(q) > v(p)$. Let $\mu(E) \in (0, 1)$ and $A = \{pEq, qEp\}$. Sender and Receiver strictly disagree on the ranking of p and q , so Sender chooses e (no information) at A . Consequently, $U(A) \in (u(q), u(p))$ and $V(A) \in (v(p), v(q))$. Let $B = \{p\}$, so that $A \cup B = \{pEq, qEp, p\}$. Then p is a Receiver-optimal act in $A \cup B$ regardless of Sender's choice of information, implying $U(A \cup B) = u(p) > U(A)$; thus, $A \not\sim A \cup B$. Since $B_\omega = \{p\} \subseteq \{p, q\} = A_\omega$ for all ω , this contradicts Preference for Statewise Flexibility. $\square$

Proof of $u \approx v \Leftrightarrow (iii)$. Suppose $u \approx v$. Then, for every A , perfect information is Sender-optimal and Receiver achieves a u -maximal lottery in A_ω for all ω . Let $\bar{u}_\omega$ denote Receiver's utility of any u -maximal lottery in A_ω . Observe that, for every σ and ω , Receiver again obtains utility $\bar{u}_\omega$ in state ω at menu σA (perfect information remains Sender-optimal and $A \subseteq \sigma A$). Thus, $A \sim \sigma A$.

Conversely, suppose $u \not\approx v$. Then there are lotteries p, q such that $u(p) > u(q)$ and $v(q) > v(p)$. Let $\mu(E) \in (0, 1)$ and $A = \{pEq, qEp\}$. As in the proof of $u \approx v \Leftrightarrow (i)$ above, this yields $U(A) \in (u(q), u(p))$. Consider σ^*A . Since $(p, \dots, p) \in \sigma^*A$, we obtain $U(\sigma^*A) = u(p) > U(A)$, so that $A \not\sim \sigma^*A$. $\square$

Proof of $u \approx v \Leftrightarrow (ii)$. Suppose $u \approx v$. Then perfect information is Sender-optimal; consequently, for all $\hat{A}$, we have

$$c(\hat{A}) = \{f \in \hat{A} : \exists \omega \text{ such that } u(f_\omega) \geq u(g_\omega) \forall g \in \hat{A}\}.$$

Suppose $B_\omega \subseteq A_\omega$ for all ω and let $f \in c(A)$. Then there is a state ω^* such that $u(f_{\omega^*}) \geq u(g_{\omega^*})$ for all $g \in A$; thus, $u(f_{\omega^*}) \geq u(p)$ for all $p \in A_{\omega^*}$. Since $B_{\omega^*} \subseteq A_{\omega^*}$, it follows that $u(f_{\omega^*}) \geq u(q)$ for all $q \in B_{\omega^*}$. Then $u(f_{\omega^*}) \geq u(g_{\omega^*})$ for all $g \in A \cup B$, so that $f \in c(A \cup B)$.

For the converse, suppose $u \not\approx v$. As in the proof of $u \approx v \Leftrightarrow (i)$ above, the menus $A = \{pEq, qEp\}$ and $B = \{p\}$ satisfy $B_\omega \subseteq A_\omega$ for all ω but lead to $c(A \cup B) = \{p\}$, so that $c(A) \not\subseteq c(A \cup B)$. $\square$

Proof of $u \approx v \Leftrightarrow (iv)$. Suppose $u \approx v$. As in the proof of $u \approx v \Leftrightarrow (ii)$ above, we have

$$c(A) = \{f \in A : \exists \omega \text{ such that } u(f_\omega) \geq u(g_\omega) \forall g \in A\}$$

for all A . Let σ denote an interior experiment. To see that $c(A) \subseteq c(\sigma A)$, first let $f \in c(A)$. Then $f \in \sigma A$ and $u(f_\omega) \geq u(p)$ for all $p \in A_\omega$. Thus, for every ω and $h \in \sigma A$, we have $u(f_\omega) \geq u(h_\omega)$ because h_ω is a mixture of lotteries in A_ω . Hence, $f \in c(\sigma A)$. To establish $c(\sigma A) \subseteq c(A)$, let $f \in c(\sigma A)$. This means there is a state ω^* such that $u(f_{\omega^*}) \geq u(h_{\omega^*})$ for all $h \in \sigma A \supseteq A$; thus, $u(f_{\omega^*}) \geq u(g_{\omega^*})$ for all $g \in A$. So, it will suffice to show that $f \in A$. Let $\omega \in \Omega$. A given lottery in $(\sigma A)_\omega \supseteq A_\omega$ is a mixture of lotteries in A_ω . If $f \in \sigma A \setminus A$ and the mixture f_ω assigns positive weight only to u -maximizers in A_ω , then σ cannot be interior: it must perfectly reveal state ω . Thus, $f \in A$.

For the converse, suppose $u \not\approx v$. Then there are lotteries p, q such that $u(p) > u(q)$ and $v(q) > v(p)$. Let $\mu(E) \in (0, 1)$ and $A = \{pEq, qEp\}$. Let $\sigma^\varepsilon = [s^\varepsilon, e - s^\varepsilon]$ where $s^\varepsilon_\omega = 1 - \varepsilon$ for $\omega \in E$ and $s^\varepsilon_\omega = \varepsilon$ for $\omega \in \Omega \setminus E$. Then

$$\begin{aligned} \sigma^\varepsilon A &= \{pEq, qEp, s^\varepsilon pEq + (e - s^\varepsilon)qEp, s^\varepsilon qEp + (e - s^\varepsilon)pEq\} \\ &= \{pEq, qEp, (1 - \varepsilon)p + \varepsilon q, \varepsilon p + (1 - \varepsilon)q\}. \end{aligned}$$

Observe that e (no information) is Sender-optimal at $\sigma^\varepsilon A$. Thus, for $\varepsilon > 0$ sufficiently small, Receiver chooses $(1 - \varepsilon)p + \varepsilon q$. Thus, $c(\sigma^\varepsilon A) \neq c(A)$. $\square$

B.3 Proof of Proposition 3

Proof of $u \approx -v \Leftrightarrow (i)$. First, suppose $u \approx -v$. Then e (no information) is Sender-optimal at every menu A , so that Receiver chooses the prior-optimal act (according to u) from A .

Thus, $\succsim$ reduces to a standard expected utility preference and, consequently, satisfies IIA.

Conversely, suppose $u \not\approx -v$. Then there are lotteries p, q such that $u(p) > u(q)$ and $v(p) > v(q)$. Let $A = \{pEq\}$ and $B = \{qEp\}$ where $1 > \mu(E) \geq \frac{1}{2}$. It follows that $U(A) = U^e(pEq) \geq U^e(qEp) = U(B)$, so that $A \succsim B$. However, at menu $A \cup B$, perfect information is Sender-optimal because it yields lottery p in every state. Thus, $U(A \cup B) = u(p) > U(A)$, so that $A \cup B \succ A$, contradicting IIA. $\square$

Proof of $u \approx -v \Leftrightarrow (iii)$. First, suppose $u \approx -v$. By Proposition 1, $\succsim$ values information: $\sigma A \succsim A$ for all σ and A . Moreover, e (no information) is Sender-optimal at all menus A , so Receiver is indifferent between A and any U^e -optimal act $f \in A$. Consider a menu $A = \{pEq, qEp\}$ where $0 < \mu(E) < 1$ and $u(p) > u(q)$. Then $\sigma^* A \sim p \succ A$, where σ^* denotes perfect information.

For the converse, suppose $\sigma A \succsim A$ for all σ, A and that there exist σ, A such that $\sigma A \succ A$. By Proposition 1, either $u \approx v$ or $u \approx -v$. By Proposition 2, $u \not\approx v$ because $\succsim$ is not indifferent to information. Thus, $u \approx -v$. $\square$

Proof of $u \approx -v \Leftrightarrow (ii)$. If $u \approx -v$, then Sender chooses e at every menu. Consequently, Receiver's choices are characterized by standard expected utility maximization and therefore satisfy WARP.

For the converse, suppose toward a contradiction that c satisfies WARP but $u \not\approx -v$. By WARP, there is a complete and transitive relation $\succsim$ such that, for all $\hat{A}$, $c(\hat{A}) = \{f \in \hat{A} : f \succsim g \ \forall g \in \hat{A}\}$. Since $u \not\approx -v$, there are lotteries p, q, r such that $u(p) > u(q) > u(r)$ and $v(p) > v(q) > v(r)$. Let E be a nonempty subset of Ω and let $f = rEp$, $g = r$, and $h = pEq$. In menu $A = \{f, g\}$, we have $c(A) = \{f, g\}$ because Sender's optimal information structure reveals whether the true state belongs to E or $\Omega \setminus E$ and both acts are chosen at states $\omega \in E$. Thus, the rationalizing preference satisfies $f \sim g$. In menu $B = \{g, h\}$, we have $c(B) = \{h\}$ because $u(p) > u(q) > u(r)$ implies act g is never chosen by Receiver. Thus, $h \succ g$. Finally, in menu $C = \{f, h\}$, we have $c(C) = \{f, h\}$ because Sender's optimal information structure reveals whether the true state belongs to E or $\Omega \setminus E$ and h is chosen for $\omega \in E$ while f is chosen for $\omega \notin E$. Thus, $f \sim h$. Combining these facts, we have $g \sim f \succ h \succ g$, a contradiction. $\square$

Proof of $u \approx -v \Leftrightarrow (iv)$. First, suppose $u \approx -v$. By Proposition 1, c satisfies Informational Sen's α . Moreover, $c(A)$ consists of all U^e -optimal acts in A because e (no information) is Sender optimal in all menus. As in the proof of part (iv) of Proposition 2, let $0 < \mu(E) < 1$ and $\sigma^\varepsilon = [s^\varepsilon, e - s^\varepsilon]$ where $s^\varepsilon_\omega = 1 - \varepsilon$ for $\omega \in E$ and $s^\varepsilon_\omega = \varepsilon$ for $\omega \in \Omega \setminus E$; note that σ^ε is interior provided $0 < \varepsilon < 1$. Then, for any $A = \{pEq, qEp\}$ where $u(p) > u(q)$, we have

$\sigma^\varepsilon A = \{pEq, qEp, (1 - \varepsilon)p + \varepsilon q, \varepsilon p + (1 - \varepsilon)q\}$. Observe that if $\varepsilon > 0$ is sufficiently close to 0, then $(1 - \varepsilon)p + \varepsilon q$ is U^ε -optimal in $\sigma^\varepsilon A$. Thus, $c(\sigma^\varepsilon A) = (1 - \varepsilon)p + \varepsilon q \notin c(A)$. $\square$

C Proofs for Sections 4.2 and 4.3

Given utility indices u and v , a pair $\{p, q\}$ of lotteries is a (u, v) -**agreement pair** if either $[u(p) \geq u(q) \text{ and } v(p) \geq v(q)]$ or $[u(q) \geq u(p) \text{ and } v(q) \geq v(p)]$.

Lemma 5. *Let $u, v, v' : X \rightarrow \mathbb{R}$ be non-constant utility indices such that $u \not\approx -v$ and $u \not\approx -v'$. Then every (u, v) -agreement pair is a (u, v') -agreement pair if and only if v' is more u -aligned than v .*

Proof. It is straightforward to show that if v' is more u -aligned than v , then every (u, v) -agreement pair is a (u, v') -agreement pair.

For the converse, let u, v, v' be non-constant utility indices such that $u \not\approx -v$ and $u \not\approx -v'$; interpret them as vectors $u = (u_1, \dots, u_N)$, $v = (v_1, \dots, v_N)$, and $v' = (v'_1, \dots, v'_N)$ in $\mathbb{R}^N$ where $X = \{x_1, \dots, x_N\}$.

Let $Z := \{\hat{u} \in \mathbb{R}^N : \sum_{n=1}^N \hat{u}_n = 0\}$; this is a hyperplane in $\mathbb{R}^N$ with normal vector $(1, \dots, 1)$. Observe that for every utility index u' , there exists $\hat{u} \in Z$ representing the same expected utility preferences; in particular, letting $B = \sum_{n=1}^N u'_n$, the index $\hat{u}$ where $\hat{u}_n = u'_n - B/N$ is a member of Z and a positive affine transformation of u' . Thus, every expected utility preference over ΔX is represented by an index in Z . Since scaling an index by a positive number does not affect expected utility preferences, we normalize the non-zero vectors in Z to length 1 by letting $C := \{\frac{\hat{u}}{\|\hat{u}\|} : 0 \neq \hat{u} \in Z\}$. Thus, every non-constant expected utility preference over ΔX is represented by a unique vector in C .¹³

Observe that if $u, v, v' \in C$, then $v' \approx \alpha u + (1 - \alpha)v$ for some $\alpha \in [0, 1]$ if and only if the vector v' belongs to the conic hull $C_{u,v} := \{\beta u + \gamma v : \beta \geq 0, \gamma \geq 0\} \subseteq Z$ of u and v . Suppose $v' \not\approx \alpha u + (1 - \alpha)v$ for all α . Since $C_{u,v} \subseteq Z$ and Z is isomorphic to $\mathbb{R}^{N-1}$, Farkas' lemma implies there is a vector $y \in Z$ such that $v' \cdot y < 0$ and $\hat{u} \cdot y \geq 0$ for all $\hat{u} \in C_{u,v}$; in particular, $u \cdot y \geq 0$ and $v \cdot y \geq 0$.

Let $q = (\frac{1}{N}, \dots, \frac{1}{N}) \in \Delta X$ (that is, $q(x) = \frac{1}{N}$ for all $x \in X$). Letting $p = q + y$, we have $\sum_{n=1}^N p_n = \sum_{n=1}^N q_n + y_n = 1$ since $q \in \Delta X$ and $y \in Z$; if necessary, replace y with δy for sufficiently small $\delta > 0$ (this does not affect any inequalities above) to ensure $p_n = q_n + y_n \in [0, 1]$ for all n . Thus, $p \in \Delta X$ as well. It follows that $0 \leq u \cdot y = u \cdot (p - q)$, so that $u \cdot p \geq u \cdot q$. Similarly, $v \cdot p \geq v \cdot q$, so $\{p, q\}$ is a (u, v) -agreement pair. However, $0 > v' \cdot y = v' \cdot (p - q)$, so that $v' \cdot p < v' \cdot q$. Since u, v and v' are linear functions and p, q

¹³This method of normalizing the set of utility indices is due to Dekel and Lipman (2012).

satisfy $u \cdot p \geq u \cdot q$ and $v \cdot p \geq v \cdot q$, we may perturb p and q to ensure $u \cdot p > u \cdot q$ (and $v \cdot p \geq v \cdot q$) without violating $v' \cdot p < v' \cdot q$. Thus, $\{p, q\}$ is not a (u, v') -agreement pair. $\square$

Lemma 6. *Suppose (μ, u, v) and (μ, u, v') represent $\succsim$ and $\succsim'$, respectively. Let $U, U' : \mathcal{A} \rightarrow \mathbb{R}$ denote the corresponding functions of the form (2) induced by these parameters. Suppose v' is more u -aligned than v . Then $U'(A) \geq U(A)$ for all A and $U'(B) = U(B)$ for all singleton menus B .*

Proof. We have $v' \approx \alpha u + (1 - \alpha)v$ for some $\alpha \in [0, 1]$. Since parameters (μ, u) are common to the representations of $\succsim$ and $\succsim'$, both involve the same set $F(A)$ of induced acts; that is, for all A ,

$$U(A) = \max_{f' \in F(A)} U^e(f') \text{ subject to } f' \in \operatorname{argmax}_{f' \in F(A)} V(f')$$

and

$$U'(A) = \max_{g' \in F(A)} U^e(g') \text{ subject to } g' \in \operatorname{argmax}_{g' \in F(A)} V'(g'),$$

where $V, V' : F \rightarrow \mathbb{R}$ are expected utility with prior μ and indices v and v' , respectively, and $U^e : F \rightarrow \mathbb{R}$ is expected utility with prior μ and index u . Clearly, then, $U'(B) = U(B)$ for all singleton menus B .

Let A be an arbitrary menu. The claim is trivial if $\alpha = 0$, so suppose $\alpha > 0$. Suppose $f \in \operatorname{argmax}_{f' \in F(A)} V(f')$ and $g \in \operatorname{argmax}_{g' \in F(A)} \alpha U^e(g') + (1 - \alpha)V(g')$; thus, g' maximizes $V' \approx \alpha U^e + (1 - \alpha)V$ on $F(A)$. Then

$$\begin{aligned} \alpha U^e(g) + (1 - \alpha)V(g) &\geq \alpha U^e(f) + (1 - \alpha)V(f) && \text{since } g \text{ maximizes } V' \text{ on } F(A) \\ &\geq \alpha U^e(f) + (1 - \alpha)V(g) && \text{since } f \text{ maximizes } V \text{ on } F(A), \end{aligned}$$

which implies $\alpha U^e(g) \geq \alpha U^e(f)$ and thus $U^e(g) \geq U^e(f)$ since $\alpha > 0$. Since f and g are arbitrary maximizers of V and V' , respectively, this implies $U'(A) \geq U(A)$. $\square$

C.1 Proof of Proposition 4

- (i) First, suppose $u \approx v$. Then perfect information is Sender-optimal at all menus B , yielding $c(B) = \{f \in B : \exists \omega \ u(f_\omega) \geq u(g_\omega) \ \forall g \in B\}$. For $B = \sigma A$, acts $f \in \sigma A$ are of the form $f = \sum_{s \in \sigma} s f^s$ ($f^s \in A$); consequently, $u(f_\omega) \geq u(g_\omega)$ for all $g \in \sigma A$ if and only if $u(\sum_{s \in \sigma} s_\omega f_\omega^s) \geq u(\sum_{s \in \sigma} s_\omega g_\omega^s) \ \forall g^s \in A$ ($s \in \sigma$). Thus, $f = \sum_{s \in \sigma} s f^s \in c(\sigma A)$ if and only if there is a state ω such that $u(f_\omega^s) \geq u(g_\omega)$ for all $s \in \sigma$ and $g \in A$. This implies $c_A(\sigma A) = c(A)$.

If $u \approx -v$, then e (no information) is Sender-optimal, yielding $c(B) = \{f \in B :$




Figure 6: Illustration of Proposition 4(i).

$U(f) \geq U(g) \forall g \in B\}$. For $B = \sigma A$, this implies $c(\sigma A) = \{f = \sum_{s \in \sigma} s f^s \in \sigma A : \forall s \in \sigma, U^s(f^s) \geq U^s(g) \forall g \in A\}$; thus, $\hat{f} \in c_A(\sigma A)$ if and only if there exists $s \in \sigma$ such that $U^s(\hat{f}) \geq U^s(g)$ for all $g \in A$. Suppose $\text{supp}(\sigma) \subseteq \text{supp}(\sigma')$ and let $\hat{f} \in c_A(\sigma A)$. Then there exists $s \in \sigma$ such that $U^s(\hat{f}) \geq U^s(g)$ for all $g \in A$. Since $\text{supp}(\sigma) \subseteq \text{supp}(\sigma')$, there exists $s' \in \sigma'$ such that s and s' yield the same Bayesian posterior. Thus, $U^{s'}(\hat{f}) \geq U^{s'}(g)$ for all $g \in A$ as well, so that $\hat{f} \in c_A(\sigma' A)$.

For the converse, suppose c is expansive in signals. Suppose toward a contradiction that $u \not\approx v$ and $u \not\approx -v$. Consider the case $|\Omega| = 2$ (this case extends to the general case via the arguments used in the proof of Proposition 1). Consider a menu $A = \{f, g, h\}$ with value function depicted in Figure 6; such a menu exists by Lemma 4. Sender is indifferent between g and h at all beliefs, strictly prefers f on $\hat{\mu}'_1 < \hat{\mu}_1 < 1$, and strictly prefers g (and h) on $0 < \hat{\mu}_1 < \hat{\mu}'$, with indifference between all three acts at $\hat{\mu}'$. Let $\sigma = e$ (so the only posterior it generates is the prior μ) and σ' be an information structure generating posteriors at $\mu, \hat{\mu}'$, and $\hat{\mu}''$. At σ , Sender chooses (due to Receiver-preferred tie breaking) additional information so as to induce posteriors at $\hat{\mu}'$ and $\hat{\mu} = 1$. Consequently, $c_A(\sigma A) = \{f, g, h\}$ because h is chosen at $\hat{\mu} = 1$ while f and g are chosen at $\hat{\mu}'$. At $\sigma' A$, Sender chooses no additional information (each posterior yields a point on the concavification, and creating spread around $\hat{\mu}'$ strictly decreases Sender's payoff); consequently, $h \notin c_A(\sigma' A)$ because no posterior induced by σ' makes Receiver choose h . Thus, $\text{supp}(\sigma) \subseteq \text{supp}(\sigma')$ but $c_A(\sigma A) \not\subseteq c_A(\sigma' A)$.

- (ii) First, suppose $u \approx v$. As demonstrated in the proof of (i), this implies $c(A) = c_A(\sigma A)$ for all A and σ . Conversely, suppose $u \not\approx v$. This implies there are lotteries p, q such that $u(p) > u(q)$ and $v(q) > v(p)$. Let $A = \{sp + (1-s)q, tp + (1-t)q\}$ be a pq -menu where Receiver has a unique prior-optimal act, say $sp + (1-s)q$. Let $\sigma = \sigma^*$ (perfect

information). Then $c(A) = sp + (1 - s)q$ but $c_A(\sigma A) = A$.

- (iii) Suppose $u \approx -v$. By (i), c is expansive in signals and by (ii), $c(A) \neq c_A(\sigma A)$ for some A and σ (in particular, u and v are non-constant, so $u \approx -v$ implies $u \not\approx v$). Conversely, (i) and (ii) imply $u \approx -v$ if c is expansive in signals and $c(A) \neq c_A(\sigma A)$ for some A and σ .

C.2 Proof of Proposition 5

- (i) Let A be a pq -menu where $u(p) > u(q)$ and $v(p) < v(q)$. Then Sender chooses no information at A , so $\sigma^* A \succ A$. Since $\dot{\succsim}$ values information more than $\succsim$, this implies $\sigma^* A \dot{\succ} A$, so $\dot{v}(p) < \dot{v}(q)$ as well. Thus, every disagreement pair for u, v is a disagreement pair for $u, \dot{v}$. By Lemma 5, this implies $\dot{v}$ is less u -aligned than v is.

To see that the converse does not hold, consider Figure 4. There, a menu $A = \{f, g, h\}$ is constructed so that Receiver chooses f on $0 \leq \hat{\mu} < \hat{\mu}'$, g on $\hat{\mu}' < \hat{\mu} < \hat{\mu}''$, and h on $\hat{\mu}'' < \hat{\mu} \leq 1$, making Sender's value function concave under $\dot{v}$ (by Lemma 4, such a menu exists). Consequently, Sender (under $\dot{v}$) chooses no information at A , so that Receiver's ex-ante value of A is given by the lower of the two red dots. Let σ be an experiment generating posteriors at $\hat{\mu}'$ and $\hat{\mu} = 1$. Again, concavity of Sender's value function under $\dot{v}$ implies no additional information is chosen; thus, $\sigma A \dot{\sim} A$. Now consider Sender utility $v \approx \alpha u + (1 - \alpha)\dot{v}$. For appropriate values of α (in the figure, $\alpha = \frac{1}{2}$), no-information remains Sender optimal at A while non-trivial information is optimal at σA . In particular, Sender's value function under v is convex around $\hat{\mu}'$, so Sender selects information $\hat{\sigma}$ generating a mean-preserving spread of the point $\hat{\mu}'$ to $\hat{\mu} = 0$ and $\hat{\mu}''$, as indicated by the arrows in the figure (additional information has no impact on the other posterior, $\hat{\mu} = 1$, generated by σ since it is a degenerate distribution). This structure is Sender-optimal at A because it raises Sender's payoff at $\hat{\mu}'$ to the point on the concavification at $\hat{\mu}'$. This information structure raises Receiver's payoff conditional at $\hat{\mu}'$ as well which, in turn, raises his overall payoff at σA to the higher of the two red dots. Thus, $\dot{v}$ is less u -aligned than v is but $\dot{\succsim}$ does not value information more than $\succsim$ because $\sigma A \succ A$ while $\sigma A \dot{\sim} A$.

- (ii) Suppose $\dot{\succsim}$ values flexibility more than $\succsim$. Let p, q be an agreement pair for u, v . If $A = \{pEq, qEp\}$ is a pq -bet, it follows that $A \succ \{pEq\}$ (where pEq is Receiver's prior-optimal act in A), and thus $A \dot{\succ} \{pEq\}$ as well. This implies the $\dot{\succ}$ -Sender chooses perfect information at A , and so p, q is an agreement pair for $u, \dot{v}$. Thus, every u, v agreement pair is a $u, \dot{v}$ agreement pair, making $\dot{v}$ more u -aligned than v (Lemma 5).

However, $\dot{\succsim}$ also values information more than $\succsim$ because $\sigma A \supseteq A$ for all A and σ . Thus, by (i), $\dot{v}$ is also less u -aligned than v . Consequently, $\dot{v} \approx v$.

C.3 Proof of Proposition 6

- (i) Suppose $\dot{v}$ is less u -aligned than v is. Since the set of induced acts only depends on μ and u , it follows from Lemma 6 that Receiver's value of A under (μ, u, v) weakly exceeds that under $(\mu, u, \dot{v})$. Thus, $f \succsim A$ implies $f \dot{\succsim} A$.

Conversely, suppose statement (i) holds and consider a pq -bet A where $u(p) > u(q)$. Let $f \in A$ denote $\dot{\succsim}$ -Receiver's prior-optimal act. Then $A \dot{\succ} f$ if and only if $\dot{v}(p) \geq \dot{v}(q)$. A similar result holds with v in place of $\dot{v}$ and $\succ$ in place of $\dot{\succ}$. Therefore, statement (i) (in contrapositive form) implies every agreement pair for $u, \dot{v}$ is an agreement pair for u, v . Lemma 5, then, implies v is more u -aligned than $\dot{v}$ is.

- (ii) Suppose $\dot{v}$ is less u -aligned than v is and that $f \sim \sigma A \succ A$ for some $f \in \sigma A$. This implies f is prior-optimal for Receiver in σA and, thus, U -minimal in $F(\sigma A)$. Since $F(\sigma A)$ only depends on μ and u , Lemma 6 implies $\dot{\succsim}$ -Sender selects a U -minimal act from $F(\sigma A)$ as well. Thus, $f \dot{\sim} \sigma A \dot{\succ} A$.

Conversely, suppose statement (ii) holds and consider a pq -bet A where $u(p) \neq u(q)$. If $f \sim \sigma^* A \succ A$, then u, v disagree on the ranking of p, q . Condition (ii) implies $\sigma^* A \dot{\succ} A$, which means $u, \dot{v}$ disagree on the ranking as well. Thus, every disagreement pair for u, v is a disagreement pair for $u, \dot{v}$; by Lemma 5, then, $\dot{v}$ is less u -aligned than v is.

C.4 Proof of Proposition 7

First, suppose $\dot{v}$ is less u -aligned than v is. If $c(A) = f$, then f is prior-optimal for Receiver, so c -Sender must be choosing e (no information) at A . Since $\dot{v}$ is less u -aligned than v is, it follows from Lemma 6 that $\dot{c}$ -Sender chooses e as well. Thus, $\dot{c}(A) = f$.

Conversely, consider a pq -menu A . If $c(A) = f$, the c -agents disagree on the ranking of p, q ; by (ii), we have $\dot{c}(A) = f$ and so the $\dot{c}$ -agents disagree as well. Thus, every disagreement pair for u, v is a disagreement pair for $u, \dot{v}$. By Lemma 5, then, $\dot{v}$ is less u -aligned than v .

C.5 Proof of Proposition 8

Clearly, $\succsim$ is sophisticated if $v' \approx v$. Conversely, suppose $\succsim$ is sophisticated. There are three cases:

1. $u \approx v$. If $v' \not\approx v$, there exist p, q such that $u(p) > u(q)$ and $v'(p) < v'(q)$; a corresponding pq -bet A yields $p \succ A$ since a Sender with preferences v' chooses no information in this case. But sophistication together with $u \approx v$ implies $A \sim f_\lambda^A = p$ for all pq -bets A where $u(p) \geq u(q)$, a contradiction.
2. $u \approx -v$. If $v' \not\approx v$, there exist p, q such that $u(p) > u(q)$ and $v'(p) > v'(q)$; a corresponding pq -bet satisfies $p \sim A$ since a Sender with preferences v' chooses perfect information in this case. But sophistication together with $u \approx -v$ implies $p \succ f_\lambda^A \sim A$ for all pq -bets A where $u(p) > u(q)$, a contradiction.
3. $u \not\approx v$ and $u \not\approx -v$. Suppose $v' \not\approx v$. If $v' \approx u$ or $v' \approx -u$, the argument from case 1 or 2 applies. If $v' \not\approx u$ and $v' \not\approx -u$, there exist p, q such that $u(p) > u(q)$, $v(p) > v(q)$, and $v'(p) < v'(q)$; now the argument from case 1 applies.

C.6 Proof of Proposition 9

We prove statement (i) (the proof for (ii) is similar). First, suppose v' is more u -aligned than v is. Let $U' : \mathcal{A} \rightarrow \mathbb{R}$ denote Receiver's value (Persuasion Representation) of $\succsim$ with parameters (μ, u, v') , and $U : \mathcal{A} \rightarrow \mathbb{R}$ Receiver's value in a Persuasion Representation with parameters (μ, u, v) . Then $U'(f) = U(f)$ for all f and, by Lemma 6, $U'(A) \geq U(A)$ for all A . Suppose $A \not\succ f_\lambda^A$. Then $U'(A) \neq U'(f_\lambda^A) = U(f_\lambda^A) = U(A)$; thus, $U'(A) > U(A) = U'(f_\lambda^A)$ and so $A \succ f_\lambda^A$.

For the converse, suppose $\succsim$ is optimistic. We show that every (u, v) -agreement pair is a (u, v') -agreement pair; the desired result then follows from Lemma 5. Suppose toward a contradiction that there exists a (u, v) -agreement pair $\{p, q\}$ that is a (u, v') -disagreement pair. Without loss of generality, suppose $u(p) > u(q)$; then $v(p) \geq v(q)$ and $v'(p) < v'(q)$. Let A be a pq -bet. Then λ -Sender chooses perfect information but $\succsim$ -Receiver expects $\succsim$ -Sender to choose no information; consequently, $A \not\succ f_\lambda^A = p$. Since $\succsim$ is optimistic, we have $A \succ f_\lambda^A = p$; but $p \succsim A$ because A is a pq -bet and $u(p) > u(q)$. Thus, $\{p, q\}$ is a (u, v') -agreement pair.

C.7 Proof of Proposition 10

Suppose $\dot{v}$ is more u -aligned than v' . Let $U, U', \dot{U}$ denote Receiver's values (utility representations of menu preferences) under parameters (μ, u, v) , (μ, u, v') , and $(\mu, u, \dot{v})$, respectively. By Lemma 6, we have $\dot{U}(A) \geq U'(A) \geq U(A)$ for all A and $\dot{U}(f) = U'(f) = U(f)$ for all f . Suppose $A \not\succ' f_\lambda^A$. Then $U'(A) > U'(f_\lambda^A) = U(f_\lambda^A) = U(A)$, so $\dot{U}(A) \geq U'(A) > U(A) = U(f_\lambda^A) = \dot{U}(f_\lambda^A)$. Thus, $A \dot{\succ} f_\lambda^A$.

For the converse, suppose $A \succ' f_\lambda^A$ implies $A \dot{\succ} f_\lambda^A$. We show that every (u, v') -agreement pair is a $(u, \dot{v})$ -agreement pair. So, let $\{p, q\}$ be a (u, v') -agreement pair; without loss of generality, assume $u(p) > u(q)$ and $v'(p) \geq v'(q)$. Let A be a pq -bet. There are two cases. First, suppose $\{p, q\}$ is a (u, v) -agreement pair. Since $\dot{\succ}$ is optimistic, $\{p, q\}$ is also a $(u, \dot{v})$ -agreement pair (see the proof of Proposition 9). Second, suppose instead that $\{p, q\}$ is a (u, v) -disagreement pair. Then $A \succ' f_\lambda^A$ because λ -Sender chooses no information but $\dot{\succ}'$ -Receiver believes $\dot{\succ}'$ -Sender chooses perfect information since $v'(p) \geq v'(q)$. By hypothesis, it follows that $A \dot{\succ} f_\lambda^A$. Since A is a pq -bet and $\{p, q\}$ is a (u, v) -disagreement pair, this implies λ -Sender chooses no information but $\dot{\succ}$ -Receiver believes $\dot{\succ}$ -Sender chooses perfect information; thus, $\dot{v}(p) \geq \dot{v}(q)$, so that $\{p, q\}$ is a $(u, \dot{v})$ -agreement pair.

D Axiomatic Foundations

This appendix develops axiomatic foundations for Persuasion Representations. Section D.1 provides axioms for menu preferences, and section D.2 adapts that approach to state-contingent random choice data.

D.1 Menu Preferences

Let $\dot{\succ}$ denote a preference on $\mathcal{A}$. The key to characterizing the model is to determine the set of induced acts. Let

$$F_{\dot{\succ}}(A) := \left\{ \sum_{s \in \sigma} s f^s : \sigma \in \mathcal{E}, f^s \in \Delta A, f^s \dot{\succ}^s g \ \forall g \in A \right\}$$

where $f \dot{\succ}^s g \Leftrightarrow s f + (1-s)h \dot{\succ} s g + (1-s)h$ for all h . The set $F_{\dot{\succ}}(A)$ is constructed using the restriction of $\dot{\succ}$ to singleton menus; when that restriction satisfies standard expected utility axioms, $F_{\dot{\succ}}(A)$ coincides with the set $F(A)$ defined in the main text using (μ, u) revealed by $\dot{\succ}$. It is then a matter of ensuring $\dot{\succ}$ has a “Strotzian” representation over $F(A)$.

Axiom 1.

- (i) $\dot{\succ}$ is complete and transitive.
- (ii) For all $A, B \in \mathcal{A}$, $\alpha \in (0, 1)$ and $p \in \Delta X$, $A \succ B \Leftrightarrow \alpha A + (1 - \alpha)p \succ \alpha B + (1 - \alpha)p$.
- (iii) For all $A \in \mathcal{A}$, the sets $\{p \in \Delta X \mid p \succ A\}$ and $\{p \in \Delta X \mid A \succ p\}$ are open in ΔX .

Axiom 1 consists of standard expected utility axioms, slightly modified to ensure comparisons between menus and lotteries (hence, comparisons of lotteries) are consistent with expected

utility. It does not ensure comparisons between individual acts are consistent with subjective expected utility. One could easily add such axioms, but it will simplify matters to assume μ is known to the analyst. Comparisons between acts are then reduced to comparisons between lotteries as follows. First, given μ and $f \in F$, let $p_f := \sum_{\omega \in \Omega} \mu_\omega f_\omega \in \Delta X$; p_f is the **lottery equivalent** of f given μ . With this construction, the following axiom ensures comparisons between acts are consistent with expected utility under prior beliefs μ .

Axiom 2. For all $f \in F$, $f \sim p_f$.

Next, let $\bar{A} = \{f \in \sigma^* A \mid \forall \omega \forall p \in A_\omega, f_\omega \succsim p\}$. Intuitively, $\bar{A}$ consists of acts that can be induced by perfect information σ^* , given that the agent chooses optimally in each state.

Axiom 3. $\bar{A} \sim \sigma^* A \succsim A \succsim f$ for all A and $f \in A$.

Axiom 3 is a “boundedness” condition; it states that the ranking of A sits between that of $\bar{A}$ and the $\succsim$ -maximal act in A . Intuitively, these bounds coincide with perfect information and no-information, respectively, for an agent who chooses optimally given available information.

The next axiom governs how the agent ranks pq -bets, henceforth denoted A_{pq} .

Axiom 4.

- (i) If A_{pq} and B_{pq} are pq -bets, then $p \succ A_{pq}$ implies $p \succ B_{pq}$.
- (ii) If $q \succ A_{pq}$ and $r \succ A_{qr}$, then $r \succ A_{pr}$.
- (iii) If $A_{pq} \succ f$ for all $f \in A_{pq}$, then $A_{pq} \sim \sigma^* A_{pq}$.

Intuitively, Axiom 4 captures the fact that at pq -bets, Sender either chooses no-information or perfect information; these cases coincide with whether Sender and Receiver rank p and q the same way. The statement $p \succ A_{pq}$ reveals that Receiver strictly prefers p to q while Sender strictly prefers q to p ; this is so because such disagreement leads Sender to choose no-information, leading Receiver to choose his prior-optimal act in A_{pq} and end up with a payoff strictly worse than p . Thus, in the representation, $p \succ A_{pq}$ means $u(p) > u(q)$ and $v(q) > v(p)$. To aid the analysis, define $\succ^*$ on ΔX by: $p \succ^* q$ if and only if $q \succ A_{pq}$. Part (i) of Axiom 4 ensures this relation is well-defined; part (ii) ensures it is transitive; and part (iii) states the only way A_{pq} is ranked preferred to its prior-optimal act is if it is equivalent to A_{pq} together with perfect public information.

Axiom 5. For all p , the sets $\{q \succ p \mid q \succ A_{pq}\}$ and $\{q \succ p \mid A_{pq} \succsim q\}$ are convex.

In terms of the relation $\succ^*$ defined above, Axiom 5 implies the sets $\{q \succ p \mid p \succ^* q\}$ and $\{q \succ p \mid p \not\succ^* q\}$ are convex. This is needed to ensure $\succ^*$ is generated by the intersection of at most two expected-utility preferences.

Finally, let $L_{\succ}(A) := \{p_f : f \in F_{\succ}(A)\}$.

Axiom 6.

- (i) If $p \in L_{\succ}(A)$ and $p \succ A$, then $p \succ A_{pq}$ for some $q \in L_{\succ}(A)$.
- (ii) If $p \in L_{\succ}(A)$ and $A \succ p$, then there exists $q \in L_{\succ}(A)$ such that $\sigma^* A_{pq} \sim A_{pq} \succ p$.

Axiom 6 makes behavior at arbitrary menus A consistent with Bayesian Persuasion. Part (i) states that if Receiver's ranking of A is lower than some induced act at A , then Sender must strictly prefer some other induced act. Intuitively, this captures the property that Sender's preferences over induced acts take priority over Receiver's. Part (ii) states that if Receiver's ranking of A is higher than some induced act, then Sender and Receiver agree that a better induced act is feasible; this captures the property that, when indifferent, Sender breaks ties in favor of Receiver's preferences.

Theorem 5. $\succsim$ satisfies Axioms 1–6 if and only if it has a Persuasion Representation.

D.2 State-Contingent Random Choices

Once again, suppose the prior μ is known to the analyst. Given state-contingent random choice data λ , consider a ranking $\succ^\lambda$ on ΔX defined as follows:

$$p \succ^\lambda q \Leftrightarrow \exists \varepsilon\text{-neighborhood } N^\varepsilon \text{ of } q \text{ such that } c(\{p, q'\}) = p \ \forall q' \in N^\varepsilon,$$

where $c(A) := \bigcup_{\omega \in \Omega} \text{supp}(\lambda_\omega^A)$. We extend $\succ^\lambda$ to a complete relation $\succsim^\lambda$ by setting $p \sim^\lambda q \Leftrightarrow q \not\succ^\lambda p$. Once again, for any act $f \in F$ let $p_f := \sum_{\omega \in \Omega} \mu_\omega f_\omega$ denote its lottery equivalent. We extend $\succsim^\lambda$ to a ranking over $\mathcal{A}$ as follows:

$$A \succsim^\lambda B \Leftrightarrow p_{f_\lambda^A} \succsim^\lambda p_{f_\lambda^B},$$

where f_λ^A is the induced act defined in section 4.3.

Definition 9. A **Generalized Persuasion Representation** of λ consists of parameters (μ, u, v) and a behavioral strategy β such that, for every A , there exists $\sigma \in \mathcal{E}^*(A)$ such that

(i) for all $s \in S$,

$$\text{supp}(\beta^{A,s}) \subseteq \underset{f}{\text{argmax}} V^s(f) \text{ subject to } f \in \underset{g \in A}{\text{argmax}} U^s(g), \text{ and}$$

(ii) for all $f \in A$ and $\omega \in \Omega$, $\lambda_\omega^A(f) = \sum_{s \in \sigma} s_\omega \beta^{A,s}(f)$.

Given the constructions above, the following result follows from Theorem 5:

Theorem 6. λ has a Generalized Persuasion Representation if and only if $\succsim^\lambda$ satisfies Axioms 1–6.

Proof of Theorem 5

It is straightforward to verify that the axioms are satisfied by a preference $\succsim$ that admits a Persuasion Representation. So, this section only establishes the converse statement. Throughout, suppose $\succsim$ satisfies Axioms 1–6 and that a full-support $\mu \in \Delta\Omega$ is given.

Lemma 7. *If A_{pq} is a pq -bet, then:*

(i) $\sigma^* A_{pq} \succsim p$ and $\sigma^* A_{pq} \succsim q$.

(ii) If $p \succsim q$, then $\sigma^* A_{pq} \sim p$; if $q \succsim p$, then $\sigma^* A_{pq} \sim q$.

(iii) If $p \sim q$, then $\sigma^* A_{pq} \sim p \sim A_{pq}$.

Proof. Let A_{pq} be a pq -bet. Since $p, q \in \sigma^* A_{pq}$ and $\sigma^* A_{pq} \succsim f$ for all $f \in \sigma^* A_{pq}$ (Axiom 3), this establishes (i). Next, observe that either $p \in \overline{\sigma^* A_{pq}}$ or $q \in \overline{\sigma^* A_{pq}}$; therefore, Axiom 3 and (i) imply $\sigma^* A_{pq} \sim p$ if $p \succsim q$, and $\sigma^* A_{pq} \sim q$ if $q \succsim p$, establishing (ii). Finally, by Axiom 2, $f \in A_{pq}$ implies $f \sim p_f \sim \alpha p + (1 - \alpha)q$ for some $\alpha \in (0, 1)$. Therefore, by (ii) and Axiom 3, $p \sim q$ implies $p \sim \sigma^* A_{pq} \succsim A_{pq} \succsim \alpha p + (1 - \alpha)q \sim q$, establishing (iii). $\square$

Lemma 8. *Let A_{pq} be a pq -bet such that $\sigma^* A_{pq} \succ A_{pq}$ (by Axioms 1, 3 and 4, this is equivalent to $A_{pq} \sim f$ for some $f \in A_{pq}$). Then $p \not\succ q$; in particular, either $p \succ A_{pq} \succ q$ or $q \succ A_{pq} \succ p$.*

Proof. Suppose $\sigma^* A_{pq} \succ A_{pq}$; as noted above, this means $A_{pq} \sim f$ for some $f \in A_{pq}$. Thus, as in the proof of Lemma 7, we have $A_{pq} \sim \alpha p + (1 - \alpha)q$ for some $\alpha \in (0, 1)$. By Lemma 7(ii), we either have $p \sim \sigma^* A_{pq}$ or $q \sim \sigma^* A_{pq}$. Thus, by Axiom 3, either $p \sim \sigma^* A_{pq} \succ A_{pq} \sim \alpha p + (1 - \alpha)q$ or $q \sim \sigma^* A_{pq} \succ A_{pq} \sim \alpha p + (1 - \alpha)q$. The first case implies $p \succ A_{pq} \succ q$ and the second $q \succ A_{pq} \succ p$. $\square$

Lemma 9. *Suppose $p \succ^* q$ (that is, $q \succ A_{pq}$). Then:*

(i) $\sigma^* A_{pq} \succ A_{pq}$ and $q \succ p$; in particular, $q \sim \sigma^* A_{pq} \succ A_{pq} \succ p$.

(ii) For all $\alpha \in (0, 1)$ and $r \in \Delta X$, $\alpha p + (1 - \alpha)r \succ^* \alpha q + (1 - \alpha)r$.

Proof. Suppose $q \succ A_{pq}$. Since $q \in \sigma^* A_{pq}$, Axiom 3 implies $\sigma^* A_{pq} \succsim q$, and so $\sigma^* A_{pq} \succ A_{pq}$. By Lemma 8 (and $q \succ A_{pq}$), this forces $q \succ A_{pq} \succ p$. By Lemma 7(ii) and $q \succ p$, we have $q \sim \sigma^* A_{pq}$, proving (i). For (ii), observe that since $q \succ p$ and $q \succ A_{pr}$, Axiom 1 implies both $\alpha q + (1 - \alpha)r \succ \alpha p + (1 - \alpha)r$ and $\alpha q + (1 - \alpha)r \succ \alpha A_{pq} + (1 - \alpha)r$. It is easy to see that $A_{p'q'} := \alpha A_{pq} + (1 - \alpha)r$ is a $p'q'$ -bet where $p' = \alpha p + (1 - \alpha)r$ and $q' = \alpha q + (1 - \alpha)r$. Thus, we have $q' \succ A_{p'q'}$, which means $\alpha p + (1 - \alpha)r \succ^* \alpha q + (1 - \alpha)r$, as desired. $\square$

Lemma 10. *Suppose $q \succ p$. Then:*

(i) $p \succ^* q$ if and only if $\sigma^* A_{pq} \succ A_{pq}$.

(ii) $p \not\succ^* q$ if and only if $\sigma^* A_{pq} \sim A_{pq}$.

Proof. Let $q \succ p$; by Lemma 7(ii), this implies $q \sim \sigma^* A_{pq}$. If $\sigma^* A_{pq} \succ A_{pq}$, then $q \sim \sigma^* A_{pq} \succ A_{pq}$ and so $q \succ A_{pq}$; that is, $p \succ^* q$. Conversely, suppose $p \succ^* q$. Then $\sigma^* A_{pq} \sim q \succ A_{pq}$, so $\sigma^* A_{pq} \succ A_{pq}$; this establishes (i). For (ii), observe that if $p \not\succ^* q$, then $A_{pq} \succsim q$ and so $A_{pq} \succsim q \sim \sigma^* A_{pq}$. Since $\sigma^* A_{pq} \succsim A_{pq}$ (Axiom 3), this implies $\sigma^* A_{pq} \sim A_{pq}$. Conversely, suppose $\sigma^* A_{pq} \sim A_{pq}$. Then $q \sim \sigma^* A_{pq} \sim A_{pq}$, so $A_{pq} \succsim q$; that is, $p \not\succ^* q$. $\square$

Lemma 11. *If $p \not\succ^* q$, then $\alpha p + (1 - \alpha)r \not\succ^* \alpha q + (1 - \alpha)r$ for all $\alpha \in (0, 1)$ and $r \in \Delta X$.*

Proof. First, observe that by Axiom 1, $A \succsim B$ implies $\alpha A + (1 - \alpha)r \succsim \alpha B + (1 - \alpha)r$ for all $\alpha \in (0, 1)$ and $r \in \Delta X$ (the axiom implies that if $\alpha A + (1 - \alpha)r \succ \alpha B + (1 - \alpha)r$ for some $\alpha \in (0, 1)$ and $r \in \Delta X$, then $A \succ B$). Therefore, if $p \not\succ^* q$, then $A_{pq} \succsim q$ and so $\alpha A_{pq} + (1 - \alpha)r \succsim \alpha q + (1 - \alpha)r$ for all $\alpha \in (0, 1)$ and $r \in \Delta X$. Clearly, $\alpha A_{pq} + (1 - \alpha)r$ is a $p'q'$ -bet where $p' = \alpha p + (1 - \alpha)r$ and $q' = \alpha q + (1 - \alpha)r$. Thus, we have $A_{p'q'} \succsim q'$, which means $\alpha p + (1 - \alpha)r \not\succ^* \alpha q + (1 - \alpha)r$, as desired. $\square$

For every lottery p , let $C_p := \{q \in \Delta X \mid q \succ p\}$, $C_p^+ := \{q \in C_p \mid p \succ^* q\}$, and $C_p^- := \{q \in C_p \mid p \not\succ^* q\}$.

Lemma 12.

(i) $C_p^+ = C_p$ for all p if and only if $C_p^+ = C_p$ for all interior p .

(ii) $C_p^- = C_p$ for all p if and only if $C_p^- = C_p$ for all interior p .

(iii) If $C_p^+ \neq \emptyset$ and $C_p^- \neq \emptyset$, then p is a boundary point of C_p^+ and of C_p^- .

Proof. For (i), suppose $C_p^+ = C_p$ for all interior p . Let $p' \in \Delta X$. If $C_{p'}^+ \neq C_{p'}$, there exists $q' \in C_{p'}$ such that $p' \not\succ^* q'$. Let r be an interior lottery and $\alpha \in (0, 1)$. By Lemma 11, $p^* = \alpha p' + (1 - \alpha)r \not\succ^* \alpha q' + (1 - \alpha)r$. But p^* is interior, so this contradicts $C_{p^*}^+ = C_{p^*}$. For (ii), the argument is similar but uses Lemma 9(ii) in place of Lemma 11. For (iii), let $q \in C_p^+$. Then $p \succ^* q$, so $p = \alpha p + (1 - \alpha)p \succ^* \alpha q + (1 - \alpha)p$ for all $\alpha \in (0, 1)$ by Lemma 9(ii). Taking a sequence $\alpha_n \rightarrow 0$ yields a sequence $q_n := \alpha_n q + (1 - \alpha_n)p$ such that $q_n \in C_p^+$ and $q_n \rightarrow p$, establishing p as a boundary point of C_p^+ . A similar argument using existence of some $q' \in C_p^-$ and Lemma 11 in place of Lemma 9(ii) establishes p as a boundary point of C_p^- . $\square$

Lemma 13. *If $C_p^+ = C_p$ for all interior p , there are non-constant indices $u \approx -v$ such that (μ, u, v) constitute a Persuasion Representation for the restriction of $\succsim$ to the domain of bets and singleton menus. If instead $C_p^- = C_p$ for all interior p , the parameters satisfy $u \approx v$.*

Proof. If $C_p^+ = C_p$ for all interior p , it follows from Lemma 12(i) that $C_p^+ = C_p$ for all p . Let A_{pq} be a pq -bet. If $p \sim q$, then $\sigma^* A_{pq} \sim p \sim A_{pq}$ by Lemma 7(iii); thus, $A_{pq} \sim p$ which is consistent with the desired Persuasion Representation since $u(p) = u(q)$ implies $U(A_{pq}) = u(p)$. If $p \not\sim q$, suppose without loss of generality that $q \succ p$. Since $C_p^+ = C_p$, we have $p \succ^* q$, hence $\sigma^* A_{pq} \succ A_{pq}$, for all $q \in C_p$; moreover, $A_{pq} \sim f$ where $f \succsim g$ for all $g \in A_{pq}$. This is consistent with the desired representation because if $u \approx -v$, Sender chooses no information and so $U(A_{pq}) = U(f)$ where $U(f) \geq U(g)$ for all $g \in A_{pq}$. Thus, $(\mu, u, v = -u)$ constitute a Persuasion Representation for the restriction of $\succsim$ to bets and singleton menus.

If instead $C_p^- = C_p$ for all interior p , we have $C_p^- = C_p$ for all p by Lemma 12(ii). Let A_{pq} be a pq -bet. The case $p \sim q$ is as above. If $p \not\sim q$, assume without loss that $q \succ p$. Then $q \sim \sigma^* A_{pq}$ by Lemma 7(ii) and $\sigma^* A_{pq} \sim A_{pq}$ by Lemma 10 and the fact that $C_p^- = C_p$ implies $p \not\succ^* q$ for all $q \in C_p$. Thus, $q \sim A_{pq}$, which is consistent with the desired representation because $u \approx v$ and $u(q) > u(p)$ induce Sender to choose perfect information at A_{pq} , yielding $U(A_{pq}) = u(q)$. $\square$

Lemma 14. *There are non-constant indices (u, v) such that (μ, u, v) constitute a Persuasion Representation for the restriction of $\succsim$ to the domain of bets and singleton menus.*

Proof. Lemma 13 establishes the desired result for the case where either (i) $C_p^+ = C_p$ for all interior p , or (ii) $C_p^- = C_p$ for all interior p . So, we assume there is an interior lottery p such that both $C_p^+ \neq \emptyset$ and $C_p^- \neq \emptyset$. By Axiom 5, C_p^+ and C_p^- are convex and by Axiom 1(iii), $C_p^+ = C_p \cap \{q \in \Delta X \mid q \succ A_{pq}\}$ is the intersection of two open sets and therefore open.

Thus, C_p^- is closed in the relative topology on C_p . By the Separating Hyperplane Theorem, there exists $v \in \mathbb{R}^X$ and $c \in \mathbb{R}$ such that $v \cdot q < c$ for $q \in C_p^+$ and $v \cdot q \geq c$ for $q \in C_p^-$. Since probabilities sum to 1, we may replace v with $v - (c, \dots, c)$ to obtain $v \cdot q < 0$ for $q \in C_p^+$ and $v \cdot q \geq 0$ for $q \in C_p^-$; thus, for all $q \in C_p$, $v \cdot q < 0$ if and only if $q \in C_p^+$. By Lemma 12(iii), p is a boundary point of C_p^+ and C_p^- ; consequently, $v \cdot p = 0$ and we have $q \in C_p^+$ if and only if $v \cdot p > v \cdot q$.

Next, let $p' \in \Delta X$ be an arbitrary lottery and $q' \in C_{p'}$. Since p is interior, there exists $p^* \in \Delta X$ and $\alpha \in (0, 1)$ such that $p = \alpha p' + (1 - \alpha)p^*$. Suppose first that $p' \succ^* q'$. Then $p = \alpha p' + (1 - \alpha)p^* \succ^* \alpha q' + (1 - \alpha)p^* := q$ by Lemma 9(ii). Thus, $q \in C_p^+$ and so $v \cdot p > v \cdot q$; that is, $v \cdot (\alpha p' + (1 - \alpha)p^*) > v \cdot (\alpha q' + (1 - \alpha)p^*)$, which implies $v \cdot p' > v \cdot q'$. If instead $p' \not\succ^* q'$, Lemma 11 implies $p = \alpha p' + (1 - \alpha)p^* \not\succ^* \alpha q' + (1 - \alpha)p^* = q$, so $q \in C_p^-$. This means $p \not\succ^* q$, so $v \cdot p \leq v \cdot q$ and so $v \cdot p' \leq v \cdot q'$. Thus, for all $p' \in \Delta X$ and $q' \in C_{p'}$, $p' \succ^* q'$ if and only if $v \cdot p' > v \cdot q'$.

Finally, we verify that (μ, u, v) constitute a Persuasion Representation for the restriction of $\succsim$ to bets. Let A_{pq} be a pq -bet. If $p \sim q$ (that is, $u(p) = u(q)$), then $\sigma^* A_{pq} \sim p \sim A_{pq}$ by Lemma 7(iii); thus, $A_{pq} \sim p$, as desired (since $u(p) = u(q)$, a Persuasion Representation induces $U(A_{pq}) = u(p)$). If instead $p \not\sim q$ (without loss of generality, suppose $q \succ p$), then either $p \succ^* q$ or $p \not\succ^* q$. If $p \succ^* q$, then $v(p) > v(q)$, $u(p) < u(q)$ and $\sigma^* A_{pq} \succ A_{pq}$ by Lemma 10(i) and $A_{pq} \sim f$ (where $f \succsim g$ for all $g \in A_{pq}$) by Axioms 3 and 4(iv). This is the desired result because any pq -bet involving $v(p) > v(q)$, $u(p) < u(q)$ results in $U(A_{pq}) = U(f)$. Finally, if $p \not\succ^* q$, we have $u(q) > u(p)$, $v(q) \geq v(p)$, and $q \sim \sigma^* A_{pq} \sim A_{pq}$ by Lemma 10(ii); again, this is the desired result because such u, v values yield $U(A_{pq}) = u(q)$ in a Persuasion Representation. $\square$

To complete the proof of Theorem 5, let A be an arbitrary menu and p a solution to

$$\max_{q \in L_{\succsim}(A)} u(q) \text{ subject to } v(q) \geq v(q') \ \forall q' \in L(A).$$

A solution exists because $F_{\succsim}(A) = F(A)$ is compact (Lemma 1). Clearly, a solution p satisfies $p \in \operatorname{argmax}_{q \in L(A)} v(q)$ and $u(p) \geq u(q)$ for all $q \in \operatorname{argmax}_{q' \in L(A)} v(q')$. We establish $A \sim p$ by way of contradiction.

First, suppose $p \succ A$. By Axiom 6(i), there exists $q \in L(A)$ such that $v(q) > v(p)$. This contradicts the fact that p maximizes v on $L(A)$. Thus, $A \succsim p$.

Next, suppose $A \succ p$. By Axiom 6(ii), there exists $q \in L(A)$ such that $\sigma^* A_{pq} \sim A_{pq} \succ p$. Then $\sigma^* A_{pq} \succ p$, so $\sigma^* A_{pq} \sim q$ by Lemma 7(ii). Consequently, $q \succ p$ so that, by Lemma 10(ii), $p \not\succ^* q$. Since v represents $\succ^*$, this means $v(q) \geq v(p)$ and $u(q) > u(p)$. Thus, both p and q maximize v on $L(A)$ but $u(q) > u(p)$, contradicting the fact that p maximizes u on

the set $\operatorname{argmax}_{q' \in L(A)} v(q')$. Thus, $A \sim p$.

References

- Ahn, D. S., R. Iijima, Y. Le Yaouanq, and T. Sarver (2019). Behavioural characterizations of naivete for time-inconsistent preferences. *The Review of Economic Studies* 86(6), 2319–2355.
- Aliprantis, C. D. and K. C. Border (2006). *Infinite Dimensional Analysis*.
- Anscombe, F. J. and R. J. Aumann (1963). A definition of subjective probability. *The annals of mathematical statistics* 34(1), 199–205.
- Azrieli, Y. and J. Rehbeck (2022). Marginal stochastic choice. *arXiv preprint arXiv:2208.08492*.
- Bénabou, R. and J. Tirole (2002). Self-confidence and personal motivation. *The quarterly journal of economics* 117(3), 871–915.
- Blackwell, D. (1951). Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, Volume 1, pp. 93–102.
- Blackwell, D. (1953). Equivalent comparisons of experiments. *The annals of mathematical statistics* 24(2), 265–272.
- Caplin, A. and M. Dean (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review* 105(7), 2183–2203.
- Carrillo, J. D. and T. Mariotti (2000). Strategic ignorance as a self-disciplining device. *The Review of Economic Studies* 67(3), 529–544.
- Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica*, 1431–1451.
- Curello, G. and L. Sinander (2022). The comparative statics of persuasion. *arXiv preprint arXiv:2204.07474*.
- Dekel, E. and B. L. Lipman (2012). Costly self-control and random self-indulgence. *Econometrica* 80(3), 1271–1302.
- Doval, L., R. Eilat, T. Liu, and Y. Zhou (2024). Revealed information. *arXiv preprint arXiv:2411.13293*.
- Ellis, A. (2018). Foundations for optimal inattention. *Journal of Economic Theory* 173, 56–94.
- Ergin, H. and T. Sarver (2010). A unique costly contemplation representation. *Econometrica* 78(4), 1285–1339.

- Falmagne, J.-C. (1978). A representation theorem for finite random scale systems. *Journal of Mathematical Psychology* 18(1), 52–72.
- Gul, F. and W. Pesendorfer (2001). Temptation and self-control. *Econometrica* 69(6), 1403–1435.
- Gul, F. and W. Pesendorfer (2006). Random expected utility. *Econometrica* 74(1), 121–146.
- Jakobsen, A. M. (2021). An axiomatic model of persuasion. *Econometrica* 89(5), 2081–2116.
- Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review* 101(6), 2590–2615.
- Kreps, D. M. (1979). A representation theorem for “preference for flexibility”. *Econometrica: Journal of the Econometric Society*, 565–577.
- Lu, J. (2016). Random choice and private information. *Econometrica* 84(6), 1983–2027.
- Rehbeck, J. (2023). Revealed bayesian expected utility with limited data. *Journal of Economic Behavior & Organization* 207, 81–95.
- Sen, A. K. (1971). Choice functions and revealed preference. *The Review of Economic Studies* 38(3), 307–317.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies* 23(3), 165–180.