

# DISCUSSION PAPER SERIES

DP16326

## **Bargaining over Treatment Choice under Disagreement**

Nabil I. Al-Najjar and Robert Gary-Bobo

**ORGANIZATIONAL ECONOMICS**

**PUBLIC ECONOMICS**

**CEPR**

# Bargaining over Treatment Choice under Disagreement

*Nabil I. Al-Najjar and Robert Gary-Bobo*

Discussion Paper DP16326

Published 05 July 2021

Submitted 28 June 2021

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

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Organizational Economics
- Public Economics

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: Nabil I. Al-Najjar and Robert Gary-Bobo

# Bargaining over Treatment Choice under Disagreement

## Abstract

A group of experts with different prior beliefs must choose a treatment. A dataset is made public and leads to revisions of beliefs. We propose a model where the experts' disagreements are resolved through bargaining, using the Nash bargaining solution. Experts bargain after disclosure of the dataset. Bargaining may lead to an inefficient use of information in a strong sense: experts receive a lower payoff in every state, and for any prior belief (i.e., inadmissibility). Bargaining exhibits under-reaction to information as compared to the normative solution in which experts bargain *ex ante* on the procedure used to exploit the data.

JEL Classification: N/A

Keywords: N/A

Nabil I. Al-Najjar - [al-najjar@northwestern.edu](mailto:al-najjar@northwestern.edu)  
*MEDS, Kellogg Business School, Northwestern University*

Robert Gary-Bobo - [robert.gary-bobo@ensae.fr](mailto:robert.gary-bobo@ensae.fr)  
*University of Paris 1 Panthéon-Sorbonne, Centre d'Economie de la Sorbonne, CREST, ENSAE and CEPR*

# Bargaining over Treatment Choice under Disagreement

Nabil I. Al-Najjar\* Robert J. Gary-Bobo†

revised version 22 June 2021

## Abstract

A group of experts with different prior beliefs must choose a treatment. A dataset is made public and leads to revisions of beliefs. We propose a model where the experts' disagreements are resolved through bargaining, using the Nash bargaining solution. Experts bargain after disclosure of the dataset. Bargaining may lead to an inefficient use of information in a strong sense: experts receive a lower payoff in every state, and for any prior belief (i.e., inadmissibility). Bargaining exhibits under-reaction to information as compared to the normative solution in which experts bargain ex ante on the procedure used to exploit the data.

---

\* Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208. [Al-Najjar's Research Page](#).

†CREST, ENSAE 5, avenue Henry Le Chatelier, 91764 Palaiseau cedex, France, and University of Paris 1 Panthéon-Sorbonne; E-mail:[robert.gary-bobo@ensae.fr](mailto:robert.gary-bobo@ensae.fr)

# 1 Introduction

This paper proposes a bargaining framework to study collective decisions when agents hold different beliefs. This study is motivated by the public policy debates where opposing sides may share similar goals, yet profoundly disagree about the best ways to achieve them. As a concrete example, consider the debate on whether capital punishment deters homicide. Donohue and Wolfers (2009) observe that “[g]iven the availability of relatively high-quality data [...] one would think that a consensus would have emerged about the answer to this ostensibly simple question.” Yet, such consensus is hard to reach because of “the large number of choices that must be made when specifying the various panel data models.” Durlauf, Fu, and Navarro (2013) provide a striking illustration of this: models that differ only in plausible econometric assumptions imply a net number of lives saved per execution that varies from 20 lives to *minus* 60 lives.<sup>1</sup>

Instances of *disagreement*—our shorthand for rational differences in prior beliefs—are, of course, not limited to the criminal justice system. There is no shortage of examples where rational agents face complex policy problems in areas ranging from macroeconomics, to climate change and education, disagree on how to frame these problems and interpret available evidence. One source of such disagreements is the radical uncertainty about key parameters, as witnessed during the 2020-21 coronavirus pandemic. Different researchers with access to identical data can and do reach different conclusions about infection and mortality rates, resulting in strikingly different policy recommendations about the effectiveness of testing, lockdowns, and hospitalization protocols. Ultimately, the source of disagreement is the difficulty of identifying causal mechanisms when—as it is often the case—policy counterfactuals cannot be observed. This difficulty, which is central to modern empirical work, implies that differences in prior beliefs can persist despite

---

<sup>1</sup>A 2012 study by the National Research Council (cited in Nagin (2013)) concludes that “Research to date on the effect of capital punishment on homicide is not informative about whether capital punishment decreases, increases, or has no effect on homicide rates.” More generally, the fact that different sets of plausible assumptions can lead to sharply different conclusions is well-recognized in applied econometrics. For example, Heckman and Vytlacil (2005) write: “Two economists analyzing the same dataset but using different valid instruments will estimate different parameters that have different economic interpretations.”

the accumulation of a large and public body of evidence.

However significant the disagreements may be, a collective decision must still be made—including the option of doing nothing. We propose a model where disagreements are resolved through bargaining. We consider an environment with a group of  $n$  “experts” who share the same utilities but hold different beliefs. The group’s decision is reached by bargaining over *treatments*—equivalent terms commonly used in the literature include *policies* and *programs*. We assume that the outcome of bargaining is determined according to the Nash bargaining solution.

A key measure of the effectiveness of a decision-making process is how well it incorporates new evidence. Adapting collective decisions to the new realities created by clinical trials, empirical studies, or intelligence reports often makes the difference between success and failure. To model the dynamic nature of the experts’ decisions, we consider bargaining over policy *rules* that condition the collective decision on publicly observed information. This richer setting is consistent with a variety of bargaining procedures that differ in the degree of commitment available to the experts.

Consider first an *ex ante bargaining procedure* where experts can commit to an enforceable rule that prescribes how the treatment choice changes in response to new information. Since the Nash bargaining solution is Pareto optimal, the chosen treatment rule is also optimal and must therefore lead to an efficient use of information.

Enforceable *ex ante* commitments require a group of experts to plan for every future signal (*i.e.*, future observations), and to credibly commit to quashing all attempts to renegotiate after a signal is observed. Since this is unrealistic in many contexts, we consider a bargaining procedure where collective decisions are made after new evidence becomes available. This *interim bargaining procedure* is less onerous and, arguably, more natural since experts engage in bargaining only at the signal that is actually observed. No advance planning for all conceivable signals is necessary. We investigate the properties of interim bargaining and show that, with sufficient diversity of beliefs, interim bargaining is generally inefficient in the sense of Pareto. But, even more radically, it can lead to collective decisions that are *inadmissible*, meaning that there exist *ex ante* treatment rules that perform better than the bargaining solution in each state—and, thus, for every possible expert

belief. A compact way to state the source of this problem is that Nash bargaining and Bayesian updating do not commute: bargaining over policy rules followed by updating always leads to Pareto-optimal allocations, while Bayesian updating, followed by bargaining result in Pareto-inefficient, if not inadmissible, collective decisions.

This loss in welfare, which can be severe in magnitude, is the result of an inefficient use of information. In Section 3, we pinpoint the form this inefficiency takes in a setting with two states. In a sense made precise, interim bargaining causes collective decisions to under-react to new public information. Roughly, collective decisions are characterized by *inertia*, because Nash bargaining tends to put more weight on the more skeptical experts whose beliefs change more slowly with the new evidence.

A by-product of our analysis is a simple, yet very useful characterization: the Nash bargaining outcome is chosen as if to maximize the expected utility of a “Nash planner” whose belief is a weighted average of the experts’ beliefs. The weights are determined by the Nash bargaining solution according to a simple and intuitive formula. Roughly, in evaluating a treatment  $t$ , the weight attached to an expert is inversely proportional to this expert’s expected surplus if  $t$  were implemented. This means that the expert who is most pessimistic about  $t$  carries a higher weight in determining whether it is chosen. Collective decisions reached through bargaining therefore tend to be more cautious.

Our interest in differences in prior beliefs has two sets of motivations. First, the assumption that agents share a common prior belief, while entrenched in economic modeling practice, has long been questioned in theoretical work. It is well-known that this assumption has weak foundational standing and leads to paradoxical implications, such as the no-trade theorems.<sup>2</sup>

Our second motivation for focusing on differences in prior beliefs is practical. One compelling argument for assuming a common prior is learning: while experts may start with different beliefs, these differences are washed out as evidence accumulates. In our context, experts eventually end up

---

<sup>2</sup>The common prior assumption has been forcefully advocated by Aumann (1987). It is controversial from a theoretical point of view, as seen, for example, in the arguments presented in Gul (1998) and Aumann (1998). Morris (1995) contains a vigorous defense of using differences in prior beliefs in economic models.

agreeing on the consequences of new treatments as data swamps differences in their priors. This idea stands in contrast with the modern empirical literature, which emphasizes the limitations of observational data, however abundant, in revealing critical policy counterfactuals.<sup>3</sup>

A useful perspective on these issues is that of *partial identification*, an empirical method that aims to limit reliance on strong non-testable assumptions (see Manski (2003)). Rather than relying on such assumptions to point-identify the “true” parameter, partial identification aims to identify bounds on this parameter, using minimal assumptions. Although one would expect rational agents to agree that the true parameter must fall within the bounds derived from the data, it is impossible, by definition, to further narrow down their beliefs without adding non-testable assumptions about which rational experts may disagree. Unless the identification bounds define a single parameter (point identification), complete agreement among rational experts would be little more than an improbable coincidence.

We relate our results to the *Bayesian model averaging* methodology in Statistics (see Brock, Durlauf, and West (2003), Steel (2019)). Nash bargaining provides a legitimate way of finding a compromise between competing models, assigning weights to different models. To the best of our knowledge, our approach has not been proposed or studied before, except in the work of Weerahandi and Zidek (1983). In this work, a group of Bayesian experts, endowed with different utilities and beliefs must select a decision rule. The authors show that the Nash solution has desirable properties in this context, mainly because of invariance with respect to linear-affine transformations of utilities<sup>4</sup>, but they did not prove or conjecture the type of results that we present below. Our approach builds on the intuition that *ex ante* Nash bargaining is likely to be the simplest and most compelling way of finding a compromise between rational agents endowed with different beliefs.

---

<sup>3</sup>For example, according to Milton and Rose Friedman: “Legalize drugs, and street crime would drop dramatically and immediately” Friedman and Friedman (1984). In the debate on whether to legalize recreational drugs, little is known about the potential impact of legalization of drugs on demand because all that is available is data collected under the current policy regime. Such data, however massive, would do little to compel rational agents to hold the same beliefs about counter-factual outcomes.

<sup>4</sup>Weerahandi and Zidek also consider the case in which a random subsample of experts bargains to choose the decision, noting that decisions are invariant under Nash bargaining, because sampling probabilities factor out.



The bargaining model synthesizes the main elements of a collective-decision problem: the different beliefs and utilities, the technological constraints, the individual rationality constraints, and pins down the weights used to compute an average of the agents' beliefs.

In the following, Section 2 presents the model and shows that it applies to collective decision-making. In Section 3, we focus on Nash bargaining; we contrast ex ante and interim bargaining situations, then present results on inertia and inadmissibility; we also show that bargaining solutions are supported by a particular weighted average of the experts' beliefs. We finally discuss applications to statistical decision problems. In Section 4, subsection 4.1 studies a class of examples called *hard choices*, in which the diversity of beliefs allows for improvements through ex ante negotiation. Subsection 4.2 states the result that interim-bargaining is always inefficient under the condition that beliefs are sufficiently dispersed. Proofs are gathered in the Appendix.

## 2 Treatment Choice under Disagreement

### 2.1 The Model

A set of  $n$  experts, indexed  $i = 1, \dots, n$ , with different beliefs, must select a *treatment*  $t$  that results in an uncertain outcome  $z \in Z$ , where  $Z$  is the set of possible outcomes.

*Uncertainty and Treatments.* We model uncertainty as a product  $\Theta \times S$  where  $\Theta$  is an index set of (unobservable) *states* of nature, or *theories*, and  $S$  is a set of publicly observed data (or *signals*). A *treatment* is any function

$$t : \Theta \times S \rightarrow \Delta(Z)$$

where  $\Delta(Z)$  denotes the set of probability distributions on  $Z$ . To fix ideas, we assume that  $Z$  is a closed and bounded subset of  $\mathbb{R}^H$ , so that outcomes may be multidimensional. Assume, for expository simplicity, that  $\Theta$  and  $S$  are finite. Let  $K$  be the number of states in  $\Theta$  and  $L$  be the number of signals in  $S$ . Then, a treatment  $t$  can be viewed as list of  $LK$  probability distributions  $t(\theta, s)$ , and  $t(\cdot, s)$  denotes a sub-vector of  $t$ , giving the probability distributions on outcomes for a given  $s$ . A treatment that is

not signal-contingent is denoted  $a = (a(\theta_1), \dots, a(\theta_K))$ , and called a basic treatment. Let  $A$  denote the feasible set of basic treatments, which we assume to be convex. Thus, feasibility requires  $t(\cdot, s) \in A$ . The set of treatments, represented as  $T = A^{|S|}$  has a natural convex structure by averaging distributions.

*Information and Likelihood.* The signal  $s$  can be any information, such as a data set in an empirical study, historical time series, or the results of an experiment. The distribution of signals conditional on the state is given by a *likelihood function*  $q(s | \theta)$ . The likelihood  $q$  is common to all experts. This is an innocuous assumption which can be satisfied without loss of generality by expanding the set of states  $\Theta$ . Differences of opinions then amount to different beliefs about the true state  $\theta$ .

*Beliefs and Disagreement.* Expert  $i$ 's model of the consequences of a treatment is represented by a probability distribution  $p_i$  on  $\Theta \times S$ , with a likelihood function  $q$ . We use  $p_i(\theta)$  and  $p_i(s)$  to denote the marginals of  $p_i$  on  $\Theta$  and  $S$ , respectively. We shall assume that  $p_i(s) > 0$  for every  $i$  and  $s$ , so beliefs can always be updated once a signal is observed.

To focus on the role of disagreement—rather than asymmetric information—we assume that the belief profile  $(p_1, \dots, p_n)$  is commonly known. By Aumann's theorem (cf. Aumann (1976)), these beliefs are not the posteriors derived from updating a common prior based on (possibly private) information. Experts therefore “agree to disagree.” Using the language of applied econometrics, even when all available data are made public, the experts' models are not fully identified. It follows that experts may legitimately use different *identifying assumptions* that cannot be proved to be false.

## 2.2 Utility

Assume a common a cardinal utility  $v$  on  $Z$ . Under this assumption, given a true state  $\theta$ , we can integrate out the outcome and the signal to obtain the *state-utility*:

$$u(t)(\theta) = \sum_{s \in S} q(s | \theta) E_{t(\theta, s)} v(z),$$

where  $E_{t(\theta,s)}v(z)$  is expected utility with respect to the distribution on outcomes  $t(\theta, s)$ , induced by the chosen treatment  $t$ . Note that  $u(t)$  doesn't depend on  $i$ 's beliefs. Now, expert  $i$ 's expected utility is by definition,

$$E_{p_i}u(t) = \sum_{\theta \in \Theta} p_i(\theta)u(t)(\theta).$$

Our bargaining model can be extended to the case of different utilities. This, however, will obscure the role of differences in beliefs—our main focus in this paper. Common utility, on the other hand, makes it possible to cleanly separate the types of uncertainty in the model. First, we have uncertainty about the outcomes and the samples, conditional on the true state. The state-utilities capture the impact of this uncertainty, which is not subject to disagreement. Second, the experts' prior beliefs  $p_i$  represent their subjective models of the likelihood of different states.

### 3 Treatment Choice as a Bargaining Problem

Like any conflict among economic agents, interactions between experts with different beliefs should be modeled as a game. The main idea of this paper is to think of the experts' collective decision as a bargaining process whose outcome is determined according to the Nash bargaining solution (cf. Nash (1950)). We comment on the interpretation of bargaining as a mechanism to resolve disagreements after a formal statement of the problem.

#### 3.1 Ex Ante Nash Bargaining

Recall the standard definition of a bargaining problem with  $n$  players as: (1) a convex set  $U \subset \mathbb{R}^n$  of feasible expected utility vectors, and (2) the status quo  $u^\circ \in U$  indicating the expected utilities if no agreement is reached. The Nash bargaining solution selects the unique vector  $u^* \in U$  that maximizes the Nash product  $\prod_{i=1}^n (u_i - u_i^\circ)$  over  $U$ .

We are interested in bargaining problems over treatment choice. A natural starting point is *ex ante bargaining* where experts bargain before any new information is received. We formally define this problem as:

$$\mathcal{B}^* = \left( T, a^\circ, (E_{p_i}u(t))_{i \in N} \right),$$

where  $T$  is a compact and convex set of feasible treatments,  $a^\circ \in T$  is the *status quo* treatment, and  $N = \{1, \dots, n\}$  is the set of players.<sup>5</sup>

Let  $t^*$  denote any treatment that produces the (necessarily unique) Nash bargaining solution vector of expected utilities for this problem. Although  $t^*$  is not in general unique,<sup>6</sup> any other such treatment must generate the same expected utility for each expert. For this reason we simplify the exposition by referring to  $t^*$  as *the* Nash bargaining solution.

To simplify notation, we define the *surplus* generated by treatment  $t$  to be the vector  $\delta(t)$  in  $\mathbb{R}^K$ , with components

$$\delta(t)(\theta) \equiv u(t)(\theta) - u(a^\circ)(\theta)$$

indicating, for each state  $\theta$ , the change in expected utility under  $t$  relative to the status quo  $a^\circ$ . We will also use the notation,  $\delta(t)(\theta, s) = E_{t(\theta, s)}v(z) - u(a^\circ)(\theta)$  to denote the surplus at point  $(\theta, s)$ .

### 3.2 Interim Nash Bargaining

Ex ante bargaining assumes that experts are bound to implement a data-contingent policy rule that is agreed to in advance. We view ex ante bargaining as an ideal, normative solution. There are certainly contexts where such commitments are reasonable. For instance, institutions, such as federal agencies and research journals, increasingly enforce *pre-analysis plans* in experimental and observational studies. See Christensen and Miguel (2018) for discussion and surveys of the use of pre-analysis plans in economics, medical science, and other fields.

In other contexts, such pre-commitments can be daunting, if not outright impossible. Researchers and policy makers may find it unrealistic to plan for all conceivable, yet to be observed, data sets. This leads us to consider treatment choices when bargaining takes place at the interim stage, after the information  $s$  has been revealed but before the consequence is realized.

---

<sup>5</sup>To map this into a bargaining problem over expected utilities, define  $U$  to be the set of vectors of expected utilities  $(E_{p_1}u(t), \dots, E_{p_n}u(t))$  obtained from each  $t \in T$  and  $u^\circ = (E_{p_1}u(a^\circ), \dots, E_{p_n}u(a^\circ))$ .

<sup>6</sup>For example, if there are more states than experts then there will necessarily be a non-zero linear subspace of state-utility vectors that generate identical expected utilities for all experts.

Define the *interim bargaining problem* as:

$$\mathcal{B}(s) = \left( A, a^\circ, (E_{p_i(\cdot|s)}u(t))_{i \in N} \right).$$

The main difference with ex ante bargaining is that experts evaluate treatments using their updated beliefs  $p_i(\cdot|s)$ . Although the outcome of a treatment at signals other than  $s$  is irrelevant, this redundancy will be convenient below.

The interim Nash bargaining solution  $\hat{t}$  is the treatment  $\hat{t} = (\hat{t}(\cdot, s))_{s \in S}$  belonging to  $A^{|S|}$  such that, for every signal  $s$ ,  $\hat{t}(\cdot, s)$  maximizes the Nash product:

$$\prod_{i=1}^n [E_{p_i(\cdot|s)}u(t) - E_{p_i(\cdot|s)}u(a^\circ)].$$

As before, we will refer to  $\hat{t}$  as *the* interim Nash bargaining solution since any two solutions must generate the same expected utility for each expert at each signal  $s$ .

We now identify a range of plausible conditions where bargaining and disagreement distort the way a group of experts uses information, and lead to overly conservative decisions that under-react to new evidence (*i.e.*, *inertia*).

### 3.3 Bargaining and the Efficient Use of Information: Inertia

To make progress and discover the structure of interim solutions, we assume here that there are two states, denoted  $\theta_0$  and  $\theta_1$ ,  $n$  experts,  $L$  signals, and a regularity condition on the set of feasible state-utility vectors  $u(A)$ . We discuss below the extent to which these assumptions may be relaxed.

The likelihood ratio of state  $\theta_1$  to state  $\theta_0$  at signal  $s$ ,

$$\ell(s) = \frac{q(s|\theta_1)}{q(s|\theta_0)}$$

defines a linear ordering on signals. A higher value of  $\ell(s)$  indicates greater support for state  $\theta_1$ . To simplify notation, for a signal  $s$  define:

$$\delta_{\theta_k}^*(s) \equiv \delta(t^*(\theta_k, s)) \quad \text{and} \quad \hat{\delta}_{\theta_k}(s) \equiv \delta(\hat{t}(\theta_k, s))$$

for the surplus in state  $\theta_k$  under the ex ante and interim solutions, respectively.

We can state the following result.

**Proposition 1.** *Suppose that:*

1. *Beliefs disagree:  $p_i \neq p_j$  for some  $i$  and  $j$ ;*
2. *The experiment  $(S, q)$  is informative in the sense that there exists a signal  $s$  such that  $q(s|\theta_0) \neq q(s|\theta_1)$  for two different states  $\theta_0, \theta_1$ ;*
3. *The set  $u(A)$  is strictly convex with differentiable boundary;*
4. *There are two states.*

*Then, there exists a positive real number  $\bar{\ell}$  such that, for any signal  $s$ ,*

$$\ell(s) > \bar{\ell} \implies \delta_{\theta_1}^*(s) > \hat{\delta}_{\theta_1}(s),$$

*and*

$$\ell(s) < \bar{\ell} \implies \delta_{\theta_0}^*(s) > \hat{\delta}_{\theta_0}(s).$$

To illustrate the result, assume that there is a signal  $\bar{s}$  such that  $\ell(\bar{s}) = \bar{\ell}$  and consider a signal  $s$  with  $\ell(s) > \ell(\bar{s})$  (the case for  $s$  with  $\ell(s) < \ell(\bar{s})$  is similar). Since  $s$  provides stronger evidence in favor of  $\theta_1$  compared to  $\bar{s}$ , the ex ante planner, who now believes  $\theta_1$  is more likely, will select a treatment  $t^*(\cdot, s)$  that does better than  $t^*(\cdot, \bar{s})$  in state  $\theta_1$ .<sup>7</sup> That is, we must have:

$$\delta_{\theta_1}^*(s) > \delta_{\theta_1}^*(\bar{s}).$$

Appendix A shows that this monotonicity in the likelihood ratio also holds for the interim solution  $\hat{t}(\cdot, s)$ . Observing an  $s$  with  $\ell(s) > \ell(\bar{s})$  causes both the ex ante and interim solutions to adjust the treatment choice towards one that does better in state  $\theta_1$  and (necessarily) worse in state  $\theta_0$ . The question is: how strong is this adjustment under the two solutions? The key assertion in the proposition,  $\delta_{\theta_1}^*(s) > \hat{\delta}_{\theta_1}(s)$ , states that the adjustment is smaller under interim bargaining than ex ante bargaining. That is, the interim solution displays inertia with respect to the information contained in  $s$ .

Inertia reflects the way information is processed under the two bargaining procedures. Both procedures update the experts' beliefs using Bayes rule.

<sup>7</sup>And does necessarily worse in state  $\theta_0$ , since  $t^*(\cdot, s)$  is on the frontier of  $A$ .

For example, when  $\ell(s)$  increases, all of the experts' posteriors put more weight on  $\theta_1$ . The difference is in the way the two procedures aggregate the posteriors. A careful analysis of interim bargaining reveals that interim bargaining places more weight on the “hold-outs,” *i.e.*, the experts with the least favorable posterior about state  $\theta_1$ .

### 3.4 Inadmissibility of Interim Bargaining

Recall that under interim bargaining, each signal gives rise to a distinct bargaining problem  $\mathcal{B}(s)$  whose solution varies in subtle ways with  $s$ .<sup>8</sup> An interesting by-product of the proof of Proposition 1 is that interim bargaining can result, not only in Pareto-inefficient choices, but, more radically, in *inadmissible* treatment rules, *i.e.*, rules that can be improved on in every state. This is always true in the two-states case.

**Proposition 2.** *Under the assumptions of Proposition 1, the interim Nash bargaining solution  $\hat{t}$  is inadmissible.*

The proofs of Propositions 1 and 2 are complicated by the fact that they require a careful analysis of how the interim bargaining solution,  $\hat{t}(\theta, s)$ , changes as beliefs are updated. We can, however, provide a rough informal argument which, despite many gaps, may help the reader develop an intuition for the result. To this end, it will be useful to state another result, which plays the role of a Lemma in the proof of the above propositions, and we will then return to our discussion of their meaning. Finally, Section 5 below provides sufficient conditions under which the interim bargaining solution is always Pareto-inefficient in the  $K$ -states, 2 experts case.

### 3.5 The Nash Planner

The bargaining solution can be described as the decision of a fictitious planner (*i.e.*, the “Nash planner”) whose belief is a particular weighted average of the expert’s beliefs. We can now state the following Proposition.

---

<sup>8</sup>This difficulty disappears when beliefs are dogmatic, because the experts’ posterior beliefs are always equal to their priors.

**Proposition 3.** *For the bargaining problem  $\mathcal{B}^*$ , with Nash bargaining solution  $t^*$ , there is a probability distribution  $\pi^*$  over  $\Theta$  such that:*

$$t^* \in \operatorname{argmax}_{t \in T} E_{\pi^*} \delta(t). \quad (1)$$

The distribution  $\pi^*$  is the unique weighted average of the experts' beliefs given by:

$$\pi^*(\theta) = \sum_{i=1}^n m_i^* p_i(\theta) \quad (2)$$

where

$$m_i^* = \frac{c}{E_{p_i} \delta(t^*)}$$

and  $c$  is the normalization constant defined by  $\sum_i m_i^* = 1$ .

Proposition 3 applies to the interim bargaining problem as well as the ex ante bargaining problem. The statement remains true if we replace  $t^*$  and  $\pi^*$  with  $\hat{t}$  and  $\hat{\pi}$ ,  $\mathcal{B}^*$  with  $\mathcal{B}(s)$  and beliefs with updated beliefs.

In the statement of Proposition 3, the probability distribution  $\pi^*$  may be interpreted as a compromise between the beliefs of the  $n$  experts. The fact that the weight  $m_i^*$  is inversely proportional to expert  $i$ 's expected surplus reflects a fundamental property of Nash bargaining as a mechanism to mediate differences in beliefs: experts who are more pessimistic about a treatment  $t$ , in the sense of a small expected surplus  $E_{p_i} \delta(t)$ , tend to carry a higher weight in determining whether  $t$  is chosen.

Once the signal  $s$  is observed, Bayes' rule requires that  $t(\cdot, s)$  maximizes expected utility with respect to the updated compromise belief:<sup>9</sup>

$$\pi^*(\theta | s) = \sum_{i=1}^n m_i^*(s) p_i(\theta | s) \quad (3)$$

where

$$m_i^*(s) = c(s) \frac{p_i(s)}{E_{p_i} \delta(t^*)} \quad (4)$$

---

<sup>9</sup>Proposition 3 tells us that  $t^*$  maximizes  $\sum_s \pi^*(s) \sum_{\theta} \pi^*(\theta | s) \delta(t)(s, \theta)$ , where  $\pi^*(s) = \sum_{\theta} q(s | \theta) \pi^*(\theta)$ . It follows that each  $t^*(\cdot, s) \in A$  must maximize  $\sum_{\theta} \pi^*(\theta | s) \delta(t)(s, \theta)$ . The rest follows easily.



and  $c(s)$  is the normalization constant defined by  $\sum_i m_i^*(s) = 1$ . It is convenient to eliminate the normalization constant by writing:

$$\frac{m_i^*(s)}{m_j^*(s)} = \underbrace{\frac{p_i(s)}{p_j(s)}}_{\text{Bayes Factor}} \times \frac{m_i^*}{m_j^*}.$$

In the statistics literature, the *Bayes factor* measures the support that new evidence lends to one model over another (cf. Wasserman (2000)). It determines how the prior odds ratio  $m_i^*/m_j^*$  of the two models changes as a result of observing the data: when the observed sample  $s$  is more likely under  $p_i$  than  $p_j$ , this is interpreted as evidence supporting  $i$ 's model, boosting its weight.

The interim bargaining process will not lead to the same compromise over updated probabilities that the Nash planner would use to implement the solution of the ex ante bargaining problem.

### 3.6 Discussion and Intuition for the Results

The Complete Class Theorem (see Appendix B) states that for the interim solution to be admissible, it must maximize expected utility with respect to some Bayesian belief, say  $\pi^{**}$ . In addition, Bayes' rule implies that the updated bargaining weights for the ex ante planner must satisfy:

$$\frac{\pi^{**}(\theta_1 | s)}{\pi^{**}(\theta_0 | s)} \propto \ell(s).$$

When the feasible set is strictly convex with differentiable boundary, then, at each signal  $s$ , the updated belief  $\pi^{**}(\theta | s)$  must coincide with the interim planner's belief (that can be expressed as a weighted average of the experts' beliefs, according to Proposition 3), that is,

$$\hat{\pi}(\theta | s) = \sum_i \hat{m}_i(s) p_i(\theta | s).$$

This, in turn, implies:<sup>10</sup>

$$\frac{\hat{\pi}(\theta_1 | s)}{\hat{\pi}(\theta_0 | s)} \propto \ell(s).$$

<sup>10</sup>The factor of proportionality is, in fact,  $\pi^{**}(\theta_1)/\pi^{**}(\theta_0)$ , a constant that does not depend on  $s$ .

A careful analysis of the interim bargaining problems shows that such a proportionality is in contradiction with the characterization of the interim compromise beliefs  $\hat{\pi}(\theta|s)$ . Intuitively, admissibility requires the interim bargaining weights to vary only with the informational content of the signal, measured by the likelihood ratio  $\ell(s)$ , while the expression of  $\hat{\pi}(\theta|s)$  shows that it must also depend on the experts' expected utilities, through the weights  $\hat{m}_i(s)$ .

To see the intuition, recall from the discussion of the Bayes factor (in Section 3.5) that ex ante bargaining selects a treatment  $t^*(s)$  based on assigning experts weights that vary proportionally with the relative accuracy of their predictions, measured by  $p_i(s)/p_j(s)$ . Contrast this with interim bargaining, where the *interim bargaining weight* of expert  $i$  relative to expert  $j$  at  $\hat{t}(\cdot, s)$  is:

$$\frac{1}{E_{p_i(\cdot|s)}\hat{\delta}(s)} \bigg/ \frac{1}{E_{p_j(\cdot|s)}\hat{\delta}(s)}.$$

Here, the relative accuracy of the predictions,  $p_i(s)/p_j(s)$ , plays no role in how the updated beliefs are weighted. Instead, the Nash bargaining solution now dictates that weights are influenced by the experts' conditional expected utilities. The injection of such non-informational consideration tends to give greater weight to the expert with the lowest expected utility conditional on  $s$ .

Next, we review the assumptions of Propositions 1 and 2. Assumption 1, stating that experts disagree, is clearly needed: under complete agreement, interim Nash bargaining is Bayesian, and therefore admissible (by the Complete Class Theorem). In addition, if the signals were perfectly informative, then all non-dogmatic experts would agree in the interim stage, and the outcome, again, would be Bayesian. The strict convexity and differentiability assumptions cannot be removed entirely, for somewhat more subtle technical reasons.<sup>11</sup>

Finally, the assumption that there are only two states is used in the proof to ensure that the likelihood ratio  $\ell(s)$  introduces a linear ordering on signals, so that a higher value of  $\ell(s)$  indicates greater support for state  $\theta_1$ . We suspect that the result would hold more generally under monotonicity assumptions about the signal structure.

---

<sup>11</sup>See a detailed discussion in the Appendix.

To sum up, with ex ante commitments, ex ante expected utilities determine the ex ante relative bargaining positions of the experts. Once these are determined, the response to new information depends only on the informational content of the observed signal, as seen in Section 3.5 above. When experts have no access to such commitments, their interim bargaining positions are calculated based on their interim expected utilities and not on the informational content of the signal alone.

Groups and organizations that lack the ability to enforce ex ante agreements will fall into a potentially inferior interim solution. In special cases, when beliefs are dogmatic, the result is complete paralysis: the status quo is sustained despite overwhelming evidence favoring change. Section 4.1 studies a class of examples, called “hard choices”, in which these problems appear. Our model provides a new perspective on the problem of commitment in organizations, namely the role of commitment in efficiently resolving differences in beliefs.

### 3.7 Regression, Prediction, and Treatment Assignments

The abstract setting above applies to collective public-policy decisions as well as a statistical decision problems. We illustrate this connection in the context of generalized regression.

*Basic Set-up.* We are given a real-valued policy-relevant consequence  $y \in Y$  and an  $l$ -dimensional vector of covariates (regressors)  $\mathbf{x} \in \mathcal{X}$ . To maintain consistency with our earlier assumptions, assume that  $\mathcal{X}$  is a finite set. The  $\mathbf{x}$ 's can, for instance, be indicator (*i.e.*, dummy) variables.

The likelihood function  $q(s | \theta)$  represents i.i.d. samples of  $m$  observations  $s \in (Y \times \mathcal{X})^m$ . The distribution of the outcome conditional on the state and a vector of covariates is assumed to follow the generalized regression equation:

$$y = F_\theta(\mathbf{x}) + \epsilon,$$

where  $F_\theta$  is a regression function and  $\epsilon$  is a random error term. For expository simplicity, the distribution of  $\epsilon$  is fully specified if  $\theta$  is given (*i.e.*, it is parameterized by  $\theta$ ). For instance, and to fix ideas, the error can be a mixture of normal distributions.

The state  $\theta$  therefore encodes all information about how samples are

generated, including the marginal distribution on covariates, possible correlations between covariates and the error term, causal interpretations of coefficients, etc. The regression function  $F_\theta$  may be linear, a more complex non-linear classifier, or a neural network. In a linear regression,  $F_\theta$  corresponds to a vector of regression parameters  $\beta$  that encodes assumptions about model specification such as exclusion restrictions.

*Prediction Problems.* Consider first the problem where experts must formulate conditional predictions of  $y$ . To this end, they choose an estimator  $\hat{f}(\mathbf{x}, s)$  of the expected value of the consequence  $y$ , given an observed vector of covariates  $\mathbf{x}$ . Any such estimator  $\hat{f} : \mathcal{X} \times \mathcal{S} \rightarrow Y$  defines a probability distribution of the prediction error  $y - \hat{f}(\mathbf{x}, s)$  given  $\theta$  and  $s$ . We interpret this prediction error as the outcome, *i.e.*,  $z = y - \hat{f}(\mathbf{x}, s)$  and its distribution for a given sample as  $t(\theta, s)$ .

The common utility assumption in this case is quite natural: experts (econometricians) equally care about reducing out-of-sample prediction errors, all the while disagreeing about the procedure to accomplish that. If, as common in practice, estimators are evaluated based on their mean square errors, then the loss function is  $z^2$  and the common utility function is  $v(z) = -z^2$ . The state-utility is then, for each model or parameter  $\theta$ ,

$$u(t)(\theta) = - \sum_{s \in \mathcal{S}} q(s|\theta) E_{t(\theta,s)} (y - \hat{f}(\mathbf{x}; s))^2.$$

*Treatment Assignment.* Next, we consider the problem of assigning a binary treatment  $D \in \{0, 1\}$  to individual cases characterized by a covariate vector  $\mathbf{x}$ . Here, samples are  $m$  observations  $s \in (\{0, 1\} \times Y \times \mathcal{X})^m$  and the regression function is  $F_\theta(D, \mathbf{x})$ , indicating the dependence of  $y$  on the treatment as well as the covariate vector. The likelihood function  $q(s | \theta)$  then represents, in addition, how treatments were assigned in-sample.<sup>12</sup>

A treatment allocation rule in this case selects  $\mathbf{d}(\mathbf{x}, s) \in \{0, 1\}$  for an out-of-sample observation  $\mathbf{x}$ , given  $s$ . This allocation rule may be constant (treat every new patient with the new drug), randomized (treat every patient with

---

<sup>12</sup>For example, in an experiment with random assignment, potential outcomes  $y_1 = F_\theta(1, \mathbf{x}) + \epsilon$ ,  $y_0 = F_\theta(0, \mathbf{x}) + \epsilon$  and treatment  $D$  would be assumed conditionally independent, given  $\mathbf{x}$ , as usual in the treatment-effects literature (see *e.g.*, Heckman and Smith (1998), Angrist and Pischke (2009)).

probability 0.5), or based on maximizing expected utility. In the latter case, the outcome is  $z = (y(\mathbf{x}))_{\mathbf{x} \in \mathcal{X}}$  where  $y(\mathbf{x}) = F_\theta(\mathbf{d}(\mathbf{x}, s), \mathbf{x}) + \epsilon$ . The utility  $v$  is some welfare function, such as a weighted average of the outcomes  $v(z) = \sum_{\mathbf{x}} \alpha(\mathbf{x})y(\mathbf{x})$ , with weights  $\alpha(\mathbf{x}) \geq 0$ .<sup>13</sup> The treatment allocation rule  $\mathbf{d}$  generates a distribution of the outcome  $y(\mathbf{x})$ , denoted  $t(\theta, s; \mathbf{d}, \mathbf{x})$ . The state-utility can now be defined as follows:

$$u(t(\cdot, \mathbf{d}))(\theta) = \sum_s q(s|\theta) \sum_{\mathbf{x}} \alpha(\mathbf{x}) E_{t(\theta, s; \mathbf{d}, \mathbf{x})} y(\mathbf{x}).$$

A single Bayesian expert with prior  $p$  over states  $\theta$  would choose  $\mathbf{d}$  to maximize  $E_p u(t(\cdot; \mathbf{d}))$ . The group of experts may be constrained to choose the rule  $\mathbf{d}$  in a certain class, for instance, in the class of rules  $d(\mathbf{x}, s) \in \{0, 1\}$  such that  $d(\mathbf{x}, s) = 1$  if and only if some estimate of the *average treatment effect* satisfies

$$ATE(\mathbf{x}; s) = \hat{f}(1, \mathbf{x}; s) - \hat{f}(0, \mathbf{x}; s) \geq c(s),$$

where  $c(s)$  is a cutoff in outcome space and  $\hat{f}(D, \mathbf{x}; s)$  is an estimator of the conditional expectation of  $y$  given  $\mathbf{x}$ . We would then study which estimator  $\hat{f}$  (and which cutoff  $c$ ) would be the result of bargaining among experts.

### 3.8 Bargaining and Bayesian Model Averaging

If we now discuss our approach in terms of Statistical Decision Theory, we see that bargaining provides a solution to the problem of model selection when there is disagreement about the correct model. In statistics—and increasingly in economics—an important methodology to deal with model uncertainty is *Bayesian model averaging*. This methodology starts with a space of possible models (different specifications, variable selections, functional forms, exclusion restrictions, etc.) and assigns to each model  $p$  in that space a weight  $m_p$  that measures the belief that the true model is  $p$ . Prediction, estimation, and policy analysis can then be conducted using the weighted average of the models.

Brock, Durlauf, and West (2003) develop persuasive arguments for using Bayesian model averaging in policy evaluation when the structure of the

---

<sup>13</sup>Manski (2004) proposes a related framework, using Minimax regret as a decision criterion.

economic environment is not known. Steel (2019) provides a survey of the vast number of applications in econometric practice. See Hoeting, Madigan, Raftery, and Volinsky (1999) and Wasserman (2000) for motivation and exposition of the underlying statistics.

Bargaining provides a better economic motivation for model averaging in treatment choice. Under bargaining, the set of models corresponds to the economic agents actually competing to achieve policy objectives. The weights derived in Proposition 3 are determined by bargaining and take into account the economically relevant information about the decision problem they face: beliefs, payoffs, feasibility, and the status quo.

In the statistical model averaging literature, weights are assigned to models based on statistical considerations such as diffuseness of priors or penalties for model complexity and overfitting. Since policy evaluation is made by the economic agents involved, diffuseness or controlling overfitting are not desirable per se, but appear only in so far as these agents believe in them or through the bargaining process itself.

Another difficulty in using statistical model averaging techniques is that they ignore payoffs. Downside risk, potential rewards, or concern about equity are obviously crucial in treatment choice. The model weights in Nash bargaining are determined not just as a function of the model space (the belief profile)  $\{p_1, \dots, p_n\}$  but also by the utility function.

While the Nash bargaining solution is extensively used in economics, our use of this concept differs from its traditional applications to distributive conflicts (*i.e.*, cake-division problems). A natural question, then, is: how appropriate is Nash bargaining as a model of resolving conflicts in beliefs?

The Nash bargaining solution may be viewed as a description of the equilibrium outcome of some underlying non-cooperative game. The literature on this topic, known as the *Nash Program*, is too large to be adequately reviewed here. See Serrano (2005) for a survey. For example, Binmore, Rubinstein, and Wolinsky (1986) show that the Nash bargaining solution emerges from an alternating offer bargaining game when the frequency of offers increases. Taking this perspective, the axioms defining the Nash bargaining solution provide a compact representation of what to expect in actual bargaining among experts.

A key feature of bargaining, compared to individual decisions, is the

pivotal role the status quo plays in treatment choice. This reflects well how real-world legislative, regulatory, and legal decisions are made (cf. Manski (2013)). New medical drugs and procedures are not approved unless they undergo stringent tests that compare them to existing ones. New legislations and regulations are proposed as alternatives to a status quo that is reverted to if no agreement is reached. The Nash bargaining solution captures this.<sup>14</sup> We return to this point in greater depth in the next section.

## 4 Hard Choices and Inefficiency

### 4.1 Hard Choices : An Example

To illustrate the impact of interim bargaining, we begin with a symmetric example. There are two states, two experts, and two treatments. Signals, payoffs, and beliefs are symmetric. The payoffs are given in the table:

$$\begin{array}{c|ccc} & a^\circ & a_1 & a_2 \\ \hline \theta_1 & 0 & \alpha & -\beta \\ \theta_2 & 0 & -\beta & \alpha \end{array} \quad \text{with } \beta > \alpha > 0.$$

The idea is that each expert can identify a treatment  $a_i$  that improves on the status quo but not by enough to out-weigh the harm caused by implementing the treatment favored by the other. The key characteristic of this environment is the potential absence of a mutually agreed way to improve on the status quo. For a concrete example, suppose that a choice has to be made whether to move from the status quo to one of two new medical procedures. There is significant uncertainty about the performance of both procedures: while  $a_i$  improves welfare by  $\alpha$  in state  $\theta_i$ , it causes greater harm  $\beta$  if the state is  $\theta_j$ .

Next consider sampling  $m$  observations from the set  $S^\circ = \{s_1, s_2\}$ .

Define  $q(s | \theta)$  as follows.

Let  $q^m$  denote the probability distribution of i.i.d. samples of  $m$  observations of  $s$ .

---

<sup>14</sup>As shown in Roth (1977), the Pareto-optimality axiom can be replaced by individual rationality, which grants each expert a veto power over treatments deemed inferior to the status quo.

	$s_1$	$s_2$	
$\theta_1$	$q$	$1 - q$	$q > 0.5.$
$\theta_2$	$1 - q$	$q$	

The two experts have symmetric beliefs: for  $i \neq j$ ,  $p_i(\theta_i) = p$  and  $p_i(\theta_j) = 1 - p$ , for some  $1 > p > \frac{\beta}{\alpha + \beta}$ . This restricts attention to the interesting case where both experts believe the status quo can be strictly improved on, but cannot agree on which alternative should be used to improve on it.

Consider first the ex ante bargaining procedure of the last section. The symmetry of the problem implies that the Nash bargaining solution must generate the same (ex ante) expected payoff to the two experts. In particular, the ex ante Nash bargaining implies equal weights on the two experts. Since beliefs are symmetric,  $\pi^*$  must satisfy  $\pi^*(\theta_1) = \pi^*(\theta_2)$ .

Bayes rule implies that, under ex ante Nash bargaining, the collective decision is determined by the likelihood ratio:

$$l(s) = \frac{\pi^*(\theta_1 | s)}{\pi^*(\theta_2 | s)} = \frac{q^m(s | \theta_1)}{q^m(s | \theta_2)}.$$

Given that the expected payoff is  $\alpha \pi^*(\theta_1 | s) - \beta \pi^*(\theta_2 | s)$  under  $\pi^*$  at sample  $s$  and treatment  $a_1$ , the bargaining outcome is (setting aside indifferences):

- ◇  $a_1$  if  $l(s) > \frac{\beta}{\alpha}$ ;
- ◇  $a_2$  if  $l(s) < \frac{\alpha}{\beta}$ ; and
- ◇  $a^\circ$  if  $l(s) \in (\frac{\alpha}{\beta}, \frac{\beta}{\alpha})$ .

To focus on the interesting case, assume that  $l(s) \notin (\frac{\alpha}{\beta}, \frac{\beta}{\alpha})$  with positive probability. This rules out that the status quo is always selected, regardless of the sample (this is a joint condition on the states,  $q$  and  $m$ ).

We turn next to interim bargaining: the two experts observe the sample  $s$ , then bargain based on their updated beliefs. If  $1 > p > \frac{\beta}{\alpha + \beta}$ , expert  $i$  ex ante believes treatment  $a_i$  is superior to the status quo and  $a_j$ , but would be willing to change his opinion if the evidence strongly suggests otherwise.

In interim bargaining, for a treatment to be selected, it has to satisfy the individual rationality constraint of both experts. So, for expert 2, say,



to agree to  $a_1$ , it is necessary that  $l(s) > (\beta/\alpha)(p/1 - p)$ .<sup>15</sup> Under ex ante bargaining, by contrast,  $a_1$  is chosen provided only that  $l(s) > \beta/\alpha$ .

To implement a treatment other than the status quo, stronger evidence in favor of that treatment is required under interim bargaining as compared to the ex ante case. The sharper the disagreement between the experts ( $p$  close to 1), the higher is  $p/(1 - p)$  and the stronger the evidence would need to be to move away from the status quo. In the extreme case where both experts are dogmatic, so  $p = 1$ , we have  $p/(1 - p) = \infty$  and experts are stuck at the status quo regardless of the strength of the evidence in favor of moving away.

Both experts are worse off under interim bargaining, as long as there is positive probability that the sample evidence is strong enough to warrant a move away from the status quo under  $\pi^*$  but not strong enough to overcome the veto power of each expert.<sup>16</sup> This inefficiency echoes long-standing ideas of the public policy literature about the status quo bias and the tyranny of the status quo (Friedman and Friedman (1984), Manski (2013)).

In Appendix B, we study the asymptotic properties of the hard-choices example, showing that when  $m \rightarrow \infty$ , the experts can approach the maximum payoff  $\alpha$  by means of ex ante bargaining.

## 4.2 Inefficiency of Interim Bargaining Solutions: A General Result

The example discussed above, although quite special, exhibits a very general property. Are interim bargaining solutions always ex ante inefficient? The answer is not obvious, because some interim solutions  $\hat{t}$  may be efficient. For instance, suppose that the signals are very powerful, in the sense that updated beliefs  $p_i(\theta|s)$  are close to each other for each  $s$ . In addition, suppose that, once  $s$  known, the increase in surplus that expert  $i$  could obtain is small when  $i$ 's favorite treatment is chosen instead of  $\hat{t}(\cdot, s)$ . Then, it may be that the interim solution is Pareto-undominated in the ex ante sense. But, in “standard cases” (a notion made precise by the proposition below),

<sup>15</sup>This follows from the fact that  $p_2(\theta_2 | s_1) = \frac{pq^m(s|\theta_2)}{(1-p)q^m(s|\theta_1) + pq^m(s|\theta_2)}$ . Thus, for  $a_1$  to be acceptable to expert 2, it must be the case that  $-\beta pq^m(s|\theta_2) + \alpha(1 - p)q^m(s|\theta_1) > 0$ .

<sup>16</sup>The probability, under  $a_1$ , that  $(\beta/\alpha)p/(1 - p) > l(s) > \beta/\alpha$ —with similar expression for  $a_2$ .

the interim solution is *always* inefficient.

To describe the interim solution  $\hat{t}$  in a convenient way, we first introduce the following notation. The treatment  $\hat{t}(\cdot, s) \in A$  yields a state-utility vector denoted  $\hat{u}(s) \in u(A)$ , where  $u(A)$  is the set of state-utility vectors generated by some basic treatment in  $A$ . Formally, we have,

$$u(A) = \{u \in \mathbb{R}^K \mid u = (E_{a(\theta_1)}v(z), \dots, E_{a(\theta_K)}v(z)), a \in A\}.$$

Let  $r_i(s)$  be expert  $i$ 's vector of conditional probabilities given  $s$ , that is,

$$r_i(s) = \begin{pmatrix} p_i(\theta_1|s) \\ \vdots \\ p_i(\theta_K|s) \end{pmatrix}.$$

For each signal  $s$ , define the feasible set  $F(s) = IR(s) \cap u(A)$ , where by definition,  $IR(s) = \bigcap_i IR_i(s)$ , and where

$$IR_i(s) = \{u \in \mathbb{R}^K \mid r_i(s) \cdot u \geq r_i(s) \cdot u(a^\circ)\},$$

is the set of individually rational state-utility vectors for expert  $i$  under signal  $s$ . Since  $\hat{u}(s)$  is a Nash bargaining solution for all  $s \in S$ , it follows that  $\hat{u}(s)$  is a Pareto-optimal point in  $F(s)$ , *i.e.*,  $\hat{u}(s) \in F(s)$  is a state-utility vector such that there doesn't exist another feasible  $u'$  with  $r_i(s) \cdot (u' - \hat{u}(s)) \geq 0$  for all  $i$  and with one strict inequality.

Recall that a state-utility vector  $u \in u(A)$  is called *admissible* if there is no other  $u' \in u(A)$  such that  $u'(\theta) \geq u(\theta)$  for all states  $\theta \in \Theta$  and with a strict inequality at least (see Ferguson (1967)). Now, define the *admissible frontier* of  $u(A)$  as the set of all admissible points in  $u(A)$ , denoted  $\mathcal{A}(u(A))$ . For each  $s$ , define the admissible and individually rational part of the frontier of  $u(A)$  as follows,

$$\mathcal{F}(s) = IR(s) \cap \mathcal{A}(u(A)).$$

Any ex-ante feasible utility vector  $u$  can be described as an array  $u = (u(s_1), \dots, u(s_L)) \in \mathbb{R}^{KL}$  where  $L$  is the number of signals, and where for each  $s$ , we have  $u(s)(\theta) = E_{t(\theta,s)}v(z)$  for each  $(\theta, s)$  and for some feasible  $t \in T$  (and therefore  $u(s) \in u(A)$ ), and  $u$  satisfies the ex ante *IR* constraints, that is, for all  $i$ ,

$$\sum_{\theta \in \Theta} \sum_{s \in S} p_i(\theta)q(s|\theta)(u(s)(\theta) - u(a^\circ)(\theta)) \geq 0.$$

An ex ante feasible  $u$  is by definition ex ante Pareto-inefficient if there exists another ex ante feasible  $u'$  such that,

$$\sum_{\theta \in \Theta} \sum_{s \in S} p_i(\theta) q(s|\theta) (u'(s)(\theta) - u(s)(\theta)) \geq 0 \quad \text{for all } i \in I,$$

with a strict inequality for at least one expert  $i$ .

Define the interim surplus  $\hat{\delta}(s) = \hat{u}(s) - u(a^\circ)$ . Recall that, by Proposition 3, the interim solution  $\hat{u}(s)$  is supported by a hyperplane with orthogonal vector  $\hat{r}(s)$  for each  $s$ , such that,

$$\hat{r}(s) = \sum_i \hat{m}_i(s) r_i(s) \quad \text{with} \quad \hat{m}_i(s) = \frac{c(s)}{r_i(s) \cdot \hat{\delta}(s)},$$

where  $c(s)$  is chosen so that  $\sum_i \hat{m}_i(s) = 1$ .

Consider now the case of two experts (i.e.,  $n = 2$ ). Consider the marginal probabilities on signals of the two experts, and the ratio  $p_1(s)/p_2(s)$ . Relabeling all signals  $s \in S$  if necessary, we can write,

$$\frac{p_1(s_1)}{p_2(s_1)} \geq \frac{p_1(s_2)}{p_2(s_2)} \geq \dots \geq \frac{p_1(s_L)}{p_2(s_L)}.$$

Define the maximal relative gap between two signals as follows,

$$\rho = \max_{(s,s') \in S^2} \left( \frac{\frac{p_1(s')}{p_2(s')} - \frac{p_1(s)}{p_2(s)}}{\frac{p_1(s)}{p_2(s)}} \right).$$

Given our relabeling of signals, we have

$$1 + \rho = \frac{p_1(s_1)}{p_2(s_1)} / \frac{p_1(s_L)}{p_2(s_L)}.$$

Proposition 4 says that if there is enough disagreement, as measured by  $\rho$ , then, the interim solutions of well-behaved bargaining problems are ex ante inefficient.

**Proposition 4.** *Let  $(\mathcal{B}(s))_{s \in S}$  be a bargaining problem.*

*Assume the following:*

- (i) there are two experts;*
- (ii)  $F(s)$  is compact with a nonempty interior for all  $s$ ;*
- (iii)  $u(A)$  is strictly convex;*
- (iv) the admissible frontier of  $u(A)$  is continuously differentiable;*
- (v) for every signal  $s$ , we have  $p_i(s) > 0$ ,  $i = 1, 2$ , and  $r_1(s) \neq r_2(s)$ , and the property, for  $i = 1, 2$ ,*

$$\min_{u \in \mathcal{F}(s)} (r_i(s) \cdot \delta) > 0.$$

*Then, there exists a threshold value  $\rho^* \geq 0$  such that, if  $\rho > \rho^*$ , the interim bargaining solution  $\hat{t}$  is ex ante Pareto-inefficient.*

*In addition, we can choose  $\rho^*$  in the following ways,*

*Case 1. If there exists a pair of signals  $(s', s'') \in S^2$  such that  $\hat{m}_1(s') < 1/2 < \hat{m}_1(s'')$ , then we can set  $\rho^* = 0$ .*

*Case 2. If  $\hat{m}_1(s) \geq 1/2$  for all  $s$ , then we can take any  $\rho^*$  such that*

$$1 + \rho^* \geq \frac{\max_{u \in \mathcal{F}(s)} (r_2(s) \cdot \delta)}{\min_{u \in \mathcal{F}(s)} (r_1(s) \cdot \delta)} \quad \text{for all } s.$$

Assumptions (ii)-(v) together define what we may call a “well-behaved” problem (in Appendix A, we show that these assumptions make sure that the bargaining problem is “standard” in a precise sense). In these problems, when experts disagree about the probability of states  $p_i$ , they also typically disagree about the marginal probability of signals  $p_i(s)$ . Note that the assumption that  $\rho$  is large is rather weak. If the set of signals  $S$  is sufficiently rich, it is quite natural, under disagreement, that  $p_1(s)/p_2(s)$  varies from 0 to infinity, or at least that this ratio has a wide range.

Intuition suggests that experts can improve efficiency by engaging in a kind of ex ante, signal-contingent trading arrangement. Remark however, that the treatment decision  $t(\theta, s)$  is like a public good for the experts. Since there is no private good, like money, that could be used to pay transfers,

it follows that transactions, in the ordinary sense, cannot be structured to compensate an expert for yielding to the demands of others. For the same reason, there are no real possibilities of ex ante insurance. Yet, in spite of these constraints creating a nontrivial problem, under the technical assumptions of the Proposition, experts can do without markets, prices, private goods or money, but they need to exploit differences in beliefs to improve on the interim bargaining solution.<sup>17</sup>

We are accustomed to think that heterogeneous beliefs are a bad thing, opening the possibility of speculation and giving rise to situations of *spurious unanimity*. This leads one to question the notion of Pareto-optimality (on this question, see *e.g.*, Mongin (1995), Mongin (2016)). A number of recent papers pointed to the adverse welfare implications of speculative behavior in private-good environments (especially financial markets). See, for example, the papers by Brunnermeier, Simsek, and Xiong (2012) and Gilboa, Samuelson, and Schmeidler (2014). The analysis of this paper offers an interesting contrast: ex ante speculation, in a collective decision context, may be essential to improving welfare. As shown above, the ex ante agreement uses the signal  $s$  as an ordeal that will determine which of the experts' theories are the most likely in an impartial way. This is as if a social planner, acting on behalf of the experts, was applying Bayes' rule before choosing the best  $t(\cdot, s)$ , once  $s$  is revealed. There is a clear sense in which any compromise is facilitated by the fact that the parties bargaining give a different meaning to words: in abstract terms, experts think that the probabilities of certain consequences of the agreed upon choices are different.

## 5 Summary and Concluding Remarks

We provide a novel application of bargaining theory to the problem of treatment choice under disagreement. In contrast to its traditional role as a mechanism for dividing private goods, we use Nash bargaining to resolve differences in prior beliefs. We studied how a group of experts chooses a treatment, or policy, by means of bargaining, when their prior beliefs (*i.e.*, their models, or theories) are different. Our model is very general and appli-

---

<sup>17</sup>For these reasons, Proposition 4 has the flavor of, but is not a restatement of Hirshleifer's paradox (see Hirshleifer (1971)).

cable to a wide variety of situations. We showed that Nash bargaining pins down the weights that a fictitious planner should use to compute a compromise belief supporting the collective decision. This approach can be used in statistical decision theory to find an average of competing models. Taking into account the fact that the disclosure of new data will lead experts to update their probabilistic beliefs, we compared the normative solution, *ex ante* bargaining, in which experts commit to a plan of action before knowing the data, and the *interim* bargaining solution, in which experts bargain once the data have been publicly revealed. Interim bargaining, which is a more realistic description of many situations, exhibits inertia, or under-reaction to new information. A more radical finding, in the case of two competing models (or two states of nature) is that interim bargaining is inadmissible, meaning that it is suboptimal for every conceivable prior belief, and hence for every Bayesian expert. More generally, in the case of  $K$  competing models (*i.e.*,  $K$  states) we found that when beliefs are sufficiently different, the interim bargaining solution is always Pareto-inefficient among experts.

One of the problems that bedevil models of heterogenous priors is that, by definition, there is no “truth” that everyone agrees on. In the extreme case of dogmatic beliefs, every expert is certain that the others are wrong. Since collective decisions must still be made despite these differences, a central question is whether information is used efficiently. Our analysis shows that *ex ante* commitments exploit differences in beliefs to achieve unambiguous, *i.e.*, prior-independent, welfare improvements and thus, a more efficient discovery of the truth. In sharp contrast, a lack of commitment leads to a distorted use of information in the form of inertia.

## A Appendix

We start with the proof of Proposition 3 because this result is used as a Lemma to prove Propositions 1 and 2. Then we prove several Lemmas leading to the proof of Proposition 2. The Proof of Proposition 1 follows. We end with the Proof of Proposition 4.

### A.1 Proof of Proposition 3

The logarithm of the Nash product in terms of the surplus in state  $k$  is:

$$\sum_{i=1}^n \log \left( \sum_{k=1}^K p_i(\theta_k) \delta(t)(\theta_k) \right).$$

The gradient of this function at  $t^*$  is the vector whose  $k$ th entry is its derivative with respect to  $\delta(t)(\theta_k)$ :

$$\sum_{i=1}^n \frac{1}{E_{p_i} \delta(t^*)} p_i(\theta_k).$$

Thus, the right hand side of (2) is proportional to the gradient of the log-Nash product at  $t^*$ , and  $\pi^*$  is its normalization to a probability distribution, obtained by dividing by  $\sum_{i=1}^n \frac{1}{E_{p_i} \delta(t^*)}$ . Since  $T$  is convex, it must be the case that  $t^*$  solves (1). ■

Proposition 3 has a partial converse: the set of treatments that solve (1) contains the set of Nash bargaining solutions. In particular, if the solution to (1) is unique, as it would be if  $T$  is strictly convex, for example, then it must coincide with the (unique) Nash bargaining solution.

### A.2 Analysis of Interim Bargaining: Preliminaries

Assume that there are two states and simplify notation by writing, for  $k = 0, 1$ :

$$\begin{aligned} p_i &= p_i(\theta_1) & i &= 1, \dots, n \\ q_k(s) &= q(s | \theta_k) \\ \hat{\delta}_{\theta_k}(s) &= \delta(\hat{t}(\theta_k, s)) \\ \delta_{\theta_k}^*(s) &= \delta(t^*(\theta_k, s)) \end{aligned}$$

and

$$\delta_{\theta_k}^* = \delta(t^*)(\theta_k) = \sum_s q_k(s) \delta_{\theta_k}^*(s),$$

where  $\hat{t} = (\hat{t}(\cdot, s))_{s \in S}$  and  $t^* = (t^*(\cdot, s))_{s \in S}$  denote interim and ex ante treatment rules, respectively.

When the feasible set of state-utility vectors  $u(A)$  is strictly convex with a differentiable frontier, we may express the admissible part of that frontier as a strictly concave and differentiable function  $\delta_1 = \phi(\delta_0)$ . Applying Proposition 3 to the interim problem at sample  $s$ , we have that  $\hat{\delta}_{\theta_0}$  and  $\hat{\delta}_{\theta_1}$  must be the unique solution to:<sup>18</sup>

$$\max_{\delta_0, \delta_1 > 0} [\hat{\pi}(\theta_0 | s) \delta_0 + \hat{\pi}(\theta_1 | s) \delta_1] \text{ subject to } \delta_1 = \phi(\delta_0). \quad (5)$$

Recall that, in general:

$$\hat{\pi}(\theta | s) = \xi(s) \sum_{i=1}^n \frac{1}{E_{p_i(\cdot | s)} \hat{\delta}(s)} p_i(\theta | s),$$

where  $\xi(s)$  is a normalization constant that depends on  $s$ . In the two-state case, Bayes rule makes it possible to simplify this expression to:

$$\begin{aligned} \hat{\pi}(\theta_1 | s) &= \xi(s) \sum_{i=1}^n \frac{p_i q_1(s) + (1 - p_i) q_0(s)}{p_i q_1(s) \hat{\delta}_{\theta_1}(s) + (1 - p_i) q_0(s) \hat{\delta}_{\theta_0}(s)} p_i(\theta_1 | s) \\ &= \xi(s) q_1(s) \sum_{i=1}^n \frac{p_i}{p_i q_1(s) \hat{\delta}_{\theta_1}(s) + (1 - p_i) q_0(s) \hat{\delta}_{\theta_0}(s)}, \end{aligned}$$

where  $p_i q_1(s) + (1 - p_i) q_0(s) = p_i(s)$  is the probability expert  $i$  assigns to observing sample  $s$ . A similar expression holds for  $\hat{\pi}(\theta_0 | s)$ . This leads to

$$\frac{\hat{\pi}(\theta_1 | s)}{\hat{\pi}(\theta_0 | s)} = \frac{q_1(s)}{q_0(s)} \frac{\sum_{i=1}^n \frac{p_i}{p_i q_1(s) \hat{\delta}_{\theta_1}(s) + (1 - p_i) q_0(s) \hat{\delta}_{\theta_0}(s)}}{\sum_{i=1}^n \frac{(1 - p_i)}{p_i q_1(s) \hat{\delta}_{\theta_1}(s) + (1 - p_i) q_0(s) \hat{\delta}_{\theta_0}(s)}} \quad (6)$$

---

<sup>18</sup> The reader may be concerned that Problem (5) appears circular: the solution is defined in terms of probabilities,  $\hat{\pi}(\theta_0 | s)$ , that are themselves defined in terms of the solution. There is no circularity here: the vector of surpluses  $\hat{\delta}_{\theta_0}, \hat{\delta}_{\theta_1}$  is *defined* as the solution to maximizing the Nash product given sample  $s$  over the convex set  $A$ . Proposition 3 simply states that this solution has the property that it also solves Problem (5).



It will be useful to define, for a real number  $\Delta > -1$ , the function:

$$\rho(\Delta) = \frac{\sum_i \frac{p_i}{p_i \Delta + 1}}{\sum_i \frac{1 - p_i}{p_i \Delta + 1}} \quad (7)$$

With this notation, we may write (6) as:

$$\frac{\hat{\pi}(\theta_1 | s)}{\hat{\pi}(\theta_0 | s)} = \frac{q_1(s)}{q_0(s)} \rho(\hat{\Delta}(s)) \quad (8)$$

where

$$\hat{\Delta}(s) = \frac{\hat{\delta}_{\theta_1}(s)q_1(s) - \hat{\delta}_{\theta_0}(s)q_0(s)}{\hat{\delta}_{\theta_0}(s)q_0(s)}.$$

Using this notation, we note, for future use, that Problem (5) can be rewritten as:

$$\max_{\delta_0, \delta_1 > 0} \left[ \delta_0 + \frac{q_1(s)}{q_0(s)} \rho(\hat{\Delta}(s)) \delta_1 \right] \text{ subject to } \delta_1 = \phi(\delta_0). \quad (9)$$

In a similar fashion, as explained in section 3.5, by Proposition 3, the ex ante solution  $\delta_{\theta_0}^*(s), \delta_{\theta_1}^*(s)$  is a solution to

$$\max_{\delta_0, \delta_1 > 0} [\pi^*(\theta_0 | s) \delta_0 + \pi^*(\theta_1 | s) \delta_1] \text{ subject to } \delta_1 = \phi(\delta_0),$$

and using expressions (3) and (4), we easily find that the ex ante solution  $\delta_{\theta_0}^*(s), \delta_{\theta_1}^*(s)$  is a solution to:

$$\max_{\delta_0, \delta_1 > 0} \left[ \delta_0 + \frac{q_1(s)}{q_0(s)} \rho(\Delta^*) \delta_1 \right] \text{ subject to } \delta_1 = \phi(\delta_0), \quad (10)$$

where

$$\Delta^* = \frac{\delta_{\theta_1}^*}{\delta_{\theta_0}^*} - 1.$$

Next, we collect some useful properties of the function  $\rho$ :

1. Note that  $\Delta > -1$  implies  $1 + p_i \Delta > 0$ , so the function  $\rho$  is always well-defined.
2. If  $p_1 = p_2 = \dots = p_n = p$ , then  $\rho(\Delta) = \frac{p}{1-p}$ .

Consider now the derivative of  $\rho$ . We can state the following result.

**Lemma A.1.** *For any  $\Delta > -1$ ,  $\rho'(\Delta) = 0$  if  $p_1 = p_2 = \dots = p_n$ , and  $\rho'(\Delta) < 0$  otherwise.*

The proof of Lemma A.1 is in Appendix B.

### A.3 Proof of Inadmissibility (Proposition 2)

The inertia and inadmissibility results are both consequences of a set of common Lemmas. More precisely, Lemmas A.1, A.2, A.3 yield Proposition 2, and Lemmas A.1, A.2, A.4, A.5 are used in the proof of Proposition 1. It follows that it is natural to prove Proposition 2 first.

Suppose, by way of contradiction, that  $\hat{t}$  is admissible. Then the Complete Class Theorem, (see Appendix B, Proposition 5), implies that  $\hat{t}$  must be optimal with respect to some Bayesian belief  $\pi^{**}$  (not necessarily related to ex ante bargaining). In this case, for every  $s$ ,  $\hat{t}(s)$  must maximize:

$$\sum_{\theta} \pi^{**}(\theta | s) \delta(a)(\theta)$$

where

$$\pi^{**}(\theta | s) = \frac{q(s | \theta) \pi^{**}(\theta)}{\sum_{\theta} q(s | \theta) \pi^{**}(\theta)}.$$

Given our assumption that the feasible set is strictly convex with smooth boundary, it follows that  $\hat{\pi}(\theta | s) = \pi^{**}(\theta | s)$  for every  $\theta$  and  $s$ . Therefore we can write,

$$\frac{\hat{\pi}(\theta_1 | s)}{\hat{\pi}(\theta_0 | s)} = \frac{q_1(s)}{q_0(s)} \frac{\pi^{**}(\theta_1)}{\pi^{**}(\theta_0)},$$

and combining (8) with this observation, it follows that  $\rho(\hat{\Delta}(s))$  must be constant in  $s$ . We show that this is not the case. We begin with a simple lemma, which we state without proof:

**Lemma A.2.** *For all  $x > 0$ , the maximization problem*

$$\max_{\delta_0, \delta_1 > 0} [\delta_0 + x \delta_1] \text{ subject to } \delta_1 = \phi(\delta_0)$$

*has a unique solution  $\bar{\delta}_0, \bar{\delta}_1$  and the ratio  $\frac{\bar{\delta}_1}{\bar{\delta}_0}$  is differentiable and strictly increasing as a function of  $x$ .*

The next lemma, which we also state without proof, applies Lemma A.1 to the interim solution:

**Lemma A.3.** *The following statements are equivalent:*

1.  $\rho(\hat{\Delta}(s))$  is constant in  $s$ .
2. For all  $s$ ,  $\hat{\Delta}(s)$  is constant in  $s$ .
3. The ratio

$$\frac{\hat{\delta}_{\theta_1}(s)q_1(s)}{\hat{\delta}_{\theta_0}(s)q_0(s)} \quad \text{is constant in } s. \quad (\text{A})$$

To complete the proof of Proposition 2, consider  $s, s'$  such that

$$\frac{q_1(s')}{q_0(s')} > \frac{q_1(s)}{q_0(s)}. \quad (11)$$

If  $\rho(\hat{\Delta}(s))$  were constant in  $s$ , then, from (8),

$$\frac{\hat{\pi}(\theta_1 | s')}{\hat{\pi}(\theta_0 | s')} > \frac{\hat{\pi}(\theta_1 | s)}{\hat{\pi}(\theta_0 | s)}. \quad (12)$$

Applying Lemma A.2 to problem 5, with  $x = \hat{\pi}(\theta_1 | s)/\hat{\pi}(\theta_0 | s)$ , it follows that

$$\frac{\hat{\delta}_{\theta_1}(s')}{\hat{\delta}_{\theta_0}(s')} > \frac{\hat{\delta}_{\theta_1}(s)}{\hat{\delta}_{\theta_0}(s)}. \quad (13)$$

From (11) and (13), it follows that:

$$\frac{\hat{\delta}_{\theta_1}(s') q_1(s')}{\hat{\delta}_{\theta_0}(s') q_0(s')} > \frac{\hat{\delta}_{\theta_1}(s) q_1(s)}{\hat{\delta}_{\theta_0}(s) q_0(s)}.$$

By the last lemma,  $\rho(\hat{\Delta}(s))$  is not constant in  $s$ . A contradiction.

*Comment on the role of differentiability in the proposition.* Suppose that the feasible set is strictly convex but its frontier contains a kink, so we lose the differentiability of the function  $\phi$ . Then (12) would still hold, indicating a change in the interim planner's beliefs. But admissibility is a property of the treatment choices; the supporting beliefs are relevant only as far as they lead to the choice of different treatments. Without differentiability, we can no longer appeal to Lemma A.2 to conclude that (13) holds with strict (rather than just weak) inequality. If (13) held with equality, then  $\hat{\pi}(\theta_1 | s)/\hat{\pi}(\theta_0 | s)$  and  $\rho(\hat{\Delta}(s))$  may vary with  $s$ , all the while that the treatments chosen remain unchanged.

#### A.4 Proof of Inertia (Proposition 1)

To prove Proposition 1, we begin with preliminary steps. Recall that the log interim Nash product evaluated at a sample  $s$  and an arbitrary vector of surpluses  $\delta_0, \delta_1$ , is:

$$\ln \mathcal{N} = \sum_i \ln \left( \frac{p_i q_1(s) \delta_1 + (1 - p_i) q_0(s) \delta_0}{p_i q_1(s) + (1 - p_i) q_0(s)} \right). \quad (14)$$

Differentiating, we obtain:

$$\frac{\partial \ln \mathcal{N}}{\partial \delta_1} = \sum_i \frac{p_i q_1(s)}{p_i q_1(s) \delta_1 + (1 - p_i) q_0(s) \delta_0}$$

and

$$\frac{\partial \ln \mathcal{N}}{\partial \delta_0} = \sum_i \frac{(1 - p_i) q_1(s)}{p_i q_1(s) \delta_1 + (1 - p_i) q_0(s) \delta_0}.$$

Next, we define  $\zeta$  as the marginal rate of substitution of the Nash product in the  $\delta_0, \delta_1$  space (so  $-\zeta^{-1}$  is the slope of the indifference curve of the Nash product). Define now  $\Delta$  as a function of  $q_1/q_0$  as:

$$\Delta = \frac{q_1 \delta_1}{q_0 \delta_0} - 1.$$

Below, we express  $\zeta$  as a function of  $\Delta$  (which is itself an increasing function of the likelihood ratio,  $q_1/q_0$ ). In the following string of equalities, we drop explicit reference to the sample  $s$  to simplify notation.

$$\begin{aligned} \zeta &= \frac{\frac{\partial \ln \mathcal{N}}{\partial \delta_1}}{\frac{\partial \ln \mathcal{N}}{\partial \delta_0}} \\ &= \frac{\sum_i \frac{p_i q_1}{p_i q_1 \delta_1 + (1 - p_i) q_0 \delta_0}}{\sum_i \frac{(1 - p_i) q_0}{p_i q_1 \delta_1 + (1 - p_i) q_0 \delta_0}} \\ &= \frac{\frac{1}{\delta_1} \sum_i \frac{p_i q_1 \delta_1}{p_i q_1 \delta_1 + (1 - p_i) q_0 \delta_0}}{\frac{1}{\delta_0} \sum_i \frac{(1 - p_i) q_0 \delta_0}{p_i q_1 \delta_1 + (1 - p_i) q_0 \delta_0}} \\ &= \frac{\delta_0 \sum_i \frac{p_i (1 + \Delta)}{p_i \Delta + 1}}{\delta_1 \sum_i \frac{(1 - p_i)}{p_i \Delta + 1}} \\ &= \frac{\delta_0}{\delta_1} (\Delta + 1) \rho(\Delta). \end{aligned}$$

**Lemma A.4.** For any fixed  $\delta_0, \delta_1 > 0$ , the function  $\zeta$  has strictly positive derivative in  $\Delta$ , and therefore also in the likelihood ratio  $q_1/q_0$ .

The proof of Lemma A.4 is given in Appendix B. We now state,

**Lemma A.5.** For any two samples  $s, s'$ ,

$$\frac{q_1(s')}{q_0(s')} > \frac{q_1(s)}{q_0(s)} \implies \frac{\hat{\delta}_{\theta_1}(s')}{\hat{\delta}_{\theta_0}(s')} > \frac{\hat{\delta}_{\theta_1}(s)}{\hat{\delta}_{\theta_0}(s)}.$$

**Proof:** Write the function  $\zeta$  as a function of the likelihood ratio and the ratio of surpluses:

$$\zeta = \zeta(q_1/q_0, \delta_1/\delta_0).$$

By Lemma A.1, for a fixed  $q_1/q_0$ ,  $\frac{\partial \zeta}{\partial (\delta_1/\delta_0)} < 0$ .

Assume by way of contradiction that  $\frac{q_1(s')}{q_0(s')} > \frac{q_1(s)}{q_0(s)}$  while  $\frac{\hat{\delta}_{\theta_1}(s')}{\hat{\delta}_{\theta_0}(s')} \leq \frac{\hat{\delta}_{\theta_1}(s)}{\hat{\delta}_{\theta_0}(s)}$ . Since  $\phi'$  is decreasing, this implies that  $\phi'(\hat{\delta}_{\theta_0}(s')) \leq \phi'(\hat{\delta}_{\theta_0}(s))$ . From the first order tangency condition at optimum,  $\phi' = -1/\zeta$ , this implies:

$$\zeta \left( \frac{q_1(s')}{q_0(s')}, \frac{\hat{\delta}_{\theta_1}(s')}{\hat{\delta}_{\theta_0}(s')} \right) \leq \zeta \left( \frac{q_1(s)}{q_0(s)}, \frac{\hat{\delta}_{\theta_1}(s)}{\hat{\delta}_{\theta_0}(s)} \right).$$

Since  $\zeta$  is decreasing in  $\delta_1/\delta_0$ , we have:

$$\zeta \left( \frac{q_1(s')}{q_0(s')}, \frac{\hat{\delta}_{\theta_1}(s)}{\hat{\delta}_{\theta_0}(s)} \right) \leq \zeta \left( \frac{q_1(s)}{q_0(s)}, \frac{\hat{\delta}_{\theta_1}(s)}{\hat{\delta}_{\theta_0}(s)} \right),$$

contradicting the conclusion of Lemma A.4. ■

**Proof of Proposition 1:** It is clear that maximizing the log of the Nash product at sample  $s$ , (14), is equivalent to maximizing

$$\max_{\delta_0, \delta_1 > 0} \sum_i \ln[p_i q_1(s) \delta_1 + (1 - p_i) q_0(s) \delta_0] \text{ subject to } \delta_1 = \phi(\delta_0).$$

The above problem yields the interim solution  $\hat{\delta}_0(s), \hat{\delta}_1(s)$ . This problem, in turn, is equivalent to

$$\max_{\delta_0, \delta_1 > 0} \sum_i \ln[p_i \ell(s) \delta_1 + (1 - p_i) \delta_0] \text{ subject to } \delta_1 = \phi(\delta_0).$$

It is easy to check that this problem has a unique solution  $\hat{\delta}_0(\ell), \hat{\delta}_1(\ell)$ , since the objective to be maximized is strictly concave, and the feasible set,  $\{\delta_1 \leq \phi(\delta_0)\}$  is strictly convex. Define  $\tilde{\Delta}(\ell) = \ell \frac{\hat{\delta}_1(\ell)}{\hat{\delta}_0(\ell)} - 1$ . By Lemma A.5,  $\tilde{\Delta}(\ell)$  is a strictly increasing function of  $\ell$ . Since  $\tilde{\Delta}(\ell)$  ranges from -1 to infinity and is continuous, by the intermediate value theorem there exists  $\ell^*$  such that

$$\tilde{\Delta}(\ell^*) = \Delta^* = \frac{\delta_{\theta_1}^*}{\delta_{\theta_0}^*} - 1.$$

Then, for any sample  $s$ :

$$\begin{aligned} \frac{q_1(s)}{q_0(s)} > \ell^* &\implies \tilde{\Delta}\left(\frac{q_1(s)}{q_0(s)}\right) > \Delta^* \\ &\implies \rho(\hat{\Delta}(s)) < \rho(\Delta^*) && \text{by Lemma A.1} \\ &\implies \frac{q_1(s)}{q_0(s)}\rho(\hat{\Delta}(s)) < \frac{q_1(s)}{q_0(s)}\rho(\Delta^*) \\ &\implies \frac{\hat{\delta}_{\theta_1}(s)}{\hat{\delta}_{\theta_0}(s)} < \frac{\delta_{\theta_1}^*(s)}{\delta_{\theta_0}^*(s)} && \text{by Lemma A.2, (9), and (10).} \end{aligned}$$

The case of  $q_1(s)/q_0(s) < \ell^*$  is similar. ■

*Connection between the proofs of Propositions 2 and 1.* Although the two propositions rely on the same set of assumptions, they deal with different properties of the interim solution  $\hat{t}$ . Proposition 1 compares the interim solution  $\hat{t}$  to a specific  $\pi^*$ -Bayesian treatment rule, where  $\pi^*$  is the planner's belief derived from ex ante bargaining. The argument used to prove the proposition is silent on whether  $\hat{t}$  might coincide with some different Bayesian treatment rule, unrelated to ex ante bargaining. At the other end, Proposition 2 shows that  $\hat{t}$  differs from any Bayesian treatment rule for *any* belief. Since inertia refers to a specific reference belief (that derived from ex ante bargaining, say), a result like Proposition 2 that considers all beliefs cannot shed light on the inertia property.

## A.5 Proof of Proposition 4

We finally prove Proposition 4 (the interim-bargaining inefficiency result). To this end we introduce a technical tool, the notion of *standard bargaining problem*. Lemma A.6 shows that under the assumptions of Proposition 4, the bargaining problems are *standard*. We then rely on this property to prove the proposition.

**Definition 1.** (*Standard bargaining problems*)

(a) A bargaining problem  $(\mathcal{B}(s))_{s \in S}$  is called *standard* at signal  $s$  if for all  $i = 1, \dots, n$ , there exist a state-utility vector  $\tilde{u}_i(s) \in F(s)$ , a conditional probability  $\tilde{r}_i(s)$  such that we have

$$\begin{aligned} p_i(s) &> 0, & \tilde{r}_i(s) &\neq \hat{r}(s), \\ r_i(s) \cdot \tilde{u}_i(s) &> r_i(s) \cdot \hat{u}(s) & \text{and} \\ \tilde{r}_i(s) \cdot \tilde{u}_i(s) &> \tilde{r}_i(s) \cdot \hat{u}(s), \end{aligned}$$

with

$$\tilde{r}_i(s) = \sum_{j \in I} \alpha_{ij}(s) r_j(s), \quad 0 < \alpha_{ij}(s) < 1, \quad \text{and} \quad \sum_{j \in I} \alpha_{ij}(s) = 1.$$

(b) The problem  $(\mathcal{B}(s))_{s \in S}$  is called *simply standard* if it is standard at every signal  $s$ .

**Remark 1:** In the case of two experts, if a problem is standard at  $s$ , then  $p_1 \neq p_2$  (i.e., there is some disagreement). *Proof of this claim:* Assume that  $p_1 = p_2$ , then,  $r_1(s) = r_2(s)$ , and it follows that  $\frac{1}{2}(r_1(s) + r_2(s))$  supports the Pareto-optimal solution  $\hat{u}(s)$  because any convex combination of  $r_1(s)$  and  $r_2(s)$  is equal to  $r_1(s)$ . In addition,  $\tilde{r}_i(s) = \hat{r}_i(s)$  for all  $i$ . It follows that the problem is nonstandard at  $s$ .

**Remark 2 (Interpretation):** The assumption of being standard at  $s$  is not very restrictive. In particular, the well-behaved feasible sets with a smooth and strictly concave admissible frontier are standard. If  $u(A)$  is the convex hull of a finite number of points, corner solutions are generic, and if  $\hat{u}(s)$  is located at a corner (or vertex), the problem is more likely to be nonstandard. In principle, we do not need to require differentiability of the frontier of  $u(A)$ , or even strict convexity, to get a standard problem.

Assume that  $n = 2$ . We can state,

**Lemma A.6.** Assume that  $p_i(s) > 0$  for all  $i$ ;  $r_1(s) \neq r_2(s)$  and  $F(s)$  has a nonempty interior at signal  $s$ . Assume in addition that  $u(A)$  is strictly convex with a nonempty interior (in the sense that for any  $u, u' \in u(A)$  and  $\lambda \in (0, 1)$ , then  $\lambda u + (1 - \lambda)u'$  belongs to the interior of  $u(A)$ ), and the admissible frontier of  $u(A)$  is differentiable, then, the problem is standard at  $s$ , and we can take any  $\alpha_{ii}(s)$  such that  $1 > \alpha_{ii}(s) > \hat{m}_i(s)$  for  $i = 1, 2$ .

**Proof of Lemma A.6:** Note first that, due to convexity and differentiability, the interim solution at  $s$  is supported by a unique hyperplane, tangent to  $F(s)$  at solution point  $\hat{u}(s)$ , with orthogonal (or normal) vector

$$\hat{r}(s) = \hat{m}_1(s)r_1(s) + \hat{m}_2(s)r_2(s),$$

where  $\hat{m}_1(s) = 1 - \hat{m}_2(s) \in (0, 1)$ . The interim solution  $\hat{u}(s)$  is Pareto-optimal given  $s$  and belongs to the admissible frontier of  $u(A)$ . Due to strict convexity, for all  $u \in F(s)$ ,  $u \neq \hat{u}(s)$ , we have  $\hat{r}(s) \cdot u < \hat{r}(s) \cdot \hat{u}(s)$ .

*Step 1.* Consider first  $i = 1$  (the reasoning is the same for  $i = 2$ ). Fix a signal  $s$ . To lighten notation, we drop ‘ $(s)$ ’ everywhere when there is no ambiguity.

First we choose  $\tilde{u}_1 \in F$  and  $\tilde{r}_1$ , as in the definition of a standard problem, such that

$$\tilde{r}_1 \cdot (\tilde{u}_1 - \hat{u}) > 0.$$

If  $\tilde{r}_1 \neq \hat{r}$ , the intersection  $F \cap \{u \mid \tilde{r}_1 \cdot (u - \hat{u}) > 0\}$  is nonempty. (If this intersection was empty, then, for all  $u \in F$ , we would have  $\tilde{r}_1 \cdot u \leq \tilde{r}_1 \cdot \hat{u}$ , implying that  $\tilde{r}_1$  defines a supporting hyperplane, tangent to  $F$  at point  $\hat{u}$ , and therefore,  $\tilde{r}_1 = \hat{r}$ , a contradiction.)

For all  $s$  and  $i$ , define the set

$$G_i(s) =$$

$$\left\{ \tilde{u} \in F(s) \mid r_i(s) \cdot (\tilde{u} - \hat{u}(s)) > 0; \tilde{r}_i(s) \cdot (\tilde{u} - \hat{u}(s)) > 0; \hat{r}(s) \cdot (\tilde{u} - \hat{u}(s)) < 0 \right\}.$$

Note that this set does not contain  $\hat{u}(s)$  and that it depends on the choice of  $\alpha_{ij}$ . Note, in addition, that the last inequality in the definition of  $G_i(s)$  must be true by definition of  $\hat{u}$ .



*Step 2.* Fix an  $s$  and take  $i = 1$ . Suppose next that we have  $\tilde{u}_1 \in G_1(s)$ , and a probability  $\tilde{r}_1 = \alpha_1 r_1 + (1 - \alpha_1) r_2$ , with  $0 < \alpha_1 < 1$ . Adding the last two inequalities in the definition of  $G_1$  yields  $(\hat{r} - \tilde{r}_1) \cdot (\tilde{u}_1 - \hat{u}) < 0$ , and we have  $(\hat{r} - \tilde{r}_1) = (\hat{m}_1 - \alpha_1)(r_1 - r_2)$ . It follows from this that we have simultaneously,

$$\begin{aligned} r_1 \cdot (\tilde{u}_1 - \hat{u}) &> 0, \quad \text{and} \\ (\hat{m}_1 - \alpha_1)(r_1 - r_2) \cdot (\hat{u} - \tilde{u}_1) &> 0. \end{aligned}$$

Assume now that  $\hat{m}_1 > \alpha_1$ . Multiply the first inequality by  $\hat{m}_1 - \alpha_1 > 0$  and add the latter two inequalities. This yields,

$$(\hat{m}_1 - \alpha_1) r_2 \cdot (\tilde{u}_1 - \hat{u}) > 0.$$

But  $r_1 \cdot (\tilde{u}_1 - \hat{u}) > 0$  and  $\hat{u}$  Pareto-optimal imply  $r_2 \cdot (\tilde{u}_1 - \hat{u}) < 0$ . This is a contradiction. It follows that  $G_1$  nonempty implies  $\alpha_1 > \hat{m}_1$ .

*Step 3.* By Step 1 above, we know that  $F \cap \{\tilde{u} \mid \tilde{r}_1 \cdot (\tilde{u} - \hat{u}) > 0\}$  is nonempty. Assume, by way of contradiction, that  $G_1 = \emptyset$ . Then, the following three inequalities must hold simultaneously for all  $\tilde{u} \in F$ :

$$r_1 \cdot (\tilde{u} - \hat{u}) \leq 0; \quad \hat{r} \cdot (\tilde{u} - \hat{u}) < 0; \quad -\tilde{r}_1 \cdot (\tilde{u} - \hat{u}) < 0.$$

Introduce nonnegative multipliers for the above three inequalities, respectively,  $(\lambda, \mu, \nu) \geq 0$  such that  $\mu + \nu > 0$ . We must have,

$$(\lambda r_1 + \mu \hat{r} - \nu \tilde{r}_1) \cdot (\tilde{u} - \hat{u}) < 0, \tag{3I}$$

and it is easy to check that,

$$\lambda r_1 + \mu \hat{r} - \nu \tilde{r}_1 = (\lambda + \mu \hat{m}_1 - \nu \alpha_1) r_1 + (\mu(1 - \hat{m}_1) - \nu(1 - \alpha_1)) r_2.$$

Now, we find positive values  $\lambda^*, \mu^*, \nu^*$  solving the following system of two linear equations,

$$\begin{cases} \lambda + \mu \hat{m}_1 - \nu \alpha_1 &= \alpha_1, \\ \mu(1 - \hat{m}_1) &= 1 - \alpha_1. \end{cases}$$

Using the second equation of the latter system, we obviously must choose

$$\mu^* = \frac{1 - \alpha_1}{1 - \hat{m}_1} > 0.$$

Then, take

$$\nu^* = \frac{\hat{m}_1 (1 - \alpha_1)}{\alpha_1 (1 - \hat{m}_1)} > 0.$$

Using the first equation of the linear system finally yields  $\lambda^* = \alpha_1 > 0$ . With these values of the multipliers, the linear combination of the three inequalities (3I) gives,

$$(\alpha_1 r_1 + (1 - \alpha_1) r_2) \cdot (\tilde{u} - \hat{u}) < \nu^* (1 - \alpha_1) r_2 \cdot (\tilde{u} - \hat{u}),$$

but the left-hand side of the latter inequality is just  $\tilde{r}_1 \cdot (\tilde{u} - \hat{u}) > 0$ . It follows that we have  $r_2 \cdot (\tilde{u} - \hat{u}) > 0$ . A contradiction will follow.

Suppose now that  $\alpha_1 > \hat{m}_1$ . Choose new values  $(\lambda^0, \mu^0, \nu^0)$  for the multipliers, namely,  $\lambda^0 = \alpha_1 - \hat{m}_1 > 0$ ,  $\mu^0 = 1$  and  $\nu^0 = 0$ . Using these new values, the linear combination of our three inequalities (3I) now yields,

$$(\alpha_1 r_1 + (1 - \hat{m}_1) r_2) \cdot (\tilde{u} - \hat{u}) < 0.$$

But we have shown above that  $r_2 \cdot (\tilde{u} - \hat{u}) > 0$  and we assume  $1 - \alpha_1 < 1 - \hat{m}_1$ . Hence,

$$(\alpha_1 r_1 + (1 - \alpha_1) r_2) \cdot (\tilde{u} - \hat{u}) < 0,$$

in other terms, we have found  $\tilde{r}_1 \cdot (\tilde{u} - \hat{u}) < 0$ , a contradiction. We conclude that:  $r_1 \neq r_2$  and  $\alpha_1 > \hat{m}_1$  implies  $G_1 \neq \emptyset$ .

To end the proof, note that the same reasoning can be applied, *mutatis mutandis*, to  $\tilde{u}_2$  and  $\tilde{r}_2$ , showing that  $\alpha_{22} > \hat{m}_2$  implies  $G_2 \neq \emptyset$ . ■

**Proof of Proposition 4.** Choose a pair of signals  $s_k$  and  $s_\ell$  with  $k < \ell$ .

Let  $\tilde{u}_i(s_h)$  be as defined in the definition of a standard problem. We construct  $u' \in F$  as follows. Let  $\epsilon_h \in (0, 1)$ , for  $h = k, \ell$ , and define,

$$\begin{aligned} u'(s_k) &= \epsilon_k \tilde{u}_1(s_k) + (1 - \epsilon_k) \hat{u}(s_k), \\ u'(s_\ell) &= \epsilon_\ell \tilde{u}_2(s_\ell) + (1 - \epsilon_\ell) \hat{u}(s_\ell), \\ u'(s_h) &= \hat{u}(s_h) \quad \text{for all } h \neq k, h \neq \ell. \end{aligned}$$

We can write,

$$\begin{aligned} & \sum_s p_i(s) r_i(s) \cdot (u'(s) - \hat{u}(s)) \\ &= p_i(s_k) r_i(s_k) \cdot (u'(s_k) - \hat{u}(s_k)) + p_i(s_\ell) r_i(s_\ell) \cdot (u'(s_\ell) - \hat{u}(s_\ell)) \\ &= \epsilon_k p_i(s_k) r_i(s_k) \cdot (\tilde{u}_1(s_k) - \hat{u}(s_k)) + \epsilon_\ell p_i(s_\ell) r_i(s_\ell) \cdot (\tilde{u}_2(s_\ell) - \hat{u}(s_\ell)) \end{aligned}$$

Now, define for convenience,

$$\begin{aligned} h_{1k} &= r_1(s_k) \cdot (\tilde{u}_1(s_k) - \hat{u}(s_k)); & h_{2k} &= r_2(s_k) \cdot (\tilde{u}_1(s_k) - \hat{u}(s_k)); \\ h_{1\ell} &= r_1(s_\ell) \cdot (\tilde{u}_2(s_\ell) - \hat{u}(s_\ell)); & h_{2\ell} &= r_2(s_\ell) \cdot (\tilde{u}_2(s_\ell) - \hat{u}(s_\ell)). \end{aligned}$$

With this notation, we have  $\sum_s p_i(s) r_i(s) \cdot (u'(s) - \hat{u}(s)) > 0$  for  $i = 1, 2$  if and only if,

$$\begin{aligned} \epsilon_k p_1(s_k) h_{1k} + \epsilon_\ell p_1(s_\ell) h_{1\ell} &> 0, \\ \epsilon_k p_2(s_k) h_{2k} + \epsilon_\ell p_2(s_\ell) h_{2\ell} &> 0. \end{aligned}$$

By Lemma A.6, under strict convexity (*i.e.*, Assumption (iii)) and differentiability (*i.e.*, Assumption (iv)) the problem is standard at  $s_k$  and  $s_\ell$ . Since the problem is standard at  $s_k$ , with  $\tilde{u}_1(s_k)$ , we have  $h_{1k} > 0$ . Similarly, since the problem is standard at  $s_\ell$  with  $\tilde{u}_2(s_\ell)$ , we have  $h_{2\ell} > 0$ . Remark that, given this, we have  $h_{2k} < 0$  since  $\hat{u}(s_k)$  is interim Pareto-optimal in  $F(s_k)$ . (Indeed, there would be a contradiction if we had  $h_{2k} \geq 0$ .) Similarly, we have  $h_{1\ell} < 0$  since  $\hat{u}(s_\ell)$  is interim Pareto-optimal in  $F(s_\ell)$ . Using these results, the above system of two strict inequalities can be rewritten,

$$\begin{aligned} \epsilon_k &> \frac{p_1(s_\ell)}{p_1(s_k)} \frac{(-h_{1\ell})}{h_{1k}} \epsilon_\ell, \\ \epsilon_\ell &> \frac{p_2(s_k)}{p_2(s_\ell)} \frac{(-h_{2k})}{h_{2\ell}} \epsilon_k. \end{aligned}$$

The set of solutions  $(\epsilon_k, \epsilon_\ell) \in (0, 1)^2$  is nonempty if we have the key inequality (KI),

$$1 > \frac{p_1(s_\ell) p_2(s_k)}{p_2(s_\ell) p_1(s_k)} \frac{|h_{2k}|}{h_{1k}} \frac{|h_{1\ell}|}{h_{2\ell}}. \quad (KI)$$

*Remark:* This is equivalent to saying that the determinant of the associated homogeneous linear system of equations in  $(\epsilon_k, \epsilon_\ell)$  is positive, namely,

$$\begin{vmatrix} p_1(s_k) h_{1k} & p_1(s_\ell) h_{1\ell} \\ p_2(s_k) h_{2k} & p_2(s_\ell) h_{2\ell} \end{vmatrix} > 0.$$

*Case 1.* In this case, there exists a pair of signals  $(s', s'') \in S^2$  such that  $\hat{m}_1(s') < 1/2 < \hat{m}_1(s'')$ . We can take  $s_k = s'$  and  $s_\ell = s''$  and assume  $k < \ell$ .

If  $k > \ell$ , we exchange the names of the experts. The above inequalities are equivalent to,

$$\hat{m}_1(s_k) < \frac{1}{2}, \quad \text{and} \quad \hat{m}_2(s_\ell) < \frac{1}{2}.$$

It follows from Lemma A.6 again that we can find a feasible  $\tilde{u}_1(s_k)$  such that

$$\left( \frac{r_1(s_k)}{2} + \frac{r_2(s_k)}{2} \right) \cdot (\tilde{u}_1(s_k) - \hat{u}(s_k)) > 0, \quad \text{and} \quad r_1(s_k) \cdot (\tilde{u}_1(s_k) - \hat{u}(s_k)) > 0,$$

(and of course we choose  $\tilde{r}_1(s_k) = \frac{1}{2}(r_1(s_k) + r_2(s_k))$ ). Similarly, we can find a feasible  $\tilde{u}_2(s_\ell)$  such that

$$\left( \frac{r_1(s_\ell)}{2} + \frac{r_2(s_\ell)}{2} \right) \cdot (\tilde{u}_2(s_\ell) - \hat{u}(s_\ell)) > 0, \quad \text{and} \quad r_2(s_\ell) \cdot (\tilde{u}_2(s_\ell) - \hat{u}(s_\ell)) > 0,$$

(and of course we choose  $\tilde{r}_2(s_\ell) = \frac{1}{2}(r_1(s_\ell) + r_2(s_\ell))$ ).

With these specific values (*i.e.*,  $\alpha_{ii} = 1/2$ ), it is easy to check that,

$$\frac{|h_{2k}|}{h_{1k}} < 1 \quad \text{and} \quad \frac{|h_{1\ell}|}{h_{2\ell}} < 1.$$

Since  $k < \ell$ , we have in addition,

$$\frac{p_1(s_\ell) p_2(s_k)}{p_2(s_\ell) p_1(s_k)} \leq 1.$$

The key inequality (*KI*) is therefore satisfied.

We conclude that there exist values  $\epsilon_h \in (0, 1)$  for  $h = k, \ell$ , such that  $\sum_s p_i(s) r_i(s) \cdot (u'(s) - \hat{u}(s)) > 0$  for  $i = 1, 2$ , showing that  $\hat{u}$  is dominated in the sense of Pareto, in the ex ante problem. We can obtain this result even if  $\rho = 0$ . This proves the result in Case 1.

*Case 2.* Assume now that for all  $s \in S$ , we have  $\hat{m}_1(s) \geq 1/2$ . Note that if  $\rho > 0$ , then

$$\frac{p_1(s_L) p_2(s_1)}{p_2(s_L) p_1(s_1)} = \frac{1}{1 + \rho} < 1.$$

We pick the two extreme signals,  $s_k = s_1$  and  $s_\ell = s_L$ . Case 2 can be divided in two complementary subcases.

*Case 2a.* In this subcase, we assume that  $\hat{m}_1(s_1) \geq 1/2$  and  $\hat{m}_1(s_L) > 1/2$ .

At signal  $s_L$ , we take a  $\tilde{u}_2(s_L)$  (as in the definition of a standard problem, since our problem is standard under Assumptions (iii) and (iv)) and  $\tilde{r}_2(s_L) = \frac{1}{2}(r_1 + r_2)$ . This is a feasible choice since  $\hat{m}_2(s_L) = 1 - \hat{m}_1(s_L) < 1/2$ . Next, we choose  $\tilde{u}_1(s_1)$  as in the definition of a standard problem, along with

$$\tilde{r}_1(s_1) = (\hat{m}_1(s_1) + \eta)r_1(s_1) + (1 - \hat{m}_1(s_1) - \eta)r_2(s_1),$$

and a small  $\eta > 0$ . From this definition of  $\tilde{r}_1(s_1)$ , we easily derive the inequality,

$$\frac{|h_{21}|}{h_{11}} \leq \frac{\hat{m}_1(s_1) + \eta}{1 - \hat{m}_1(s_1) - \eta}.$$

Since  $-h_{1L}/h_{2L} < 1$ , it is always possible to choose  $\eta > 0$  small enough so that,

$$\frac{|h_{1L}|}{h_{2L}} \frac{(\hat{m}_1(s_1) + \eta)}{(1 - \hat{m}_1(s_1) - \eta)} \leq \frac{\hat{m}_1(s_1)}{(1 - \hat{m}_1(s_1))}.$$

It follows that the product of ratios appearing in  $(KI)$  satisfies,

$$\frac{p_1(s_L)}{p_2(s_L)} \frac{p_2(s_1)}{p_1(s_1)} \frac{|h_{1L}|}{h_{2L}} \frac{|h_{21}|}{h_{11}} \leq \left( \frac{1}{1 + \rho} \right) \frac{\hat{m}_1(s_1)}{(1 - \hat{m}_1(s_1))}.$$

Define

$$H(s) = \left( \frac{\max_{u \in \mathcal{F}(s)} (r_2(s) \cdot \delta)}{\min_{u \in \mathcal{F}(s)} (r_1(s) \cdot \delta)} \right).$$

Now, we have, by definition of  $\hat{m}_i(s)$ , and with  $\delta = u - u(a^\circ)$ ,

$$\frac{\hat{m}_1(s_1)}{(1 - \hat{m}_1(s_1))} = \frac{r_2(s_1) \cdot \hat{\delta}(s_1)}{r_1(s_1) \cdot \hat{\delta}(s_1)} \leq H(s_1).$$

It follows that if we set  $1 + \rho > H(s_1)$ , the key inequality  $(KI)$  holds true again, showing that there exists  $(\epsilon_1, \epsilon_L) \in (0, 1)^2$  such that  $u'$  dominates  $\hat{u}$ . As a consequence, in this case too,  $\hat{u}$  is dominated by  $u'$  in the sense of *ex ante* Pareto-optimality.

*Case 2b.* Finally, we treat the only remaining case in which  $\hat{m}_1(s_1) \geq \hat{m}_1(s_L) = 1/2$ .

In this latter subcase, we can take  $\tilde{u}_1(s_1)$  with  $\eta_1 > 0$  small, such that

$$\tilde{r}_1(s_1) = (\hat{m}_1(s_1) + \eta_1)r_1(s_1) + (1 - \hat{m}_1(s_1) - \eta_1)r_2(s_1).$$

This yields,

$$\frac{-h_{21}}{h_{11}} < \frac{(\hat{m}_1(s_1) + \eta_1)}{(1 - \hat{m}_1(s_1) - \eta_1)}.$$

We take  $\tilde{u}_2(s_L)$  under signal  $s_L$ , and choose

$$\tilde{r}_2(s_L) = (\hat{m}_1(s_L) - \eta_2)r_1(s_L) + (1 - \hat{m}_1(s_L) + \eta_2)r_2(s_L),$$

with a small  $\eta_2 > 0$ . Since  $\hat{m}_1(s_L) = 1/2$ , this yields,

$$\frac{-h_{1L}}{h_{2L}} < \frac{(1 - \hat{m}_1(s_L) + \eta_2)}{(\hat{m}_1(s_L) - \eta_2)} = \frac{(1 + 2\eta_2)}{(1 - 2\eta_2)},$$

the right-hand side being as close to 1 as desired for sufficiently small  $\eta_2 > 0$ .

It follows that the key product in  $(KI)$  can be written

$$\left( \frac{1}{1 + \rho} \right) \frac{(1 + 2\eta_2)}{(1 - 2\eta_2)} \frac{(\hat{m}_1(s_1) + \eta_1)}{(1 - \hat{m}_1(s_1) - \eta_1)}.$$

But we have,

$$\frac{\hat{m}_1(s_1)}{(1 - \hat{m}_1(s_1))} = \frac{r_2(s_L) \cdot \hat{\delta}(s_L)}{r_1(s_L) \cdot \hat{\delta}(s_L)} \leq H(s_L).$$

We conclude that with  $\eta_1$  and  $\eta_2$  sufficiently small and  $1 + \rho > H(s_L)$ , the key product is smaller than 1 and  $(KI)$  holds true. Again, this implies that  $\hat{u}$  is Pareto-dominated by  $u'$  in the ex ante bargaining problem. ■

## References

- ANGRIST, J. D., AND J.-S. PISCHKE (2009): *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- AUMANN, R. J. (1976): "Agreeing to Disagree," *The Annals of Statistics*, 4(6), 1236–1239.
- AUMANN, R. J. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55(1), 1–18.
- AUMANN, R. J. (1998): "Common Priors: A Reply to Gul," *Econometrica*, 66, 929–38.
- BERGER, J. (1985): *Statistical Decision Theory and Bayesian Analysis*. Springer.
- BINMORE, K., A. RUBINSTEIN, AND A. WOLINSKY (1986): "The Nash Bargaining Solution in Economic Modelling," *The RAND Journal of Economics*, pp. 176–188.
- BROCK, W. A., S. N. DURLAUF, AND K. D. WEST (2003): "Policy Evaluation in Uncertain Economic Environments," in *Brookings Papers on Economic Activity*, vol. 1, pp. 235–322. Brookings Institution.
- BRUNNERMEIER, M., A. SIMSEK, AND W. XIONG (2012): "A Welfare Criterion for Models with Distorted Beliefs," Princeton University.
- CHRISTENSEN, G., AND E. MIGUEL (2018): "Transparency, Reproducibility, and the Credibility of Economics Research," *Journal of Economic Literature*, 56(3), 920–80.
- DONOHUE, J. J., AND J. WOLFERS (2009): "Estimating the Impact of the Death Penalty on Murder," *American Law and Economics Review*, 11(2), 249–309.
- DURLAUF, S. N., C. FU, AND S. NAVARRO (2013): "Capital Punishment and Deterrence: Understanding Disparate Results," *Journal of Quantitative Criminology*, 29(1), 103–121.

- FERGUSON, T. S. (1967): *Mathematical Statistics: A Decision Theoretic Approach*. Academic Press, New York.
- FRIEDMAN, M., AND R. D. FRIEDMAN (1984): *Tyranny of the Status Quo*. Harcourt, Brace, Jovanovitch.
- GILBOA, I., L. SAMUELSON, AND D. SCHMEIDLER (2014): “No-Betting Pareto Dominance,” *Econometrica*, 82, 1405–42.
- GUL, F. (1998): “A Comment on Aumann’s Bayesian View,” *Econometrica*, 66(4), 923–927.
- HECKMAN, J., AND J. SMITH (1998): “Evaluating the Welfare State,” in *Econometrics and Economic Theory in the 20th Century. The Ragnar Frisch Centennial Symposium*, ed. by S. Strom, chap. 8, pp. 235–322. Cambridge University Press.
- HECKMAN, J. J., AND E. VYTLACIL (2005): “Structural Equations, Treatment Effects, and Econometric Policy Evaluation,” *Econometrica*, 73(3), 669–738.
- HIRSHLEIFER, J. (1971): “The Private and Social Value of Information and the Reward to Inventive Activity,” *American Economic Review*, 61(4), 561–574.
- HOETING, J. A., D. MADIGAN, A. E. RAFTERY, AND C. T. VOLINSKY (1999): “Bayesian Model Averaging: a Tutorial,” *Statistical Science*, pp. 382–401.
- MANSKI, C. F. (2003): *Partial Identification of Probability Distributions*. Springer Science & Business Media.
- (2004): “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, 72(4), 1221–1246.
- (2013): “Status Quo Deference and Policy Choice under Ambiguity,” *Journal of Institutional and Theoretical Economics*, pp. 116–128.
- MONGIN, P. (1995): “Consistent Bayesian Aggregation,” *Journal of Economic Theory*, 66(2), 511–532.



- (2016): “Spurious Unanimity and the Pareto Principle (original version 1997),” *Economics & Philosophy*, 32, 511–532.
- MORRIS, S. (1995): “The Common Prior Assumption in Economic Theory,” *Economics and Philosophy*, 11, 227–253.
- NAGIN, D. S. (2013): “Deterrence: A Review of the Evidence by a Criminologist for Economists,” *Annual Review of Economics*, 5(1), 83–105.
- NASH, J. F. (1950): “The Bargaining Problem,” *Econometrica*, 18(2), 155–162.
- ROTH, A. E. (1977): “Individual Rationality and Nash’s Solution to the Bargaining Problem,” *Mathematics of Operations Research*, 2(1), 64–65.
- SERRANO, R. (2005): “Fifty years of the Nash program, 1953-2003,” *Investigaciones Economicas*, 29(2), 219–258.
- STEEL, M. F. (2019): “Model Averaging and its Use in Economics,” *arXiv preprint arXiv:1709.08221*.
- WASSERMAN, L. (2000): “Bayesian Model Selection and Model Averaging,” *Journal of Mathematical Psychology*, 44(1), 92–107.
- WEERAHANDI, S., AND J. V. ZIDEK (1983): “Elements of Multi-Bayesian Decision Theory,” *Annals of Statistics*, 11(4), 1032–1046.

## B Online Appendix (Not for Publication)

### B.1 A Simple Illustrative Example: Robust Randomizations

Proposition 3 makes it possible to think about Nash bargaining as a mechanism to mediate differences in opinion. Figure 1 illustrates this in a two-state, two-expert example. The horizontal axis measures the utility of a treatment in state  $\theta_1$ , while the vertical axis measures utility in state  $\theta_2$ .

Treatments are identified with the corresponding vectors of state-utilities. For the status quo,  $u(a^\circ)$  defines the point where the individual rationality constraints are binding for both experts. The set of individually rational state-utility vectors  $IR$  is the cone with vertex  $u(a^\circ)$ , and boundaries orthogonal to the experts' beliefs,  $p_1$  and  $p_2$ .

The curves in the figure represent level sets of the Nash product. It is evident that the Nash planner is not a Bayesian decision maker, since this would have required these level sets to be linear and  $IR$  to be a half space. Clearly, this is not true here, unless the experts' beliefs agree. The figure also shows the normalized gradient of the Nash product at  $u(t)$ , denoted  $\pi(t)$ . In light of Proposition 3,  $\pi(t)$  may be interpreted as the planner's belief at  $t$ .

Compare the two treatments  $t_1$  and  $t_2$  in the figure. Expert 2 holds the most pessimistic belief about  $t_1$ , with an expected surplus  $E_{p_2}\delta(t_1)$  close to zero. From Proposition 3, expert 2's belief is weighted more heavily when evaluating  $t_1$  relative to expert 1. In fact, the planner's belief  $\pi(t_1)$  nearly coincides with that of expert 2. The converse holds for  $t_2$ . This illustrates an important implication of Nash bargaining in our context, namely that, in comparing alternative treatments, the opinions of the most pessimistic experts are given greater weight. This may be interpreted as displaying caution when a treatment is chosen under disagreement.

A related consequence of the Nash bargaining solution is robust randomizations. Suppose that  $A = \text{co}\{a^\circ, t_1, t_2\}$  is the feasible set. Since the Nash bargaining solution is Pareto optimal, it must select a treatment  $t^*$  that lies on the Pareto frontier of this set, shown as the dotted line connecting  $u(t_1)$  and  $u(t_2)$ . It is clear that neither treatment  $t_1$  nor  $t_2$  can be the Nash bargaining solution for this problem. Rather,  $t^*$  is a convex combination of the two. How should this be interpreted? One possibility is that either




Figure 1: Planner's "Indifference Curves"

$t_1$  or  $t_2$  is chosen based on a coin flip (with the appropriate probability). Justifications for such randomization appear in Manski (2004), based on a non-Bayesian decision criterion. Alternatively, convex combinations may correspond to proportions of the population assigned to a particular treatment. Informally, we assume an underlying continuum of individuals, and interpret a convex combination  $\lambda t_1 + (1 - \lambda)t_2$  as applying treatment  $t_1$  to a fraction  $\lambda$  and treatment  $t_2$  to the rest. No randomization is invoked under this interpretation.

Convex combinations in a bargaining model are robust. In the example, if we interpret  $t^*$  as a coin flip between  $t_1$  or  $t_2$ , then the randomization remains (possibly with different weights) for small changes in utility, beliefs, or the feasible set. For a Bayesian planner, by contrast, a randomized decision arises only in knife-edge cases of indifference and would disappear with minute changes in the parameters of the problem.

## B.2 Expanding State Space

If experts  $i$  have different likelihood functions  $q_i$ , we can reformulate the problem, by adding states, in such a way that a likelihood function  $q$  becomes common to all experts. For each  $\theta$ , consider all experts  $i$  with  $p_i(\theta) > 0$ . For any such expert, define  $\theta_i$  such that: (1)  $p_i(\theta_i) = p_i(\theta)$ ; and (2)  $t(\theta_i, s) = t(\theta, s)$  for all  $t$  and  $s$ . Define  $q(s|\theta_i) = p_i(s|\theta)$ . Replace  $\theta$  by the  $\theta_i$ 's constructed above.

## B.3 Extension to Different Utility Functions

We have assumed that experts share the same utility function in order to focus on the role of bargaining in mediating differences in beliefs. This is appropriate in collective decision problems where experts agree on the objectives but have different opinions on how to achieve them.

A bargaining framework can still be used when experts have different utilities as well as different beliefs. The bargaining solution  $t^*$  continues to be well-defined but the analysis becomes less transparent: since utilities now depend on  $i$ , the log Nash product is a function of  $n \times K$  terms of the form  $\delta_i(t)(\theta_k)$ , so we cannot characterize  $t^*$  in terms of a planner's belief. While we suspect that our results on commitment, inadmissibility, and inertia would continue to hold when experts have different utilities, new proofs are needed

since the present proofs rely on the characterization of the planner’s belief derived in Proposition 3 in a fundamental way.

In summary, the broader point of this paper is the use of Nash bargaining to study collective decisions under disagreement. This requires neither a common utility nor concordant beliefs. The narrower path we pursue in this paper, on the other hand, makes it possible to obtain sharper results about bargaining under disagreement.

#### B.4 Asymptotic Behavior of Hard Choices Example

We provide here additional details relative to the example presented in subsection 4.2.

In Figure 2, the set of feasible treatments is the triangle representing the convex hull of the state-utilities of the feasible treatments,  $co\{u(a^\circ), u(a_1), u(a_2)\}$ . The set of individually rational state-utilities,  $IR$ , is shown as the shaded cone with vertex  $u(a^\circ)$ . The status quo is the only point that is both feasible and individually rational at the interim stage, and is therefore (trivially) the interim Nash bargaining solution for this problem. We have shown how the bargaining deadlock can be broken by conditioning on the outcome of a statistical experiment.

Data takes the form of observations  $s = (s^1, \dots, s^m)$  that are conditionally i.i.d. given the state. The experts agree on the econometric model  $q$  generating the data but disagree on the probabilities of states  $\theta$ . We interpret  $m$  as sample size. For  $m = 1$ , the set of signals is  $S = \{s_0, s_1\}$  and  $q(s_1 | \theta_1) = q > 1/2$ ,  $q(s_1 | \theta_2) = 1 - q$  as presented in the main body of the paper. For  $m > 1$ , define  $S^m = \{s_0, s_1\}^m$  and let  $q^m(\cdot | \theta)$  be the product of the  $q(s^h | \theta)$ ,  $h = 1, \dots, m$ . Now we show that, as the number of observations  $m \rightarrow \infty$ , the ex ante expected utility for both experts approaches the maximum payoff  $u(t^{max})$ , corresponding to perfect information.

Define  $k(s) = |\{j | s^j = s_1\}|$  the number of times  $s_1$  has been drawn in  $s$ . For any  $m$ , consider a simple treatment rule  $t_m(k(s))$  that depends on  $s$  only through  $k(s)$ . For any such treatment rule  $t$  in the hard choices problem,




Figure 2: Hard Choices

and for any  $m$ , the expected utility of expert  $i$  is:

$$\begin{aligned}
 E_{p_i} u(t) &= p_i(\theta_1) \sum_{s \in \mathcal{S}^m} q^{k(s)} (1-q)^{(m-k(s))} u(t(s)) \\
 &\quad + p_i(\theta_2) \sum_{s \in \mathcal{S}^m} (1-q)^{k(s)} q^{(m-k(s))} u(t(s)).
 \end{aligned}$$

More specifically, for  $t_m(k)$ , we have,

$$\begin{aligned}
 E_{p_i} u(t_m) &= p_i(\theta_1) \sum_{k=0}^m \binom{m}{k} q^k (1-q)^{(m-k)} u(t_m(k)) \\
 &\quad + p_i(\theta_2) \sum_{k=0}^m \binom{m}{k} (1-q)^k q^{(m-k)} u(t_m(k)).
 \end{aligned}$$

Let  $\lfloor m/2 \rfloor$  be equal to  $m/2$  if  $m$  is even, and equal to  $(m-1)/2$  if  $m$  is odd. Consider next the particular treatment rule:

$$t_m(k) = \begin{cases} a_1 & \text{if } k \geq \lfloor m/2 \rfloor \\ a_2 & \text{if } k < \lfloor m/2 \rfloor; \end{cases}$$

Then expected utility can be rewritten as:

$$\begin{aligned} E_{p_i} u(t_m) = & p_i(\theta_1) \left[ \alpha \Pr \left( k \geq 1 + \left\lfloor \frac{m}{2} \right\rfloor \mid \theta_1 \right) - \beta \Pr \left( k \leq \left\lfloor \frac{m}{2} \right\rfloor \mid \theta_1 \right) \right] \\ & + p_i(\theta_2) \left[ \alpha \Pr \left( k \leq \left\lfloor \frac{m}{2} \right\rfloor \mid \theta_2 \right) - \beta \Pr \left( k \geq 1 + \left\lfloor \frac{m}{2} \right\rfloor \mid \theta_2 \right) \right]. \end{aligned}$$

Now consider a sequence of even values of  $m$ , that is,  $m = 2r$ ,  $r$  increasing without bound. Since  $q > 1/2$ , by the Law of Large Numbers, when  $m \rightarrow +\infty$  we have,

$$\begin{aligned} \Pr \left[ k \leq \frac{m}{2} \mid \theta_1 \right] &= \Pr \left[ \frac{k}{m} \leq \frac{1}{2} \mid \theta_1 \right] \rightarrow 0, \\ \Pr \left[ k \geq 1 + \frac{m}{2} \mid \theta_1 \right] &= \Pr \left[ \frac{k}{m} \geq \frac{1}{2} + \frac{1}{m} \mid \theta_1 \right] \rightarrow 1. \end{aligned}$$

We find similar results for  $m$  odd and when the conditioning state is  $\theta_2$ . We then easily find that for any belief  $p$ , the expected payoffs converge, that is,  $E_p u(t_m) \rightarrow \alpha$ . We conclude that the Nash planner cannot achieve a smaller expected utility with the optimal ex ante solution, as  $m$  grows without bound. The expert's payoffs must be approaching  $\alpha$  too.

## B.5 Pareto Optimality and Admissibility

In assessing the optimality of a treatment rule, it is natural to consider the Pareto criterion:

**Definition 2.** *Given a profile of beliefs  $\{p_1, \dots, p_n\}$ , a treatment rule  $t$  Pareto dominates another treatment rule  $t'$  if  $E_{p_i} u(t) \geq E_{p_i} u(t')$  for each expert  $i$ , with at least one strict inequality.*

*Treatment rule  $t$  is Pareto optimal relative to a feasible set  $T$  if it is not Pareto dominated by any other treatment rule  $t' \in T$ .*

A number of authors questioned the appropriateness of the Pareto criterion when agents have different beliefs. See, for example, Mongin (2016), Brunnermeier, Simsek, and Xiong (2012), and Gilboa, Samuelson, and Schmeidler (2014). This suggests *admissibility* as an attractive alternative:

**Definition 3.** A treatment rule  $t'$  dominates another treatment rule  $t$  if  $u(t')(\theta) \geq u(t)(\theta)$  for each state  $\theta$ , with at least one strict inequality.

Treatment rule  $t$  is admissible relative to a feasible set  $T$  if it is not dominated by any other treatment rule  $t' \in T$ .

Admissibility is an appealing criterion, commonly used in statistical decision theory and in the treatment choice literature (see, for example, Berger (1985)).<sup>19</sup> We first observe the following:

**Fact:** A treatment rule  $t$  that is Pareto optimal relative to a feasible set  $T$  is admissible.<sup>20</sup>

A key advantage of admissibility is that it is belief-free—in contrast to Pareto optimality which depends on the experts' profile of beliefs. To say that a treatment rule is inadmissible is an unambiguous judgement about its inefficiency since that rule can be improved on in every state and, therefore, in expectation for *any* belief.

We conclude by recalling a well-known result, the Complete Class Theorem (cf. Ferguson (1967)), which we use extensively in this paper, and which characterizes admissible rules as those that are Bayesian:

**Definition 4.** A treatment rule  $t$  is Bayesian if it maximizes expected utility with respect to some prior  $p$  on  $\Omega$ .

**Proposition 5.** A Bayesian treatment rule with respect to a full-support prior  $p$  is admissible. Conversely, an admissible treatment rule must be Bayesian.

## B.6 Proof of Lemma A.1

**Proof:**

$$\begin{aligned} \rho'(\Delta) &= \frac{\sum_i \frac{-p_i^2}{(p_i\Delta+1)^2}}{\sum_i \frac{1-p_i}{p_i\Delta+1}} + \frac{\sum_i \frac{p_i}{p_i\Delta+1}}{\left(\sum_i \frac{1-p_i}{p_i\Delta+1}\right)^2} \left(\sum_i \frac{(1-p_i)p_i}{(p_i\Delta+1)^2}\right) \\ &= B \times A \end{aligned}$$

---

<sup>19</sup>The definition of admissibility in our bargaining context coincides with that in the statistics literature under the assumptions of common values and when the experts share the likelihood function  $q(s|\theta)$ .

<sup>20</sup>It is easy to see that the converse is not true in general.



where

$$B = \left( \sum_i \frac{1 - p_i}{p_i \Delta + 1} \right)^{-2} > 0 \quad (15)$$

and

$$A = \left[ \left( \sum_i \frac{p_i}{p_i \Delta + 1} \right) \left( \sum_i \frac{(1 - p_i)p_i}{(p_i \Delta + 1)^2} \right) - \left( \sum_i \frac{p_i^2}{(p_i \Delta + 1)^2} \right) \left( \sum_i \frac{1 - p_i}{p_i \Delta + 1} \right) \right]. \quad (16)$$

We show that  $A < 0$ :

$$\begin{aligned} A &= \sum_i \sum_j \frac{p_i p_j (1 - p_j)}{(p_i \Delta + 1)(p_j \Delta + 1)^2} - \sum_i \sum_j \frac{p_i^2 (1 - p_j)}{(p_i \Delta + 1)^2 (p_j \Delta + 1)} \\ &= \sum_i \sum_j \frac{p_i p_j - p_i p_j^2}{(p_i \Delta + 1)(p_j \Delta + 1)^2} - \sum_j \sum_i \frac{p_j^2 (1 - p_i)}{(p_j \Delta + 1)^2 (p_i \Delta + 1)} \\ &= \sum_i \sum_j \frac{p_i p_j - p_j^2}{(p_i \Delta + 1)(p_j \Delta + 1)^2} \\ &= \sum_i \sum_{\substack{j \\ j \neq i}} \frac{p_j (p_i - p_j)(p_i \Delta + 1)}{(p_i \Delta + 1)^2 (p_j \Delta + 1)^2} \\ &= \sum_i \sum_{\substack{j \\ j < i}} \frac{p_j (p_i - p_j)(p_i \Delta + 1) + p_i (p_j - p_i)(p_j \Delta + 1)}{(p_i \Delta + 1)^2 (p_j \Delta + 1)^2} \quad (17) \\ &= - \sum_i \sum_{\substack{j \\ j < i}} \frac{(p_i - p_j)^2}{(p_i \Delta + 1)^2 (p_j \Delta + 1)^2}, \end{aligned}$$

where the last equality follows from cancellation of the numerator in (17) :

$$\begin{aligned} p_j (p_i - p_j)(p_i \Delta + 1) + p_i (p_j - p_i)(p_j \Delta + 1) &= p_j p_i^2 \Delta + p_j p_i - p_j^2 p_i \Delta \\ &\quad - p_j^2 + p_i p_j^2 \Delta + p_i p_j \\ &\quad - p_i^2 p_j \Delta - p_i^2 \\ &= -(p_i - p_j)^2. \end{aligned}$$

■

## B.7 Proof of Lemma A.4

**Proof:** From Lemma A.1, the function  $\rho$  is strictly decreasing in  $\Delta$ :

$$\rho'(\Delta) = BA < 0,$$

where the terms  $A, B$  are defined in (16) and (15). Next, using (7), we write  $\rho$  as:

$$\rho(\Delta) = BC,$$

where

$$C = \left( \sum_i \frac{p_i}{p_i\Delta + 1} \right) \left( \sum_i \frac{1 - p_i}{p_i\Delta + 1} \right).$$

Thus, the derivative of  $\zeta$  with respect to  $\Delta$  at a fixed  $(\delta_0, \delta_1)$  may be written as:

$$\begin{aligned} \zeta' &= \frac{\delta_0}{\delta_1} [\rho'(\Delta)(1 + \Delta) + \rho(\Delta)] \\ &= \frac{\delta_0}{\delta_1} B[A(1 + \Delta) + C]. \end{aligned}$$

Since both  $B$  and  $\delta_0/\delta_1$  are positive, the sign of  $\zeta'$  is the same as that of  $A(1 + \Delta) + C$ .

Recall from Lemma A.1 that

$$A = - \sum_i \sum_{\substack{j \\ j < i}} \frac{(p_i - p_j)^2}{(p_i\Delta + 1)^2(p_j\Delta + 1)^2}.$$

Furthermore,

$$\begin{aligned} C &= \sum_i \sum_j \frac{p_i(1 - p_j)}{(p_i\Delta + 1)(p_j\Delta + 1)} \\ &= \sum_i \sum_{\substack{j \\ j \neq i}} \frac{(p_i - p_i p_j)(p_i\Delta + 1)(p_j\Delta + 1)}{(p_i\Delta + 1)^2(p_j\Delta + 1)^2} + \sum_k \frac{p_k(1 - p_k)}{(p_k\Delta + 1)^2} \\ &= \sum_i \sum_{\substack{j \\ j < i}} \frac{(p_i - 2p_i p_j + p_j)(p_i\Delta + 1)(p_j\Delta + 1)}{(p_i\Delta + 1)^2(p_j\Delta + 1)^2} + \sum_k \frac{p_k(1 - p_k)}{(p_k\Delta + 1)^2}. \end{aligned}$$

Since the second term in the last equality is obviously positive, the desired conclusion obtains if:  $\tilde{A}_{ij}(\Delta + 1) + \tilde{C}_{ij} > 0$  for every  $i$  and  $j$ , where

$$\tilde{A}_{ij} = -(p_i - p_j)^2,$$

and

$$\tilde{C}_{ij} = (p_i - 2p_i p_j + p_j)(p_i \Delta + 1)(p_j \Delta + 1).$$

We find that: test

$$\begin{aligned} \tilde{A}_{ij}(\Delta + 1) + \tilde{C}_{ij} &= 4p_i p_j \Delta + p_i(1 - p_j) + p_j(1 - p_i) \\ &\quad + p_i p_j \Delta [p_i \Delta - 2p_i p_j \Delta + p_j \Delta - 2p_i - 2p_j] \\ &= 4p_i p_j \Delta + p_i(1 - p_j) + p_j(1 - p_i) \\ &\quad + p_i p_j \Delta^2 [p_i - 2p_i p_j + p_j] - 2p_i p_j \Delta [p_i + p_j] \\ &= p_i(1 - p_i) + p_j(1 - p_j) \\ &\quad + p_i p_j \Delta^2 [p_i(1 - p_j) + p_j(1 - p_i)] + 4p_i p_j \Delta \left[ 1 - \frac{p_i + p_j}{2} \right]. \end{aligned}$$

Clearly, each of the four terms in the last expression is positive, proving that  $\zeta' > 0$ . ■