Information Selling under Prior Disagreement

Ernesto Rivera Mora* Nicolás Rodríguez†

April 2024

Abstract

This paper studies monopolistic information selling in environments in which (1) the seller has limited commitment power, and (2) the buyer and the seller hold different beliefs about the state of the world. We show that in environments with a common prior, there is no advantage to selling information sequentially; the seller cannot achieve higher revenue than by offering an experiment that fully reveals the state in one period. We find that if, on the other hand, the agents *agree to disagree* about their prior beliefs, the seller achieves a strictly higher revenue by gradually selling information over multiple periods. Moreover, increasing the number of periods of the protocol strictly increases the seller's expected revenue. In addition, in some environments, it is optimal for the seller to first offer a *free sample test*, i.e., an experiment that partially reveals information, at no charge.

*Cowles Foundation, ernesto.riveramora@yale.edu
†University of Arizona, nirodriguezgo@arizona.edu

1 Introduction

Information buyers and sellers do not always agree on the value of advice. Individuals often undervalue recommendations provided by professionals such as lawyers, physicians, dietitians, technicians, and financial advisors. As a result, these professionals frequently employ strategies that encourage buyers to reassess the value of their advice. A common tactic is the provision of "complimentary consultations" which tends, on average, to increase the buyers' willingness to pay for additional information. Essentially, complimentary consultations serve as "hooks," persuading buyers to accept higher fees they might not have initially considered.

This paper sheds light on this practice. We study monopolistic information markets in which (1) the seller has limited commitment, and (2) the seller and the buyer hold different prior beliefs regarding the state of the world. Differences in prior beliefs are often derived from differences in information. We depart from models with private information, focusing on settings in which the common prior assumption is dropped, and agents *agree to disagree* about their beliefs regarding the state of the world. Prior disagreement, for instance, may stem from *overconfidence* (Grubb, 2009), differences in *opinions* (Che and Kartik, 2009), or simply from different *views of the world* (Alonso and Câmara, 2016).

We introduce a general monopolistic framework in which an information seller interacts with an information buyer. The buyer faces a decision problem and hence is willing to pay for experiments that reveal information about the state of the world. Before the buyer selects an action in the decision problem, the seller sequentially offers Blackwell experiments to the buyer. The seller can implement experiments without incurring any cost but has limited commitment: at each period, the seller commits to honor the experiment she offers, but cannot further commit to transfers or offers after the signal of the experiment is realized.

To illustrate our main insight, we introduce a two-period example in which the seller and the buyer disagree in their prior beliefs. Consider a manager (the information buyer) who is *overconfident* in the security of her firm's data, believing the data is well-protected with high probability.[1] A technician (the information seller) has more cautious beliefs, and can offer tests that reveal whether the data is vulnerable or not. The agents' beliefs are transparent to them. Disagreement in prior beliefs leads to a disagreement in how the agents value information. The buyer, confident in the data's security, has minimal interest in purchasing information. The seller disagrees, appreciating how the information could mitigate costly errors for the buyer. The seller cannot prove the value of her information through one-period selling protocols, hence obtaining low revenue from disclosing it. Two-period selling schemes, however, can boost revenue. Consider, for instance, the following two-period scheme. In the first period, the seller offers a free sample test that partially reveals the state of the world at no charge. In the second period, the seller offers a fully revealing test by charging a high fee. We show that the seller strictly benefits from using the free sample as it, on average, reduces the buyer's confidence that the data is safe, thereby enhancing the buyer's perceived value of further information. Moreover, we show that offering an initial free sample test plus a subsequent fully-revealing test at a high price is the unique optimal two-period selling strategy.

[1] By *overconfident* we mean prior disagreement in the sense of Grubb (2009).

Our main result characterizes the agents' equilibrium payoffs for (1) any decision problem that the buyer faces, (2) any prior beliefs of the agents, and (3) any number of periods that the agents are allowed to trade. If the agents share a common prior, then there is no advantage in selling information sequentially. The seller gets the same expected revenue by either selling all information in one period or by slowly selling information over multiple periods. Under prior disagreement, however, selling all information in one period is strictly suboptimal. Intuitively, selling information gradually allows the seller to tailor an experiment that *drives* the buyer's posterior belief towards paths where the seller expects higher future payments. The drift induced by prior disagreement creates a non-trivial trade-off between maximizing present and future revenue. Our key insight is that, at each history before the last period, the trade-off is optimized by selling *some* information, but not all. Moreover, by using standard arguments of ultimatum-style games, we show that this is the only behavior consistent with subgame perfect equilibrium.

Our second main result characterizes the marginal revenue of time. While the seller strictly benefits from having extra periods of trade, we show that the marginal value of an extra period sharply decreases over time. Moreover, the seller's total revenue is bounded regardless of the number of periods of the interaction. So, even if the seller has an arbitrarily large number of periods to trade, the seller's revenue remains bounded. While a sequential interaction allows the seller to steer the buyer's beliefs toward high-revenue paths, the seller ends up selling most of the 'stock' of information at early periods, making long-run periods almost irrelevant.

Related Literature The common prior assumption has long played a prominent role in economic theory. (See Harsanyi (1968); Aumann (1976); Halpern (2002).) Nevertheless, models without a common prior are fully consistent with rationality, especially if beliefs are interpreted through a "personalistic" or "subjectivist" Bayesian lens. (See Savage (1972); Morris (1995).) Our paper demonstrates that dropping the common prior assumption has important consequences for the behavior of information monopolists.

This paper contributes to the literature on principals with limited commitment. Most of the literature primarily focuses on common-prior environments in which some agents have private information. (See Acharya and Ortner (2017); Bester and Strausz (2001); Krishna and Morgan (2008); Doval and Skreta (2022).) By contrast, we explore the implications of limited commitment with belief heterogeneity, absent any private information.

Our analysis combines tools from the literatures on dynamic programming (Stokey, 1989; Miao, 2020) and information design (Kamenica and Gentzkow, 2011; Rayo and Segal, 2010), and draws particularly from the literature on dynamic information design (Ely, Frankel, and Kamenica, 2015; Ely, 2017; Renault, Solan, and Vieille, 2017; Ely and Szydlowski, 2020; Bizzotto, Rüdiger, and Vigier, 2021; Escudé and Sinander, 2023). Our paper closely relates to Che et al. (2023), which explores a dynamic information design setting under limited commitment. We contribute to these literatures by characterizing the *"drift"* of the agents' posterior beliefs under prior disagreement. This allows us to reduce the seller's dynamic problem into a static Bayesian persuasion problem with prior disagreement (Alonso and Câmara, 2016).

This paper fits into a broad literature that analyzes information markets. Initiated by the seminal work of Arrow (1973); Admati and Pfleiderer (1986) the literature has been recently extended by Hörner and Skrzypacz (2016); Bergemann, Bonatti, and Smolin (2018); Bergemann and Bonatti (2019); Ichihashi (2021); Ali, Haghpanah, Lin, and Siegel (2022); Zhong (2022); Bergemann, Bonatti, and Gan (2022). Among these, the paper closest to this one is Hörner and Skrzypacz (2016). As in their paper, the seller's optimal selling scheme is a sequential procedure that gradually sells imperfect signals. However, there are important differences. Their results apply to a model where (1) the buyer's decision problem has two actions and two states; (2) the seller has private information; (3) the agents share a common prior; and (4) the seller has preferences about the action that the buyer takes. In contrast, our paper studies settings in which: (1) the buyer faces an arbitrary decision problem; (2) there is no private information; (3) there may be no common prior; and (4) the action taken by the buyer has no impact on the seller's payoffs.

Lastly, our paper is related to the literature that studies information monopolist offering free-samples of information and data. Zheng and Chen (2021) study optimal free-sampling strategies in settings with two periods. Our paper differs in that it allows for general T-period settings and does not impose the free-sample restriction. (The seller can charge positive prices at any period.) We show that in our main example free-sampling is the unique equilibrium strategy of the seller. So, rather than exogenously imposing this restriction, free-sampling endogenously emerges. Drakopoulos and Makhdoumi (2023) analyze a continuous-time environment in which (1) the state is normally distributed, (2) the seller offers signals that follow an exogenous normal distribution, and (3) the buyer communicates his private information to the seller. They restrict to a class of selling strategies that offer (potentially free) signals at a constant rate and sells the entire data set at the end of the interaction. Our paper differs in that (1) beliefs are arbitrary, (2) does not restrict the structure of the signals, and (3) there is no private information.

Organization of the paper The remainder of the paper is organized as follows. In the next section, we introduce the leading example. Section 3 lays out the model. In Section 4, we represent the game as a dynamic programming problem, and characterize the equilibrium payoffs. In Section 5, we show that the seller strictly benefits from interacting over more periods, but that the revenue is asymptotically bounded. Finally, Section 6 discusses some of our modeling assumptions. All proofs are collected in the appendix.

2 Example

An information buyer (a manager of a firm) faces a decision problem under uncertainty. There are two possible states: $\underline{\theta}$ (the firm's data is vulnerable) and $\overline{\theta}$ (the firm's data is safe). Denote the state space as $\Theta = \{\underline{\theta}, \overline{\theta}\}$. The set of actions is $A = \{\mathcal{U}, \mathcal{N}\}$, where $\mathcal{U}$ denotes updating the firm's firewall and $\mathcal{N}$ denotes not updating the firewall.

Assume that the information buyer has a budget of 1 to cover the cost of a security update. If the buyer updates the firewall, no breach occurs, and their final payoff is zero. If the buyer does not update the firewall and the state is $\overline{\theta}$, there is no breach, and the buyer retains their budget of 1. However, if the buyer does not update and the state is $\underline{\theta}$, a breach occurs, resulting in a cost of 2 and a final payoff of -1. Thus, $\mathcal{U}$ is the right action for $\underline{\theta}$ and $\mathcal{N}$ is the right action for $\overline{\theta}$.

The buyer's payoff function $u : A \times \Theta \rightarrow \mathbb{R}$ is summarized in the table below.

u	$\mathcal{U}$	$\mathcal{N}$
$\overline{\theta}$	0	1
$\underline{\theta}$	0	-1

Table 1 Buyer's utility function

The buyer is uncertain about the true state. Let $\nu_b \in (0, 1)$ be the buyer's prior belief of state $\overline{\theta}$. Notice that, absent any information, the buyer prefers $\mathcal{N}$ if $\nu_b \geq \frac{1}{2}$ and $\mathcal{U}$ if $\nu_b \leq \frac{1}{2}$. Figure 2.1 describes the buyer's expected utility from choosing optimally.

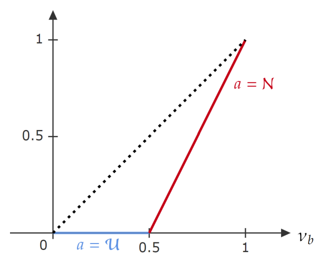

Figure 2.1 Expected utility of the buyer

For each prior belief $\nu_b \in (0, 1)$, we can compute the buyer's value of observing the state. This value is the difference between the dotted line and the buyer's expected utility represented by the blue and red lines in Figure 2.1. So, the buyer's value of fully observing the state is given by:

$$V(\nu_b) = \begin{cases} \nu_b & \text{if } \nu_b \leq \frac{1}{2} \\ 1 - \nu_b & \text{otherwise} \end{cases}$$

Figure 2.2 illustrates the function $V(\cdot)$ in terms of the prior belief ν_b. Notice, the buyer places a higher value for priors that are closer to $\frac{1}{2}$, where the buyer is more uncertain about the state of the world.

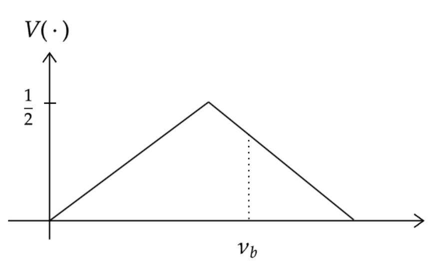

Figure 2.2 Buyer's value of observing the state for prior belief ν_b

A monopolistic information seller has access to "experiments" or "tests" that can fully or partially reveal the state of the world. The seller has a prior belief $\nu_s \in (0, 1)$ that the state is $\overline{\theta}$. Importantly, we allow ν_s to be different from ν_b. The agents' prior beliefs are transparent to them and there is no private information.

Fix a set of signals M. An experiment is a stochastic mapping $\pi : \Theta \to \Delta M$ that describes the probability of each signal conditional on each state. We assume that both agents agree about the probabilities described by π.

The agents interact in two periods. In each period t, the seller offers an experiment π^t at a fixed price $p^t \geq 0$. If the buyer purchases the experiment, she pays the price p^t to the seller, and both agents observe the realized signal $m^t \in M$. This gives rise to some posterior beliefs that are consistent with Bayes' rule.

At each period t, the seller commits to honor the priced experiment (π^t, p^t) offered to the buyer. That is, the seller commits to charge p^t and truthfully reveal the realized signal m^t of the experiment π^t. The seller cannot further commit to transfers that are contingent upon the state or the signal realization, or to provide any priced experiment at a future period.[2]

2.1 Common Prior

We start our analysis with the benchmark case where the buyer and the seller agree about their prior beliefs regarding the state.

Proposition 2.1. *If $\nu_b = \nu_s$, then the seller's maximum expected revenue is $V(\nu_b)$. In particular, it is optimal to offer an experiment that completely reveals the state in the first period.*

A fully revealing experiment in the first period gives the seller a maximum payoff of $V(\nu_b)$. Moreover, under a common prior, the seller cannot exceed this revenue by selling information sequentially. To see this, suppose that the seller's expected revenue is higher than $V(\nu_b)$. Since the agents share the same prior, the agents agree about the probability of any outcome of the experiment. Consequently, if the seller expects to receive a higher revenue than $V(\nu_b)$ then the buyer expects to pay more than $V(\nu_b)$, a payment higher than her willingness to pay for full disclosure. Consequently, the buyer is better off by not participating in the seller's scheme.

2.2 Prior Disagreement

We now consider a setting in which the agents have prior disagreement stemming from overconfidence in the sense of Grubb (2009). Consider the case in which the seller's and the buyer's prior belief about $\overline{\theta}$ are $\nu_s = 0.5$ and $\nu_b = 0.9$, respectively. Notice that the agents not only disagree about their beliefs, they also disagree about the value of information. From the seller's point of view, the value of information is $V(0.5) = 0.5$. However, the buyer is overconfident and is willing to pay only $V(0.9) = 0.1$ for fully observing the state. As a result, the maximum price the seller can charge in one period is $p = 0.1$, even though the seller believes that the information has a higher value.

[2]See Section 6.2 for a discussion of this assumption.

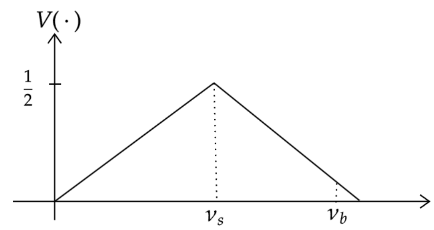

Figure 2.3 Value of information under prior disagreement

We show that the seller can strictly increase her revenue by sequentially selling information in two periods. The first experiment partially reveals the state and has zero price, i.e., it is a free sample. This experiment is described by a set of signals $M = \{\underline{m}, \overline{m}\}$ and signal mapping $\pi : \Theta \to \Delta(M)$ given by the table below:

$\pi(m \mid \theta)$	$\underline{m}$	$\overline{m}$
$\overline{\theta}$	$\frac{1}{9}$	$\frac{8}{9}$
$\underline{\theta}$	1	0

Using Bayes' rule, the posterior probability that the buyer assigns to state $\overline{\theta}$ after each signal is

$$\mu_b(\overline{m}) = \frac{0.9 \cdot \frac{8}{9}}{0.9 \cdot \frac{8}{9} + 0.1 \cdot 0} = 1 \qquad \text{and} \qquad \mu_b(\underline{m}) = \frac{0.9 \cdot \frac{1}{9}}{0.9 \cdot \frac{1}{9} + 0.1 \cdot 1} = \tfrac{1}{2}.$$

The experiment is tailored so that it maximizes the probability that the buyer has the posterior probability $\mu_b = 0.5$, which leads to the highest valuation of information. An illustration of the "posterior spread" of the free sample is given in Figure 2.4 below.

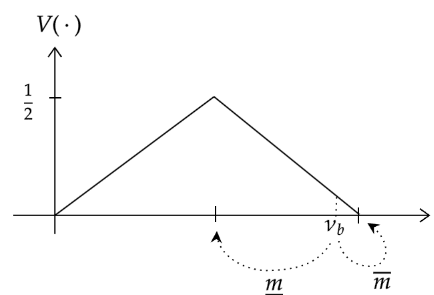

Figure 2.4 buyer's new beliefs after the free sample experiment

Note, after receiving signal $\overline{m}$, the buyer becomes certain that the state is $\overline{\theta}$, eliminating the need to purchase additional information. However, after receiving signal $\underline{m}$, the buyer's uncertainty increases. Moreover, after signal $\underline{m}$, the buyer is willing to buy a second experiment that fully reveals the state for $p = \frac{1}{2}$.

Notice that the agents disagree about the probability of the outcomes of the experiment. While the buyer believes that the probability of $\underline{m}$ is $\frac{1}{10} \cdot 1 + \frac{9}{10} \cdot \frac{1}{9} = \frac{1}{5}$, the seller believes that such probability is $0.5 \cdot 1 + 0.5 \cdot \frac{1}{9} = \frac{5}{9}$. After such a signal, the seller charges a fully revealing experiment at price $\frac{1}{2}$. Hence, the seller's expected revenue under the dynamic selling scheme is $\frac{5}{9} \cdot \frac{1}{2} = \frac{5}{18} \approx 0.27$, which is higher than the revenue from selling the information in a single period, $V(0.9) = 0.1$.

The increase in revenue can be interpreted as an exploitation of the buyer's "incorrect" prior belief.[3] When the information buyer is overconfident, he underestimates the probability of state $\underline{\theta}$ and, consequently, the probability of receiving a low signal $\underline{m}$. Thus, the free sample moves the buyer's beliefs closer to $\frac{1}{2}$, increasing (on average) the buyer's value for fully observing the state. This, in turn, allows to raise the price of full disclosure and increases expected revenue to $\frac{5}{18}$. Our main results imply that $\frac{5}{18}$ is the maximum revenue that can be achieved in two periods, and that revenue increases by selling information in three or more periods.

3 Model

Throughout the paper, take the following conventions. Endow a compact metric space C with its Borel sigma-algebra. Denote by ΔC the set of probability measures on C and endow ΔC with the topology of weak convergence. Denote the interior of ΔC by $\operatorname{int} \Delta C$. For each $c \in C$ write $\delta_c \in \Delta C$ for the probability measure that assigns probability one to the singleton $\{c\}$.

3.1 Environment

There are two agents, the buyer (denoted by b) and the seller (denoted by s). The buyer faces an individual decision problem described by a finite state space Θ, a compact set of actions A, and a continuous utility function $u : A \times \Theta \to \mathbb{R}$. We assume that there are $\theta, \theta' \in \Theta$ such that

$$\arg\max_{a \in A} u(a, \theta) \bigcap \arg\max_{a \in A} u(a, \theta') = \emptyset.$$

In this sense, the problem is not trivial, and the buyer strictly benefits from observing the state.

Each agent $i \in \{s, b\}$ has a prior belief $\nu_i \in \operatorname{int} \Delta\Theta$ about the state. The agents' prior beliefs are *transparent* to them. That is, the agents' beliefs are described by a type structure $(\Theta, (\mathcal{T}_i, \beta_i)_{i \in \{s,b\}})$, where (1) each type set $\mathcal{T}_i$ is a singleton, and (2) each belief mapping $\beta_i : \mathcal{T}_i \to \Delta(\Theta \times \mathcal{T}_{-i})$ satisfies $\operatorname{marg}_{\Theta} \beta_i(t_i) = \nu_i$. Note, in particular, that the agents do not have private information. So, if $\nu_b = \nu_s$ the agents share a *common prior*. If $\nu_b \neq \nu_s$, then the agents *agree to disagree* about their priors.

3.2 Interaction

We consider a dynamic game in which the seller sequentially offers information to a buyer. Information is revealed by implementing the experiments (in the sense of Blackwell (1951, 1953)) that the buyer chooses to purchase.

[3] The buyer's beliefs are "incorrect" from the seller's subjective point of view. In the same way as Grubb (2009) the seller's payoffs are computed using the seller's prior belief.

The timing is as follows. Nature first chooses a state of the world $\theta \in \Theta$. Agents do not directly observe the state during the interaction. After the state is realized, there is a finite sequence of periods that are indexed (backward) by $t = T, T-1, \ldots, 1, 0$. The number of periods T is exogenous. In each period $t > 0$ the agents trade information and in the last period $t = 0$, the buyer takes an action. So, t reflects the number of periods left before the buyer faces her decision problem.

We assume that there is a finite set of signals M that satisfies $|M| \geq |\Theta|$.[4] At each period $t > 0$, the seller makes a take-it-or-leave-it offer to the buyer, consisting of a priced experiment $E^t = (\pi^t, p^t)$, where $\pi^t : \Theta \to \Delta M$ is an experiment, and p^t is the experiment's price. The experiment π^t describes how signals are distributed conditional on the state. Importantly, the agents agree about the conditional distribution of signals π^t. The buyer observes E^t and decides whether to accept or reject it. If E^t is accepted, the buyer pays the price p^t to the seller, the experiment π^t is implemented, both agents observe the realized signal $m^t \in M$, and period $t-1$ starts. If E^t is rejected, no transfer is made, the experiment is not implemented, and period $t-1$ starts. In the last period, $t = 0$, the buyer decides what action $a \in A$ to take. Once the game is over, the state θ is revealed and the buyer receives utility $u(a, \theta)$.

To isolate the strategic effects of selling information, we assume that (1) there is no time-discounting, (2) the seller can implement any sequence of experiments without incurring any cost, and (3) the seller's payoff does not depend either on the state or on the action the receiver takes.[5] Thus, the seller only cares about maximizing her total revenue given by $\sum_{t=1}^{T} p^t \mathbb{1}\{E^t \text{ is accepted}\}$. The buyer has quasilinear preferences regarding the decision problem and the payments made to the seller. So, at the end of the interaction, the buyer's total utility is $u(a,\theta) - \sum_{t=1}^{T} p^t \mathbb{1}\{E^t \text{ is accepted}\}$.

The seller has limited commitment power. Within each period t, the seller commits to honor the priced experiment $E^t = (\pi^t, p^t)$, provided that the buyer accepts it. That is, the seller commits to charge p^t and truthfully reveal the realized signal of the experiment π^t. The seller cannot further commit to transfers or future priced experiments that are contingent upon the state θ or the signal realization m^t.[6]

3.3 Histories, Strategies, and Equilibrium

Write $\mathcal{E}$ for the set of all priced experiments. A history for the seller at period $T > t \geq 1$ is a sequence $h_s^t = \{(E^{t'}, c^{t'}, m^{t'}(c^{t'}))\}_{t'=T}^{t+1}$, where $E^{t'} \in \mathcal{E}$ is the experiment offered at time t', and $c^{t'} \in \{\text{accept, reject}\}$ is the buyer's choice. The entry $m^{t'}(c^{t'}) \in \text{Supp}\,(\pi^{t'})$ is the signal generated by $\pi^{t'}$ provided that $c^{t'}$ =accept, and $m^{t'}(c^{t'}) = \emptyset$ otherwise. A history for the buyer at period $T \leq t \leq 1$ is a sequence $h_b^t = \{h_s^t, E^t\}$, describing the seller's history h_s^t up to that period and the priced experiment E^t that the seller offers at period t. Write H_i^t for i's set of histories at period $t \geq 1$ and set $H_s^T = \{\emptyset\}$. Write $\mathcal{H}_i := \bigcup_{t=T}^{1} H_i^t$ for i's set of histories prior to the last period $t = 0$. A **trading history** is a sequence $h^0 = \{(E^{t'}, c^{t'}, m^{t'}(c^{t'}))\}_{t'=T}^{1}$. Denote by $\mathcal{H}^0$ the set of all trading histories.

[4]Assuming that M is finite simplifies the analysis by avoiding updating in zero probability events. Assuming that M is finite does not affect the results. (See Part (i) of Lemma 4.3.)

[5]Section 6.3 discusses the environment in which agents discount the future. Section 6.4 discusses the environment in which the seller faces non-zero costs for implementing the experiments.

[6]See Section 6.2 for a discussion of this assumption.

A **behavior strategy for the seller** is a mapping $\sigma : \mathcal{H}_s \to \Delta\mathcal{E}$ that associates a priced experiment with each history in $\mathcal{H}_s$. A **behavior strategy for the buyer** is a pair (c, α) where $c : \mathcal{H}_b \to \Delta\{\text{accept, reject}\}$ is an acceptance rule and $\alpha : \mathcal{H}^0 \to \Delta A$ is a final action rule. So, for each buyer's history $h_b^t = (h_s^t, E^t) \in \mathcal{H}^b$, $c(h_b^t)$ prescribes whether to accept or reject the priced experiment E^t given the seller's history h_s^t; and, for each $h^0 \in \mathcal{H}^0$, $\alpha(h^0)$ prescribes the action to take after acquiring the information induced by the history h^0.

In this game the agents have no private information and all signals are public. Hence, we employ subgame perfect equilibrium (SPE) as solution concept. Discussion 6.1 discusses how this solution concept is equivalent to strong perfect Bayesian equilibrium (SPBE).

4 A Dynamic Programming Approach

To analyze equilibrium behavior, we employ a dynamic programming approach that characterizes (1) how the agents' beliefs evolve after any sequence of experiments and (2) how equilibrium payoffs and behavior depend on such beliefs at any history.

4.1 Posterior Dynamics

Before describing the strategic behavior of the agents, we first describe the dynamics of the agents' posteriors after observing any sequence of signal realizations.

For each period $t \in \{T, \ldots, 1\}$, write $\mu_i^t \in \Delta\Theta$ for i's beliefs at the beginning of period t. Hence, each experiment π^t and signal $m^t \in \operatorname{Supp} \pi^t$ induce an agent i's posterior belief μ_i^{t-1}, following Bayesian updating.

We first describe how the agents' posterior beliefs evolve. To do so, for each belief $\mu \in \Delta\Theta$, write $\mathrm{PS}[\mu] := \{\tau \in \Delta(\Delta\Theta) : \mathbb{E}_\tau[\mu'] = \mu\}$ for the set of *posterior spreads* of μ. Assuming that agent i has belief μ_i^t at time t, each experiment $E^t = (\pi^t, p^t)$ induces a posterior spread $\tau_i \in \mathrm{PS}[\mu_i^t]$ for agent i. Moreover, each posterior spread $\tau_i \in \mathrm{PS}[\mu_i^t]$ such that $|\operatorname{Supp}(\tau_i)| \leq M$ is induced by some experiment $\pi^t : \Theta \to \Delta(M)$. (See Kamenica and Gentzkow (2011).)

For each $\theta \in \Theta$, write $r(\theta) = \frac{\nu_s(\theta)}{\nu_b(\theta)}$ for the agents' likelihood ratio of state θ. Notice that, given the agents' prior beliefs are in the interior of the simplex, the vector of likelihood ratios $r := (r(\theta))_{\theta \in \Theta}$ is well defined. Let $g : \Delta\Theta \to \Delta\Theta$ be given by

$$g(\mu)(\theta) := \frac{r(\theta)\mu(\theta)}{r \cdot \mu}.$$

Notice, the mapping g is the identity if and only if the agents share a common prior. This mapping describes the relation of the agents' beliefs along the entire interaction. To see this, consider an experiment $E^T = (\pi^T, p^T)$ in the first period T. The mapping g links the seller's posterior with the buyer's posterior. That is, after any signal realization $m^T \in \operatorname{Supp} \pi^T$, the agents' posterior beliefs satisfy $\mu_s^{T-1} = g(\mu_b^{T-1})$. (See Proposition 1 in Alonso and Câmara (2016).) Furthermore, since any sequence of experiments is itself an experiment, the relation $\mu_s^t = g(\mu_b^t)$ continues to hold at each subsequent period $t \geq 0$ after any sequence of experiments and realized signals.

The mapping g allows to characterize the behavior and the outcomes of both agents in terms of the posterior of just one agent. We pick the buyer's belief as a state variable and study how this belief influences the agents' behavior and payoffs.[7]

4.2 Posterior Drifts

Describing the agents' equilibrium behavior requires understanding not only how beliefs differ after each signal, but also how the agents disagree about the likelihood of the realization of such posterior beliefs. Assume that, at some period $t > 0$ the agents have beliefs (μ_b^t, μ_s^t) with $\mu_s^t = g(\mu_b^t)$. Notice, after the realization of experiment π^t, the agents not only disagree about the posterior beliefs ($\mu_s^{t-1} \neq \mu_b^{t-1}$) at period $t-1$ but also about the likelihood of the realization of each pair $(\mu_b^{t-1}, \mu_s^{t-1})$. That is, before the implementation of π^t, the seller believes that some posterior pairs $(\mu_b^{t-1}, \mu_s^{t-1})$ are more likely than what the buyer believes.

To describe this disagreement, fix a buyer's belief μ_b^t and define $\rho(\cdot \mid \mu_b^t) : \Delta\Theta \to \mathbb{R}$ by

$$\rho(\mu_b^{t-1} \mid \mu_b^t) := \frac{r \cdot \mu_b^{t-1}}{r \cdot \mu_b^t},$$

where r is the vector of likelihood ratios given by the priors (ν_s, ν_b). The following lemma shows that the mapping $\rho(\cdot \mid \mu_b^t)$ captures the agents' disagreement about their posteriors at period $t-1$, provided that the buyer's belief at period t is μ_b^t.

Lemma 4.1. *Assume that at period t the agents have prior beliefs (μ_b^t, μ_s^t) with $\mu_s^t = g(\mu_b^t)$. If π^t is an experiment such that for each agent $i \in \{s, b\}$ induces a posterior spread $\tau_i \in \mathrm{PS}[\mu_i^t]$, then for each posterior $\mu_b^{t-1} \in \Delta\Theta$,*

$$\tau_s(g(\mu_b^{t-1})) = \tau_b(\mu_b^{t-1})\rho(\mu_b^{t-1} \mid \mu_b^t).$$

Lemma 4.1 describes the agents' disagreement about the likelihood of their posterior beliefs. To illustrate this, fix an experiment π^t and let $\tau_b \in \mathrm{PS}[\mu_b^t]$ be the buyer's posterior spread it induces.

Fix a belief $\mu_b^t \in \Delta\Theta$ and write $H_{\rho>1}(\mu_b^t) := \{\mu_b^{t-1} \in \Delta\Theta \mid \rho(\mu_b^{t-1} \mid \mu_b^t) > 1\}$ for the posterior beliefs that the seller deems more likely as compared to the buyer. Notice, since the function $\rho(\mu_b^{t-1} \mid \mu_b^t)$ is linear in μ_b^{t-1}, the set $H_{\rho>1}(\mu_b^t)$ is an open half-space. Figure 4.1 illustrates this open half-space as a shaded area of the simplex. Intuitively, if an experiment induces a posterior spread $\tau_b \in \mathrm{PS}[\mu_b^t]$ such that $\mathrm{Supp}\,(\tau_b) \cap H_{\rho>1}(\mu_b^t) \neq \emptyset$, then the seller expects the average buyer's posterior to "drift" towards the region $H_{\rho>1}(\mu_b^t)$. This has an important implication: while the buyer's Bayesian plausibility condition holds from the buyer's perspective, it does not hold from the seller's perspective. The next sections will show how this drift affects equilibrium payoffs and behavior.

[7] The analysis is analogous if the analyst chooses the seller's belief as the state variable.

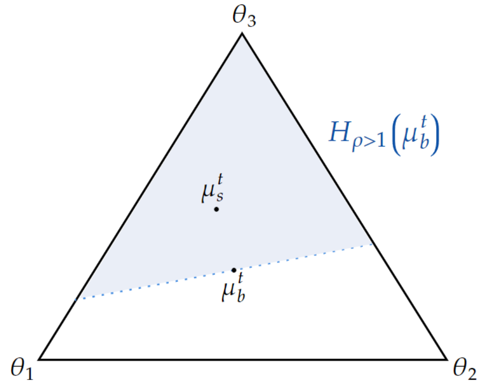

Figure 4.1 Illustration of the set $H_{\rho>1}(\mu_b^t)$.

4.3 Dynamic Programming: Last Period

We proceed by backward induction to characterize payoffs and behavior in the sequential game. We start by finding the experiment that the seller offers at period $t = 1$ after any history.

The buyer's value of each posterior belief is key in characterizing the equilibrium payoffs. With this in mind, write $\mu_b^0 \in \Delta\Theta$ for the buyer's posterior belief at period $t = 0$ and write

$$U(\mu_b^0) = \max_{a \in A} \mathbb{E}_{\mu_b^0}[u(a, \theta)]$$

for the buyer's expected payoff when he has posterior belief μ_b^0.

Lemma 4.2. *The function $U(\cdot)$ is continuous and convex.*

Following Blackwell (1951, 1953), Lemma 4.2 implies that the buyer is weakly better off under any posterior spread and, hence, is weakly better off by obtaining more information. Assume that the buyer has belief μ_b^1 at period $t = 1$. Notice, the buyer's willingness to pay for a posterior spread $\tau_b \in \mathrm{PS}[\mu_b^1]$ is non-negative and is given by $\mathbb{E}_{\tau_b}[U(\mu_b^0)] - U(\mu_b^1)$. Therefore,

$$V^1(\mu_b^1) := \sup_{\tau_b \in \mathrm{PS}[\mu_b^1]} \mathbb{E}_{\tau_b}[U(\mu_b^0)] - U(\mu_b^1)$$

is the seller's maximum revenue that can be achieved in a single period, provided that the buyer has belief $\mu_b^1 \in \Delta\Theta$. Moreover, since U is convex, the supremum defining V^1 is achieved by the posterior spread of a fully revealing experiment. Thus,

$$V^1(\mu_b^1) = \bar{U} \cdot \mu_b^1 - U(\mu_b^1),$$

where $\bar{U} = (U(\delta_\theta))_{\theta \in \Theta}$ denotes the vector of maximum utilities at each state. Note, Lemma 4.2 implies that V^1 is a continuous, non-negative, and concave mapping such that $V^1(\delta_\theta) = 0$ for each $\theta \in \Theta$. Moreover, because the buyer's decision problem is not trivial, $V^1(\mu_b^t) > 0$ for each $\mu_b^t \in \operatorname{int} \Delta\Theta$.

Notice, given a buyer's belief μ_b^1, the agents' behavior in the last period is strategically equivalent to a setting with only one period in which the buyer has prior belief $\nu_b = \mu_b^1$. The following result characterizes behavior and payoffs for one-period protocols.

Theorem 4.1. *Assume there is only one period. There exists an SPE that satisfies the following:*

(i) The seller offers a fully revealing experiment,

(ii) the buyer accepts the offer with probability one,

(iii) the seller has expected payoff $V^1(\nu_b)$, and

(iv) the buyer has expected payoff $U(\nu_b)$.

Moreover, each SPE satisfies properties (ii)-(iv).[8]

Theorem 4.1 characterize the unique equilibrium payoffs of the one-period game in terms of the model's primitives (A, u) and ν_b.[9] Notice that, in the one-period game, the seller's prior does not play a role. Hence, belief disagreements become relevant only in sequential-selling protocols.

The proof of Theorem 4.1 follows from a standard argument used in ultimatum-style games. In any SPE, the seller offers the best experiment at the maximum price that the buyer would accept. Each strategy profile in which there is positive probability that the buyer does not accept the offer is not an SPE. Under such a strategy profile, the seller would have incentives to deviate to a fully disclosing experiment at a slightly lower price.

4.4 Dynamic Programming: Previous Periods

This section characterizes equilibrium payoffs by using a backward induction argument. Fix a period $t > 1$, a belief $\mu_t \in \Delta\Theta$, and a posterior spread $\tau_b \in \mathrm{PS}[\mu_t]$. Write $\tilde{\tau}_b \in \Delta(\Delta\Theta)$ for the distribution of beliefs such that $\tilde{\tau}_b(\mu_b^{t-1}) := \tau_b(\mu_b^{t-1})\rho(\mu_b^{t-1} \mid \mu_b^t)$. By Lemma 4.1, $\tilde{\tau}_b$ describes the seller's ex-ante beliefs about the buyer's posterior μ_b^{t-1} after observing some experiment π^t. Call $\tilde{\tau}_b$ the **posterior spread τ_b from the seller's perspective**.

We inductively define a sequence of real mappings $(V^t)_{t\in\mathbb{N}}$ defined on $\Delta\Theta$. Assume that $V^{t-1} : \Delta\Theta \to \mathbb{R}$ is well-defined and write

$$V^t(\mu_b^t) := \sup_{\tau_b \in \mathrm{PS}[\mu_b^t]} \left(\mathbb{E}_{\tau_b}[U(\mu_b^{t-1})] - U(\mu_b^t) + \mathbb{E}_{\tilde{\tau}_b}[V^{t-1}(\mu_b^{t-1})] \right),$$

where $\widetilde{\tau}_b$ is the posterior spread τ_b from the seller's perspective.

We will show that the mapping V^t captures the seller's maximum total revenue that she can extract when there are $t > 1$ periods remaining and the buyer has belief $\mu_b^t \in \Delta\Theta$. (Recall that the seller has belief $\mu_s^t = g(\mu_b^t)$.) To see this, assume that the seller offers an experiment that induces a posterior spread $\tau_b \in \mathrm{PS}[\mu_b^t]$. The first component, $\mathbb{E}_{\tau_b}[U(\mu_b^{t-1})] - U(\mu_b^t)$, captures the *seller's present revenue:* the buyer's willingness to pay for an experiment that induces τ_b at the current belief μ_b^t. This relies on the fact that each agent anticipates that the buyer's continuation value

[8] Notice, if there are some $\theta, \theta' \in \Theta$ such that $\arg\max_{a\in A} u(a, \theta) = \arg\max_{a\in A} u(a, \theta')$, then there exists an optimal experiment that pools states θ and θ' into a single signal $m \in M$. Observe that, conditional on signal m, the buyer has no value from observing the true state. Therefore, condition (i) does not hold for some SPE.

[9] Indeed, the functions U and V^1 are defined by the decision problem (A, u).

for the next period $t-1$ is $U(\mu_b^{t-1})$. The second component, $\mathbb{E}_{\tilde{\tau}_b}[V^{t-1}(\mu_b^{t-1})]$, captures the *seller's future revenue*: the sum of expected transfers from optimally selling information in the remaining $t-1$ periods. Importantly, while the expectation of the first component is based on the buyer's posterior spread τ_b, the expectation of the second component is based on $\tilde{\tau}_b$—which captures the buyer's posterior spread from the seller's perspective. The difference between $\tilde{\tau}_b$ and τ_b is the key driver of our results.

Notice that in principle, the supremum defining V^t may not be attained at some period t. We show that this is not the case. Moreover, we show that the optimization problem associated to $V^t(\mu_b^t)$ can be written as a standard Bayesian persuasion problem (Kamenica and Gentzkow, 2011). To do so, fix $\mu_b^t \in \Delta\Theta$ and define the auxiliary mapping $\Lambda^t(\cdot \mid \mu_b^t) : \Delta\Theta \to \mathbb{R}$ as

$$\Lambda^t(\mu_b^{t-1} \mid \mu_b^t) := U(\mu_b^{t-1}) + V^{t-1}(\mu_b^{t-1})\rho(\mu_b^{t-1} \mid \mu_b^t).$$

The following Lemma describes the value $V^t(\mu_b^t)$ as a standard concavification problem for the objective $\Lambda^t(\cdot \mid \mu_b^t)$.

Lemma 4.3. *For each $t > 1$ and each $\mu_b^t \in \Delta\Theta$,*

$$V^t(\mu_b^t) = \sup_{\tau_b \in \mathrm{PS}[\mu_b^t]} \mathbb{E}_{\tau_b}[\Lambda^t(\mu_b^{t-1} \mid \mu_b^t)] - U(\mu_b^t).$$

Moreover,

(i) For each $\mu_b^t \in \Delta\Theta$, the supremum defining $V^t(\mu_b^t)$ is achieved for some posterior spread $\tau_b \in \mathrm{PS}[\mu_b^t]$ that has at most $|\Theta|$ elements in its support.

(ii) The mapping $V^t(\cdot)$ is continuous.

(iii) The mapping $V^t(\cdot)$ satisfies $V^{t+1}(\cdot) \geq V^t(\cdot)$ and $V^t(\delta_\theta) = 0$ for each $\theta \in \Theta$.

Lemma 4.3 states that $V^t(\mu_b^t)$ is the value of a well-defined finite-dimensional Bayesian persuasion problem. Thus, the supremum defining $V^t(\mu_b^t)$ can be found by computing the concave envelope of $\Lambda^t(\cdot \mid \mu_b^t)$ evaluated at the belief $\mu_b^{t-1} = \mu_b^t$. Moreover, part (i) indicates that the supremum is achieved by an experiment with at most $|\Theta|$ signals. So, provided that $|M| \geq |\Theta|$, there is some experiment $\pi^t : \Theta \to \Delta M$ that induces the optimal posterior spread. Part (ii) shows that V^t is continuous and therefore bounded. Part (iii) shows that the seller weakly benefits from having extra periods and that the future revenue is zero if the buyer becomes certain about the state.

Theorem 4.2. *Assume there are T periods. There exists an SPE that satisfies the following:*

(i) On the equilibrium path the buyer accepts each offer with probability one,

(ii) the seller has expected payoff $V^T(\nu_b)$, and

(iii) the buyer has expected payoff $U(\nu_b)$.

Moreover, each SPE satisfies these properties.

Theorem 4.2 characterizes the unique equilibrium payoffs of the T-period game in terms of the agent's prior beliefs. The proof follows an inductive argument reminiscent of those used in ultimatum-style games. Fix a history that induces buyer's belief μ_b^t at period t. An inductive argument shows that in each SPE the agents anticipate that the buyer's continuation value is given by $U(\mu_b^{t-1})$. So, the seller takes this buyer's outside option as given and offers an experiment that maximizes current and future payments described by the expression defining $V^t(\mu_b^t)$. Overall, this results in an expected payoff $V^T(\nu_b)$ for the seller and an expected payoff of $U(\nu_b)$ for the buyer. Notice, as in standard ultimatum-style games, the buyer accepts all offers on the equilibrium path. If, with positive probability, the buyer does not accept the offer, then the seller would have incentives to deviate to a lower price which the buyer accepts with probability one.

Observe that each history $h_s^t \in \mathcal{H}_s$ induces an information-selling game of t periods. Hence, the results of Theorem 4.2 extend to all such induced subgames, even those outside the equilibrium path. So, each history $h_s^t \in \mathcal{H}_s$ in which the buyer's initial belief is μ_b^t, the buyer accepts the equilibrium seller's offer with probability one, and the seller's and buyer's continuation expected payoffs are $V^t(\mu_b^t)$ and $U(\mu_b^t)$, respectively.

4.5 Example Revisited

We now apply the dynamic programming approach to analyze the example of Section 2. Recall that in this example the agents' priors are $(\nu_b(\overline{\theta}), \nu_s(\overline{\theta})) = (0.9, 0.5)$ and the value function U is piecewise linear.

We first characterize the case in which the agents have two periods to trade information. In the first period $t = 2$, the seller seeks an experiment that in expectation maximizes the expected value of the objective

$$\Lambda^2(\cdot \mid \nu_b) := U(\cdot) + V^1(\cdot)\rho(\cdot \mid \nu_b).$$

For each $x \in [0, 1]$, write μ_x for the belief such that $\mu_x(\overline{\theta}) = x$. Notice that for these prior beliefs $\rho(\mu_x \mid \nu_b) > 1$ if and only if $x < 0.9$. So, the seller believes that the buyer's posteriors μ_x with $x < 0.9$ are more likely relative to the buyer. Hence, the seller can tailor an experiment that drives the average buyer's posterior towards posteriors μ_x with $x < 0.9$. Figure 4.2 plots the mapping $\Lambda^2(\cdot \mid \nu_b)$ (in blue) and its concave envelope (in red dashed lines).

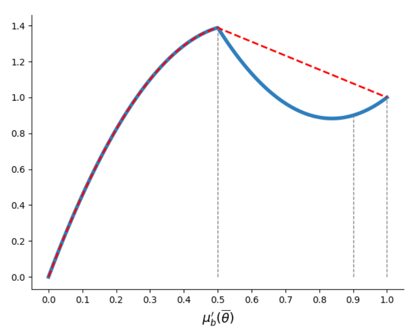

Figure 4.2 Mapping $\Lambda^2(\cdot \mid \nu_b)$ and its concave envelope.

Notice, the posterior spread $\tau \in \text{PS}[\nu_b]$ that maximizes the objective $\mathbb{E}_\tau[\Lambda^2(\mu_b^1 \mid \nu_b)]$ is such that $\text{Supp}\,(\tau) = \{\mu_{0.5}, \mu_1\}$. Moreover, the maximum price that the seller can achieve for the associated experiment is $p^1 = \mathbb{E}_\tau[U(\mu_b^1)] - U(\nu_b) = 0$. (Recall that $U(\cdot)$ is linear from $\mu_{0.5}$ to μ_1.) Therefore, provided there are two periods, the selling scheme with free sample described in Section 2 maximizes the seller's revenue and provides an expected revenue of $V^2(\nu_b) = \frac{5}{18}$.

One can apply the dynamic programming approach to analyze the example for the case in which there are three periods. In this case, the optimal information-selling scheme has the following features: In the first period, the seller provides a free sample experiment that induces a posterior spread $\tau_b^3 \in \mathrm{PS}[\nu_b]$ such that $\mathrm{Supp}\,(\tau_b^3) = \{\mu_{0.5}, \mu_1\}$. If the posterior $\mu_{0.5}$ is realized in the second period, then the seller offers an experiment inducing a posterior spread $\tau_b^2 \in \mathrm{PS}[\mu_{0.5}]$ such that $\mathrm{Supp}\,(\tau_b^2) = \{\mu_x, \mu_1\}$ for some $x \approx 0.3$. If the posterior μ_x is realized in the last period, the seller finally offers an experiment that fully reveals the state of the world. Notice, if at some period the posterior μ_1 is realized, then the buyer does not purchase further information.

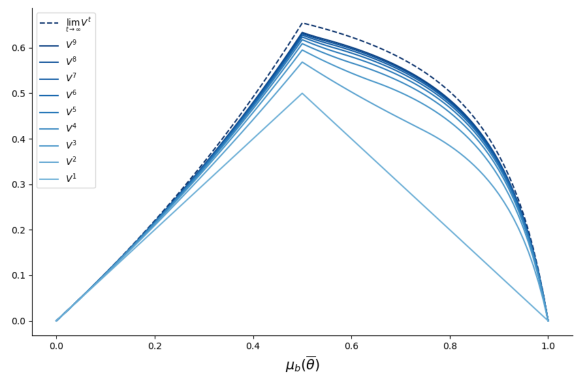

Figure 4.3 Mappings V^t in terms of $\mu_b(\overline{\theta})$ given that priors are $(\nu_b, \nu_s) = (0.9, 0.5)$.

The example can be easily extended to a sequential game with an arbitrary finite number of periods. By applying the concave envelope approach on the associated function Λ^t, one can iteratively compute the value functions $(V^t)_{t\in\mathbb{N}}$ and the optimal selling scheme for an arbitrary number of periods. Figure 4.3 illustrates the mappings $V^t(\cdot)$ for $t = 1\ldots, 9$. The asymptotic limit $\lim_{t\to\infty} V^t(\cdot)$ is plotted in dashed lines. There are two important observations. First, at each period $t > 0$ and each belief $\mu_b \in \mathrm{int}\,\Delta\Theta$, $V^{t+1}(\mu_b) > V^t(\mu_b)$. So, provided that there is information left to offer, the seller strictly benefits from having more time to interact with the buyer. Second, the value functions $(V^t)_{t\in\mathbb{N}}$ are bounded and converge to a continuous function. Section 5 shows that these two phenomena emerge in all environments with prior disagreement regardless of the buyer's decision problem (Θ, A, u), and the agents' prior beliefs (ν_b, ν_s).

5 Main Results

Theorem 4.2 describes the seller's payoffs in terms of the mapping V^T. This section explores the properties of the mappings $(V^t)_{t\in\mathbb{N}}$ as a way to describe the payoffs and behavior of the sequential game with an arbitrary number of periods.

Write $\mathcal{C}$ for the set of real continuous functions defined in $\Delta\Theta$ and write

$$\mathcal{F} = \{V \in \mathcal{C} : \ V \text{ is non-negative and } V(\delta_\theta) = 0 \text{ for each } \theta \in \Theta\}.$$

Note that Lemma 4.3 and 4.2 imply that $V^t \in \mathcal{F}$ for each $t \in \mathbb{N}$. Define the functional $\phi : \mathcal{F} \to \mathcal{F}$ as follows:

$$\phi(V)(\mu_b) := \sup_{\tau \in \mathrm{PS}[\mu_b]} \mathbb{E}_\tau[U(\mu_b') + V(\mu_b')\rho(\mu_b' \mid \mu_b)] - U(\mu_b).$$

The functional ϕ is well-defined because $\phi(V) \in \mathcal{F}$ for each $V \in \mathcal{F}$. (See Lemma A.2 in the appendix). This functional identifies the mappings $(V^t)_{t\in\mathbb{N}}$ in the sense that $V^{t+1} = \phi^t(V^1)$ for each $t \in \mathbb{N}$. Consequently, features of the agents' payoffs and behavior are derived from the monotonicity properties of the functional ϕ. (Lemma A.3 in the appendix describes these properties.)

5.1 Impact of Time

Theorem 4.2 shows that, while the buyer's expected utility remains unaffected, the seller's expected revenue may increase with the length of the interaction. Moreover, in the example of Section 2, the seller strictly benefits from having an additional period to trade information. This section explores the benefits of extra trading periods for general environments.

Write $\mathcal{D}^+ := \{\mu_b \in \Delta\Theta : \mu_b \neq g(\mu_b) \text{ and } V^1(\mu_b) > 0\}$ for the set of beliefs where the agents disagree and the buyer has a positive value for information. Lemma A.6 in the appendix shows that (1) if $\nu_s \neq \nu_b$, then $\operatorname{int}\Delta\Theta \subseteq \mathcal{D}^+$, and (2) if $\nu_s = \nu_b$ then $\mathcal{D}^+ = \emptyset$.[10] Moreover. Lemma A.7 shows that $\mathcal{D}^+ = \{\mu_b \in \Delta\Theta : V^2(\mu_b) > V^1(\mu_b)\}$. Our main result shows that $\mathcal{D}^+$ characterizes the beliefs for which the mappings $V^t(\cdot)$ strictly increase in t.

Theorem 5.1. *The following holds for each $t \in \mathbb{N}$ and each $\mu_b \in \Delta\Theta$:*

(i) If $\mu_b \notin \mathcal{D}^+$, then $V^{t+1}(\mu_b) = V^t(\mu_b)$.

(ii) If $\mu_b \in \mathcal{D}^+$, then $V^{t+1}(\mu_b) > V^t(\mu_b)$.

Theorem 5.1 follows from the monotonicity properties of the functional ϕ. Such properties imply that $\mathcal{D}^+ = \{\mu_b \in \Delta\Theta : \phi(V^2)(\mu_b) > \phi(V^1)(\mu_b)\}$. An inductive argument shows the result holds for all $t \in \mathbb{N}$.

Observe that Theorem 5.1 characterizes the seller's equilibrium payoffs for each belief μ_b, including the exogenous prior belief ν_b. Moreover, it provides a sharp difference between the case of common prior and the case of prior disagreement. Notice, under a common prior $\mathcal{D}^+ = \emptyset$. Consequently, $V^{T+1}(\nu_b) = V^T(\nu_b) = \ldots = V^1(\nu_b)$, implying that no selling strategy surpasses an experiment that fully reveals the state in the first period. By contrast, under prior disagreement $\nu_b \in \mathcal{D}^+$. As a result, $V^{T+1}(\nu_b) > V^T(\nu_b) > \ldots > V^1(\nu_b) > 0$, showing that the seller's expected revenue strictly increases in the total number of periods, T.

This result not only characterizes the equilibrium payoffs but also identifies features of the experiments that the seller offers at any history in the game. Note, since fully revealing the state at any

[10]Notice that the set $\mathcal{D}^+$ depend on both, the decision problem (Θ, A, u) and the prior beliefs (ν_b, ν_s).

stage yields a payoff of $V^1(\mu_b^t)$, it is strictly suboptimal to fully reveal the state whenever $\mu_b^t \in \mathcal{D}^+$ and $t > 1$. This implies the following corollary.

Corollary 5.1. *Assume that $\nu_s \neq \nu_b$ and $T > 1$. At each SPE the first experiment offered by the seller is not fully revealing. Moreover, if at some history at period $t > 1$, the buyer has belief $\mu_b^t \in \mathcal{D}^+$, then the seller does not offer a fully revealing experiment.*

Intuitively, Theorem 5.1 and Corollary 5.1 follow from the interplay between selling information in the current period versus future periods. On the one hand, the sequential framework allows the seller to steer the buyer's beliefs toward paths in which the seller expects larger future payments. On the other hand, this drift comes with the cost of depleting the 'stock' of information available for future interactions. So, revealing too much information diminishes the available information for future sales, and withholding too much information decreases the posterior drift. Consequently, the optimal revenue is not achieved at extremes—neither through full disclosure nor complete withholding—but rather through an experiment that partially reveals information.

5.2 Asymptotic Payoffs

This section identifies the marginal effect of time on the seller's revenue for environments in which the agents interact in an asymptotically large number of periods.

In dynamic settings, long-run outcomes are usually characterized by the convergence properties of Banach contractions. By contrast, here the functional ϕ is not a Banach contraction.[11] Despite this, we show that the set of fixed points of ϕ characterizes the agents' asymptotic payoffs.

Lemma 5.1. *Fix $\hat{\theta} \in \arg\max_{\theta\in\Theta} r(\theta)$ and $B : \Delta\Theta \to \mathbb{R}$ be defined by $B(\mu_b) := V^1(\mu_b)\rho(\delta_{\hat{\theta}} \mid \mu_b)$. The following statements hold:*

(i) The mapping B is a fixed point of ϕ.

(ii) For each $\mu_b \in \Delta\Theta$, $B(\mu_b) \geq V^1(\mu_b)$.

Lemma 5.1 constructs a fixed point of ϕ that (1) has a closed-form expression in terms of primitives, and (2) dominates the value function V^1. Moreover, due to the monotonicity properties of ϕ, the fact that B dominates V^1 implies that B dominates all the mappings $(V^t)_{t\in\mathbb{N}}$. (See Lemma A.3.) This shows that the benefits from using a sequential scheme are bounded by the multiplicative factor $\rho(\delta_{\hat{\theta}} \mid \nu_b)$.[12] Moreover, this result implies that for each $\mu_b \in \Delta\Theta$, the sequence $V^t(\mu_b)$ is monotone and bounded, and thus, its limit exists. The following result shows that the point-wise limit of the sequence $(V^t)_{t\in\mathbb{N}}$ defines a well-behaved mapping.

Theorem 5.2. *There exists a mapping $V^\infty : \Delta\Theta \to \mathbb{R}$ so that, for each $\mu_b \in \Delta\Theta$,*

$$V^\infty(\mu_b) = \lim_{t\to\infty} V^t(\mu_b).$$

Moreover, $V^\infty \in \mathcal{F}$ and V^∞ is a fixed point of ϕ.

[11] The mapping ϕ has multiple fixed points and hence is not a Banach contraction. For instance, let $B \in \mathcal{F}$ be the mapping described by Lemma 5.1. For each $\lambda \geq 1$, the mapping λB is a fixed point of ϕ.

[12] Notice that $\rho(\delta_{\hat{\theta}} \mid \nu_b) > 1$ if and only if the agents have prior disagreement.

Theorem 5.2 states that the seller's asymptotic revenue is captured by V^{∞}, a bounded and continuous mapping. This result implies that the seller's marginal revenue with respect to T sharply decreases. Intuitively, the optimal strategy reveals most of the information in the early periods, depleting the stock of information available for the final extra periods. Consequently, when T is large, an additional period generates almost no extra value for the seller.

6 Discussion

6.1 Subgame Perfect Equilibrium

As described in Section 3, the agents in this game have no private information. Moreover, since prior beliefs are transparent and all signals are public, the agents' posterior beliefs remain transparent after any sequence of experiments.

We impose the restriction that agents "cannot signal what they do not know." That is, the actions of the co-player do not convey information about the state θ. At each history, the agents' beliefs about the state depend solely on the signals selected by chance. More specifically, given a stream of experiments $\boldsymbol{\pi} = (\pi^t)_{t\in\mathcal{T}}$ purchased by the buyer at periods $\mathcal{T} \subseteq \{1, 2, \ldots, T\}$ the conditional distribution of the stream of signals $\boldsymbol{m} = (m^t)_{t\in\mathcal{T}}$ is given by

$$\mathbb{P}_{\boldsymbol{\pi}}[\boldsymbol{m} \mid \theta] := \prod_{t\in\mathcal{T}} \pi^t(m^t \mid \theta).$$

Therefore, since i has prior $\nu_i \in \operatorname{int}\Delta\Theta$, i's posterior beliefs at period $\hat{t} < \min(\mathcal{T})$ are given by

$$\mu_i^{\hat{t}}(\theta) = \frac{\nu_i(\theta)\mathbb{P}_{\boldsymbol{\pi}}[\boldsymbol{m} \mid \theta]}{\sum_{\theta'\in\Theta} \nu_i(\theta')\mathbb{P}_{\boldsymbol{\pi}}[\boldsymbol{m} \mid \theta']}. \tag{1}$$

Absent any experiment, the belief of each agent i remains fixed at the prior ν_i, even after a deviation of the co-player.

In this game, the agents' posterior beliefs are transparent after any history. Hence, subgame perfect equilibrium is equivalent to strong perfect Bayesian equilibrium (SPBE) under the following requirements: First, at each history that implements some experiments, the agents' beliefs are described by Equation (1). Second, at each history in which the buyer does not purchase any experiment, the belief of agent i equals its prior ν_i.

6.2 Commitment Power

The paper explores a setting in which the seller is endowed with limited commitment power. At each period t the seller commits to honor the priced experiment (π^t, p^t). So, if the buyer accepts to buy π^t at price p^t, the seller commits to truthfully reveal the signal m^t that π^t generates. The seller has no further commitment after the signal m^t is revealed. In particular, the seller cannot further commit to transfers or future priced experiments $(\pi^{t'}, p^{t'}) \in \mathcal{E}$ that are contingent on the state θ or the signal realization m^t.

Offering a fixed price for a "consultation" is a widespread practice for experts offering advice. Technicians, dietitians, attorneys, financial advisors, and other consultants typically set their fees prior the consultations, and the fees are not influenced by the specific nature of the advice and information dispensed. Moreover, consistent with our findings, in some instances professionals offer "complimentary consultations" as a way to engage with clients.

The widespread adoption of fixed-price consultations establishes the commitment assumptions studied here as a natural benchmark. Changing these assumptions induces different results. For instance, if (1) the seller could commit to transfers that depend on the signal realization, and (2) the amount of these transfers could be arbitrarily large, then the seller can achieve unbounded (subjective) expected revenue. Arbitrarily large levels of revenue can be achieved by designing a contingent contract that allows the agents to bet about the state at arbitrarily high stakes. Nevertheless, such contracts require the seller paying to the buyer a large transfer after some realizations of the test, behavior that is not usual for professionals selling advice.

6.3 Discounting the Future

This paper investigates the effects of belief disagreement in information markets. Our main finding reveals that the seller benefits from longer interactions. To isolate the effects of prior disagreement, we abstract from other factors that could influence our dynamic setting, such as time discounting and costs of experiments.

Incorporating a discounting parameter ($\delta \in (0, 1)$) makes deferring the payments less attractive for the seller. As a result, the seller refrains from exploiting trading opportunities in the long run and instead, opts to disclose more information in early stages. At one extreme, when the seller is highly impatient ($\delta \approx 0$), the optimal strategy entails disclosing all information within a single period. At the opposite extreme, with a sufficiently patient seller ($\delta \approx 1$), our characterization closely approximates the equilibrium payoffs and behavior of the agents.

6.4 Costly experiments

This paper assumes that the seller faces no cost for executing experiments. This benchmark covers multiple economic interactions in which the marginal cost per experiment is negligible. For instance, software firms incur zero marginal costs for running antivirus tests.

We anticipate that adding costs to the experiments will the experiments that the seller offers in equilibrium. If the seller faces a small fixed cost per experiment, the seller will reveal more information (in comparison with no-cost environment) with the goal of decreasing the expected number of experiments executed. If the cost is sufficiently big, then the seller will opt to offer only fully-revealing experiments.

References

Avidit Acharya and Juan Ortner. Progressive learning. *Econometrica*, 85(6):1965–1990, 2017.

Anat R Admati and Paul Pfleiderer. A monopolistic market for information. *Journal of Economic Theory*, 39(2):400–438, 1986.

S Nageeb Ali, Nima Haghpanah, Xiao Lin, and Ron Siegel. How to sell hard information. *The Quarterly Journal of Economics*, 137(1):619–678, 2022.

Ricardo Alonso and Odilon Câmara. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706, 2016.

Kenneth Joseph Arrow. *Information and economic behavior*, volume 28. Federation of Swedish Industries Stockholm, 1973.

Robert J Aumann. Agreeing to disagree. *The Annals of Statistics*, pages 1236–1239, 1976.

Dirk Bergemann and Alessandro Bonatti. Markets for information: An introduction. *Annual Review of Economics*, 11:85–107, 2019.

Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The design and price of information. *American Economic Review*, 108(1):1–48, 2018.

Dirk Bergemann, Alessandro Bonatti, and Tan Gan. The economics of social data. *The RAND Journal of Economics*, 53(2):263–296, 2022.

Helmut Bester and Roland Strausz. Contracting with imperfect commitment and the revelation principle: the single agent case. *Econometrica*, 69(4):1077–1098, 2001.

Jacopo Bizzotto, Jesper Rüdiger, and Adrien Vigier. Dynamic persuasion with outside information. *American Economic Journal: Microeconomics*, 13(1):179–194, 2021.

David Blackwell. Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, volume 2, pages 93–103. University of California Press, 1951.

David Blackwell. Equivalent comparisons of experiments. *The annals of mathematical statistics*, pages 265–272, 1953.

Yeon-Koo Che and Navin Kartik. Opinions as incentives. *Journal of Political Economy*, 117(5): 815–860, 2009.

Yeon-Koo Che, Kyungmin Kim, and Konrad Mierendorff. Keeping the listener engaged: a dynamic model of bayesian persuasion. *Journal of Political Economy*, 131(7):000–000, 2023.

Laura Doval and Vasiliki Skreta. Mechanism design with limited commitment. *Econometrica*, 90 (4):1463–1500, 2022.

Kimon Drakopoulos and Ali Makhdoumi. Providing data samples for free. *Management Science*, 69 (6):3536–3560, 2023.

Jeffrey Ely, Alexander Frankel, and Emir Kamenica. Suspense and surprise. *Journal of Political Economy*, 123(1):215–260, 2015.

Jeffrey C Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.

Jeffrey C Ely and Martin Szydlowski. Moving the goalposts. *Journal of Political Economy*, 128(2): 468–506, 2020.

Matteo Escudé and Ludvig Sinander. Slow persuasion. *Theoretical Economics*, 18(1):129–162, 2023.

Michael D. Grubb. Selling to overconfident consumers. *American Economic Review*, 99(5):1770–1807, 2009.

Joseph Y Halpern. Characterizing the common prior assumption. *Journal of Economic Theory*, 106 (2):316–355, 2002.

John C Harsanyi. Games with incomplete information played by "bayesian" players part ii. bayesian equilibrium points. *Management Science*, 14(5):320–334, 1968.

Johannes Hörner and Andrzej Skrzypacz. Selling information. *Journal of Political Economy*, 124 (6):1515–1562, 2016.

Shota Ichihashi. The economics of data externalities. *Journal of Economic Theory*, 196:105316, 2021.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6): 2590–2615, 2011.

Vijay Krishna and John Morgan. Contracting for information under imperfect commitment. *The RAND Journal of Economics*, 39(4):905–925, 2008.

Jianjun Miao. *Economic dynamics in discrete time*. MIT press, 2020.

Stephen Morris. The common prior assumption in economic theory. *Economics & Philosophy*, 11 (2):227–253, 1995.

Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of political Economy*, 118(5): 949–987, 2010.

Jérôme Renault, Eilon Solan, and Nicolas Vieille. Optimal dynamic information provision. *Games and Economic Behavior*, 104:329–349, 2017.

Leonard J Savage. *The foundations of statistics*. Courier Corporation, 1972.

Nancy L Stokey. *Recursive methods in economic dynamics*. Harvard University Press, 1989.

Shuran Zheng and Yiling Chen. Optimal advertising for information products. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 888–906, 2021.

Weijie Zhong. Optimal dynamic information acquisition. *Econometrica*, 90(4):1537–1582, 2022.

A Appendix

A.1 Proofs of Section 2

Proof of Proposition 2.1. The statement directly follows from Theorems 4.2 and 5.1 part (i). ■

A.2 Proofs of Section 4

Proof of Lemma 4.1. Assume that at the beginning of period t, the agents' beliefs are (μ_b^t, μ_s^t) with $\mu_s^t = g(\mu_b^t)$. Fix an experiment π^t. Let $\tau_i \in \mathrm{PS}[\mu_i]$ be the associated agent i's posterior spread. Write $M[\mu_i^{t-1} \mid \mu_i^t] \subseteq M$ for the set of messages that induce posterior μ_i^{t-1} on i given that i has prior μ_i^t. Then, notice that, for each i

$$\tau_i(\mu_i^{t-1}) = \sum_{m \in M[\mu_i^{t-1}|\mu_i^t]} \sum_{\theta \in \Theta} \pi^t(m \mid \theta)\mu_i^t(\theta).$$

Moreover, for each posterior belief μ_b^{t-1}, $M[g(\mu_b^{t-1})|\mu_b^t] = M[\mu_s^{t-1}|\mu_s^t]$. Notice, Bayes rule states that for each $m \in M[\mu_i^{t-1} \mid \mu_i^t]$,

$$\pi^t(m \mid \theta)\mu_i^t(\theta) = \mu_i^{t-1}(\theta)\left(\sum_{\theta \in \Theta} \pi^t(m \mid \theta)\mu_i^t(\theta)\right)$$

In addition, recall that

$$\mu_s^{t-1}(\theta) = g(\mu_b^{t-1})(\theta) = \frac{r(\theta)\mu_b^{t-1}(\theta)}{r \cdot \mu_b^{t-1}}. \tag{2}$$

Notice, since $\operatorname{Supp}\mu_s^t = \operatorname{Supp}\mu_b^t$, it follows that

$$\begin{aligned}
\tau_s(\mu_s^{t-1}) &= \sum_{m \in M[\mu_s^{t-1}|\mu_s^t]} \sum_{\theta \in \operatorname{Supp}\mu_s^t} \pi^t(m \mid \theta)\mu_s^t(\theta) \\
&= \sum_{m \in M[\mu_s^{t-1}|\mu_s^t]} \sum_{\theta \in \operatorname{Supp}\mu_s^t} \left(\frac{\mu_b^{t-1}(\theta)}{\mu_b^t(\theta)} \sum_{\theta' \in \Theta} \pi^t(m \mid \theta')\mu_b^t(\theta')\right)\mu_s^t(\theta) \\
&= \sum_{\theta \in \operatorname{Supp}\mu_s^t} \frac{\mu_b^{t-1}(\theta)\mu_s^t(\theta)}{\mu_b^t(\theta)} \left(\sum_{m \in M[\mu_s^{t-1}|\mu_s^t]} \sum_{\theta' \in \Theta} \pi^t(m \mid \theta')\mu_b^t(\theta')\right) \\
&= \sum_{\theta \in \operatorname{Supp}\mu_s^t} \frac{\mu_b^{t-1}(\theta)\mu_s^t(\theta)}{\mu_b^t(\theta)} \tau_b(\mu_b^{t-1}) \\
&= \sum_{\theta \in \operatorname{Supp}\mu_s^t} \frac{r(\theta)\mu_b^{t-1}(\theta)}{r \cdot \mu_b^t} \tau_b(\mu_b^{t-1}) \\
&= \frac{r \cdot \mu_b^{t-1}}{r \cdot \mu_b^t} \tau_b(\mu_b^{t-1}),
\end{aligned}$$

where the fourth equality follows from Equation (2). ■

Proof of Lemma 4.2. Fix an action $a \in A$. Notice that the function $u_a : \Delta\Theta \to \mathbb{R}$ defined by $u_a(\mu) = \sum_{\theta \in \Theta} u(a, \theta)\mu(\theta)$ is linear, which implies that it is convex and continuous. Hence, observe that $U(\mu) = \max_{a \in A} u_a(\mu)$, so U is convex and continuous. ■

Proof of Theorem 4.1. First, we show existence. Let $E^1 = (\pi^1, p^1)$ be a priced experiment that fully reveals the state at price $p^1 = V^1(\nu_b) = \overline{U} \cdot \nu_b - U(\nu_b)$. Write σ for the seller's strategy that selects E^1 at the root of the game. Write $c : \mathcal{E} \to \Delta\{\text{accept}, \text{reject}\}$ for the buyer's strategy such that satisfies the following: for each priced experiment $E' = (\pi', p') \in \mathcal{E}$,

$$c(E')(\text{accept}) = \begin{cases} 1 & \text{if } p' \leq \mathbb{E}_{\tau'}[U(\mu_b^0)] - U(\nu_b) \\ 0 & \text{otherwise,} \end{cases}$$

where $\tau' \in \text{PS}[\nu_b]$ is the posterior spread induced by π'. Finally, write α for the strategy profile such that prescribes an optimal action given the buyer's beliefs. That is, for each history $h^0 \in \mathcal{H}^0$,

$$\text{Supp}\,(\alpha(h^0)) \subseteq \arg\max_{a \in A} \mathbb{E}_{\mu_b^0}[u(a, \theta)], \tag{3}$$

where μ_b^0 is the buyer's belief at history h^0. Notice, by construction, under the profile $(\sigma, (c, \alpha))$ the buyer and the seller have no profitable deviation at any history, so it is an SPE. Moreover, note that $(\sigma, (c, \alpha))$ satisfies properties (i)-(iv).

Now, we show that each SPE satisfies properties (ii)-(iv). Fix an SPE $(\sigma, (c, \alpha))$. Notice, each history $h^0 \in \mathcal{H}^0$ must satisfy Equation (3). Hence, after any history, the buyer's value from accepting an experiment $E = (\pi, p) \in \mathcal{E}$ is $\mathbb{E}_\tau[U(\mu_b^0)] - U(\nu_b)$, where $\tau \in \text{PS}[\nu_b]$ is the posterior spread induced by π. Consequently, it must be that $c(E) = 1$ if $p < \mathbb{E}_\tau[U(\mu_b^0)] - U(\nu_b)$, and $c(E) = 0$ if $p > \mathbb{E}_\tau[U(\mu_b^0)] - U(\nu_b)$.

In the case where $p = \mathbb{E}_\tau[U(\mu_b^0)] - U(\nu_b)$, the buyer is indifferent between accepting and rejecting, so any randomization $c(E) \in \Delta\{\text{accept}, \text{reject}\}$ is optimal. We will show that the buyer must accept at least one fully revealing experiment at a price $V^1(\nu_b)$ with probability one. Suppose, by way of contradiction, that he rejects with positive probability all such priced experiments. Consequently, the seller cannot achieve an expected revenue of $V^1(\nu_b)$. However, he can achieve any strictly lower payoff since the buyer would accept with probability one any fully revealing experiment at price p, for any $p \in \mathbb{R}_+$ with $p < V^1(\nu_b)$. Therefore, the seller has no optimal choice, which contradicts that the strategy profile $(\sigma, (c, \alpha))$ is a SPE.

As a result, the buyer must accept at least one fully revealing experiment at a price $V^1(\nu_b)$ with probability one. Since such a priced experiment yields the highest possible expected revenue, then σ must prescribe choosing one of such experiments. In conclusion, any SPE satisfies properties (ii)-(iv). ■

Lemma A.1. *Fix a mapping $f : \Delta\Theta \times \Delta\Theta \to \mathbb{R}$ and let $F : \Delta\Theta \to \mathbb{R}$ be defined by $F(\mu) = \sup_{\tau \in \text{PS}[\mu]} \mathbb{E}_{\mu' \sim \tau}[f(\mu', \mu)]$. If f is continuous, then F is continuous.*

Proof. Fix $\mu \in \Delta\Theta$. We show that F is continuous at μ. Fix a sequence $(\mu_k)_{k \in \mathbb{N}}$ such that $\mu_k \in \Delta\Theta$

and $\lim \mu_k = \mu$. We show $\lim_{k\to\infty} F(\mu_k) = F(\mu)$. We divide the proof into two steps. Step one shows that $\limsup_{k\to\infty} V(\mu_k) \le V(\mu)$ and step two shows that $\liminf_{k\to\infty} V(\mu_k) \le V(\mu)$.

Step 1. Notice, since f is continuous, there is some $\tau \in \mathrm{PS}[\mu]$ such that $\mathbb{E}_\tau[f(\mu', \mu)] = F(\mu)$ (See Kamenica and Gentzkow (2011).) Moreover, there exist an affine mapping $L : \Delta\Theta \to \mathbb{R}$ such that

(i) $L(\mu) = \mathbb{E}_\tau[f(\mu', \mu)] = F(\mu)$.

(ii) $L(\mu') \ge f(\mu', \mu)$ for each $\mu' \in \Delta(\Theta)$.

Notice, since the sets $\Delta\Theta$ and $\Delta\Theta \times \Delta\Theta$ is compact, the mappings L and f are uniformly continuous. Hence, there is some $\delta > 0$ such that $||\mu - \mu_k||_\infty < \delta$ implies that for each $\mu' \in \Delta\Theta$, $|f(\mu', \mu_k) - f(\mu', \mu)| < \frac{\varepsilon}{2}$ and $|L(\mu_k) - L(\mu)| < \frac{\varepsilon}{2}$. So,

$$\begin{aligned}
F(\mu_k) &= \sup_{\tau' \in \mathrm{PS}[\mu_k]} \mathbb{E}_{\mu' \sim \tau'}[f(\mu', \mu_k)] \\
&\le \sup_{\tau' \in \mathrm{PS}[\mu_k]} \mathbb{E}_{\mu' \sim \tau'}[L(\mu') + \tfrac{\varepsilon}{2}] \\
&= L(\mu_k) + \tfrac{\varepsilon}{2} \\
&< L(\mu) + \varepsilon \\
&= F(\mu) + \varepsilon.
\end{aligned}$$

Note, since $\varepsilon > 0$ is arbitrary and $\lim_{k\to\infty} \mu_k = \mu$, it follows that $\limsup_{k\to\infty} F(\mu_k) \le F(\mu)$.

Step 2. Write $\tau \in \mathrm{PS}[\mu]$ for the posterior spread that satisfies $F(\mu) = \mathbb{E}_{\mu' \sim \tau}[f(\mu', \mu)]$. Note, by Kamenica and Gentzkow (2011), there is some finite message space M with $|M| \le |\Theta|$ and some experiment $\pi : \Theta \to \Delta(M)$ such that π induces τ. Let $\tau_k \in \mathrm{PS}[\mu_k]$ be the posterior spread induced by π under prior belief μ_k.

For each $m \in M$ and each prior belief μ' write $\mathbb{P}_\pi[m \mid \mu']$ for the probability of m under prior belief $\mu' \in \Delta\Theta$. Likewise, write $\mathcal{P}_m(\mu') \in \Delta\Theta$ for the posterior belief induced by a message $m \in M$ and prior belief $\mu' \in \Delta\Theta$. Notice that $\mathcal{P}_m(\mu')$ and $\mathbb{P}_\pi[m \mid \mu']$ are continuous at μ. Hence, for each $m \in M$, $\lim_{k\to\infty} \mathbb{P}_\pi[m \mid \mu_k] = \mathbb{P}_\pi[m \mid \mu]$ and $\lim_{k\to\infty} \mathcal{P}_m(\mu_k) = \mathcal{P}_m(\mu)$. Moreover, since f is continuous,

$$\lim_{k\to\infty} \mathbb{P}_\pi[m \mid \mu_k] f(\mathcal{P}_m(\mu_k), \mu_k) = \mathbb{P}_\pi[m \mid \mu] f(\mathcal{P}_m(\mu), \mu).$$

Therefore,

$$\begin{aligned}
\lim_{k\to\infty} \mathbb{E}_{\mu' \sim \tau_k}[f(\mu', \mu_k)] &= \lim_{k\to\infty} \sum_{\mu' \in \mathrm{Supp}\,(\tau_k)} \tau_k(\mu') f(\mu', \mu_k) \\
&= \lim_{k\to\infty} \sum_{m \in M} \mathbb{P}_\pi[m \mid \mu_k] f(\mathcal{P}_m(\mu_k), \mu_k) \\
&= \sum_{m \in M} \mathbb{P}_\pi[m \mid \mu] f(\mathcal{P}_m(\mu), \mu) \\
&= F(\mu).
\end{aligned}$$

Finally, notice that for each $k \in \mathbb{N}$,

$$F(\mu_k) = \sup_{\tau' \in \text{PS}[\mu_k]} \mathbb{E}_{\mu' \sim \tau'}[f(\mu', \mu_k)] \geq \mathbb{E}_{\mu' \sim \tau_k}[f(\mu', \mu_k)]$$

Thus, it follows that $\liminf_{k\to\infty} F(\mu_k) \geq F(\mu)$, as desired. ■

Proof of Lemma 4.3. Fix $t > 1$ and $\mu_b^t \in \Delta\Theta$. Lemma 4.1 ensures that for any posterior spread $\tau_b \in \text{PS}[\mu_b^t]$ the associated seller's beliefs about the buyer's posterior μ_b^{t-1} are given by $\tilde{\tau}_b(\mu_b^{t-1}) = \tau_b(\mu_b^{t-1})\rho(\mu_b^{t-1} \mid \mu_b^t)$. As a result,

$$\mathbb{E}_{\tilde{\tau}_b}[V^{t-1}(\mu_b^{t-1})] = \sum_{\mu_b^{t-1} \in \text{Supp}\, \tilde{\tau}_b} V^1(\mu_b^{t-1})\tau(\mu_b^{t-1})\rho(\mu_b^{t-1} \mid \mu_b^t) = \mathbb{E}_{\tau_b}[V^1(\mu_b^{t-1})\rho(\mu_b^{t-1} \mid \mu_b^t)].$$

Therefore,

$$\begin{aligned}\mathbb{E}_{\tau_b}[U(\mu_b^{t-1})] - U(\mu_b^t) + \mathbb{E}_{\tilde{\tau}_b}[V^{t-1}(\mu_b^{t-1})] &= \mathbb{E}_{\tau_b}[U(\mu_b^{t-1}) + V^1(\mu_b^{t-1})\rho(\mu_b^{t-1} \mid \mu_b^t)] - U(\mu_b^t) \\ &= \mathbb{E}_{\tau_b}[\Lambda^t(\mu_b^{t-1} \mid \mu_b^t)] - U(\mu_b^t).\end{aligned}$$

From here we conclude that

$$V^t(\mu_b^t) = \sup_{\tau_b \in \text{PS}[\mu_b^t]} \mathbb{E}_{\tau_b}[\Lambda^t(\mu_b^{t-1} \mid \mu_b^t)] - U(\mu_b^t).$$

In addition, notice that condition (i) follows from Proposition 9 in the working paper version of Kamenica and Gentzkow (2011). As for condition (ii), observe that Lemmas A.1 and 4.2 ensure that $\sup_{\tau_b \in \text{PS}[\mu_b^t]} \mathbb{E}_{\tau_b}[\Lambda^t(\mu_b^{t-1} \mid \mu_b^t)]$ and $U(\mu_b^t)$ are continuous on μ_b^t. Hence, $V^t(\mu_b^t)$ is a continuous mapping.

Finally, to prove condition (iii), we first show that $V^{t+1} \geq V^t$ by arguing that a seller with $t+1$ periods can offer a non-informative experiment in period $t+1$ and then behave optimally from period t on. Second, we show that $V^t(\delta_\theta) = 0$ for each $t \in \mathbb{N}$ and each $\theta \in \Theta$ through an inductive argument.

First, fix $\mu \in \Delta\Theta$ and $t \geq 1$. Notice that $\Lambda^{t+1}(\mu \mid \mu) = U(\mu) + V^t(\mu)\rho(\mu \mid \mu) = U(\mu) + V^t(\mu)$. In addition, observe that $\delta_\mu \in \text{PS}[\mu]$. As a result,

$$\begin{aligned}V^{t+1}(\mu) &= \sup_{\tau \in \text{PS}[\mu]} \mathbb{E}_\tau[\Lambda^{t+1}(\mu' \mid \mu)] - U(\mu) \\ &\geq \mathbb{E}_{\delta_\mu}[\Lambda^{t+1}(\mu' \mid \mu)] - U(\mu) \\ &= \Lambda^{t+1}(\mu \mid \mu) - U(\mu) \\ &= V^t(\mu).\end{aligned}$$

Second, fix $\theta \in \Theta$. We will show that $V^1(\delta_\theta) = 0$ for each $t \in \mathbb{N}$. We proceed by induction. Notice that $V^1(\delta_\theta) = \overline{U} \cdot \delta_\theta - U(\delta_\theta) = 0$. Now, fix $t \geq 1$. Assume that $V^t(\delta_\theta) = 0$. Since δ_θ is an extreme point of $\Delta\Theta$, it cannot be written as a non-trivial convex combination of elements of $\Delta\Theta$.

Hence, $\text{PS}[\delta_\theta] = \{\delta_\theta\}$. As a result,

$$\begin{aligned}V^{t+1}(\delta_\theta) &= \sup_{\tau \in \text{PS}[\mu]} \mathbb{E}_\tau[\Lambda^{t+1}(\mu' \mid \mu)] - U(\mu) \\ &= \Lambda^t(\delta_\theta \mid \delta_\theta)] - U(\delta_\theta) \\ &= V^t(\delta_\theta) \\ &= 0.\end{aligned}$$

In conclusion, $V^t(\delta_\theta) = 0$ for each $t \in \mathbb{N}$. ■

Proof of Theorem 4.2. We divide the proof into two steps. Step 1 shows existence of a SPE that satisfies conditions (i)-(iii). Step 2 shows that each SPE satisfies conditions (i)-(iii).

Step 1. We proceed by induction on T. Notice that the base case ($T = 1$) is shown in Theorem 4.1.

Now we show the inductive step. Fix $T \geq 1$. Assume that, for every pair of agents' beliefs (μ_b^T, μ_s^T) with $\mu_s^T = g(\mu_b^T)$ at the beginning of period T, there exists an SPE $(\sigma', (c', \alpha'))$ as described in Theorem 4.2 in the game with T periods. We will use these strategies to construct the SPE in the game of $T + 1$ periods. For any seller's (non-initial) history at period $T + 1 > t \geq 1$, $h_s^t = \{(E^{t'}, c^{t'}, m^{t'}(c^{t'})\}_{t'=T+1}^{t+1}$, write $\hat{h}_s^t = \{(E^{t'}, c^{t'}, m^{t'}(c^{t'})\}_{t'=T}^{t+1}$ for the associated pruned seller's history, which describes the history of play in the subgame that starts after $\{\emptyset, (E^{T+1}, c^{T+1}, m^{T+1}(c^{T+1})\}$. Similarly, for any buyer's history h_b^t, write $\hat{h}_b^t$ for the associated pruned buyer's history that describes the history of play in the subgame that starts after $\{\emptyset, (E^{T+1}, c^{T+1}, m^{T+1}(c^{T+1}), E_T\}$. Last, for every trading history h^0, write $\hat{h}^0$ for the associated pruned trading history that describes the history of play in the subgame that starts after $\{\emptyset, (E^{T+1}, c^{T+1}, m^{T+1}(c^{T+1})\}$.

Now, let $\tilde{\pi}^{T+1}$ be the experiment inducing an optimal posterior spread $\tilde{\tau}$ such that (1) $|\text{Supp}\,(\tilde{\tau})| \leq |\Theta|$ and (2) $V^{t+1}(\nu_b) = \mathbb{E}_\tau[\Lambda^{t+1}(\mu' \mid \nu_b)] - U(\nu_b)$. Set $\tilde{p}^{T+1} = \mathbb{E}_{\tilde{\tau}}[U(\mu_b^T)] - U(\nu_b)$, and $\tilde{E}^{T+1} = (\tilde{\pi}^{T+1}, \tilde{p}^{T+1})$. Consider the seller's strategy given by:

$$\sigma(h_s^t) = \begin{cases} \widetilde{E}^{T+1} & \text{if } t = T+1 \\ \sigma'(\hat{h}_s^t) & \text{if } t \leq T. \end{cases}$$

The buyer's acceptance rule defined by:

$$c(h_b^t)(\text{accept}) = \begin{cases} 1 & \text{if } h_b^t = \{(\pi^{T+1'}, p^{T+1'})\} \text{ and } p^{T+1'} \leq \mathbb{E}_{\tau'}[U(\mu_b^T)] - U(\nu_b) \\ 0 & \text{if } h_b^t = \{(\pi^{T+1'}, p^{T+1'})\} \text{ and } p^{T+1'} > \mathbb{E}_{\tau'}[U(\mu_b^T)] - U(\nu_b) \\ c'(\hat{h}_b^t) & \text{if } t \leq T. \end{cases}$$

where $c'(\hat{h}_b^t)$ is the acceptance rule in the SPE associated with the initial buyer's belief induced by the subgame that starts after $\{\emptyset, (E^{T+1}, c^{T+1}, m^{T+1}(c^{T+1})\}$. Likewise, the buyer's final action rule given by $\alpha(h^0) = \alpha'(\hat{h}^0)$ where $\alpha'(\hat{h}^0)$ is the final choice rule in the SPE associated to the initial buyer's belief induced by the subgame that starts after $\{\emptyset, (E^{T+1}, c^{T+1}, m^{T+1}(c^{T+1})\}$.

To show that $(\sigma, (c, \alpha))$ is an SPE, we appeal to the one-shot deviation principle. Since $(\sigma', (c', \alpha'))$ is an SPE for any game of T periods, then no single-deviation is optimal at any history after period T.

Now, a buyer's history at period $T+1$ is characterized by an initial priced experiment (π^{T+1}, p^{T+1}) offered. In the case that $p^{T+1} \leq \mathbb{E}_\tau[U(\mu_b^T)] - U(\nu_b)$, a buyer's deviation would imply a positive probability of rejection, which implies that the buyer's continuation payoffs will be a convex combination of $U(\nu_b)$ and $\mathbb{E}_\tau[U(\mu_b^T)] - p^{T+1}$, which will be weakly smaller than the payoff from sticking to c, $\mathbb{E}_\tau[U(\mu_b^T)] - p^{T+1}$. In the case that $p^{T+1} > \mathbb{E}_\tau[U(\mu_b^T)] - U(\nu_b)$, a buyer's deviation would imply a positive probability of accepting the experiment, which implies that the buyer's continuation payoffs will be a convex combination between $U(\nu_b)$ and $\mathbb{E}_\tau[U(\mu_b^T)] - p^{T+1}$, which will be weakly smaller than the payoff from sticking to c, $U(\nu_b)$.

Finally, a one-shot deviation for the seller at the root of the game would imply choosing other priced experiment than $\tilde{E}^{T+1}$. A deviation to the priced experiment (π^{T+1}, p^{T+1}) would imply a seller's expected revenue of either $p^{T+1} + \mathbb{E}_{\tilde{\tau}}[V^T(\mu_b^T)] \leq \mathbb{E}_\tau[U(\mu_b^T)] - U(\nu_b) + \mathbb{E}_{\tilde{\tau}}[V^T(\mu_b^T)]$ or $V^T(\nu_b^T)$, which, either way, is lower than $V^{T+1}(\nu_b) = \sup_{\tau_b \in \mathrm{PS}[\nu_b]} \mathbb{E}_{\tau_b}[U(\mu_b^T)] - U(\nu_b) + \mathbb{E}_{\tilde{\tau}_b}[V^T(\mu_b^T)]$. Hence, the seller has no incentives to deviate once at the root and then conform back to σ.

Step 2. Fix a SPE $(\sigma, (c, \alpha))$. To see that $(\sigma, (c, \alpha))$ satisfies condition (i), first notice that in any SPE, it must be that $\mathrm{Supp}\, \alpha(h^0) \subseteq \arg\max_{a \in A} \mathbb{E}_{\mu_b^0}[u(a, \theta)] = \arg\max_{a \in A} U(\mu_0^b)$ where μ_b^0 is the buyer's belief at history h^0. Therefore, at any trading history $h_b^t = \{h_s^t, E^t\} \in \mathcal{H}_b$ in which the buyer's initial belief is μ_b^t, the buyer's value of accepting the priced experiment $E^t = (\pi^t, p^t)$ is $\mathbb{E}_\tau[U(\mu_b^{t-1})] - U(\mu_b^t)$, where τ is the posterior spread induced by π^t. As a result, it must be that

$$c(h_b^t)(\text{accept}) = \begin{cases} 1 & \text{if } p^t < \mathbb{E}_\tau[U(\mu_b^{t-1})] - U(\mu_b^t) \\ 0 & \text{if } p^t > \mathbb{E}_\tau[U(\mu_b^{t-1})] - U(\mu_b^t) \end{cases}.$$

In the case that $p^t = \mathbb{E}_\tau[U(\mu_b^{t-1})] - U(\mu_b^t)$, the receiver is indifferent between accepting and rejecting the offer, so any randomization is optimal. We will show that the buyer must accept with probability one at least one experiment inducing the posterior spread τ_b that defines $V^t(\mu_b^t)$ at a price $p^t = \mathbb{E}_{\tau_b}[U(\mu_b^{t-1})] - U(\mu_b^t)$. Suppose, by contradiction, that the buyer rejects all such priced experiments with positive probability. Hence, the seller cannot achieve an expected revenue stream of $V^t(\mu_b^t)$. However, he can achieve any strictly lower payoff stream since the buyer would accept with probability one any priced experiment inducing posterior spread τ_b at price p, for any $p \in \mathbb{R}$ with $p < \mathbb{E}_{\tau_b}[U(\mu_b^{t-1})] - U(\mu_b^t)$. Therefore, the seller has no optimal choice, which contradicts that the strategy profile $(\sigma, (c, \alpha))$ is a SPE.

As a result, the buyer must accept at least one such priced experiment with probability one. Since this experiment yields the highest possible expected revenue stream, then σ must prescribe choosing one of such experiments. In conclusion, any SPE satisfies property (i). Notice that this implies that the seller's expected revenue stream at any given history h_s^t is given by $V^t(\mu_b^t)$ where μ_b^t is the buyer's initial belief at period t in such a history. In particular, the seller's ex-ante expected revenue is $V^T(\nu_b)$, so property (ii) holds. In addition, observe that at any on-path buyer's history h_b^t with initial buyer's belief μ_b^t, we have that the buyer accepts the seller's offered experiment at a price $p^t = \mathbb{E}_{\tau_b}[U(\mu_b^{t-1})] - U(\mu_b^t)$. This implies that the buyer's continuation payoffs at such a history are $U(\mu_b^t)$. In particular, the ex-ante buyer's expected payoff is $U(\nu_b)$. In other words, any SPE satisfies property (iii). ■

A.3 Proofs of Section 5

Lemma A.2. *Assume that $V \in \mathcal{F}$. Then, the mapping $\phi(V) : \Delta\Theta \to \mathbb{R}$ is continuous and satisfies $\phi(V)(\delta_\theta) = 0$ for each $\theta \in \Theta$.*

Proof. Notice that $V \in \mathcal{F}$ implies V is continuous. Hence, the mapping defined by $f(\mu', \mu) = U(\mu') + V(\mu')\rho(\mu' \mid \mu)$ is continuous. Thus, the mapping $\phi(V)(\cdot) : \Delta\Theta \to \mathbb{R}$ is continuous. (See Lemma A.1.)

Fix $\theta \in \Theta$. Notice $\tau \in \text{PS}[\delta_\theta]$ if and only if $\text{Supp}\,(\tau) = \{\delta_\theta\}$. Hence

$$\phi(V)(\delta_\theta) = \mathbb{E}_\tau[U(\mu') + V(\mu')\rho(\mu' \mid \delta_\theta)] = U(\delta_\theta) + V(\delta_\theta) - U(\delta_\theta) = 0.$$

Therefore, $\phi(V) \in \mathcal{F}$. ■

Lemma A.3. *Fix a pair of mappings $V, W \in \mathcal{F}$ such that $W(\mu) \geq V(\mu)$ for each $\mu \in \Delta\Theta$, and write $\mathcal{S} = \{\mu \in \Delta\Theta : V(\mu) = W(\mu)\}$. The following properties hold:*

(i) For each $\mu \in \Delta\Theta$, $\phi(V)(\mu) \geq V(\mu)$.

(ii) For each $\mu \in \Delta\Theta$, $\phi(W)(\mu) \geq \phi(V)(\mu)$.

(iii) If $\phi(W)(\mu) = \phi(V)(\mu)$, then there is some $\tau \in \text{PS}[\mu]$ with $\text{Supp}\,(\tau) \subseteq \mathcal{S}$ such that

$$\phi(V)(\mu) = \mathbb{E}_\tau[U(\mu') + V(\mu')\rho(\mu' \mid \mu)] - U(\mu) = \phi(W)(\mu) \tag{4}$$

Proof. We first show part (i). Fix $\mu \in \Delta\Theta$ and $V \in \mathcal{F}$. Since $\delta_\mu \in \text{PS}[\mu]$, then

$$\begin{aligned}
\phi(V)(\mu) &\geq \mathbb{E}_{\delta_\mu}[U(\mu') + V(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&= U(\mu) + V(\mu)\rho(\mu \mid \mu) - U(\mu) \\
&= V(\mu).
\end{aligned}$$

We now show part (ii). Fix $\mu \in \Delta\Theta$ and notice that

$$\begin{aligned}
\phi(W)(\mu) &= \sup_{\tau' \in \text{PS}[\mu]} \mathbb{E}_{\tau'}[U(\mu') + W(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&\geq \sup_{\tau' \in \text{PS}[\mu]} \mathbb{E}_{\tau'}[U(\mu') + V(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&= \phi(V)(\mu).
\end{aligned}$$

Now we show part (iii). Assume that $\phi(W)(\mu) = \phi(V)(\mu)$. Notice, since $\phi(V)$ is continuous, then $U(\mu') + \phi(V)(\mu')\rho(\mu' \mid \mu)$ is continuous in μ'. Thus, there is some $\tau \in \text{PS}[\mu]$ with finite support such that

$$\phi(V)(\mu) = \max_{\tau' \in \mathrm{PS}[\mu]} \mathbb{E}_{\tau'}[U(\mu') + V(\mu')\rho(\mu' \mid \mu)] = \mathbb{E}_{\tau}[U(\mu') + V(\mu')\rho(\mu' \mid \mu)] \tag{5}$$

Thus,

$$\begin{aligned}
\phi(V)(\mu) &= \mathbb{E}_{\tau}[U(\mu') + V(\mu')\rho(\mu' \mid \mu)] \\
&\leq \mathbb{E}_{\tau}[U(\mu') + W(\mu')\rho(\mu' \mid \mu)] \\
&= \max_{\tau' \in \mathrm{PS}[\mu]} \mathbb{E}_{\tau'}[U(\mu') + W(\mu')\rho(\mu' \mid \mu)] \\
&\leq \phi(W)(\mu) \\
&= \phi(V)(\mu),
\end{aligned}$$

which implies Equation (4). Moreover, if $\phi(V)(\mu') > V(\mu')$ for some $\mu' \in \operatorname{Supp}(\tau)$, then we obtain $\phi^2(V)(\mu) > \phi(V)(\mu)$, a contradiction. Thus, $\phi(V)(\mu') = V(\mu')$ for each $\mu' \in \operatorname{Supp}\tau$, as desired. ■

Lemma A.4. *Fix $\mu \in \Delta\Theta$. The following statements are equivalent:*

(i) $g(\mu) = \mu$.

(ii) $r(\theta) = r(\theta')$ *for each* $\theta, \theta' \in \operatorname{Supp}\mu$.

(iii) $\rho(\mu' \mid \mu) = 1$ *for each* $\mu' \in \Delta\Theta$ *with* $\operatorname{Supp}\mu' \subseteq \operatorname{Supp}\mu$.

(iv) $g(\mu') = \mu'$ *for all* $\mu' \in \Delta\Theta$ *with* $\operatorname{Supp}\mu' \subseteq \operatorname{Supp}\mu$.

Proof. Fix $\mu \in \Delta\Theta$. We first prove that condition (i) is equivalent to condition (ii). Assume that $g(\mu) = \mu$. Let $\theta \in \operatorname{Supp}\mu$, then $g(\mu)(\theta) = \frac{r(\theta)\mu(\theta)}{r \cdot \mu} = \mu(\theta) > 0$. This implies that $r(\theta) = r \cdot \mu$, which is independent of θ. Conversely, assume that for each $\theta \in \operatorname{Supp}\mu$, $r(\theta) = c$ is constant. Then, $g(\mu)(\theta) = \frac{r(\theta)\mu(\theta)}{r \cdot \mu} = \frac{c\mu(\theta)}{c} = \mu(\theta)$.

Now, we prove that condition (ii) implies conditions (iv) and (iii). Assume that $r(\theta) = c$ is constant for all $\theta \in \operatorname{Supp}\mu$. Fix $\mu' \in \Delta\operatorname{Supp}\mu$. Then, $\operatorname{Supp}\mu' \subset \operatorname{Supp}\mu$. Therefore, for each $\theta \in \operatorname{Supp}\mu' \subset \operatorname{Supp}\mu$, we conclude that $r(\theta) = c$. Hence, $g(\mu') = \mu'$ given the equivalence between conditions (i) and (ii). Moreover, $\rho(\mu' \mid \mu) = \frac{r \cdot \mu'}{r \cdot \mu} = \frac{c}{c} = 1$.

Next, we show that condition (iii) implies condition (ii). Assume that $\rho(\mu' \mid \mu) = 1$ for each $\mu' \in \Delta\Theta$ with $\operatorname{Supp}\mu' \subset \operatorname{Supp}\mu$. Then, for any $\theta \in \operatorname{Supp}\mu$, we have that $1 = \rho(\delta_\theta \mid \mu) = \frac{r(\theta)}{r \cdot \mu}$. Thus, $r(\theta) = r \cdot \mu$ is constant as it is independent of θ.

Finally, notice that condition (iv) implies condition (i) since $\operatorname{Supp}\mu \subset \operatorname{Supp}\mu$, so $g(\mu) = \mu$. ■

Lemma A.5. *Fix $\mu \in \Delta\Theta$. The following statements are equivalent:*

(i) $V^1(\mu) = 0$.

(ii) $\arg\max_{a \in A} \mathbb{E}_\mu[u(a, \theta)] \subseteq \arg\max_{a \in A} u(a, \theta)$ *for each* $\theta \in \operatorname{Supp}\mu$.

(iii) $V^1(\mu') = 0$ *for every* $\mu' \in \Delta\Theta$ *with* $\operatorname{Supp}\mu' \subseteq \operatorname{Supp}\mu$.

Proof. Fix $\mu \in \Delta\Theta$. We first prove that parts (i) and (ii) are equivalent. First, suppose that $V^1(\mu) = \overline{U} \cdot \mu - U(\mu) = 0$. Fix $\theta \in \operatorname{Supp}\mu$ and $a' \in A$. Assume, by contradiction, that $a' \in$

$\arg\max_{a\in A} \mathbb{E}_\mu[u(a,\theta)]$ but $a' \notin \arg\max_{a\in A} u(a,\theta)$ for some $\theta \in \operatorname{Supp}\mu$. Hence, $U(\delta_\theta) > U(\mu)$. Since $U(\delta_{\theta'}) \geq U(\mu)$, then $\overline{U}\cdot\mu > U(\mu)$, a contradiction.

Conversely, assume that $\arg\max_{a\in A} \mathbb{E}_\mu[u(a,\theta)] \subset \arg\max_{a\in A} u(a,\theta)$ for each $\theta \in \operatorname{Supp}\mu$. Hence, $U(\mu) = U(\delta_\theta)$ for all $\theta \in \operatorname{Supp}\mu$. Finally, $\overline{U}\cdot\mu = U(\mu)$ which is to say $V^1(\mu) = 0$.

Now, we prove that part (ii) implies part (iii). Assume that for each $\theta \in \operatorname{Supp}\mu$ the following is satisfied: $\arg\max_{a\in A} \mathbb{E}_\mu[u(a,\theta)] \subset \arg\max_{a\in A} u(a,\theta)$. Fix $\mu' \in \Delta\Theta$ with $\operatorname{Supp}\mu' \subset \operatorname{Supp}\mu$, and $a' \in \arg\max_{a\in A} \mathbb{E}_\mu[u(a,\theta)]$. Therefore, $u(a',\theta) = U(\delta_\theta)$ for each $\theta \in \operatorname{Supp}\mu$. As a result, $\overline{U}\cdot\mu' \geq U(\mu') \geq \mathbb{E}_{\mu'}[u(a',\theta)] = \mathbb{E}_{\mu'}[U(\delta_\theta)] = \overline{U}\cdot\mu'$. In conclusion, $V^1(\mu') = 0$.

Last, we prove that part (iii) implies part (i). Note, since $\operatorname{Supp}\mu' \subseteq \operatorname{Supp}\mu$, then $V^1(\mu) = 0$. ■

Lemma A.6.

(i) If $\nu_s = \nu_b$, then $\mathcal{D}^+ = \emptyset$.

(ii) If $\nu_s \neq \nu_b$, then $\Delta\Theta \subseteq \mathcal{D}^+$.

Proof. First assume that $\nu_s = \nu_b$. This implies that $g(\mu_b) = \mu_b$ for each μ_b. Thus $\mathcal{D}^+ = \emptyset$. Now assume that $\nu_s = g(\nu_b) \neq \nu_b$ and fix $\mu_b \in \operatorname{int}\Delta\theta$. Notice, by Lemma A.4, $g(\mu_b) \neq \mu_b$ for each $\mu_b \in \operatorname{int}\Delta\Theta$. Notice, since the buyer's decision problem is not trivial, there are some $\theta, \theta' \in \Theta$ such that $\arg\max_{a\in A} u(a,\theta) \cap \arg\max_{a\in A} u(a,\theta') = \emptyset$. Since $\operatorname{Supp}(\mu_b) = \Theta$, it follows that $V^1(\mu_b) > 0$. (See Lemma A.5.) Therefore, $\operatorname{int}\Theta \subseteq \mathcal{D}^+$. ■

Lemma A.7. *Fix $\mu \in \Delta\Theta$. Then, $\mu \in \mathcal{D}^+$ if and only if $V^2(\mu) > V^1(\mu)$.*

Proof. Let $\mu \in \Delta\Theta$. First, assume that $\mu \in \mathcal{D}^+$, which means that $\mu \neq g(\mu)$ and $V^1(\mu) > 0$. We first show that there is some $\tau \in \mathrm{PS}[\mu]$ such that $\mathbb{E}_\tau[\Lambda^2(\mu' \mid \mu)] > \mathbb{E}_\tau[\overline{U}\cdot\mu']$. To show this, we will show the following:

(i) There is some $\mu' \in \operatorname{int}(\Delta\operatorname{Supp}\mu)$, such that $\Lambda^2(\mu' \mid \mu) > \overline{U}\cdot\mu'$.

(ii) For each $\theta \in \Theta$ it follows that $\Lambda^2(\delta_\theta \mid \mu) = \overline{U}\cdot\delta_\theta$.

So, if $\tau \in \mathrm{PS}[\mu]$ is such that $\operatorname{Supp}\tau = \{\mu'\} \cup \{\delta_\theta : \theta \in \Theta\}$, then

$$\begin{aligned} V^2(\mu) &= \sup_{\tau'\in\mathrm{PS}[\mu]} \mathbb{E}_{\tau'}[\Lambda^2(\mu' \mid \mu)] - U(\mu) \\ &\geq \mathbb{E}_\tau[\Lambda^2(\mu' \mid \mu)] - U(\mu) \\ &> \overline{U}\cdot\mu - U(\mu) \\ &= V^1(\mu), \end{aligned}$$

as desired.

To show condition (i), notice that $\Delta\operatorname{Supp}\mu \not\subset H_{\rho=1}(\mu)$. (See Lemma A.4.) Moreover, for every $\theta \in \Theta$ for which $g(\mu)(\theta) > \mu(\theta)$, it holds that $\delta_\theta \in H_{\rho>1}$. Hence, by linearity of ρ , it follows that $H_{\rho>1} \cap \Delta\operatorname{Supp}\mu$ is non-empty and open relative to $\Delta\operatorname{Supp}\mu$. Fix $\mu' \in H_{\rho>1} \cap \Delta\operatorname{Supp}\mu$. Notice, that $V^1(\mu') > 0$. (See Lemma A.5.) Thus,

$$\Lambda^2(\mu' \mid \mu) = U(\mu') + V^1(\mu')\rho(\mu' \mid \mu) > U(\mu') + V^1(\mu') = \overline{U}\cdot\mu'.$$

To show condition (ii), fix $\theta \in \Theta$ and notice that

$$\Lambda^2(\delta_\theta \mid \mu) = U(\delta_\theta) + V^1(\delta_\theta)\rho(\delta_\theta \mid \mu) = U(\delta_\theta) = \overline{U} \cdot \delta_\theta.$$

Now we show the converse. Assume, by contrapositive, that $\mu \notin \mathcal{D}^+$. So, either $V^1(\mu) = 0$ or $g(\mu) = \mu$. First consider the case $g(\mu) = \mu$. Lemma A.4 shows that for for each $\tau \in \mathrm{PS}[\mu]$ and $\mu' \in \operatorname{Supp} \tau$, it follows that $\rho(\mu' \mid \mu) = 1$. Hence,

$$\begin{aligned}
V^2(\mu) &= \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + V^1(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&= \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + V^1(\mu)] - U(\mu) \\
&= \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_\tau[\overline{U} \cdot \mu'] - U(\mu_b) \\
&= V^1(\mu).
\end{aligned}$$

Consider the case $V^1(\mu) = 0$. Then, Lemma A.5 shows that, for each $\tau \in \mathrm{PS}[\mu]$ and each $\mu' \in \operatorname{Supp} \tau$ it follows that $V^1(\mu') = 0$. Consequently,

$$\begin{aligned}
V^2(\mu) &= \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + V^1(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&= \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu')] - U(\mu) \\
&= V^1(\mu).
\end{aligned}$$

Therefore, $\mu \notin \mathcal{D}^+$ implies $V^2(\mu) = V^1(\mu)$ as desired. ∎

Proof of Theorem 5.1. We proceed by induction on $t \in \mathbb{N}$. The base case $t = 1$ follows directly from Lemma A.7. Fix $t > 1$ and assume that for each $t \geq t' \geq 1$.

(i) $V^{t'+1}(\mu') = V^{t'}(\mu')$ for each $\mu' \notin \mathcal{D}^+$, and

(ii) $V^{t'+1}(\mu') > V^{t'}(\mu')$ for each $\mu' \in \mathcal{D}^+$.

We will show that these two statements hold for $t + 1$. Notice that $V^{t+2}(\mu') \geq V^{t+1}(\mu')$ for each $\mu' \in \Delta\Theta$ (See Lemma 4.3). Thus, it suffices to show that that $V^{t+2}(\mu') > V^{t+1}(\mu')$ if and only if $\mu' \in \mathcal{D}^+$.

First consider the case $\mu \notin \mathcal{D}^+$. Note that $\operatorname{Supp}(\tau) \cap \mathcal{D}^+ = \emptyset$ for each $\tau \in \mathrm{PS}[\mu]$. (See Lemmata A.4 and A.5). Therefore, for each $\tau \in \mathrm{PS}[\mu]$ and each $\mu' \in \operatorname{Supp}(\tau)$, $V^{t+1}(\mu') = V^t(\mu')$. This implies that

$$\begin{aligned}
V^{t+2}(\mu) &= \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + V^{t+1}(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&= \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + V^t(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&= V^{t+1}(\mu),
\end{aligned}$$

as desired.

Now consider the case $\mu \in \mathcal{D}^+$. We show that $V^{t+2}(\mu) > V^{t+1}(\mu)$ by contradiction. Suppose $\phi(V^{t+1})(\mu) = V^{t+2}(\mu) = V^{t+1}(\mu) = \phi(V^t)(\mu)$. Then, by Lemma A.3, there is some $\tau \in \text{PS}[\mu]$ such that

$$V^{t+1}(\mu) = \mathbb{E}_\tau[U(\mu') + V^t(\mu')\rho(\mu' \mid \mu)] - U(\mu) = V^{t+2}(\mu),$$

and $\text{Supp}\,(\tau) \subseteq \{\mu' : V^{t+1}(\mu') = V^t(\mu')\}$. Moreover, by conditions (i) and (ii), it follows that

$$\text{Supp}\,(\tau) \subseteq \Delta\Theta \backslash \mathcal{D}^+ = \{\mu' : V^t(\mu') = V^{t-1}(\mu')\}.$$

Therefore,

$$\begin{aligned}
V^{t+1}(\mu) &= \mathbb{E}_\tau[U(\mu') + V^t(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&= \mathbb{E}_\tau[U(\mu') + V^{t-1}(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&\leq \sup_{\tau \in \text{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + V^{t-1}(\mu')\rho(\mu' \mid \mu)] - U(\mu) \\
&= V^t(\mu).
\end{aligned}$$

Which contradicts condition (ii). Therefore, we conclude that $V^{t+2}(\mu) > V^{t+1}(\mu)$, as desired. ■

Lemma A.8. *Let $\hat{\theta} \in \arg\max_{\theta\in\Theta}\{r(\theta)\}$. Then, $\rho(\delta_{\hat{\theta}} \mid \mu_b) \geq 1$ for each $\mu_b \in \Delta(\Theta)$.*

Proof. Note that, by definition of ρ, for each $\mu_b \in \Delta\Theta$, $\rho(\delta_{\hat{\theta}} \mid \mu_b)\rho(\mu_b \mid \delta_{\hat{\theta}}) = 1$. Hence, it suffices to show that $\rho(\mu_b \mid \delta_{\hat{\theta}}) \leq 1$ for each $\mu_b \in \Delta\Theta$. Notice, for each $\theta \in \Theta$, $\rho(\delta_\theta \mid \delta_{\hat{\theta}}) = \frac{r(\theta)}{r(\hat{\theta})} \leq 1$. Thus, since $\rho(\mu_b \mid \delta_{\hat{\theta}})$ is linear in μ_b, if follows that $\rho(\mu_b \mid \delta_{\hat{\theta}}) \leq 1$ for each $\mu_b \in \Delta\Theta$. ■

Proof of Lemma 5.1. We first show part (i). Fix $\mu_b \in \Delta(\Theta)$. First we prove that for the mapping

$$\Lambda_B(\mu'_b \mid \mu_b) := U(\mu'_b) + B(\mu'_b) \cdot \rho(\mu'_b \mid \mu_b)$$

is weakly concave in μ'_b. To show this, note that

$$\begin{aligned}
\Lambda_B(\mu'_b \mid \mu_b) &= U(\mu'_b) + B(\mu'_b) \cdot \rho(\mu'_b \mid \mu_b) \\
&= U(\mu'_b) + V^1(\mu'_b) \cdot \rho(\delta_{\hat{\theta}} \mid \mu'_b) \cdot \rho(\mu'_b \mid \mu_b) \\
&= U(\mu'_b) + V^1(\mu') \cdot \rho(\delta_{\hat{\theta}} \mid \mu_b) \\
&= U(\mu'_b) + \left(\lambda \sum_\theta \mu'_b(\theta) \cdot U(\delta_\theta) - U(\mu')\right) \cdot \rho(\delta_{\hat{\theta}} \mid \mu_b) \\
&= -U(\mu'_b)\left(\rho(\delta_{\hat{\theta}} \mid \mu_b) - 1\right) + \rho(\delta_{\hat{\theta}} \mid \mu_b) \sum_{\theta\in\Theta} \mu'_b(\theta) \cdot U(\delta_\theta).
\end{aligned}$$

In addition, note that $\rho(\delta_{\hat{\theta}} \mid \mu_b) \geq 1$. (See Lemma A.8.) Hence, for each $\mu_b \in \Delta(\Theta)$, the mapping $\Lambda_B(\cdot \mid \mu_B)$ is the sum of a weakly concave function and a linear function. Hence, it is weakly concave. Thus, for each μ_b,

$$\sup_{\tau \in \mathrm{PS}[\mu_b]} \mathbb{E}_\tau[\Lambda_B(\mu_b' \mid \mu_b)] = \Lambda_B(\mu_b \mid \mu_b).$$

We now show that B is a fixed point of ϕ. Fix $\mu_b \in \Delta\Theta$ and note that

$$\begin{aligned}\phi(B)(\mu_b) &= \sup_{\tau \in \mathrm{PS}[\mu_b]} \mathbb{E}_\tau[\Lambda_B(\mu_b' \mid \mu_b)] - U(\mu_b) \\ &= \Lambda_B(\mu_b \mid \mu_b) - U(\mu_b) \\ &= B(\mu_b) \cdot \rho(\mu_b \mid \mu_b) \\ &= B(\mu_b),\end{aligned}$$

where the third equality equation follows from definition of Λ_B. This shows that B is a fixed point of ϕ.

We now show part (ii). Fix $\mu_b \in \Delta(\Theta)$ and notice that $\rho(\delta_{\hat{\theta}} \mid \mu_b) \geq 1$. (See Lemma A.8). Hence, $B(\mu_b) = V^1(\mu_b) \cdot \rho(\delta_{\hat{\theta}} \mid \mu_b) \geq V^1(\mu_b)$, as desired. ■

Proof of Theorem 5.2. First we show that $V^t(\mu) \leq B(\mu)$ for each $t \in \mathbb{N}$ and each belief $\mu \in \Delta(\Theta)$. We proceed by induction. Notice that $V^1(\cdot) \leq B(\cdot)$ (See Lemma 5.1). Now, assume that $V^t(\cdot) \leq B(\cdot)$ for $k \geq 1$. Therefore, by monotonicity of ϕ, for each $\mu \in \Delta\Theta$, it follows that

$$V^{k+1}(\mu) = \phi(V^t)(\mu) \leq \phi(B)(\mu) = B(\mu). \tag{6}$$

(See Lemma A.3.)

This implies that $(V^t(\mu))_{t \in \mathbb{N}}$ is a bounded increasing sequence and hence it has a limit. Write $V^\infty(\mu) := \lim_{t\to\infty} V^t(\mu)$.

We first will show that V^∞ satisfies the following: for each $\mu_b \in \Delta\Theta$,

$$V^\infty(\mu) = \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + V^\infty(\mu') \cdot \rho(\mu' \mid \mu)] - U(\mu). \tag{7}$$

To show this, first fix $\tau \in \mathrm{PS}[\mu]$ Notice, since $V^1 \leq V^\infty$, it follows that

$$\mathbb{E}_\tau[U(\mu') + V^t(\mu') \cdot \rho(\mu' \mid \mu)] \leq \mathbb{E}_\tau[U(\mu') + V^\infty(\mu') \cdot \rho(\mu' \mid \mu)]. \tag{8}$$

In addition, notice that for each $t \in \mathbb{N}$, the mapping $U(\mu') + V^t(\mu') \cdot \rho(\mu' \mid \mu)$ is bounded by the continuous mapping $U(\mu') + B(\mu') \cdot \rho(\mu' \mid \mu)$. (See Equation 6.) Then, by the Dominated Convergence Theorem,

$$\mathbb{E}_\tau[U(\mu') + \lim_{t\to\infty} V^t(\mu') \cdot \rho(\mu' \mid \mu)] = \lim_{t\to\infty} \mathbb{E}_\tau[U(\mu') + V^t(\mu') \cdot \rho(\mu' \mid \mu)] \tag{9}$$

Finally, notice that, by definition of V^{t+1}.

$$\mathbb{E}_\tau[U(\mu') + V^t(\mu')\cdot\rho(\mu' \mid \mu)] - U(\mu) \leq V^{t+1}(\mu) \tag{10}$$

Therefore,

$$\begin{aligned}
V^\infty(\mu) &= \lim_{t\to\infty} V^{t+1}(\mu) \\
&= \lim_{t\to\infty} \sup_{\tau\in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + V^t(\mu')\cdot\rho(\mu' \mid \mu)] - U(\mu) \\
&\leq \sup_{\tau\in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + V^\infty(\mu')\cdot\rho(\mu' \mid \mu)] - U(\mu) \\
&= \sup_{\tau\in \mathrm{PS}[\mu]} \mathbb{E}_\tau[U(\mu') + \lim_{t\to\infty} V^t(\mu')\cdot\rho(\mu' \mid \mu)] - U(\mu) \\
&= \sup_{\tau\in \mathrm{PS}[\mu]} \lim_{t\to\infty} \mathbb{E}_\tau[U(\mu') + V^t(\mu')\cdot\rho(\mu' \mid \mu)] - U(\mu) \\
&\leq \sup_{\tau\in \mathrm{PS}[\mu]} \lim_{t\to\infty} V^{t+1}(\mu) \\
&= V^\infty(\mu),
\end{aligned}$$

where the first inequality follows from Equation(8), the fourth equality follows from Equation (9), and the last inequality follows from Equation (10). This shows Equation (7).

Notice that $V^\infty(\delta_\theta) = 0$ for each $\theta \in \Theta$. Thus, to show that $V^\infty \in \mathcal{F}$ is suffices to show that V^∞ is continuous.

Fix $\mu \in \Delta\Theta$ and let $(\mu_k)_{k\in\mathbb{N}}$ be a sequence such that $\mu_k \in \Delta\Theta$ and $\lim \mu_k = \mu$. We show $\lim V^\infty(\mu_k) = V^\infty(\mu)$. We divide the proof into two steps. Step one shows that $\limsup_{k\to\infty} V^\infty(\mu_k) \leq V(\mu)$ and step two shows that $\liminf_{k\to\infty} V^\infty(\mu_k) \leq V^\infty(\mu)$.

Step 1. Let $f : \Delta\Theta \times \Delta\Theta \to \mathbb{R}$ given by $f(\mu', \mu) = U(\mu') + V^\infty(\mu')\rho(\mu' \mid \mu)$. Notice, by Equation (7)

$$V^\infty(\mu) = \sup_{\tau\in \mathrm{PS}[\mu]} \mathbb{E}_\tau[f(\mu', \mu)] - U(\mu)$$

Notice that f is bounded since $U(\mu') + V^\infty(\mu')\rho(\mu' \mid \mu) \leq U(\mu') + B(\mu')\rho(\mu' \mid \mu)$. Thus, there exist an affine mapping $L : \Delta\Theta \to \mathbb{R}$ such that

(i) $L(\mu') \geq f(\mu', \mu)$ for each $\mu' \in \Delta(\Theta)$.

(ii) $L(\mu) = \sup_\tau \mathbb{E}_\tau[f(\mu', \mu)]$.

Let $M > 0$ be a bound of $V^\infty(\cdot)$. Since the set $\Delta\Theta \times \Delta\Theta$ is compact, the mapping ρ is uniformly continuous. Hence, there is some $\delta > 0$ such that $||\mu - \mu_k||_\infty < \delta$ implies that for each $\mu' \in \Delta\Theta$,

$$\rho(\mu' \mid \mu_k) < \rho(\mu' \mid \mu) + \tfrac{\varepsilon}{2M},$$

Then, for each $\mu' \in \Delta\Theta$,

$$V^\infty(\mu')\rho(\mu' \mid \mu_k) < V^\infty(\mu')\rho(\mu' \mid \mu) + \tfrac{\varepsilon}{2}.$$

Thus,

$$\begin{aligned} f(\mu', \mu_k) &= U(\mu') + V^{\infty}(\mu')\rho(\mu' \mid \mu_k) \\ &< U(\mu') + V^{\infty}(\mu')\rho(\mu' \mid \mu) + \tfrac{\epsilon}{2} \\ &= f(\mu', \mu) + \tfrac{\epsilon}{2} \\ &\leq L(\mu') + \tfrac{\epsilon}{2}. \end{aligned}$$

Consequently,

$$\begin{aligned} V^{\infty}(\mu_k) &= \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_{\tau}[f(\mu', \mu_k)] - U(\mu_k) \\ &\leq \sup_{\tau \in \mathrm{PS}[\mu]} \mathbb{E}_{\tau}[L(\mu') + \tfrac{\varepsilon}{2}] - U(\mu) + \tfrac{\varepsilon}{2} \\ &= L(\mu) - U(\mu) + \varepsilon \\ &= \sup_{\tau} \mathbb{E}_{\tau}[f(\mu', \mu)] - U(\mu) + \varepsilon \\ &= V^{\infty}(\mu) + \varepsilon. \end{aligned}$$

This implies that $\limsup_{k \to \infty} V^{\infty}(\mu_k) \leq V^{\infty}(\mu) + \varepsilon$. Moreover, since $\varepsilon > 0$ is arbitrary it follows that $\limsup_{k \to \infty} V^{\infty}(\mu_k) \leq V^{\infty}(\mu)$.

Step 2. Fix $\varepsilon > 0$. Notice, there is some $K \in \mathbb{N}$ such that $V^K(\mu) > V^{\infty}(\mu) + \frac{\varepsilon}{2}$. Moreover, there is some $\delta > 0$ such that $||\mu_k - \mu||_{\infty} < \delta$ implies $V^K(\mu_k) > V^K(\mu) + \frac{\varepsilon}{2}$. Therefore, if $||\mu_k - \mu||_{\infty} < \delta$, then

$$V^{\infty}(\mu_k) \geq V^K(\mu_k) \geq V^K(\mu) - \tfrac{\varepsilon}{2} \geq V^{\infty}(\mu) - \epsilon.$$

Therefore, $\liminf_{k \to \infty} V^{\infty}(\mu_k) \geq V^{\infty}(\mu) - \varepsilon$. Moreover, since the $\varepsilon >$ is arbitrary, it follows that $\liminf_{k \to \infty} V^{\infty}(\mu_k) \geq V^{\infty}(\mu)$. ■