

DISCUSSION PAPER SERIES

DP12482

ON THE VALUE OF PERSUASION BY EXPERTS

Ricardo Alonso and Odilon Câmara

INDUSTRIAL ORGANIZATION



ON THE VALUE OF PERSUASION BY EXPERTS

Ricardo Alonso and Odilon Câmara

Discussion Paper DP12482
Published 03 December 2017
Submitted 04 December 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Ricardo Alonso and Odilon Câmara

ON THE VALUE OF PERSUASION BY EXPERTS

Abstract

We consider a persuasion model in which a sender influences the actions of a receiver by selecting an experiment (public signal) from a set of feasible experiments. We ask: does the sender benefit from becoming an expert — observing a private signal prior to her selection? We provide necessary and sufficient conditions for a sender to never gain by becoming informed. Our key condition (sequential redundancy) shows that the informativeness of public experiments can substitute for the sender's expertise. We then provide conditions for private information to strictly benefit or strictly hurt the sender. Expertise is beneficial when the sender values the ability to change her experimental choice according to her private information. When the sender does not gain from expertise, she is strictly hurt when different types cannot pool on an optimal experiment.

JEL Classification: D83

Keywords: information design, Bayesian persuasion, experts.

Ricardo Alonso - r.alonso@lse.ac.uk
London School of Economics and CEPR

Odilon Câmara - ocamara@marshall.usc.edu
University of Southern California

Acknowledgements

We thank Emir Kamenica, Anton Kolotilin, Alex Smolin, and participants of the 2017 SAET Conference for their comments and suggestions

On the Value of Persuasion by Experts ^{*}

RICARDO ALONSO[†]

ODILON CÂMARA[‡]

London School of Economics and CEPR

University of Southern California

September 15, 2017

Abstract

We consider a persuasion model in which a sender influences the actions of a receiver by selecting an experiment (public signal) from a set of feasible experiments. We ask: does the sender benefit from becoming an expert — observing a private signal prior to her selection? We provide necessary and sufficient conditions for a sender to never gain by becoming informed. Our key condition (sequential redundancy) shows that the informativeness of public experiments can substitute for the sender’s expertise. We then provide conditions for private information to strictly benefit or strictly hurt the sender. Expertise is beneficial when the sender values the ability to change her experimental choice according to her private information. When the sender does not gain from expertise, she is strictly hurt when different types cannot pool on an optimal experiment.

JEL classification: D83.

Keywords: Information design, Bayesian persuasion, experts.

^{*}We thank Emir Kamenica, Anton Kolotilin, Alex Smolin, and participants of the 2017 SAET Conference for their comments and suggestions.

[†]LSE, Houghton Street, London WC2A 2AE, United Kingdom. R.Alonso@lse.ac.uk

[‡]USC FBE Dept, 701 Exposition Blvd, Ste. 231, HOH-231, MC-1422, Los Angeles, CA 90089-1422. ocamara@marshall.usc.edu.

1 Introduction

A prosecutor would like to persuade a judge to convict a defendant. She can submit to the court the testimony from one of several expert witnesses, who vary in their expertise in evaluating the available evidence. Before choosing which testimony to submit, she has private access to some information relevant to the case (e.g., she may be able to privately ask the witnesses exploratory questions, or privately observe unofficial reports from law enforcement). Can she increase the chances of a conviction by accessing such information? That is, is an “informed” prosecutor a more successful persuader?

We investigate this question in the broader setting of a sender (she) who can affect the decisions of a receiver (he) by controlling his information environment — as in Kamenica and Gentzkow (2011) (KG henceforth).¹ The receiver chooses the action a that maximizes his utility $u_R(a, \theta)$, where θ is an unknown state of the world. The sender wants to maximize her utility $u_S(a, \theta)$, and can influence the receiver’s action by providing a public signal (an experiment) whose outcome is correlated with θ . We expand on the KG model in two ways. First, the sender in our model might be constrained in her choice of an experiment; she must choose one experiment π from a given set Π , but she can garble its outcome.² For instance, the prosecutor can frame the questions to the expert witness in a way that coarsens the informativeness of his testimony, or she may refrain from asking certain questions altogether. Second, the sender in our model privately observes the realization of an exogenous signal π_e before committing to an experiment π . The sender in KG commits to an experiment π prior to observing any private information. We contrast this uninformed-sender case with the case in which the sender is an expert.

Does the sender benefit from becoming an expert? That is, does she prefer to observe π_e before choosing π , or does she prefer to commit to a public signal without observing π_e ? The answers to these questions depend on the informational content of π_e relative to experiments in Π . We say that experiment π_e is *redundant given* Π if, for every experiment

¹See, also, Brocas and Carrillo (2007), Rayo and Segal (2010), Boleslavsky, Cotton, and Gurnani (2017), Alonso and Câmara (2016a,b,c) and Bergemann and Morris (2016).

²Note that if the set Π has a signal π_{FI} that is fully informative of the state θ , then the sender in our model effectively has access to the same signals as in KG.

$\pi \in \Pi$, there is an experiment $\pi' \in \Pi$ such that π' is at least as Blackwell-informative as jointly observing the outcome of π_e and π . In effect, when the private signal is redundant, disclosing its outcome alongside the outcome of any other experiment in Π does not generate a more informative signal than what is already available through experiments in Π .

It is easy to see how a sender may benefit from privately observing a non-redundant signal: if she could credibly disclose its outcome, she could then convey more information to the receiver than if uninformed.

While a sender may benefit from observing non-redundant information, redundancy by itself does not guarantee that the sender cannot gain from becoming an expert. In some cases, even if π_e is redundant given Π , she can use its outcome to revert to an experiment that is more likely to induce the desired behavior in the receiver, as in the following example.

Example 1: π_e redundant given Π , but positive value of expertise. A prosecutor wants to persuade a judge that a defendant is corrupt. The prosecutor has evidence that the defendant has a secret offshore account, but she is not sure if it is in country A or country B. The money in this account is either legal (the defendant is innocent (I)) or the result of corruption (the defendant is guilty (G)). Consequently, there are four possible states, $\{AI, AG, BI, BG\}$, representing where the defendant has an account and whether or not the source of the money is legal. The prosecutor and the judge have uniform prior beliefs.

The judge will convict the defendant if and only if he learns that the defendant is guilty for sure. The prosecutor receives payoff 1 if the defendant is convicted, and zero otherwise. The prosecutor has to choose one of two possible experiments. Experiment π_A is a full investigation of banks in country A, and reveals all information about country A but no information about B. That is, it reveals partitions $\{AI\}$, $\{AG\}$, and $\{BI, BG\}$. Experiment π_B is a full investigation of country B and reveals partitions $\{AI, AG\}$, $\{BI\}$, and $\{BG\}$.

Without private information, the prosecutor's optimal strategy is to select either investigation and fully disclose its results. The conviction probability is only 25% in this case, since the prosecutor might investigate the wrong country — i.e., the one in which the defendant does not have an account. Now suppose that the prosecutor has access to a private source of information π_e that simply reveals the country of the account, resulting in partitions

$\{AI, AG\}$ and $\{BI, BG\}$. Note that π_e is redundant given Π .³ Nevertheless, the prosecutor strictly benefits from this redundant private signal. The expert prosecutor launches the investigation of the country in which she knows the defendant has an account, increasing the conviction probability to 50%. \square

To understand how expertise is beneficial in Example 1, note that experiments in Π differ in their informativeness, and the expert's redundant signal can pinpoint which one would ex-post reveal more information about the state. Our first result shows that if this is not the case, then a sender cannot benefit from observing redundant information. More precisely, Proposition 1 establishes that an uninformed sender is ex-ante (weakly) better off than an expert for every $u_S(a, \theta)$ and $u_R(a, \theta)$ if and only if π_e is *sequentially redundant* given Π . Sequential redundancy implies that, for *every* possible selection rule in which the sender first observes the result of π_e and then selects an experiment, there exists an available experiment π' that is at least as Blackwell-informative as the sender's sequential experimentation. In other words, sequential redundancy ensures that an uninformed sender can always replicate both the expert's private signal and her ensuing choice of experiment by garbling an experiment in Π .

An important corollary of Proposition 1 is that if Π contains a fully informative experiment, then a sender can *never* benefit from privately observing π_e , independently of the utility functions $u_R(a, \theta)$ and $u_S(a, \theta)$ and the informational content of π_e . This result shows that the informativeness of available experiments can substitute for the sender's expertise. In fact, a necessary condition for a sender to benefit from becoming an expert is that no experiment in Π fully reveals the state.

While sequentially redundant private information can never benefit the sender, redundant private information can strictly benefit or strictly hurt the sender. Sections 4 and 5 provide sufficient conditions for these two cases. In Section 6, we apply our results to an important economic phenomenon: the strategic use of *real information* in marketing (see Johnson and Myatt 2006). We show that a retailer who can only offer a limited set of experiments to consumers strictly benefits from a salesperson that is an expert and can, consequently,

³Actually, π_e is redundant in a stronger sense: learning the outcome of π_e after observing the outcome of either π_A or π_B would not change the posterior beliefs of the sender.

select the best experiment for each consumer. Interestingly, the same expert salesperson will strictly hurt a retailer that has access to a fully informative signal. This happens because the expert salesperson destroys the company’s ability to strategically garble the experiments and hide information from the consumer. Therefore, this retailer benefits from strategic ignorance: hiring non-experts and limiting the training she provides to her salesforce.

Our paper is related to the recent literature that studies the strategic design of a public signal by an informed sender. Gill and SgROI (2008, 2012) consider a privately-informed principal who can subject herself to a test that is informative of her type, and can optimally choose the test’s difficulty. Li and Li (2013) study a privately-informed candidate who can choose the accuracy of a costly public signal (campaign) about the qualifications of the politicians competing for office. Degan and Li (2016) study a persuasion model in which the sender privately knows the realized binary state and can provide a public signal to the receiver — the cost of the signal to the sender increases in its precision. Rosar (2017) studies test design by a principal who seeks to learn the binary quality of an imperfectly informed agent, when test-taking by the agent is voluntary. Perez-Richet (2014) considers an informed sender who might be constrained in her choice of a signal. In his model, the receiver can only take two actions (validation or non-validation), and there are only two types of senders, both of whom receive the same net payoff from validation. In contrast to that paper, we consider any finite set of types for the sender and finite action space for the receiver, and we allow for general sender and receiver utility functions. Hedlund (2017) considers an informed sender who has access to any signal that is correlated with the state (so the space of signals is the same as in KG). He compares the equilibrium payoffs of a game in which the sender’s type is private information and a game in which the sender’s type is public information. Both Perez-Richet (2014) and Hedlund (2017) focus on characterizing the properties of equilibria, and explore how different refinements narrow the equilibrium predictions. Our focus, however, is on understanding the value of expertise in persuasion games. That is, we compare the payoff of a privately informed sender and the payoff of an uninformed sender. In particular, by looking at senders’ payoffs attainable in a Perfect Bayesian equilibrium, we provide an upper bound on the value of expertise.

2 The Model

Our model features a game between a sender (she) and a receiver (he). The receiver chooses an action that affects the utility of both players. The sender can influence this choice by supplying the receiver with a public signal (experiment) that is correlated with the state. We contrast two cases. In the first case, the sender has no private information about the state, or, equivalently, she can commit to the signal before becoming privately informed (as in KG). In the second case, the sender has private information about the state and cannot commit to the public signal before becoming informed.

Preferences and Prior Beliefs: All players are expected utility maximizers and process information according to Bayes' rule. The receiver selects an action a from a finite set A , which has at least two actions. We will relax this assumption and allow A to be a compact set for some of our results. The sender and the receiver have preferences over actions characterized by continuous von Neumann-Morgenstern utility functions $u_S(a, \theta)$ and $u_R(a, \theta)$, with $\theta \in \Theta$ and Θ a finite state space. Players share a common prior belief p belonging to the interior of the simplex $\Delta(\Theta)$.

Private Information: An experiment π is a Z_π -valued random variable that depends on the state, with a finite realization space Z_π . The sender privately observes the realization of experiment π_e . Let the sender's type $t \in \Delta(\Theta)$ represent her interim belief after observing $z_{\pi_e}(t) \in Z_{\pi_e}$ —i.e., $\Pr[\theta | z_{\pi_e}(t)] = t_\theta$, and $\beta(t)$ the probability of t . Let T be the (finite) set of possible interim beliefs induced by π_e , where, for simplicity, we assume that different realizations result in different types. Bayes' rule requires that $E_\beta[t] = \sum_{t \in T} \beta(t)t = p$. The set T and the probabilities $\beta(t)$ are common knowledge. Throughout the paper, we contrast two cases: T has a single element (the sender is uninformed) and T has at least two elements (the sender is privately informed).

Feasible Experiments: After observing her private signal, but before the receiver chooses his action, the sender supplies an experiment π formed from a finite set of feasible experiments Π . All experiments π in Π carry the same cost to the sender, which we assume to be zero. When comparing the informational content of different experiments, we will consider the Blackwell information order \preceq_B (Blackwell 1953).

We impose no special structure on the joint distribution of signal realizations of π_e and experiments in Π . For instance, we could have $\pi_e \in \Pi$, so that the sender could certify her type by supplying experiment π_e . We also denote by π_{FI} a fully informative experiment — i.e., an experiment such that the posterior belief of a Bayesian decision maker puts non-zero probability in at most one state.

Finally, we assume that the sender can costlessly garble any experiment π and select arbitrary mixtures of experiments. A garbling of experiment π is an experiment whose realizations are independent of θ and π_e conditional on the realization of experiment π . Note that by allowing for garblings implicitly allows the sender to engage in other forms of communication. For instance, changing the labels associated to realizations in Z_π would mimic cheap talk communication by the sender. A mixture λ is an experiment with realization space $\Pi \times \tilde{Z}$, where $\tilde{Z} = \times_{\pi \in \Pi} \tilde{Z}_\pi$, and \tilde{Z}_π the space of realizations of a garbling of π , in which $\lambda(\pi)$ is the probability of selecting experiment π and observing the realization of its garbling $\tilde{z}_\pi \in \tilde{Z}_\pi$. As the sender mixes among experiments without observing their actual realizations, we assume that the choice of experiment is independent of these realizations given their type, i.e. $\Pr_\lambda [(\pi, \tilde{z})|t] = \Pr(\tilde{z}|t) \lambda(\pi|t)$, $\pi \in \Pi, \tilde{z} \in \tilde{Z}, t \in T$. We denote by $\Gamma(\Pi)$ the set of all possible mixtures of garblings of experiments in Π , so that the sender supplies the receiver an experiment $\pi \in \Gamma(\Pi)$. Following Blackwell (1953), this implies that the sender also has access to any experiment that is less informative than any $\pi \in \Pi$. In particular, if the sender has a fully informative signal available, then she can choose *any experiment* that is correlated with the state (as in KG).

Timing: The sender privately learns her type t and then chooses an experiment $\pi \in \Gamma(\Pi)$. The receiver simultaneously observes π and its realization $z_\pi \in Z_\pi$, updates his beliefs, and then chooses an action $a \in A$. Payoffs are then realized.

Sender's Equilibrium Payoff: We consider Perfect Bayesian equilibria (PBE). We will henceforth use the term *equilibrium* to refer to a PBE. After observing experiment π and realization z_π , the receiver updates his information consistently, taking into account equilibrium strategies and the informational content of $\{\pi, z_\pi\}$. Following Perez-Richet (2014), off the equilibrium path, if the hard information $\{\pi, z_\pi\}$ is inconsistent with equilibrium strategies,

then the hard information has preeminence on the receiver's belief updating.

If experiment π_e is uninformative, then the sender is uninformed, and we can use the results from KG to compute the sender's maximum expected equilibrium payoff V_U . We refer to the informed sender's game to denote the case in which experiment π_e is informative. Let $w = (w(t))_{t \in T}$, where $w(t)$ is the expected payoff of type t , and $W \subset \mathbb{R}^{card(T)}$ is the set of type-dependent equilibrium payoffs of the sender. Let $V_I = \sup_{w \in W} \sum_{t \in T} \beta(t)w(t)$. That is, V_I is the sender's maximum ex-ante expected utility in the informed sender game. We refer to V_U as the value of persuasion by an uninformed sender and V_I as the value of persuasion by an expert.

2.1 Definitions

We now introduce some properties of the information environment that allows us to compare the information that can be conveyed by an informed and an uninformed sender. For an arbitrary set of experiments Π , we say that experiment π_e is *redundant* given Π if for every $\pi \in \Pi$, there exists $\pi' \in \Pi$ such that $\{\pi_e, \pi\} \preceq_B \pi'$, where $\{\pi_e, \pi\}$ refers to the experiment in which the decision maker observes the realizations of both π_e and π . In other words, π_e is redundant given Π if observing the outcome of π_e in addition to the outcome of some experiment π cannot generate more information than what is already available through experiments in Π . Note that redundancy is different from the notion of mutual information (see Cover and Thomas 1991). For instance, if $\Pi = \{\pi_{FI}\}$ then we trivially have $\{\pi_e, \pi_{FI}\} \preceq_B \pi_{FI}$ for any choice of π_e and yet knowing the state does not always allow the decision maker to predict the outcome of π_e .

Likewise, we say that experiment π_e is *strongly redundant* given Π if for every $\pi \in \Pi$, $\{\pi_e, \pi\} \preceq_B \pi$. In this case, a decision maker who observes π would never change his beliefs if he then observes the realization of π_e . In statistical terms, π_e is strongly redundant given Π if, for each $\pi \in \Pi$, π is a sufficient statistic for $\{\pi_e, \pi\}$ (DeGroot 1970).

The receiver may not be able to infer the sender's type from the realization of π , even when π_e is strongly redundant given π . We say that π_e *can be replicated* with π if there

exists a T -valued garbling of π , denoted by g_π , such that

$$\Pr [\pi_e = g_\pi \circ \pi] = 1. \tag{1}$$

Therefore, if π is available, the sender can offer a garbling of π that certifies her type with probability 1.

Finally, we say that experiment π_e is *sequentially redundant* given Π if for every z_{π_e} -contingent selection of experiments $\pi(z_{\pi_e}) \in \Pi$, where $\pi(z_{\pi_e})$ is selected whenever z_{π_e} occurs, there exists $\pi' \in \Pi$ such that $\{\pi_e, \pi(z_{\pi_e})\} \preceq_B \pi'$. Trivially, every π_e that is sequentially redundant given Π must also be redundant given Π .

3 Non-positive Value of Expertise

How can a sender benefit from gathering some information prior to choosing an experiment? In the absence of cost differences among experiments, an informed sender may be able to revert to an experiment that she believes more likely to induce the desired behavior in the receiver. However, such interim information will not confer an advantage to the sender if her private signal is sequentially redundant.

Proposition 1 *We have that $V_U \geq V_I$ for all $u_S(a, \theta)$ and $u_R(a, \theta)$ if and only if π_e is sequentially redundant given $\Gamma(\Pi)$.*

The proposition clarifies that the informativeness of available experiments substitutes for the sender's expertise whenever these experiments make such expertise sequentially redundant. Conversely, if π_e is not sequentially redundant given $\Gamma(\Pi)$, then an informed sender can convey more information than an uninformed sender by a judicious choice of experiment following each realization of π_e . It is then easy to think of situations in which this could be beneficial to the sender. For example, suppose that sender and receiver share the same preferences, and an uninformed sender offers experiment π . Then, if the sender is privately informed, she cannot be made worse off by credibly signaling the realization of π_e and offering π , so $V_U \leq V_I$ — and, in many situations, she can improve by adapting her choice of experiment, so that $V_U < V_I$.

The proof of the proposition hinges on the ability of an uninformed sender to replicate through experiments in $\Gamma(\Pi)$ *both* the informed sender’s private signal π_e and her outcome-contingent choice of experiment. For instance, if π_e is redundant, then for each experiment $\pi \in \Gamma(\Pi)$, an uninformed sender can find an experiment that replicates the information revealed by an expert who discloses the outcomes of π and her private signal π_e . Redundancy, however, does not guarantee that an uninformed sender can also replicate an expert’s *choice* of experiment. As Example 1 in the Introduction shows, an expert with redundant private information can generate more informative experiments by conditioning her choice of a signal on her type. To guarantee that the uninformed sender can replicate both π_e and the informed sender’s choice, her private signal must be sequentially redundant.

The main insight of Proposition 1 is that the informativeness of the public experiments available to a sender can substitute for a sender’s lack of expertise when persuading a receiver. In fact, an important implication of Proposition 1 is that a sender with access to a fully informative experiment can never benefit from becoming informed, regardless of the correlation of her private signal with experiments in Π .

Corollary 1 *Suppose that A is a compact set and $\pi_{FI} \in \Pi$. Then, $V_U \geq V_I$.*

As an illustration of the corollary, consider a prosecutor persuading a judge to convict a defendant. If the prosecutor could submit *any* number of expert witness testimonies to the judge, and could commit to a garbling of the submitted testimonies, she could never benefit from observing some (or all) of the actual findings.

To help with applications, we next provide a characterization of sequentially redundant private signals. We will work with a minimal representation of available experiments to the sender. We say that a set of experiments $\Pi_B = \{\pi_i\}_{i \in I_B}$ is *linearly independent* if for any mixture λ such that $\pi_{i'} \preceq_B \sum_{i \in I_B} \lambda_i \pi_i$ we must have $\lambda_i = 0$ for $i \neq i'$.

Our characterization is framed in terms of the posterior beliefs induced by realizations of different experiments. Let $Q(\pi_i)$ be the set of posterior beliefs induced by experiment π_i , with $Q^{ij} = Q(\pi_i) \cup Q(\pi_j)$, and $Q_{ext}^{ij} = \text{ext}(\text{conv}(Q^{ij}))$ the posterior beliefs that are extreme points of the convex hull of Q^{ij} . Note that, whenever π_e is strongly redundant given Π_B , any posterior induced after observing z_{π_e} and the outcome of π_i must belong to $Q(\pi_i)$.

Proposition 2 Let $\Pi_B = \{\pi_i\}_{i \in I_B}$ be a finite set of linearly independent experiments.

(a) If π_e is sequentially redundant given $\Gamma(\Pi_B)$, then (i) π_e is strongly redundant given Π_B , and (ii) for each $i, j \in I_B$ there exist non-negative numbers $\alpha_{ij}(z_{\pi_e}) \geq 0$, with

$$\sum_{z_{\pi_e} \in Z_{\pi_e}} \Pr(z_{\pi_e}) \alpha_{ij}(z_{\pi_e}) = 1,$$

so that for every $q \in Q_{ext}^{ij}$,

$$\Pr_i(q|z_{\pi_e}) - \Pr_j(q|z_{\pi_e}) = \alpha_{ij}(z_{\pi_e}) (\Pr_i(q) - \Pr_j(q)). \quad (2)$$

(b) Let $Q = \cup_{i \in I_B} Q(\pi_i)$ and $Q_{ext} = ext(conv(Q))$. Suppose that $Q(\pi_i) \subset Q_{ext}$ for every $i \in I_B$. Then, π_e is sequentially redundant given $\Gamma(\Pi_B)$ if and only if there exist non-negative numbers $\alpha(z_{\pi_e}) \geq 0$, with $\sum_{z_{\pi_e} \in Z_{\pi_e}} \Pr(z_{\pi_e}) \alpha(z_{\pi_e}) = 1$, so that for every $i, j \in I_B$ and $q \in Q$,

$$\Pr_i(q|z_{\pi_e}) - \Pr_j(q|z_{\pi_e}) = \alpha(z_{\pi_e}) (\Pr_i(q) - \Pr_j(q)). \quad (3)$$

A property of the extreme beliefs Q_{ext} is that whenever $q \in Q_{ext}$ is an outcome of a mixture of garblings of experiments in Π_B , q can only be induced from realizations of experiments π_i that induced posterior q . While different experiments may still induce different conditional distributions over beliefs in Q_{ext} after observing z_{π_e} , condition (2) shows that their difference must always be proportional to the unconditional difference. This property is necessary for an uninformed sender to replicate the distribution over posterior beliefs of any z_{π_e} -contingent choice of experiment through appropriate mixtures of experiments in Π_B .

The next corollary provides an easy-to-verify necessary condition for sequential redundancy.

Corollary 2 Let $\Pi_B = \{\pi_i\}_{i \in I_B}$ be a finite linearly independent set of experiments. Suppose that π_e is sequentially redundant given $\Gamma(\Pi_B)$ and for some $\pi_i, \pi_j \in \Pi_B$ and posteriors $q, q' \in Q_{ext}$, we have $q \in Q(\pi_i)$, $q' \in Q(\pi_j)$ but $q \notin Q(\pi_j)$ and $q' \notin Q(\pi_i)$. Then, for all z_{π_e} we have

$$\frac{\Pr_i[q|z_{\pi_e}]}{\Pr_i[q]} = \frac{\Pr_j[q'|z_{\pi_e}]}{\Pr_j[q']}. \quad (4)$$

The ratio $\Pr_i[q|z_{\pi_e}]/\Pr_i[q]$ represents the *pointwise mutual information* of the pair of outcomes z_{π_e} of π_e and q of π_i —how much the likelihood of q under experiment π_i is revised

after observing z_{π_e} . Suppose that experiments π_i and π_j induce different extreme beliefs q and q' . The corollary shows that the pointwise mutual information conveyed by z_{π_e} about q and q' must be identical: experiment π_e cannot lead to different relative occurrences of posteriors q and q' . Otherwise, knowledge of z_{π_e} can be used to generate q and q' at different relative frequencies.

As an application of this corollary, consider a linearly independent set Π_B of partitional experiments π_1 and π_2 , and a partitional π_e (as in Example 1 in the Introduction). Then, it is immediate that π_e is strongly redundant given Π if and only if the partition induced by π_e is coarser than the one induced by every π_i , $i = \{1, 2\}$. However, for π_e to be sequentially redundant, it must be that there exists at most one realization z_{π_e} such that the restriction of experiments π_i to z_{π_e} are distinct.

4 When does Redundant Expertise Benefit the Sender?

Proposition 1 reveals that the sender cannot benefit if her expertise is sequentially redundant. Alternatively, if her expertise is not redundant, then her choice of signal could be used to reveal her non-redundant information to the receiver. Consequently, if the receiver benefits from acquiring more information and players' preferences are sufficiently aligned, then the sender can strictly benefit from non-redundant expertise. In this section, we focus on the intermediate case: when can the sender strictly benefit from redundant, but not sequentially redundant, information?

We can exploit the concavification argument from KG to compute V_U . However, computing V_I can be a much harder task. To overcome this problem, we first consider a simpler game in which the signal π_e is publicly observed. We then provide conditions such that the sender's payoff in this simpler game is a lower bound for V_I .

Formally, consider an alternative game in which players publicly observe the realization of π_e before the sender chooses experiment π . After observing realization $z_{\pi_e} \in Z_{\pi_e}$, players update their beliefs to $q(z_{\pi_e})$. The sender then chooses the signal $\pi^*(z_{\pi_e})$ that maximizes her expected payoff — that is, she optimally selects a signal $\pi^* \in \Gamma(\Pi)$. Let $\Pi_{Pub}^* \equiv \{\pi^*(z_{\pi_e})\}_{z_{\pi_e} \in Z_{\pi_e}}$ be the set of optimal signals selected by the sender in equilibrium. Let $V_{z_{\pi_e}}$

be the sender's expected equilibrium payoff after players publicly observe z_{π_e} and the sender optimally selects $\pi^*(z_{\pi_e})$. The sender's ex ante expected payoff in this game is

$$V_{Pub} \equiv \sum_{z_{\pi_e} \in Z_{\pi_e}} Pr[z_{\pi_e}] V_{z_{\pi_e}}.$$

Notice that computing V_{Pub} is typically a much simpler task than computing V_I , as it only requires repetitive use of the arguments in KG to solve for the optimal signals — we do not have to worry about the receiver's interim beliefs about the sender's private information. Consequently, it is often simpler to verify if the sender benefits from public information (verify if $V_{Pub} > V_U$) than it is to verify if the sender benefits from private information (verify if $V_I > V_U$).

If π_e is privately observed by the sender, however, then she may not achieve V_{Pub} in equilibrium. Indeed, a strategy for the sender that selects $\pi^*(z_{\pi_e}(t))$ may not constitute a separating equilibrium if a type t prefers the receiver's choice under $\pi^*(z'_{\pi_e})$ when the receiver interprets signal $\pi^*(z'_{\pi_e})$ as being offered by type t' . The next proposition provides sufficient conditions for the existence of an equilibrium in which each privately informed sender chooses the same signal as the publicly informed sender, which implies that $V_I \geq V_{Pub}$. The proposition exploits a property of strongly redundant experiments: if the sender offers experiments that makes π_e strongly redundant, then the receiver would not revise his beliefs if she were to observe the actual realization of π_e .

Assumption (A1) (Monotone Preferences) For all $a, a' \in A$ and $\theta, \theta' \in \Theta$

$$(u_S(a', \theta) - u_S(a, \theta)) (u_S(a', \theta') - u_S(a, \theta')) \geq 0.$$

To simplify notation, let $A \subset \mathbb{R}$ and $u_S(a', \theta) \geq u_S(a, \theta)$ for $a' > a$ and $\theta \in \Theta$.

Proposition 3 Suppose (A1) holds.

(i) If there exists a selection of public optimal signals $\pi^*(z_{\pi_e})$, $z_{\pi_e} \in Z_{\pi_e}$, such that π_e is strongly redundant given $\Pi_{Pub}^* \equiv \{\pi^*(z_{\pi_e})\}_{z_{\pi_e} \in Z_{\pi_e}}$, then $V_I \geq V_{Pub}$.

(ii) If π_e can be replicated by each $\pi \in \Pi$, then $V_I \geq V_{Pub}$. In particular, if π_e and all signals in Π are partitional, with π_e coarser than each $\pi \in \Pi$, then $V_I \geq V_{Pub}$.

To prove the first part of the proposition, we construct a separating equilibrium in which each type t sender selects an experiment $\pi^*(z_{\pi_e}(t))$ that would be optimal if $z_{\pi_e}(t)$ were publicly observed. Strong redundancy implies that, on the equilibrium path, no sender benefits from mimicking another type's choice. In a sense, by offering experiments that make her private information strongly redundant, the sender is “letting the evidence speak for itself” — the receiver's interim belief after observing the choice of signal $\pi^* \in \Pi_{Pub}^*$ does not affect his posterior belief after observing the realization z_{π^*} of π^* .

Assumption **(A1)** allows us to form inferences off-the-equilibrium-path that discipline the sender to avoid signals not used in equilibrium. When the receiver observes realization $z_{\tilde{\pi}}$ of a signal $\tilde{\pi} \notin \Pi_{Pub}^*$, the receiver's interim belief about the sender is such that it leads him to choose the worst action from the point of view of all senders (cf. Assumption **(A1)**). Therefore, offering $\tilde{\pi}$ off-the-equilibrium leads to (weakly) lower actions than if $\tilde{\pi}$ were offered by type t when the realization of π_e were public. Therefore, type t cannot gain from offering $\tilde{\pi} \notin \Pi_{Pub}^*$ instead of $\pi^*(z_{\pi_e}(t))$.

The proof of part (ii) shows that if the sender's private signal can be replicated with each $\pi \in \Pi$, then there is a set of optimal signals that make her private signal strongly redundant. For instance, for a fixed z_{π_e} , she can construct a bidimensional garbling with one dimension that is perfectly correlated with her type t while the other dimension provides the outcome of experiment $\pi^*(z_{\pi_e})$. In particular, if π_e and all signals in Π are partitional, with π_e coarser than Π , then π_e can be replicated with each $\pi \in \Pi$ and we must have $V_I \geq V_{Pub}$.

We can then use Proposition 3 as a sufficient condition for the sender to strictly benefit from redundant information.

Corollary 3 *Suppose **(A1)** holds and there exists a selection of optimal experiments $\pi^*(z_{\pi_e})$ such that π_e is strongly redundant given $\Pi_{Pub}^* \equiv \{\pi^*(z_{\pi_e})\}_{z_{\pi_e} \in Z_{\pi_e}}$. If the sender strictly benefits from publicly observing π_e , then she strictly benefits from privately observing π_e , $V_{Pub} > V_U \Rightarrow V_I > V_U$.*

For instance, in Example 1 in the Introduction, if we assume that the realization of π_e is public, then it is easy to verify that $V_{Pub} > V_U$ — the sender strictly benefits from the ability to adapt her choice of experiment to the actual realization z_{π_e} . Moreover, **(A1)** holds

and π_e and all signals in Π are partitional, with π_e coarser than Π . Therefore, Proposition 3(ii) implies that $V_I \geq V_{Pub}$ and the sender strictly benefits from privately observing π_e .

5 When does Redundant Expertise Hurt the Sender?

Proposition 1 shows that a sender cannot benefit from observing a sequentially redundant private signal. We now study cases in which the sender is actually hurt by this interim information—that is, situations in which $V_U > V_I$. The fact that expertise can be detrimental resonates with some applications in which information limits the ability to persuade decision makers. For instance, failure to take a polygraph often leads to a negative update on a defendant’s innocence. Furthermore, in some cases, failure to submit to DNA testing in paternity lawsuits automatically assigns paternity to the non-compliant.

To obtain a sharp characterization, we restrict our attention to the following case:

Assumption (A2) $\Pi = \{\hat{\pi}\}$ and π_e can be replicated with $\hat{\pi}$.

As Π consists of a single experiment $\hat{\pi}$ and π_e can be replicated with $\hat{\pi}$, Assumption (A2) implies that π_e is also sequentially redundant and Proposition 1 establishes that $V_U \geq V_I$. One important case that satisfies (A2) is the case of partitional experiments π_e and $\hat{\pi}$, with $\hat{\pi}$ corresponding to a finer partition than π_e .

We allow for the possibility that the uninformed sender’s game admits multiple optimal experiments, and we let Π_U^* be the set of such optimal experiments. When studying the set of type-dependent equilibrium payoffs, we show that one can without loss restrict attention to pooling equilibria if (A2) holds. That is, the ability to replicate one’s type allows a sender to sustain any vector of equilibrium payoffs by pooling on a single experiment.

While the restriction to pooling equilibria resembles the inscrutability principle of Myerson (1983), our result is based on the statistical properties of available experiments relative to π_e . Following Myerson (1983), we could allow ex-post communication by the sender and posit that all sender types select a single experiment and then communicate the information about their type revealed in an equilibrium. In principle, this would be possible if, for instance, π_e is sequentially redundant, as one could always find an experiment $\pi \in \Gamma(\Pi)$ that

replicates the same distribution over realizations as the one induced by any given equilibrium (cf. Proposition 1). However, sequential redundancy is not sufficient for the distribution over receiver’s actions *conditional on each sender’s type* to be the same as in the given equilibrium. As we show in the proof of Proposition 4, a pooling equilibrium can replicate both the sender’s and receiver’s equilibrium payoffs if, instead, **(A2)** holds.

We then show that a sender is strictly worse off when privately informed if and only if, for every optimal experiment $\pi_U^* \in \Pi_U^*$, she cannot pool in equilibrium and offer π_U^* . Let $v_{\pi_U^*}^*(t)$ be the interim expected utility of a type t sender when all types pool on experiment π_U^* . From Section 4, recall that $V_{z_{\pi_e}(t)}$ is the value of persuasion when sender and receiver publicly observed the realization $z_{\pi_e}(t)$.

Proposition 4 *Suppose that **(A1)** and **(A2)** hold. Then, $V_U > V_I$ if and only if*

$$\min_{\pi_U^* \in \Pi_U^*} \max_{t \in T} \left[V_{z_{\pi_e}(t)} - v_{\pi_U^*}^*(t) \right] > 0. \quad (5)$$

To understand the proposition, we first compute the lowest interim expected utility for a type t in any equilibrium of the informed sender game. In the proof of the proposition, we show that, when the receiver updates in the most adverse way following the experiment’s realization of an out-of-the-equilibrium deviation, the sender can always improve her payoffs from such deviation by simultaneously replicating her type. This implies that the minimum expected utility that a type t can guarantee herself in any equilibrium is $V_{z_{\pi_e}(t)}$.

We can now interpret (5). Condition (5) implies that, for every optimal experiment $\pi_U^* \in \Pi_U^*$, there is some type t' with $V_{z_{\pi_e}(t')} > v_{\pi_U^*}^*(t')$. Therefore, pooling on π_U^* cannot be an equilibrium of the informed-sender game. To wit, the informed sender is hurt by her expertise if, for every optimal pooling experiment, some type would prefer to offer an experiment that both “certifies” her type and is an optimal experiment when her type is public.

Under the conditions in Proposition 4, pooling is sustained in equilibrium because possible deviations are made unprofitable by eliciting the receiver’s most adverse update. The ability of the sender to certify her type, coupled with the most adverse update by the receiver, allows us to characterize the lowest individually rational payoff that each type can obtain in any equilibrium. Considering refinements of signaling games that may rule out certain inferences would certainly make pooling less likely and would widen the gap between the uninformed

and the informed senders' payoffs. Situations where the sender cannot replicate her type may in turn allow the sender to sustain payoff profiles that are not attainable through pooling equilibria. We leave these investigations for future research.

6 Application: Persuading Consumers

We next present an application that illustrates our results. To persuade a consumer (receiver), the seller (sender) can design a public signal (test of the product or marketing campaign) that allows the consumer to learn about his true valuation of the product — see Johnson and Myatt (2006) for many examples of how a firm can provide real information to consumers, allowing them to learn of their personal match with a product. Our application captures many important market features: (i) the consumer often faces a menu with several different possible options, (ii) he is uncertain about how important for him are the many different features of the products, and (iii) the seller has some control over what the consumer can learn, by strategically designing product tests and the informational content of marketing campaigns.

Formally, consider a consumer who must choose which product to buy. For concreteness, suppose it is a smartphone. The consumer can buy one phone from brands A, B or C, or the consumer can choose not to buy a phone. Brand C is a more expensive and advanced phone, while brands A and B are cheaper but have very distinctive features. The consumer is not familiar with the different brands and types of phones, so he is uncertain about which phone is the best match for his needs. This uncertainty is captured by the unknown state $\theta \in \{AH, AL, BH, BL, C\}$, and players have a uniform prior belief over these states. State $\theta = C$ implies that brand C is the best match and the consumer should buy phone C. State $\theta = AH$ implies that brand A is the best match and it has a high consumption value (above price), so the consumer should buy A. State $\theta = AL$ implies that brand A is the best match but it has a low consumption value (below price), so the consumer should not buy a phone. Similarly, state $\theta = BH$ implies that the consumer should buy B, while $\theta = BL$ implies that he should not buy a phone.

Let q be the consumer's posterior belief. To streamline the presentation, we assume that the consumer's optimal action follows a simple rule. The consumer buys phone C if and

only if he is certain that this is the best option, $Pr(\theta = C) = 1$. The consumer buys phone A if and only if $Pr(\theta = C) + Pr(\theta = AH) \geq 0.8$ and $Pr(\theta = BH) + Pr(\theta = BL) = 0$. That is, the consumer has great disutility if he buys phone A when his type is B, but he continues to value phone A if his type is C. Similarly, the consumer buys phone B if and only if $Pr(\theta = C) + Pr(\theta = BH) \geq 0.8$ and $Pr(\theta = AH) + Pr(\theta = AL) = 0$. The consumer does not buy a phone in the remaining cases.

A retailer profits from selling the phones. The retailer's payoff from selling a C phone is 12, while her payoff from selling an A or B phone is 10. The retailer receives zero if she does not sell. We next consider two types of retailers.

Constrained Retailer: Consider a retailer that is constrained on the experiments that she can offer to a consumer. She only has access to two partitional experiments, $\Pi = \{\pi_A, \pi_B\}$. Experiment π_A reveals partitions $\{AH\}$, $\{AL\}$, $\{BH, BL\}$ and $\{C\}$. Experiment π_B reveals partitions $\{AH, AL\}$, $\{BH\}$, $\{BL\}$ and $\{C\}$. That is, both experiments can easily identify if the consumer's type is A, B, or C. However, for types A and B, the retailer needs to use the targeted test π_A or π_B to differentiate between a high and a low value from consumption. This captures the natural assumption that a more specific experiment is needed to test the consumer's valuation of the distinctive features of each brand.

If the retailer has no private information, then the following is an optimal experiment. The retailer garbles π_A and designs a test with two realizations, $S = \{s_A, s_0\}$. States C and AH induce realization s_A with probability one, while partition $\{BH, BL\}$ induces realization s_0 with probability one. State AL induces realizations s_A and s_0 with equal probability. Upon observing s_A , the consumer's posterior belief becomes $Pr(\theta = C) = 0.4$, $Pr(\theta = AH) = 0.4$, $Pr(\theta = AL) = 0.2$, and $Pr(\theta = BH) = Pr(\theta = BL) = 0$. The consumer then chooses to buy A. The consumer does not buy a phone if she observes s_0 . Figure 1(a) illustrates how this constrained and uninformed retailer bundles the different states into the different recommendations to the consumer. The retailer's expected payoff is $(0.2 + 0.2 + \frac{1}{2} \times 0.2) \times 10 + 0 = 5$. Note that this retailer does not find it optimal to sell the more expensive phone C. It is more profitable to bundle type C and type A consumers. The same optimal payoff can be attained by a similar garbling of π_B .

Now suppose that the retailer can acquire private information. For example, the retailer




Figure 1: Retailer's Equilibrium Recommendations to the Consumer

can hire an expert salesperson that is trained to quickly identify the consumer's type. The retailer's private signal π_e identifies partitions $\{AH, AL\}$, $\{BH, BL\}$ and $\{C\}$. In this case, the privately informed retailer can no longer bundle consumers with type C . In equilibrium, if the retailer learns that the consumer is C , she will always provide this information and the consumer will buy phone C . If the retailer privately observes partition $\{AH, AL\}$, then she garbles π_A as to maximize the probability of the consumer buying phone A , given that the consumer understands that the state is not C . For instance, upon observing partition $\{AH, AL\}$, the retailer implements a binary signal $S = \{s_A, s_0\}$. State AH induces realization s_A with probability one; state AL induces realization s_A with probability 0.25 and s_0 with probability 0.75; the remaining states induce s_0 with probability one. A similar garble of π_B is optimal when the retailer learns that the state is in partition $\{BH, BL\}$. Figure 1(b) illustrates how, in equilibrium, the informed retailer bundles the states into different recommendations. The retailer's expected payoff is $0.2 \times 12 + (0.2 + 0.25 \times 0.2) \times 10 + (0.2 + 0.25 \times 0.2) \times 10 = 7.4$.

Note that our application satisfies the conditions of Proposition 3(ii) and Corollary 3, so that our constrained sender strictly benefits from becoming an expert.

Unconstrained Retailer: Now suppose that the retailer has access to a fully informative signal, $\Pi = \{\pi_{FI}\}$. If she does not have private information, then the following is an optimal experiment. She implements an experiment with three realizations, $S = \{s_A, s_B, s_0\}$. State C induces realizations s_A and s_B with probability 0.5 each. For $j \in \{A, B\}$, state jH induces realization s_j with probability one; state jL induces realization s_j with probability

0.375 and realization s_0 with probability 0.625. Consequently, the consumer buys a phone j when he observes s_j , and she does not buy a phone when she observes s_0 . Figure 1(c) illustrates how this unconstrained and uninformed retailer bundles the states into different recommendations. The retailer's expected payoff is $(0.1 + 0.2 + 0.375 \times 0.2) \times 10 + (0.1 + 0.2 + 0.375 \times 0.2) \times 10 + 0 = 7.5$. Note that the uninformed retailer prefers to bundle type C with types A and B .

Finally, suppose that the retailer can acquire private information π_E — she can hire the same expert salesperson from the previous example, who can quickly identify partitions $\{AH, AL\}$, $\{BH, BL\}$ and $\{C\}$. In this case, the informed sender can no longer commit to bundle a type C with the other types. In our application, the optimal experiment of the unconstrained informed sender results in the same distribution over buyers' actions as in the previous example with a constrained informed sender — See Figure 1(b). In both cases, the retailer's expected payoff is 7.4.

Note that our application satisfies the conditions in Proposition 4, so that our unconstrained sender strictly loses from becoming an expert.

In summary, our application captures the fact that consumers are often overwhelmed by a long menu of choices. If the retailer does not have access to a fully informative test, then she may benefit from an expert salesperson who is able to select the best experiment for each particular consumer. Hence, we can think about the case of companies strategically training their salesforce, or the firm's executives gathering information prior to designing their marketing strategies. However, if the retailer has access to a fully informative test, then she might be worse off if workers and executives have private information, as they might be unable to withhold disclosure of this information to the consumers.

7 Conclusion

When is an expert a more effective persuader? If a sender has access to a set of certifiable public signals (experiments), observing a private signal prior to choosing an experiment may allow her to revert to an experiment that elicits a more desired behavior in a receiver. However, we show that the informativeness of public experiments substitutes for the sender's

expertise: an uninformed sender can always achieve the payoffs of an expert if she has access to experiments that are sufficiently informative. Our key condition (sequential redundancy) ensures that an uninformed sender can always replicate both the expert’s private signal and her ensuing choice of experiment. Perhaps surprisingly, redundant private information may still be valuable to the sender when it allows her to choose between experiments that carry different information.

An important implication of our analysis is that a sender can never benefit from becoming an expert when a fully informative public experiment is available. We then show that expertise may be detrimental to a sender if pooling on the uninformed sender’s optimal experiment is not an equilibrium on the informed-sender game. In these situations, the sender could benefit from strategic ignorance — taking steps to guarantee to the receiver that she did not acquire private information. For instance, the prosecutor might prefer not to meet a particular witness, so that the judge knows that she did not ask the witness exploratory questions before the trial. Similarly, a retailer with access to a fully informative experiment might prefer to hire uninformed salespeople, while a constrained retailer might prefer to hire expert salespeople.

8 Appendix

Proof of Proposition 1: Sufficiency: We prove sufficiency without requiring A to be finite. Furthermore, we will show that every joint distribution over payoffs and the state achieved in equilibrium by an informed sender can be replicated by an uninformed sender. This trivially implies that $V_U \geq V_I$.

Consider an equilibrium of the informed sender game in which the sender selects an experiment according to the mixing $\sigma(\pi|t)$ with support $\Pi_t \subset \Gamma(\Pi)$, and denote by π_σ the corresponding sequential experiment induced in equilibrium, where the receiver observes first the chosen experiment π and then its realization. Sequential redundancy implies that there exists a mixture $\{\lambda_i\}_{i \in I}$ so that

$$\pi_\sigma \preceq_B \sum_{i \in I} \lambda_i \pi_i.$$

Thus, there exists a garbling of $\sum_{i \in I} \lambda_i \pi_i$ that generates the same joint distribution over posterior beliefs and the state as π_σ (Blackwell 1953).

Necessity: As in Perez-Richet (2016), let $\Phi(\pi)$ denote the set of distributions over $A \times \Theta$ induced by decision rules based on the outcome of experiment π , and note that $\Phi(\pi)$ is compact and convex in $[0, 1]^{|A| \times |\Theta|}$. Clearly, for any garbling $\tilde{\pi}$ of π we have $\Phi(\tilde{\pi}) \subseteq \Phi(\pi)$. If one considers mixtures over experiments in Π , then the set of distributions over $A \times \Theta$ induced by decision rules based on arbitrary mixtures of experiments in Π coincides with $\Phi(\Gamma(\Pi)) = \text{conv}(\cup_{\pi \in \Pi} \Phi(\pi))$, which is clearly convex. Finiteness of Π implies that $\Phi(\Gamma(\Pi))$ is compact.

Suppose that π_e is not sequentially redundant given $\Gamma(\Pi)$. Then, there exists a sequential experiment π_σ described by a randomized selection rule $\sigma(\cdot | z_{\pi_e})$, with support $\Pi_{z_{\pi_e}} \subset \Gamma(\Pi)$, so that $\sigma(\pi | z_{\pi_e})$ is the probability of choosing experiment $\pi \in \Pi_{z_{\pi_e}}$ after realization z_{π_e} , and such that for every mixture λ we have $\pi_\sigma \not\preceq_B \sum_{i \in I} \lambda_i \pi_i$. Therefore, there exists a joint distribution $\varphi(a, \theta)$ induced by selecting actions according to the outcome of π_σ such that $\varphi(a, \theta) \notin \Phi(\Gamma(\Pi))$ (Perez-Richet 2016). The separating hyperplane theorem ensures the existence of a payoff function $u(a, \theta)$ such that

$$\sum_{A \times \Theta} \varphi(a, \theta) u(a, \theta) > W(\Pi, u) = \max_{\varphi' \in \Phi(\Gamma(\Pi))} \sum_{A \times \Theta} \varphi'(a, \theta) u(a, \theta),$$

where $W(\Pi, u)$ denotes the maximum expected payoff of a decision maker with utility $u(a, \theta)$ when making decisions based on mixtures of experiments in Π . Consider, then, a sender-receiver game in which $u_R(a, \theta) = u_S(a, \theta) = u(a, \theta)$. A type t sender can select after each $z_{\pi_e}(t)$ a mixture $\sigma(\pi | z_{\pi_e}(t))$ that has the same distribution over states and outcomes as π_σ (by relabeling, if necessary, the outcomes of π_σ). This constitutes a separating equilibrium of this persuasion game in which the sender credibly signals her type and thus $V_I > V_U$.

■

Proof of Corollary 1: It follows from the sufficiency proof of Proposition 1, which does not require A to be finite, and the fact that every π_e is sequentially redundant given $\Gamma(\Pi)$ if $\pi_{FI} \in \Pi$.

Lemma A1: Suppose that $\pi_a \preceq_B \sum_{j \in I_B} \lambda_j^a \pi_j$ and $\pi_b \preceq_B \sum_{j \in I_B} \lambda_j^b \pi_j$. Then, for $\alpha \in [0, 1]$, we have $\alpha \pi_a + (1 - \alpha) \pi_b \preceq_B \sum_{j \in I_B} (\alpha \lambda_j^a + (1 - \alpha) \lambda_j^b) \pi_j$.

Proof of Lemma A1: Let M_j^k be a Markov matrix representing the garbling of experiment π_j such that, for $k \in \{a, b\}$, the mixture $\sum_{j \in I_B} \lambda_j^k \pi_j$ followed by the garbling M_j^k of π_j ,

$j \in I_B$, has the same distribution over outcomes as π_k . Now consider the garbling of the mixture $\sum_{j \in I_B} (\alpha \lambda_j^a + (1 - \alpha) \lambda_j^b) \pi_j$ where, for any j such that $\alpha \lambda_j^a + (1 - \alpha) \lambda_j^b > 0$, the outcome of experiment π_j is garbled according to the Markov matrix

$$\frac{\alpha \lambda_j^a}{\alpha \lambda_j^a + (1 - \alpha) \lambda_j^b} M_j^a + \frac{(1 - \alpha) \lambda_j^b}{\alpha \lambda_j^a + (1 - \alpha) \lambda_j^b} M_j^b.$$

It is then immediate to verify that such experiment generates the same distribution as $\alpha \pi_a + (1 - \alpha) \pi_b$. ■

Proof of Proposition 2: Part (a)(i)- Sequential redundancy implies that, for each $i \in I_B$, there exists a mixture λ^i such that $\{\pi_e, \pi_i\} \preceq_B \sum_{j \in I_B} \lambda_j^i \pi_j$. But, then, $\pi_i \preceq_B \sum_{j \in I_B} \lambda_j^i \pi_j$ and linear independence implies $\lambda_j^i = 0$ for $i \neq j$, so that $\{\pi_e, \pi_i\} \preceq_B \pi_i$. As π_e is strongly redundant, the set of posterior beliefs induced by jointly observing the realization of π_e and π_i , $i \in I_B$, is contained in $Q(\pi_i)$.

Part (a)(ii)- If Π_B has a single element then (2) is trivially satisfied.⁴ Suppose, then, that Π_B has at least two elements. To simplify notation, let τ_i and $\tau_i(z_{\pi_e})$ denote the distribution over posterior beliefs, both with support on $Q = \cup_{i \in I_B} Q(\pi_i)$, induced by experiment π_i and by experiment π_i conditional on observing realization z_{π_e} —so $\tau_i(q) = \Pr_i(q)$ and $\tau_i(z_{\pi_e})(q) = \Pr_i(q|z_{\pi_e})$. Define

$$\begin{aligned} \Delta_{i,j}(q) &= \tau_i(q) - \tau_j(q), \\ \Delta_{i,j}(z_{\pi_e})(q) &= \tau_i(z_{\pi_e})(q) - \tau_j(z_{\pi_e})(q). \end{aligned} \tag{6}$$

The proposition then states that there is $\alpha_{ij}(z_{\pi_e}) \geq 0$ so that

$$\Delta_{i,j}(z_{\pi_e})(q) = \alpha_{ij}(z_{\pi_e}) \Delta_{i,j}(q), \tag{7}$$

for every $q \in Q_{ext}^{ij}$. Multiplying both sides of (7) by $p(z_{\pi_e})$ and adding over $z_{\pi_e} \in Z_{\pi_e}$ implies that $\sum_{z_{\pi_e} \in Z_{\pi_e}} \Pr(z_{\pi_e}) \alpha_{ij}(z_{\pi_e}) = 1$.

We now prove (7). Consider the experiment $\tilde{\pi}_{ij, z_{\pi_e}}$ constructed as follows: if z_{π_e} is realized then it selects experiment π_i , otherwise it selects experiment π_j , $j \neq i$. Denote by $\tilde{\tau}_{ij, z_{\pi_e}}$ the

⁴This would be the case for a general set of experiments Π if Π admits a most informative experiment according to the Blackwell information order.

induced distribution on Q . Sequential redundancy implies that there are weights $\mu_{ij,z_{\pi_e}}^k \geq 0$ with $\sum_{k \in I_B} \mu_{ij,z_{\pi_e}}^k = 1$ and

$$\tilde{\pi}_{ij,z_{\pi_e}} \preceq_B \sum_{k \in I_B} \mu_{ij,z_{\pi_e}}^k \pi_k.$$

Let $R = \text{card}(Z_{\pi_e})$ be the number of outcomes of experiment π_e . Then we have

$$\sum_{z_{\pi_e} \in Z_{\pi_e}} \tilde{\tau}_{ij,z_{\pi_e}} = \sum_{z_{\pi_e} \in Z_{\pi_e}} (\tau_j + p(z_{\pi_e})\Delta_{i,j}(z_{\pi_e})) = R\tau_j + (\tau_i - \tau_j),$$

so that

$$\sum_{z_{\pi_e} \in Z_{\pi_e}} \frac{1}{R} \tilde{\tau}_{ij,z_{\pi_e}} = \frac{1}{R} \tau_i + \left(\frac{1}{R} - 1 \right) \tau_j.$$

Repeated application of Lemma A1 to the mixture of experiments $\sum_{z_{\pi_e} \in Z_{\pi_e}} \frac{1}{R} \tilde{\pi}_{ij,z_{\pi_e}}$ establishes that

$$\frac{1}{R} \sum_{z_{\pi_e} \in Z_{\pi_e}} \tilde{\pi}_{ij,z_{\pi_e}} = \frac{1}{R} \pi_i + \left(\frac{1}{R} - 1 \right) \pi_j \preceq_B \sum_{z_{\pi_e} \in Z_{\pi_e}} \sum_{k \in I_B} \frac{\mu_{ij,z_{\pi_e}}^k}{R} \pi_k.$$

Linear independence of Π_B then implies

$$\sum_{z_{\pi_e} \in Z_{\pi_e}} \sum_{k \neq i,j} \mu_{ij,z_{\pi_e}}^k = 0.$$

As $\mu_{ij,z_{\pi_e}}^k \geq 0$, we must then have $\mu_{ij,z_{\pi_e}}^k = 0$ for each $k \neq i, j$ and $z_{\pi_e} \in Z_{\pi_e}$. Therefore, the experiment $\tilde{\pi}_{ij,z_{\pi_e}}$ can only be represented through a garbling of mixtures involving only experiments π_i and π_j .

Consider now a belief $q \in Q_{ext}^{ij}$. The probability of posterior q under experiment $\tilde{\pi}_{ij,z_{\pi_e}}$ is

$$\tilde{\tau}_{ij,z_{\pi_e}}(q) = \tau_j(q) + p(z_{\pi_e})\Delta_{i,j}(z_{\pi_e})(q).$$

As q is an extreme point of $Q^{ij} = Q(\pi_i) \cup Q(\pi_j)$, it can only be induced by the same belief induced after realizations of π_i and π_j . Sequential redundancy implies $\tilde{\pi}_{ij,z_{\pi_e}} \preceq_B \alpha_{ij}(z_{\pi_e})\pi_i + (1 - \alpha_{ij}(z_{\pi_e}))\pi_j$ for some $\alpha_{ij}(z_{\pi_e})$ so that, for all $q \in Q_{ext}^{ij}$,

$$\alpha_{ij}(z_{\pi_e})\tau_i(q) + (1 - \alpha_{ij}(z_{\pi_e}))\tau_j(q) = \tau_j(q) + p(z_{\pi_e})\Delta_{i,j}(z_{\pi_e})(q),$$

which implies (7).

Part (b). Necessity. As all posteriors in Q are extreme, $Q = Q_{ext}$, then Part (a)(ii) already showed $\Delta_{i,j}(z_{\pi_e}) = \alpha_{ij}(z_{\pi_e})\Delta_{i,j}$. We now show that $\alpha_{ij}(z_{\pi_e})$ is constant across pairs i ,

$j \in I_B$. Consider an experiment $\tilde{\pi}$ characterized by the collection of mixtures $\{\lambda(z_{\pi_e})\}_{z_{\pi_e} \in Z_{\pi_e}}$, with $\lambda_i(z_{\pi_e})$ the probability that experiment $\pi_i \in \Pi_B$ is selected after observing z_{π_e} , and $\tilde{\tau}$ the induced distribution on Q . For a fixed $l \in I_B$, we have

$$\begin{aligned}
\tilde{\tau} &= \sum_{z_{\pi_e} \in Z_{\pi_e}} p(z_{\pi_e}) \left(\sum_{i \in I_B} \lambda_i(z_{\pi_e}) \tau_i(z_{\pi_e}) \right) = \\
&= \sum_{z_{\pi_e} \in Z_{\pi_e}} p(z_{\pi_e}) \left(\sum_{i \in I_B} \lambda_i(z_{\pi_e}) (\tau_l(z_{\pi_e}) + \alpha_{il}(z_{\pi_e}) \Delta_{i,j}) \right) \\
&= \sum_{z_{\pi_e} \in Z_{\pi_e}} p(z_{\pi_e}) \left(\tau_l(z_{\pi_e}) + \sum_{i \in I_B} \lambda_i(z_{\pi_e}) \alpha_{il}(z_{\pi_e}) \Delta_{i,j} \right) \\
&= \tau_l + \sum_{z_{\pi_e} \in Z_{\pi_e}} \sum_{i \in I_B} p(z_{\pi_e}) \lambda_i(z_{\pi_e}) \alpha_{il}(z_{\pi_e}) \Delta_{i,j}.
\end{aligned}$$

Setting

$$\gamma_{il} = \sum_{z_{\pi_e} \in Z_{\pi_e}} p(z_{\pi_e}) \lambda_i(z_{\pi_e}) \alpha_{il}(z_{\pi_e}),$$

and noting that $\gamma_{ll} = 0$, we have

$$\tilde{\tau} = \tau_l + \sum_{i \in I_B} \gamma_{il} (\tau_i - \tau_l) = \sum_{i \in I_B, i \neq l} \gamma_{il} \tau_i + \left(1 - \sum_{i \in I_B, i \neq l} \gamma_{il} \right) \tau_l.$$

Since π_e is sequentially redundant given Π_B , and all posterior beliefs in Q are extreme points, we must have that $\tilde{\tau}$ belongs to the convex hull of $\{\tau_i\}_{i \in I_B}$. This requires that, for each possible choice of $\{\lambda(z_{\pi_e})\}_{z_{\pi_e} \in Z_{\pi_e}}$, we must have $\sum_{i \in I_B, i \neq l} \gamma_{il} \leq 1$. Note that

$$\sum_{i \in I_B, i \neq l} \gamma_{il} = \sum_{z_{\pi_e} \in Z_{\pi_e}} p(z_{\pi_e}) \sum_{i \in I_B, i \neq l} \lambda_i(z_{\pi_e}) \alpha_{il}(z_{\pi_e}).$$

For each realization z_{π_e} , select $i_{z_{\pi_e}}^*$ so that $\alpha_{i_{z_{\pi_e}}^* l}(z_{\pi_e}) = \max_{j \in I_B} \alpha_{jl}(z_{\pi_e})$. The program

$$V_l = \max_{\lambda(z_{\pi_e})} \sum_{i \in I_B, i \neq l} \gamma_{il}, \text{ s.t. } \lambda_i(z_{\pi_e}) \geq 0, \sum_{i \in I_B, i \neq l} \lambda_i(z_{\pi_e}) \leq 1,$$

is maximized by setting $\lambda_j(z_{\pi_e}) = 1$ iff $j = i_{z_{\pi_e}}^*$. In this case, we have

$$V_l = \sum_{z_{\pi_e} \in Z_{\pi_e}} p(z_{\pi_e}) \max_{j \in I_B} \alpha_{jl}(z_{\pi_e}) \geq \sum_{z_{\pi_e} \in Z_{\pi_e}} p(z_{\pi_e}) \alpha_{kl}(z_{\pi_e}) = 1$$

for every $k \in I_B, k \neq l$, where the inequality is strict whenever $\max_{j \in I} \alpha_{jl}(z_{\pi_e}) > \alpha_{kl}(z_{\pi_e})$ for some z_{π_e} and k . Therefore, we must have that $\alpha_{kl}(z_{\pi_e}) = \max_{j \in I} \alpha_{jl}(z_{\pi_e}) = \alpha_l(z_{\pi_e})$.

Since $\alpha_{kl}(z_{\pi_e}) = \alpha_{lk}(z_{\pi_e})$, we then must have that $\alpha_{lk}(z_{\pi_e})$ is constant across all z_{π_e} for each $l, k \in I_B$.

Sufficiency: It suffices to show that (3) implies that π_e is sequentially redundant given mixtures in Π_B . Consider an experiment $\hat{\pi}$ characterized by the collection of mixtures $\{\lambda(z_{\pi_e})\}_{z_{\pi_e} \in Z_{\pi_e}}$ of experiments in Π_B , that induces a distribution $\hat{\tau}$ on Q . Then with $\Delta_{i,j} = \tau_i - \tau_j$,

$$\begin{aligned} \hat{\tau} &= \sum_{z_{\pi_e} \in Z_{\pi_e}} p(z_{\pi_e}) \left(\sum_{i \in I_B} \lambda_i(z_{\pi_e}) \tau_i(z_{\pi_e}) \right) = \\ &= \sum_{z_{\pi_e} \in Z_{\pi_e}} p(z_{\pi_e}) \left(\sum_{i \in I_B} \lambda_i(z_{\pi_e}) (\tau_k(z_{\pi_e}) + \alpha(z_{\pi_e}) \Delta_{ik}) \right) \\ &= \tau_k + \sum_{z_{\pi_e} \in Z_{\pi_e}} \sum_{i \in I_B} p(z_{\pi_e}) \lambda_i(z_{\pi_e}) \alpha(z_{\pi_e}) \Delta_{ik}. \end{aligned}$$

Define

$$\begin{aligned} \kappa_i &= \sum_{z_{\pi_e} \in Z_{\pi_e}} \lambda_i(z_{\pi_e}) p(z_{\pi_e}) \alpha(z_{\pi_e}), i \neq k \\ \kappa_k &= 1 - \sum_{i \in I_B} \kappa_i. \end{aligned}$$

Note that $\{\kappa_i\}_{i \in I_B}$ satisfies $\kappa_i \geq 0$, $\sum_{i \in I_B} \kappa_i = 1$ and $\hat{\tau} = \sum_{i \in I_B} \kappa_i \tau_i$. Therefore, for each sequential experiment $\hat{\pi}$, we can find a mixture $\{\kappa_i\}_{i \in I}$ of experiments in Π_B with

$$\hat{\pi} \preceq_B \sum_{i \in I_B} \kappa_i \pi_i,$$

implying that π_e is sequentially redundant given Π_B . ■

Proof of Corollary 2: Follows immediately from (2) by setting for each $q, q' \in Q_{ext}$

$$\frac{\Pr_i(q|z_{\pi_e}) - \Pr_j(q|z_{\pi_e})}{\Pr_i(q) - \Pr_j(q)} = \alpha_{ij}(z_{\pi_e}) = \frac{\Pr_i(q'|z_{\pi_e}) - \Pr_j(q'|z_{\pi_e})}{\Pr_i(q') - \Pr_j(q')}.$$

■

Proof of Proposition 3: Part (i): Let $\hat{a}(z_{\pi}, z_{\pi_e})$ be the receiver's equilibrium action after publicly observing z_{π_e} and the realization of π . We show that the sender's strategy $\pi^*(t) = \pi^*(z_{\pi_e}(t))$ and receiver's choice $a(z_{\pi^*(t)}) = \hat{a}(z_{\pi^*(t)}, z_{\pi_e}(t))$ after observing $\pi^*(t) \in \Pi_{Pub}^*$ and its realization forms a separating equilibrium of the informed sender game. First, consider

deviations off-the-equilibrium path to an experiment $\tilde{\pi} \in \Gamma(\Pi)$ with $\tilde{\pi} \notin \Pi_{Pub}^*$, and suppose that the receiver's posterior action after observing $z_{\tilde{\pi}}$ satisfies $a(z_{\tilde{\pi}}) = \min_{t' \in T} \hat{a}(z_{\tilde{\pi}}, z_{\pi_e}(t'))$. That is, the receiver's update off-the-equilibrium path assigns probability 1 to a type t' that leads to the lowest possible action consistent with the realization $z_{\tilde{\pi}}$ of $\tilde{\pi}$. Since $\hat{a}(z_{\tilde{\pi}}, z_{\pi_e}(t)) \geq a(z_{\tilde{\pi}})$, then experiment $\tilde{\pi}$ induces a pointwise lower action off-the-equilibrium path than if the sender's type was observed by the receiver. Since each type t has available experiment $\tilde{\pi}$ when π_e is publicly observed, then type t cannot profit from offering $\tilde{\pi}$ when privately informed.

Second, consider on-the-equilibrium path deviations so that type t offers $\pi^*(t')$ instead of $\pi^*(t)$. The strong redundancy assumption $\{\pi_e, \pi^*(z_{\pi_e})\} \preceq_B \pi^*(z_{\pi_e})$ implies that the receiver's posterior belief after observing $z_{\pi^*(t')}$ is independent of his interim belief over T . That is, type t cannot gain by mimicking another type t' when the set of type-dependent optimal experiments makes the private signal strongly redundant. As $\pi^*(t')$ is available to type t when types are public, then she cannot gain by selecting $\pi^*(t')$ instead of $\pi^*(t)$.

Part (ii): Let I index the experiments in Π . We now show that if π_e can be replicated with π_i , $i \in I$, then there exists a selection of experiments $\hat{\pi}_{z_{\pi_e}}, z_{\pi_e} \in Z_{\pi_e}$, that makes π_e strongly redundant given $\Pi_{Pub}^* \equiv \{\hat{\pi}_{z_{\pi_e}}\}_{z_{\pi_e} \in Z_{\pi_e}}$. Fix a realization z_{π_e} and an optimal public experiment $\pi^*(z_{\pi_e})$. Suppose that $\pi^*(z_{\pi_e})$ is generated by a mixture μ_{z_e} of garblings g_{i,z_e}^* , $\pi^*(z_{\pi_e}) = \sum_{i \in I} \mu_{i,z_e} (g_{i,z_e}^* \circ \pi_i)$. Since π_e can be replicated with π_i , $i \in I$, there exists a garbling g_i of π_i satisfying (1). The bidimensional garbling \tilde{g}_{i,z_e} of π_i with outcome $(g_{i,z_e}^* \circ \pi_i, g_i \circ \pi_i)$ also replicates π_e . Then, the bidimensional experiment $\hat{\pi}_{z_{\pi_e}}$,

$$\hat{\pi}_{z_{\pi_e}} = \sum_{i \in I} \mu_{i,z_e} (g_{i,z_e}^* \circ \pi_i, g_i \circ \pi_i),$$

which is a garbling of the mixture μ_{z_e} , satisfies $\{\pi_e, \hat{\pi}_{z_{\pi_e}}\} \preceq_B \hat{\pi}_{z_{\pi_e}}$.

Finally, suppose that π_e and all signals in Π are partitionial, with π_e coarser than each $\pi \in \Pi$. Then each $\pi \in \Pi$ can replicate π_e and, combining parts (i) and (ii) of Proposition 3, we have $V_I \geq V_{Pub}$. ■

Proof of Proposition 4: Step 1) We show that, for any equilibrium of the informed-sender game, **(A2)** guarantees the existence of a pooling equilibrium with the same type-dependent payoffs. To see this, consider an equilibrium in which the sender selects an experiment according to the mixing $\sigma(\pi|t)$ with support $\Pi_t \subset \Gamma(\Pi)$ and let $\Pi' = \cup_{t \in T} \Pi_t$.

Let $\{1, \dots, J\}$ index the set Π' . As $\Pi = \{\hat{\pi}\}$, every $\pi_j \in \Pi'$, $j \in \{1, \dots, J\}$, is a garbling of the same experiment $\hat{\pi}$.

Let $\tilde{\pi}_e$ be a garbling of $\hat{\pi}$ that replicates $\tilde{\pi}_e$. Consider the multidimensional experiment $\pi' = (\pi_1, \dots, \pi_J, \tilde{\pi}_e)$ which is a garbling of $\hat{\pi}$. We now construct an experiment π_P from π' that can be supported in a pooling equilibrium and it induces the same type-dependent payoffs as the original equilibrium. From π' , experiment π_P reveals $z_{\pi_P} = (j, z_{\pi_j})$ with probability $\sigma(\pi_j|\tilde{t})$, where \tilde{t} is the type associated to the realization of $z_{\tilde{\pi}_e}$. That is, π_P selects an experiment in (π_1, \dots, π_J) according to the mixing in the original equilibrium where, instead of type t , it uses the outcome of the experiment $\tilde{\pi}_e$.

First, as $\tilde{\pi}_e$ replicates π_e , we have that for each type t , $\Pr[z_{\tilde{\pi}_e} = t|t] = 1$. Therefore, the conditional distribution over experiments in the original equilibrium, given by $\sigma(\pi|t)$, is the same as the conditional distribution $\sigma(\pi_j|\tilde{t})$ of the pooling equilibrium given the realization of $\tilde{\pi}_e$. Thus, it is a sequentially rational response for the receiver to select the same actions after observing (j, z_{π_j}) as in the original equilibrium.

Second, since $\Pr[z_{\tilde{\pi}_e} = t|t] = 1$ and all experiments are garblings of $\hat{\pi}$, we must then have that for each experiment π_j , $\Pr[z_{\pi_j}|t] = \Pr[z_{\pi_j}|z_{\tilde{\pi}_e} = t]$. Therefore, the type-dependent payoffs are the same when pooling on π_P as in the original equilibrium, and any deviation from pooling can be made unprofitable by the same inference off-the-equilibrium as in the original equilibrium. Thus, pooling on π_P is an equilibrium of the informed sender game.

Step 2) (Sufficiency) We prove the contrapositive. Thus, suppose that $V_U = V_I$. Then, there is an equilibrium of the informed sender game that achieves the same ex-ante expected payoff as the equilibrium if uninformed. From Step 1, we must have an equilibrium in which all types pool on some experiment π_p with $\pi_p \in \Pi_U^*$. As π_p is an equilibrium, each type cannot profit from certifying her type, so

$$\max_{t \in T} [V_{z_{\pi_e}(t)} - v_{\pi_p}^*(t)] \leq 0,$$

thus violating (5).

Step 3) (Necessity) We first show that the minimum payoff a sender can obtain by deviating from a pooling equilibrium is weakly higher if it also replicates her type. Let $a(z_\pi, \mu)$ be the receiver's action after observing realization z_π when he assigns probability μ_t to type t ,

and let $\underline{a}(z_\pi) = \min_{\mu \in \text{conv}(T)} a(z_\pi, \mu)$. Suppose that type t deviates from a pooling equilibrium by offering experiment π' which is a garbling of $\hat{\pi}$. Consider the bidimensional experiment $(\pi', \tilde{\pi}_e)$ which is again a garbling of $\hat{\pi}$. As $(\pi', \tilde{\pi}_e)$ allows the sender to replicate its type, then we must have

$$\underline{a}(z_{\pi'}) \leq \underline{a}(z_{\pi'}, z_{\tilde{\pi}_e}),$$

where $z_{\tilde{\pi}_e} = t$ with probability 1. Therefore, the minimum payoff that a sender can obtain in any equilibrium is achieved when she “certifies” her type. This is formally equivalent to $V_{z_{\pi_e}(t)}$.

We prove that if (5) does not hold, so that for some $\pi_U^* \in \Pi_U^*$ we have

$$V_{z_{\pi_e}(t)} \leq v_{\pi_U^*}^*(t), t \in T, \quad (8)$$

then pooling on π_U^* is an equilibrium of the informed sender game and $V_U = V_I$. Suppose that, after observing any deviation from π_U^* to an alternative experiment π' and observing its realization $z_{\pi'}$, the receiver selects the action $\underline{a}(z_{\pi'})$. Then, the maximum payoff that a type t would obtain from such deviation is $V_{z_{\pi_e}(t)}$. Given (8), no type would profit from deviating, and pooling on π_U^* is an equilibrium. ■

References

- [1] ALONSO, R., AND O. CÂMARA. 2016a. “Bayesian Persuasion with Heterogenous Priors.” *Journal of Economic Theory* 165, 672-706.
- [2] ALONSO, R., AND O. CÂMARA. 2016b. “Persuading Voters.” *American Economic Review* 106(11), 3590-3605.
- [3] ALONSO, R., AND O. CÂMARA. 2016c. “Political Disagreement and Information in Elections.” *Games and Economic Behavior* 100, 390-412.
- [4] BERGEMANN, D., AND S. MORRIS. 2016. “Bayes correlated equilibrium and the comparison of information structures in games.” *Theoretical Economics* 11(2), 487-522.
- [5] BLACKWELL, D. 1953. “Equivalent comparisons of experiments.” *The annals of mathematical statistics* 24(2), 265-272.
- [6] BOLESLAVSKY, R., C. COTTON, AND H. GURNANI. 2017. “Demonstrations and Price Competition in New Product Release.” *Management Science* 63(6), 2016-2026.

- [7] BROCAS, I., AND J. CARRILLO. 2007. "Influence through Ignorance." *Rand Journal of Economics* 38(4), 931-947.
- [8] COVER, T.M., AND J. A. THOMAS. 1991. *Elements of Information Theory*, Wiley Interscience, NY.
- [9] DEGAN, A., AND M. LI. 2016. "Persuasion with Costly Precision," *mimeo*.
- [10] DEGROOT, M.H. 1970. *Optimal Statistical Decisions*, McGraw-Hill.
- [11] GILL, D., AND D. SGROI. 2008. "Sequential Decisions with Tests." *Games and Economic Behavior* 63(2), 663-678.
- [12] GILL, D., AND D. SGROI. 2012. "The Optimal Choice of Pre-Launch Reviewer." *Journal of Economic Theory* 147(3), 1247-1260.
- [13] HEDLUND, J. 2017. "Bayesian persuasion by a privately informed sender." *Journal of Economic Theory* 167, 229-268.
- [14] JOHNSON, J. P., AND D. P. MYATT. 2006. "On the simple economics of advertising, marketing, and product design." *American Economic Review* 96(3), 756-784.
- [15] KAMENICA, E., AND M. GENTZKOW. 2011. "Bayesian Persuasion," *American Economic Review* 101, 2590-2615.
- [16] LI, H. AND LI, W. 2013. "Misinformation." *International Economic Review* 54(1), 253-277.
- [17] MYERSON, R. B. 1983. "Mechanism design by an informed principal." *Econometrica* 51(6), 1767-1797.
- [18] PEREZ-RICHET, E. 2014. "Interim Bayesian Persuasion: First Steps." *American Economic Review: Papers & Proceedings* 104(5), 469-74.
- [19] PEREZ-RICHET, E. 2016. "A Proof of Blackwell's Theorem," *Mimeo*.
- [20] RAYO, L., AND I. SEGAL. 2010. "Optimal Information Disclosure." *Journal of Political Economy* 118(5), 949-987.
- [21] ROSAR, F. 2017. "Test design under voluntary participation," *Games and Economic Behavior* 104, 632-655.
- [22] SHAKED, M., AND J. G. SHANTHIKUMAR. 2007. *Stochastic Orders*, Springer.