Compromise, Don't Optimize: Generalizing Perfect Bayesian Equilibrium to Games with Ambiguity

Karl Schlag and Andriy Zapechelnyuk

Abstract. We introduce a solution concept for extensive-form games of incomplete information in which players need not assign likelihoods to what they do not know. This is embedded in a model in which players can hold a set of beliefs. Players make choices by looking for compromises that yield a good performance under each of their beliefs. Our solution concept is called perfect compromise equilibrium. It generalizes perfect Bayesian equilibrium. We show how it deals with uncertainty without using probabilities in Cournot and Bertrand markets, Spence's job market signaling, as well as in bilateral trade with common value.

JEL Classification: D81, D83

Keywords: compromise, multiple priors, loss, robustness, perfect Bayesian equilibrium, perfect compromise equilibrium, solution concept

Date: October 31, 2020.

Schlag: Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria. *E-mail:* karl.schlag@univie.ac.at.
Zapechelnyuk: School of Economics and Finance, University of St Andrews, Castlecliffe, the Scores, St Andrews KY16 9AR, UK. *E-mail:* az48@st-andrews.ac.uk.

We are grateful to Pierpaolo Battigalli, Simon Grant, and Clara Ponsati for their comments.

1. Introduction

Modeling lack of information is at the center stage of economics. Agents might not know the previous choice of someone else. Or they might not know the type of their opponent. Or they might not know their own payoffs.

The concept of *perfect Bayesian equilibrium* (PBE) forces an agent to specify precise likelihoods of everything they are uncertain about. However, being uncertain seems to contradict their ability to assign such probabilities. Players might be ambiguous and not willing or able to specify any probabilities. They might only be able to specify bounds on what can happen.

In this paper we introduce a solution concept for extensive-form games that does not force players to formulate priors about what they do not know. Our solution concept generalizes PBE and hence also allows players to have priors. These two extreme settings, players with priors and players without probability assessments, are included in the same model by assuming that each player makes choices using a set of beliefs. A traditional player is incorporated by having a unique belief. A player without probability assessments is incorporated by considering a set of degenerate beliefs, each of which puts all mass on a single event. Intermediate cases of ambiguity are included too, where a player has a better but not complete understanding of the environment, by assigning to this player a set of beliefs.

Our solution concept is called *perfect compromise equilibrium* (PCE). It applies to extensive-form games of incomplete information where players are allowed to be ambiguous about what they do not know. As motivated above, ambiguity is modeled by allowing players to hold multiple beliefs at each of their information sets. The solution concept is readily defined once we have resolved the following two issues. How to learn from the past? How to model decision making under ambiguity?

To capture learning from the past, a set of beliefs at a given information set has to include the beliefs derived by Bayes' rule whenever possible. There are no other constraints on the set of beliefs. This makes our concept flexible. On the one hand, one can limit attention to the updated beliefs to model a player who only learns from the past. On the other hand, one can include other beliefs to take the possibility of mistakes by other players into account. Thus, our concept bridges the discontinuous jump in PBE, where beliefs are completely arbitrary off the equilibrium path, but Bayes' rule pins down a unique belief on the equilibrium path, even if that equilibrium path is extremely unlikely.

Decision making under ambiguity is modeled as follows. A player strives to find the *best compromise*. This is a choice that balances the loss of not making the optimal decision for each of her beliefs. This criterion collapses to expected utility maximization if there is only one belief or if there is a dominant action at the given information set. The concept of best compromise follows the tradition of minimax regret and is founded on many pillars. It has an axiomatic foundation. It is close to classic expected utility maximization when there is little ambiguity, in the sense that all beliefs are similar. It can be used to justify behavior in front of people with different preferences. It formalizes a criterion that can be interpreted as the everyday notion of making a compromise.

Formally, a PCE specifies the action each player chooses at each of her information sets, together with a set of beliefs over which decision node she is at within this information set.

We are particularly interested in modeling players who are extremely ambiguous and only focus on which states are possible, without assessing their likelihoods. For instance, it seems beyond any realism to assume that firms conjecture a specific probability distribution when they think about what demand they will be facing. Yet it seems plausible that agents might use simple functional forms to put bounds on their uncertain demand. Modeling situations like this is possible within our framework by letting the set of priors consist only of degenerate priors. We call this *genuine ambiguity*. In particular, PCE can be used instead of PBE when players have difficulty forming priors. Modeling a game of incomplete information without using priors comes with numerous advantages in comparison to PBE. Solutions are often easier to obtain. They are more parsimonious as they do not change with a prior. Solutions can be more intuitive as they are simple and depend on observables and not on fictitious distributions. These advantages are demonstrated in four salient economic examples.

We investigate four examples. We consider Cournot competition with unknown demand, where firms postulate bounds on the true demand. We consider Bertrand competition where firms assess lower and upper bounds on the costs of their rivals. We consider Spence's job market where employers do not know the cost of education and the productivity of workers and they impose bounds on these parameters. In the fourth example we consider bilateral trade under common value where each party knows an interval that contains the true value.

These examples highlight the value of the PCE concept in terms of realism, good compromises, and new insights. All our examples are arguably more realistic than those found in the literature, as we do not have to confine ourselves to parametric models of uncertainty or to models with two states (high and low). Our examples involve strategic decision making under rich uncertainty where the PBE analysis is intractable. Compromise values are small in the Cournot and Bertrand competition settings. In these contexts it makes little sense to think in more detail about which state is really the true one, as payoffs would only be slightly higher in some states but could be substantially lower in other states. New insights appear. We find that adding uncertainty makes firms more competitive under Cournot competition and less competitive under Bertrand competition. In the separating equilibrium of Spence's job market signaling game, better educated workers are not necessarily more productive, unlike in the classic model with two types (Spence, 1973). In bilateral trade with common value, we find that trade is possible, as opposed to the famous no-trade theorem for PBE (Milgrom and Stokey, 1982). Under PCE the possibility that the trading partners have different valuations leads to trade with positive probability, as ignoring this possibility generates losses that the traders want to minimize. Under PBE there is no trade as the trading partners always agree on the expected valuation of the good.

Related Literature. Our paper contributes to the literature on robustness and ambiguity in games of incomplete information.

A paper that at a glance may seem very similar to ours is Hanany, Klibanoff and Mukerji (2018). They also consider general extensive form games with incomplete information. Their players have smooth ambiguous preferences (see also Klibanoff, Marinacci and Mukerji, 2005). Specifically, a player combines or aggregates different possible priors into a single belief using a distribution over these priors and a concave aggregator function. This aggregated belief is updated over time in a dynamically consistent fashion. Thus, a player has a very detailed understanding of how the different priors should be weighted. In contrast, the different priors in our model remain conceptually separated. The inability or unwillingness to combine priors is at the heart of our approach. Compromises are chosen as a way to resolve the conflict of having different possible understandings of the environment. In fact, one of the emphases of our paper is that it offers a means to get away from probability assessments. As such we explain how our approach can handle a player who wishes to capture uncertainty by assessing a set of possible states, without any use of priors.

An important ingredient of our solution concept is our use of compromise for making choices when the true state is unknown. A popular alternative approach in the literature on ambiguity is maximin preferences (Wald, 1950; Gilboa and Schmeidler, 1989). These preferences have been brought to simultaneous-move games with incomplete information and multiple priors by Epstein and Wang (1996), Kajii and Morris (1997), Kajii and Ui (2005), and Azrieli and Teper (2011). While this approach can be suitable in applications where players are pessimistic and care about payoff guarantees in the worst case instead of compromises, it leads to unintuitive results in our examples. For instance, in Bertand duopoly with ambiguity about the rival's cost, maximin utility leads firms to shut down. To obtain nontrivial results, additional structural assumptions need to be added, such as assuming knowledge of the mean state. Another approach found in literature is Knightian uncertainty with incomplete preferences. This has been used by Chiesa et al. (2015) to model bidding in auctions.

Our idea of best compromise has origins in minimax regret (Savage, 1951) and connects to approximate optimality. Our optimization criterion differs from minimax regret as evaluation occurs at each information set, while minimax regret traditionally evaluates regret ex-post. Furthermore, PCE retains the strategic reasoning of PBE, as players have certainty about each others' strategies. For an investigation of minimax regret under strategic uncertainty see Linhart and Radner (1989), and under partial strategic uncertainty see Renou and Schlag (2010).

In simultaneous-move games, PCE can be considered as a generalization of ex-post Nash equilibrium (Cremer and McLean, 1985). It can be thought of as an ε-ex-post Nash equilibrium in which the smallest possible value of ε is chosen for each player. In the context of ε-Nash equilibrium (Radner, 1980) the value of ε is interpreted a minimal level of improvement necessary to trigger a deviation. Our interpretation is different. The value of ε measures the compromise needed to accommodate all beliefs. In particular, the threshold ε is endogenous in a PCE.

PCE can be interpreted as a robust version of PBE where robustness in the sense of Huber (1965) means to make choices that also perform well if the model is slightly misspecified. Being a compromise, our suggested strategies perform well under each prior given how others make their choices, never doing too badly relative to what could be achieved under that prior. Stauber (2011) analyzes the local robustness of PBE to small degrees of ambiguity about player's beliefs. In particular, players do not adjust their play to this ambiguity, unlike our paper.

We proceed as follows. In Section 2 we introduce our solution concept. In Section 3 we illustrate PCE in four self-contained examples. Section 4 concludes. All proofs are in Appendix A. Some additional examples are in Appendix B.

2. Perfect Compromise Equilibrium

We introduce a solution concept called *perfect compromise equilibrium (PCE).* A formal definition is presented in Section 2.1 below. A reader who wishes to be spared with the formalities and seeks to understand the essence of PCE and its applicability can jump to Section 3 that presents self-contained examples.

2.1. Formal Setting.

Consider a finite extensive-form game described by $(N, \mathcal{G}, \Omega, \Pi, u)$, where $N = \{1, ..., n\}$ is a set of players, $\mathcal{G}$ is a finite game tree, Ω is a finite set of states, $\Pi = (\Pi_1, ..., \Pi_n)$ is a profile of sets of priors, where $\Pi_i \subset \Delta(\Omega)$ is a set of priors of player i, and $u = (u_1, ..., u_n)$ is a profile of payoff functions.

Game tree $\mathcal{G}$ describes the order of players' moves, players' information sets, and actions that are available at each information set. It is defined by a set of linked decision nodes and terminal nodes that form a tree. Each decision node is assigned three elements: a player i, an information set ϕ_i, and a set of actions available to player i at that information set. Information set ϕ_i is a set of all the decision nodes that player i cannot distinguish. Information sets and action sets satisfy the standard assumptions of games with perfect recall. Let ϕ_0 be the initial decision node of the game, let Φ_i be the set of all information sets of player i for each $i \in N$, and let $\mathcal{T}$ be the set of terminal nodes of the game. Let $A(\phi_i)$ be a finite set of actions available at an information set ϕ_i.

In the spirit of Harsanyi (1967), all incomplete information is captured by a move of nature at the beginning of the game. Nature moves only once, at the initial decision node ϕ_0. An action of nature ω is called *state* and is chosen from the set of states Ω.

The game terminates after finitely many moves at some terminal node, and players obtain payoffs. A payoff function of each player $i \in N$ specifies the payoff $u_i(\tau)$ of player i at each terminal node $\tau \in \mathcal{T}$.

A strategy of player $i \in N$ prescribes a mixed action $s_i(\phi_i)$ for each information set $\phi_i \in \Phi_i$, so $s_i(\phi_i) \in \Delta(A(\phi_i))$. A strategy profile s describes the behavior of all players throughout the game.

Like in Bayesian games, we specify not only strategies, but also beliefs of the players at each of their information sets. The crucial difference from a Bayesian game is that, in our setting, the players may have multiple beliefs in each of their information sets. For each player i and each an information set $\phi_i \in \Phi_i$, a belief β_i specifies a probability distribution over the decision nodes in ϕ_i, so $\beta_i \in \Delta(\phi_i)$. Let $B_i(\phi_i)$ be a nonempty set of beliefs that player i holds at information set ϕ_i. We will refer to elements of $B_i(\phi_i)$ as *speculated beliefs* at ϕ_i. With a slight abuse of notation, let us consider the set of priors of a player i as her speculated beliefs at node ϕ_0, so $B_i(\phi_0) = \Pi_i$. Let $B = (B_1, ..., B_N)$ be the profile of speculated beliefs for all players at all the information sets.

Like in PBE, we will require consistency of beliefs. Let $\phi_i \in \Phi_i$. We say that player i's information set ϕ_i' *is preceding* ϕ_i if ϕ_i' is the last information set where player i moved before ϕ_i.[1] If player i moves at ϕ_i for the very first time, then $\phi_i' = \phi_0$.

Definition 1. A set of speculated beliefs $B_i(\phi_i)$ at an information set ϕ_i is called *consistent* under a strategy profile s if $B_i(\phi_i)$ contains every distribution over the nodes in ϕ_i that is derived by Bayes' rule from some belief β_i' that is contained in the set of speculated beliefs at the preceding information set ϕ_i', so $\beta_i' \in B_i(\phi_i')$.

A profile B of speculated beliefs is *consistent* with a strategy profile s if for each $i \in N$ and each $\phi_i \in \Phi_i$ the set of speculated beliefs $B_i(\phi_i)$ at ϕ_i is consistent under s.

Note that our definition of consistency does not impose any discipline on the out-of-equilibrium beliefs. If an information set ϕ_i cannot be reached from the preceding information set ϕ_i' under a given strategy profile s, then every nonempty set of beliefs at ϕ_i is consistent under s. Of course, not every choice of out-of-equilibrium beliefs can be sensible in applications. This is the very same problem that emerged in the context of PBE and gave rise to a vast literature on PBE refinements. This problem is of equally high importance for PCE. However, addressing this problem would take us away from the main messages of this paper. Neither the idea of PCE, nor its properties in the examples considered in this paper change if additional assumptions about out-of-equilibrium beliefs are made. So, we leave this question for future research.

Next we define how decisions are made at an information set ϕ_i. We fix a strategy profile s and determine how to make a choice at ϕ_i, while keeping choices at all other information sets fixed. The difficulty of making a decision at ϕ_i is that the player does not know which belief in the set of speculated beliefs $B_i(\phi_i)$ should be used to evaluate the expected payoff. We follow that minimax regret approach.

[1]Perfect recall implies that there is at most one preceding information set.

Roughly speaking, the player chooses a compromise action that is never too far from the best action under each belief in $B_i(\phi_i)$.

Formally, consider a profile (s, B) of strategies and sets of speculated beliefs. Denote by $\bar{u}_i(x_i|\phi_i, s, \beta_i)$ the expected payoff of player i from choosing a mixed action $x_i \in \Delta(A(\phi_i))$ in an information set ϕ_i under the speculated belief β_i over the decision nodes in ϕ_i, assuming that the play is given by s elsewhere in the game. The payoff difference

$$\sup_{a_i \in A(\phi_i)} \bar{u}_i(a_i|\phi_i, s, \beta_i) - \bar{u}_i(x_i|\phi_i, s, \beta_i)$$

is called player i's *loss* from choosing mixed action x_i at information set ϕ_i given speculated belief β_i. It describes how much better off player i could have been at this information set given this belief if, instead of choosing x_i, she had chosen the best action, assuming that the actions in all other information sets are prescribed by s. The *maximum loss* of player i from choosing a mixed action x_i in an information set ϕ_i under (s, B) is given by

$$l(x_i|\phi_i, s, B) = \sup_{\beta_i \in B_i(\phi_i)} \left(\sup_{a_i \in A(\phi_i)} \bar{u}_i(a_i|\phi_i, s, \beta_i) - \bar{u}_i(x_i|\phi_i, s, \beta_i) \right).$$

So the maximum (supremum) is sought over all speculated beliefs of player i at ϕ_i.

The player making a decision seeks to minimize the maximum loss. Such a choice is called a *best compromise.* Formally she chooses

$$s_i(\phi_i) \in \operatorname*{arg\,min}_{x_i \in \Delta(A(\phi_i))} l(x_i|\phi_i, s, B), \tag{1}$$

We now formulate our equilibrium concept that is based on the ideas of best compromises and consistent beliefs.

Definition 2. A profile (s, B) is called a *perfect compromise equilibrium* if
(a) each player chooses a best compromise in each of her information sets;
(b) profile B of speculated beliefs is consistent under strategy profile s.

Before proceeding any further, let us establish the existence of PCE.

Theorem 1. *A perfect compromise equilibrium exists.*

The proof is in Appendix A.1.

Remark 1. In some applications, it is unrealistic to assume that players can choose mixed actions. Our definition of PCE can be easily adjusted if players are only allowed to use pure actions. In this case, each player minimizes her maximal loss among her pure actions, so instead of (1) we require

$$s_i(\phi_i) \in \operatorname*{arg\,min}_{a_i \in A(\phi_i)} l(a_i|\phi_i, s, B). \tag{1'}$$

Remark 2. In applications, there can be a continuum of strategies and states, and game trees can be infinite. The definition of PCE readily extends to such settings, but some additional assumptions have to be made to ensure existence.

2.2. **Genuine Ambiguity.** We are particularly interested in understanding strategic play under uncertainty without using probabilities. This is possible within the above framework as we now demonstrate. Consider players who do not use mixed strategies and who cannot or are unwilling to assess the likelihood of different states and of decision nodes within information sets. We call this *genuine ambiguity*. Formally, players are limited to pure strategies as outlined in Remark 1, they can only have degenerate priors, and they can assign only degenerate speculated beliefs to decision nodes in each of their information sets. All our examples presented in Section 3 deal with genuine ambiguity.

In the presence of genuine ambiguity, it is natural to offer an equivalent representation of our concepts that do not involve distributions. There are two small changes. First, we consider PCE where players are only allowed to choose among their pure actions. Second, we consider PCE in which all speculated beliefs are degenerate. Instead of talking about speculated beliefs, we can then talk about speculated decision nodes. A decision node is called speculated at a given information set if the degenerate belief that puts all weight on this node is speculated. Note that Bayes' rule applied to a degenerate prior and a pure strategy profile will generate a degenerate belief at an information set that can be reached. Hence the consistency requirement given in Definition 1 is well defined.

2.3. **Discussion of Perfect Compromise Equilibrium.** We highlight some properties of PCE.

Best Compromise. Our decision making criterion for how to make choices at a given information set captures the intuitive notion of making a compromise. As a compromise, the performance should be satisfactory in all potential situations, as opposed to being best under some and, possibly, very bad under others. The concept of best compromise identifies the smallest maximal distance from first best as a measure of how large the compromise has to be. Compromises are valuable when decisions have to be justified in front of others who have heterogeneous perceptions about the environment.

The concept of a best compromise follows the tradition of decision making under minimax regret, thus having an axiomatic underpinning (Milnor, 1954; Stoye, 2011). Traditionally, minimax regret is evaluated ex-post after all uncertainty is resolved. In contrast, to model a compromise in the face of several beliefs, we measure loss ex-ante for a given belief. Stoye's (2011) axioms continue to hold from this ex-ante viewpoint. Furthermore, our concept retains the strategic reasoning of PBE, as players know each others' strategies. This is unlike Linhart and Radner (1989) who reduce the game to an individual decision problem, where the behavior of the others is a part of unknown nature.

Clearly, instead of best compromise, any other decision making criterion under ambiguity could be used for determining choices at information sets. For instance, the maximin utility criterion can be used to model pessimism or cautiousness, a world in which the player always anticipates the worst outcome.

PCE vs PBE. Our definition of PCE generalizes the concept of PBE to games where some players may be ambiguous about what they do not know. When there is no ambiguity, so there is a single speculated belief at each information set, then our setting describes a standard game of incomplete information. In this case, the loss minimization objective, as described in (2), reduces to the standard utility maximization objective. So, an action minimizes the maximum loss of a player if and only if it is a best response. Moreover, whenever there is only a single speculated belief, the consistency requirement introduced in Definition 1 reduces to the standard Bayesian consistency of beliefs. Hence, PCE becomes PBE.

The difference between PCE and PBE emerges in models where some players are ambiguous about the occurrence of an event. The standard PBE approach forces players to quantify the uncertainty by specifying a unique belief at each information set, and then assuming that the players optimize with respect to these beliefs. Our approach sidesteps this issue by letting the players have multiple beliefs at each information set and find compromises with respect to these beliefs.

Ex-post Nash equilibrium. In simultaneous move games there is a relationship of PCE to the concept of an ex-post Nash equilibrium. Ex-post Nash equilibria are profiles that are Nash equilibria in the game in which the state is observed by all players at the outset of the game. This means that the maximum loss of each player at her single information set is equal to zero. Consequently, any ex-post Nash equilibrium is also a PCE. Note, however, that ex-post Nash equilibria often do not exist.

Dominance. A PCE survives the elimination of strictly dominated strategies, as we now demonstrate. We say that an action $a_i \in A(\phi_i)$ at an information set ϕ_i is *strictly dominated* for player i if there exists a mixed action $x_i \in \Delta(A(\phi_i))$ such that player i's payoff from choosing a_i is strictly worse than that from choosing x_i, regardless of the state $\omega \in \Omega$ and of the choices of other players at any of their information sets. Iterated dominance is defined as usual. After having excluded actions that were strictly dominated in previous rounds, one checks the dominance condition w.r.t. the remaining actions of each player. Now observe that if an action a_i at some information set ϕ_i is strictly dominated, then it cannot be a best compromise at this information set. This is because the (mixed) action that strictly dominates a_i will achieve a strictly lower loss for each speculated belief, and hence its maximal loss will be strictly smaller. Thus, a strictly dominated action cannot be a part of any PCE. This argument can be iterated, so any iterated strictly dominated action cannot be a part of any PCE.

3. Examples

We illustrate our solution concept in a few applications that are prominent in the literature. We consider Cournot and Bertrand duopoly, Spence's job market signaling, and bilateral trade with common value. The examples presented in this section are self-contained as they do not require knowledge of the formalities presented in Section 2. Further examples on public good provision and forecasting can be found in Appendix B.

In the applications that we consider, uncertainty is traditionally incorporated in a very simple fashion, often only considering two states, high and low. We consider richer sets of uncertain events in order to capture more realistic uncertainty. In our examples we specify uncertainty in terms of bounds on what the players do not know. Probability distributions do not play a role. Players do not have beliefs. Instead, they speculate about which state is true or about what decision node within an information set they are at. Moreover, players do not use mixed strategies. They search among their pure strategies for a best compromise. Thus we perform a strategic analysis without using probabilities, which is referred to as genuine ambiguity in Section 2.2.

To maintain an intuitive appeal, we describe the speculation of the players about what they do not know in terms of what is payoff relevant to them. Thus we avoid dealing with details included in the description of states and decision nodes that are not relevant for the players' decision making.

3.1. **Cournot Duopoly with Unknown Demand.** We investigate how two firms compete in quantities when neither firm knows the demand.

There are two firms that produce a homogeneous good. For clarity of exposition, we assume that there are no costs of production. Each firm $i = 1, 2$ chooses a number of units $q_i \geq 0$ to produce. Choices are made simultaneously. The firms face an inverse demand function given by $P(q_1 + q_2)$. Each firm i's profit is given by

$$u_i(q_i, q_{-i}; P) = P(q_i + q_{-i})q_i, \quad i = 1, 2.$$

Neither firm knows the inverse demand P, but they know that it belongs to a set $\mathcal{P}$ given as follows. Let

$$\underline{P}(q) = \underline{a} - \underline{b}q \quad \text{and} \quad \bar{P}(q) = \bar{a} - \bar{b}q, \quad \text{where} \quad \bar{a} \geq \underline{a} > 0 \quad \text{and} \quad \bar{a}/\bar{b} \geq \underline{a}/\underline{b} > 0.$$

Let $\mathcal{P}$ be the set of inverse demand functions that satisfy

$$\begin{aligned} &P(q) \text{ is continuously differentiable in } q, \\ &\underline{P}(q) \leq P(q) \leq \bar{P}(q) \quad \text{and} \quad \underline{P}'(q) \leq P'(q) \leq \bar{P}'(q). \end{aligned} \tag{2}$$

A firm i's *maximum loss* of choosing quantity q_i when the other firm chooses quantity q_{-i} is given by

$$l_i(q_i, q_{-i}) = \sup_{P \in \mathcal{P}} \left(\sup_{q_i' \geq 0} u_i(q_i', q_{-i}; P) - u_i(q_i, q_{-i}; P) \right).$$

The maximum loss describes how much more profit firm i could have obtained if it had known the inverse demand P when anticipating the other firm to produce q_{-i}. Firm i's *best compromise* given a choice q_{-i}^* of the other firm is a quantity q_i^* that achieves the lowest maximum loss, so

$$q_i^* \in \arg\min_{q_i \geq 0} l_i(q_i, q_{-i}).$$

A strategy profile (q_1^*, q_2^*) is a *perfect compromise equilibrium* if each firm chooses a best compromise given the choice of the other firm.

Proposition 1. *There exists a unique perfect compromise equilibrium. In this PCE, the strategy profile (q_1^*, q_2^*) is given by*

$$q_i^* = \frac{1}{3\left(\sqrt{\underline{b}} + \sqrt{\bar{b}}\right)} \left(\frac{\underline{a}}{\sqrt{\underline{b}}} + \frac{\bar{a}}{\sqrt{\bar{b}}} \right), \quad i = 1, 2. \tag{3}$$

The associated maximum losses are

$$l_i(q_i^*, q_{-i}^*) = \frac{(\underline{a}\bar{b} - \bar{a}\underline{b})^2}{4\underline{b}\bar{b}\left(\sqrt{\underline{b}} + \sqrt{\bar{b}}\right)^2}, \quad i = 1, 2. \tag{4}$$

The proof is in Appendix A.2.

Let us discuss the strategic concerns underlying the PCE in this game. Each firm i, when facing unknown inverse demand and deciding about the quantity to produce, worries about two possible situations. It could be that the inverse demand is actually very high, so the firm is losing profit by producing too little. The greatest such loss occurs when the inverse demand is the highest, so $P = \bar{P}$. Alternatively, it could be that the inverse demand is actually very low, so the firm is losing profit by producing too much. The greatest such loss occurs when the inverse demand is the lowest, so $P = \underline{P}$. The firm thus chooses the best compromise q_i^* that balances these two losses, assuming that the other firm follows its equilibrium strategy q_{-i}^*.

Remark 3. It is generally intractable to find a PBE in this game with such a rich set of possible inverse demand functions. It can only be done under very specific priors about the inverse demand. For example, PBE can be found if a prior describes the uncertainty about the parameters of the linear inverse demand function $P(q) = a - bq$ (Vives, 1984).

Remark 4. Our equilibrium analysis can shed light on how the firms' behavior changes in response to increasing uncertainty. For comparative statics, let us consider as a benchmark a linear inverse demand function $P_0(q) = a_0 - b_0 q$. We normalize constants a_0 and b_0 so that the monopoly profit is equal to 1, that is,

$$\sup_{q \geq 0}(a_0 - b_0 q)q = \frac{a_0^2}{4b_0} = 1.$$

Suppose that there is a small uncertainty. Specifically, for $\varepsilon > 0$ let $P(q)$ satisfy (2) where

$$\underline{P}(q) = \left(1 - \frac{\varepsilon}{2}\right) a_0 - \left(1 + \frac{\varepsilon}{2}\right) b_0 q \quad \text{and} \quad \bar{P}(q) = \left(1 + \frac{\varepsilon}{2}\right) a_0 - \left(1 - \frac{\varepsilon}{2}\right) b_0 q.$$

Denote by $q^\varepsilon = (q_1^\varepsilon, q_2^\varepsilon)$ the strategies of the PCE as given by Proposition 1. We then obtain

$$\frac{\mathrm{d}q_i^\varepsilon}{\mathrm{d}\varepsilon} = \frac{2\varepsilon}{3a_0} + O(\varepsilon^3) > 0.$$

So the firms optimally respond to a growing uncertainty about the demand by increasing their output, and do so at an increasing rate as ε grows. Next, consider the associated maximum losses as shown in (4). Then

$$l_i(q_i^\varepsilon, q_{-i}^\varepsilon) = \varepsilon^2 + O(\varepsilon^4), \quad i = 1, 2.$$

So the maximum losses in the PCE increase very slowly as uncertainty increases. Moreover, if $\varepsilon = 0.1$, then $l_i(q_i^\varepsilon, q_{-i}^\varepsilon) \approx 0.01$. So the firms lose no more than about 1% of the maximum profit due to not knowing the demand.

3.2. **Bertrand Duopoly with Private Costs.**

We now consider how two firms compete in prices when the cost of the rival firm is unknown.

There are two firms that produce a homogeneous good. Each firm $i = 1, 2$ chooses a price p_i. Choices are made simultaneously. The consumers only buy from the firm that offers a lower price. In particular, the quantity that firm i sells is given by

$$q_i(p_i, p_{-i}) = \begin{cases} Q(p_i), & \text{if } p_i < p_{-i}, \\ Q(p_i)/2, & \text{if } p_i = p_{-i}, \\ 0, & \text{if } p_i > p_{-i}, \end{cases}$$

where $Q(p)$ is the demand function. For clarity of exposition we assume that the demand function is given by

$$Q(p) = \max\left\{\frac{a-p}{b}, 0\right\}$$

The cost of producing q_i units is $c_i q_i$. Each firm i's profit is given by

$$u_i(p_i, p_{-i}; c_i) = (p_i - c_i) q_i(p_i, p_{-i}), \quad i = 1, 2.$$

Each firm i knows her own marginal cost but not that of the other firm, and it is common knowledge that

$$c_1, c_2 \in [\underline{c}, \bar{c}], \quad \text{where } 0 \le \underline{c} \le \bar{c} \le a/2.$$

A firm i's pricing strategy $s_i(c_i)$ describes its choice of the price given its marginal cost c_i.

For each marginal cost c_i, firm i's *maximum loss* of choosing a price p_i when facing pricing strategy s_{-i} of the other firm is given by

$$l_i(p_i, s_{-i}; c_i) = \sup_{c_{-i} \in [\underline{c}, \bar{c}]} \left(\sup_{p_i' \ge 0} u_i(p_i', s_{-i}(c_{-i}); c_i) - u_i(p_i, s_{-i}(c_{-i}); c_i) \right).$$

The maximum loss describes how much more profit i could have obtained if it had known the other firm's marginal cost c_{-i}, anticipating the other firm to follow the pricing strategy s_{-i}. Firm i's *best compromise* given c_i is a pricing strategy $s_i^*(c_i)$ that achieves the lowest maximum loss for a given strategy s_{-i}^* of the other firm:

$$s_i^*(c_i) \in \arg\min_{p_i \ge 0} l_i(p_i, s_{-i}^*; c_i).$$

A strategy profile (s_1^*, s_2^*) is a *perfect compromise equilibrium* if each firm i chooses a best compromise given its marginal cost c_i when facing the strategy s_{-i}^* of the other firm.

Proposition 2. *There exists a unique perfect compromise equilibrium. In this PCE, the pricing strategies are given by*

$$s_i^*(c_i) = \frac{1}{2}\left(a + c_i - \sqrt{(a-\bar{c})^2 + (\bar{c} - c_i)^2}\right), \quad i = 1, 2. \tag{5}$$

The associated maximum losses are

$$l_i(s_i^*(c_i), s_{-i}^*, c_i) = \frac{(a-\bar{c})(\bar{c}-c_i)}{2} \leq \frac{(a-\bar{c})(\bar{c}-\underline{c})}{2}, \quad i = 1, 2. \tag{6}$$

The proof is in Appendix A.3.

Let us discuss the strategic concerns underlying the PCE in this game. Each firm i, when deciding about the price $p_i > c_i$ and facing an unknown cost of the other firm, worries about two possible situations. It could be that the other firm chooses a weakly lower price $p_{-i} \leq p_i$. Thus, firm i could have obtained more profit by undercutting p_{-i}. The greatest such loss occurs when the other firm's price marginally undercuts p_i. Alternatively, it could be that the other firm chooses a higher price, $p_{-i} > p_i$. Thus, unless p_i is the profit maximizing price for the monopoly, firm i is losing profit by charging too little. The greatest such loss occurs when the other firm's cost is the highest possible, $\bar{c}$. The firm thus chooses the best compromise $s_i^*(c_i)$ that balances these two losses, assuming that the other firm follows its equilibrium strategy.

We find that the PCE price $s_i^*(c_i)$ is strictly increasing in c_i and lies strictly above the marginal cost c_i whenever $c_i < \bar{c}$. Moreover, $s_i^*(\bar{c}) = \bar{c}$. So, any sale with the cost below $\bar{c}$ leads to a positive profit. The fact that the equilibrium price cannot not lie above $\bar{c}$ is intuitive. It is common knowledge that the costs are at most $\bar{c}$. So if a firm charges a price above $\bar{c}$, the other firm would undercut it. Note also that the largest equilibrium price cannot lie below $\bar{c}$. This is because a firm with cost $\bar{c}$ will never charge a price below $\bar{c}$.

Note that the lowest equilibrium price $s_i^*(\underline{c})$ is strictly positive, even if $\underline{c} = 0$. This is because when the price is very low, then the potential loss due to not undercutting the other firm is small, while the potential loss due to not setting a price much higher is large. This has an upward effect on prices.

Remark 5. It is generally intractable to find a PBE in this application under any reasonable prior, even in this simplest setting with linear demand and constant marginal costs. The PBE strategy profile for this simplest setting is implicitly defined by a differential equation with no closed form solution (see Spulber, 1995).

Remark 6. As in Section 3.1, our equilibrium analysis can shed light on how the firms' behavior changes in response to increasing uncertainty. For comparative statics, let us consider as a benchmark marginal cost $c_0 = a/4$ (recall that we require $0 \leq c_i \leq a/2$, so $c_0 = a/4$ is the midpoint). We normalize the constants

a and b of the demand function $Q(p) = (a - p)/b$ so that the monopoly profit is equal to 1, that is,

$$\sup_{p\geq 0}(p - c_0)\frac{a-p}{b} = \frac{(a-c_0)^2}{4b} = 1.$$

Suppose that there is a small uncertainty. Specifically, for $0 < \varepsilon < 1$ let $c_i \in [\underline{c}, \bar{c}]$, $i = 1, 2$, where

$$\underline{c} = \left(1 - \frac{\varepsilon}{2}\right) c_0 \quad \text{and} \quad \bar{c} = \left(1 + \frac{\varepsilon}{2}\right) c_0.$$

Denote by $s^\varepsilon = (s_1^\varepsilon, s_2^\varepsilon)$ the PCE strategy profile as given by Proposition 2. We then obtain

$$\frac{\mathrm{d}s_i^\varepsilon(c_i)}{\mathrm{d}\varepsilon} = \frac{(a + c_i - 2\bar{c})c_0}{4\sqrt{(a-\bar{c})^2 + (\bar{c} - c_i)^2}} > 0,$$

because, using our assumptions on the parameters,

$$a + c_i - 2\bar{c} \geq a - 2\bar{c} = 1 - 2\left(1 + \frac{\varepsilon}{2}\right) c_0 = \frac{1}{4}(2 - \varepsilon) > 0.$$

So the firms optimally respond to the growing uncertainty about the demand by increasing their prices. They become less competitive. Next, consider the associated maximum losses as shown in (6). Then

$$l_i(s_i^\varepsilon(c_i), s_{-i}^\varepsilon, c_i) \leq \frac{3\varepsilon}{32} - \frac{\varepsilon^2}{64}, \quad i = 1, 2.$$

So the maximum losses are small. For example, if $\varepsilon = 0.1$, then the maximum losses are bounded by 0.01. So the firms lose no more than about 1% of the maximum profit due to not knowing the cost of the other firm.

3.3. **Job Market Signaling.** Here we investigate Spence's job market signaling (Spence, 1973) when the worker's productivity and cost of education are unknown to the firms.

There is a single worker and two firms. The worker has productivity θ with $\theta \in [0, 1]$. The worker publicly chooses a level of education e, either low (e_L) or high (e_H), to signal her productivity to the firms. The cost of low education is zero. The cost of high education is c with $c \geq 0$. The firms observe the worker's education level e and simultaneously offer wages w_1 and w_2. The worker chooses the better of the two wages. Her payoff is given by

$$v(w_1, w_2, e; \theta, c) = \max\{w_1, w_2\} - \begin{cases} 0, & \text{if } e = e_L, \\ c, & \text{if } e = e_H. \end{cases}$$

Each firm i's payoff is given by

$$u_i(w_i, w_{-i}; \theta) = \begin{cases} \theta - w_i, & \text{if } w_i > w_{-i}, \\ (\theta - w_i)/2, & \text{if } w_i = w_{-i}, \\ 0, & \text{if } w_i < w_{-i}. \end{cases}$$

The worker knows her productivity type θ and her cost of high education c. The firms know neither. They only know that the worker can have any productivity θ in $[0, 1]$ and that her cost of high education c lies between two linearly decreasing functions of θ. Specifically, c is between $1 - b\theta$ and $1 - b\theta + \delta$, where b and δ are parameters that satisfy $0 \leq \delta \leq b \leq 1$. Formally, the firms know that (θ, c) belongs to the set Ω given by

$$\Omega = \{(\theta, c) : \theta \in [0, 1] \text{ and } c \in [1 - b\theta, 1 - b\theta + \delta].\} \tag{7}$$

The worker's strategy $e^*(\theta, c)$ describes her choice of the education level for each pair $(\theta, c) \in \Omega$. Each firm i's strategy $w_i^*(e)$ describes its wage offer conditional on each education level $e \in \{e_L, e_H\}$.

Consider how a firm makes inference from the observed level of education of the worker. This is formalized with the notion of speculated states. These are the pairs (θ, c) that a firm thinks are possible after observing the education level of the worker. The set of speculated states is denoted by $S_i(e)$. This set is *consistent* with the worker's equilibrium strategy e^* if it includes all pairs (θ, c) under which the worker chooses $e \in \{e_L, e_H\}$, so $(\theta, c) \in S_i(e)$ if $e^*(\theta, c) = e$.

For each education level e, firm i's *maximum loss* of choosing wage w_i when the other firm chooses the wage according to its strategy w_{-i}^* is given by

$$l_i(w_i, w_{-i}^*; e) = \sup_{(\theta,c) \in S_i(e)} \left(\sup_{w_i' \geq 0} u_i(w_i', w_{-i}^*(e); \theta) - u_i(w_i, w_{-i}^*(e); \theta) \right).$$

The maximum loss describes how much more profit firm i could have obtained if it had known the true productivity and cost of education of the worker, anticipating that the other firm follows its strategy w_{-i}^*. Firm i's *best compromise* given e is a wage $w_i^*(e)$ that achieves the lowest maximum loss for a given strategy w_{-i}^* of the other firm:

$$w_i^*(e) \in \arg\min_{w_i \geq 0} l_i(w_i, w_{-i}^*; e). \tag{8}$$

Observe that the worker has complete information. There is no need for a compromise. So, the worker simply chooses a best-response:

$$e^*(\theta, c) \in \arg\max_{e \in \{e_L, e_H\}} v(w_1^*(e), w_2^*(e), e; \theta, c). \tag{9}$$

A profile $(e^*, w_1^*, w_2^*, S_1, S_2)$ of strategies and speculated sets is a *perfect compromise equilibrium* (PCE) if two conditions hold. First, the strategies satisfy (8) and (9), so each firm i chooses a best compromise, and the worker chooses a best response to the strategies of the others. Second, the firms' sets of speculated states are consistent with the worker's strategy e^*.

A PCE is *pooling* if the worker chooses the same level of education for all $(\theta, c) \in \Omega$. A PCE is *separating* if the set Ω can be partitioned into two subsets such that worker types belonging to the same subset choose the same level of education, but these levels differ between the two subsets.

Proposition 3. *(i) There exists a pooling PCE in which the worker chooses low education, so*

$$e^*(\theta, c) = e_L \quad \text{for all } (\theta, c) \in \Omega,$$

and the firms' wages are given by

$$w_i^*(e_H) = w_i^*(e_L) = \frac{1}{2}, \quad i = 1, 2.$$

After each observed education level e, each firm i's set of speculated states $S_i(e)$ contains all states.

(ii) If $\delta \geq 2b^2 - b$, then a separating PCE does not exist.

(iii) If $\delta < 2b^2 - b$, then there exists a separating PCE in which the worker chooses high education if and only if her cost c is at most $\frac{1}{2b}(b-\delta)$, so for all $(\theta, c) \in \Omega$

$$e^*(\theta, c) = \begin{cases} e_H, & \text{if } c \leq \frac{1}{2b}(b-\delta), \\ e_L, & \text{if } c > \frac{1}{2b}(b-\delta), \end{cases}$$

and the firms' wages are given by

$$w_i^*(e_H) = \frac{1}{2} + \frac{b+\delta}{4b^2} \quad \text{and} \quad w_i^*(e_L) = \frac{\delta}{2b} + \frac{b+\delta}{4b^2}, \quad i = 1, 2. \tag{10}$$

After each observed education level e, each firm i's set of speculated states $S_i(e)$ contains each state $(\theta, c) \in \Omega$ that satisfies

$$\theta \in \left[0, \frac{b+\delta}{2b^2} + \frac{\delta}{b}\right] \quad \text{if } e = e_L, \quad \text{and} \quad \theta \in \left[\frac{b+\delta}{2b^2}, 1\right] \quad \text{if } e = e_H. \tag{11}$$

The proof is in Appendix A.4.

Let us discuss the strategic concerns underlying these PCE. Each firm i, when deciding about the wage offer w_i and facing unknown productivity of the worker, worries about two possible situations. It could be that the productivity is high, so offering a wage that is marginally greater than that of the competitor would improve profit. The greatest such loss occurs when the productivity is the highest possible. Alternatively, it could be that the productivity is low, so offering a wage that is smaller than the competitor's would eliminate the loss. The greatest such loss occurs when the productivity is the lowest possible. The firm thus offers the best compromise wage that balances these two losses, assuming that the other firm follows its equilibrium strategy. In equilibrium, both firms offer the same wage, so each of them has probability 1/2 to hire the worker. Hiring with probability 1/2 is the best compromise between not hiring a productive worker and hiring an unproductive worker.

An essential detail in the above considerations is that the greatest and smallest productivities are now endogenous and can depend on the level of education e that the worker chooses. In the pooling equilibrium, $e = e_L$ does not provide any useful information, so all productivity types are possible. However, in the separating equilibrium, the firms believe that the productivity belongs to a different interval when observing a different level of education. For example, if $b = 1$ and $\delta = 1/4$, then the firms believe that $\theta \in [0, 7/8]$ if the education is low, and that $\theta \in [5/8, 1]$ if the education is high.

Observe that, among the workers with productivity $\theta \in [5/8, 7/8]$, some choose low education, while others choose high education. This overlap is due to the richness of the state space. The same productivity type θ can have different costs of education c that can fall below or above the threshold at which high education is profitable. Clearly, this result cannot emerge in the traditional setting where the workers are differentiated only by their productivity.

The parameter δ captures the firms' uncertainty about the worker's cost of high education given her productivity type. As δ goes up, this range of costs increases. When δ is sufficiently large, education signaling is not very informative. A costly signal cannot be used to differentiate high and low productivity types, and the separating PCE does not exist.

3.4. **Bilateral Trade with Common Value.** We now examine bilateral trade with common value. In this example we show that trade can occur when traders follow a PCE. This is in stark contrast to the no-trade theorem under common values as predicted by PBE (Milgrom and Stokey, 1982).

A seller wants to sell an indivisible good to a buyer. The value v of the good is the same for each of them. If the good is traded at some price p, then the buyer obtains $v - p$ and the seller obtains $p - v$. If the good is not traded, then both traders obtain zero.[2]

Neither trader knows v. Before the trade takes place, the traders privately consults independent experts to obtain some information about v. Each expert provides an interval of possible values, from the most pessimistic to the most optimistic assessment of the true value. Specifically, the seller privately learns that $v \in [x_0, x_1]$ and the buyer privately learns that $v \in [y_0, y_1]$.

The traders commonly know the lower and upper bounds of the value v. These bounds are normalized to be 0 and 1, so $v \in [0, 1]$. In addition, the traders commonly know that the experts cannot be wrong, so

$$v \in [x_0, x_1] \cap [y_0, y_1]. \tag{12}$$

We do not impose constraints on how precise or imprecise the experts' information is. We allow $[x_0, x_1]$ and $[y_0, y_1]$ to be arbitrary intervals contained in $[0, 1]$ that satisfy (12).

We consider a take-it-or-leave-it protocol in which the seller is the proposer. The protocol is as follows. First, the traders observe their private information $[x_0, x_1]$ and $[y_0, y_1]$. Then the seller asks a price $p \in [0, 1]$. Finally, the buyer decides whether to accept or to reject the seller's asked price.

Let us describe the traders' strategies. Let $p^*(x_0, x_1)$ be the seller's asked price given her information $[x_0, x_1]$. Let $\alpha^*(p, y_0, y_1)$ be the buyer's decision whether to accept or to reject the asked price p given the buyer's private information $[y_0, y_1]$, where $\alpha^*(p, y_0, y_1) = 1$ means to buy, and $\alpha^*(p, y_0, y_1) = 0$ means not to buy.

Next we describe how the buyer makes inference from the price asked by the seller. This is formalized with the concept of speculated values. These are values for v that the buyer thinks are possible after he observes the price asked by the

[2]The same analysis applies if the seller obtains p when the good is sold and v when the good is not sold.

seller. Let $V_b(p, y_0, y_1)$ be the buyer's set of speculated values when the seller asks price p. Clearly, the buyer rules out the values outside of $[y_0, y_1]$, so $V_b(p, y_0, y_1) \subset [y_0, y_1]$. But some values in $[y_0, y_1]$ may be ruled out too, because $p = p^*(x_0, x_1)$ depends on x_0 and x_1, and the buyer knows that $v \in [x_0, x_1] \cap [y_0, y_1]$.

The buyer's maximum loss from his choice $\alpha \in \{0, 1\}$, given the asked price p and his set of speculated values $V_b(p, y_0, y_1)$, is

$$l_b(\alpha; p, y_0, y_1) = \sup_{v \in V_b(p,y_0,y_1)} \big(\max\{v - p, 0\} - (v - p)\,\alpha \big).$$

It describes how much more the buyer could have obtained if he knew the true value v. The seller's maximum loss of asking price p, given the buyer's acceptance strategy α^*, is

$$l_s(p; x_0, x_1) = \sup_{\substack{(v,y_0,y_1)\in[0,1]^3:\\ v\in[x_0,x_1]\cap[y_0,y_1]}} \left(\sup_{p'\in[0,1]} (p' - v)\,\alpha^*(p', y_0, y_1) - (p - v)\,\alpha^*(p, y_0, y_1) \right).$$

It describes how much more the seller could have obtained if she knew both v and the buyer's private information $[y_0, y_1]$, anticipating that the buyer would follow his strategy α^*. Each trader's *best compromise* is a choice that achieves the lowest maximum loss for a given strategy of the other trader. A strategy profile (p^*, α^*) is a *perfect compromise equilibrium* (PCE) if each trader chooses a best compromise given the strategy of the other trader.

Proposition 4. *A perfect compromise equilibrium is as follows. The seller asks*

$$p^*(x_0, x_1) = \max\left\{ \frac{x_0 + x_1}{2} + \frac{1 - x_1}{4}, \frac{1}{2} \right\}. \tag{13}$$

If the seller is asks $p \geq \frac{1}{2}$, then the buyer speculates that $v \in [\max\{y_0, 2p - 1\}, y_1]$ and accepts this price if and only if

$$p \leq \frac{y_0 + y_1}{2}.$$

If the seller asks $p < \frac{1}{2}$, then the buyer speculates that $v \in \{y_0\}$ and accepts this price if and only if $p \leq y_0$.

The formal proof is in Appendix A.5. Here we sketch the arguments that lead to this proposition.

Consider first how the buyer makes his choice when the seller asks p. To build the intuition, let us first assume that the buyer makes no inference from the value of the asked price. So the buyer speculates that $v \in [y_0, y_1]$ and compares her maximal losses when buying and not buying the good. The maximal loss of buying is attained when $v = y_0$, giving the loss of $p - y_0$. The maximal loss of not buying is attained when $v = y_1$, giving the loss of $y_1 - p$. The best compromise between these two situations is for the buyer to buy if and only if $p \leq (y_0 + y_1)/2$.

Now consider the inference about v that the buyer makes from the asked price p. When $p \geq 1/2$, the buyer concludes that v cannot be below $2p - 1$. This weakly increases the lower bound on v to $\max\{y_0, 2p - 1\}$. When $2p - 1 \leq y_0$, the inference from observing p is not useful. So the buyer behaves as described above. When $2p-1 > y_0$, the maximal loss from buying is larger than that from not buying. So the buyer does not buy. Notice that $2p-1 > y_0$ implies $p > (y_0+y_1)/2$, and hence the rule described above continues to apply. In summary, the buyer behaves as if she ignores how the seller chooses the price when $p \geq 1/2$.

Alternatively, suppose that $p < 1/2$. This cannot happen in equilibrium, so the buyer can have any speculated beliefs. Assume that the buyer speculates that $x_0 = x_1 = y_0$. So, the buyer speculates that $v = y_0$. Clearly, it is then best to buy the good if and only if $p \leq y_0$.

Consider now how the seller chooses the price when anticipating the buyer's equilibrium behavior. Observe that p should be at least 1/2. This is because if $p < 1/2$, then the buyer accepts p if and only if $p \leq y_0$. So the seller knows that she will only sell the good if it its value is above its price. Thus, choosing $p < 1/2$ is dominated by choosing $p = 1$.

To understand how a price $p \geq 1/2$ should be chosen, consider briefly an alternative setting where it is common knowledge that $v \in [x_0, x_1]$. So the buyer has the same information as the seller. Then the seller will ask $p = (x_0 + x_1)/2$, as this is the highest price that the buyer is willing to accept, and any higher price leads to no sale with the same maximal loss.

Now return to our model. Assume that the buyer does not buy at price p. The maximal loss is attained when the buyer would have bought at a marginally lower price and the value of the good is x_0. So the maximal loss equals $p - x_0$. Now assume that the buyer buys at price p. The maximal loss is attained when the buyer is extremely optimistic and believes that $v \in [y_0, y_1] = [x_1, 1]$. This buyer will also accept the price $(x_1 + 1)/2$. So the maximal loss equals $(x_1 + 1)/2 - p$. The seller chooses a best compromise price that balances these two losses, and hence sets $p = \frac{1}{2}(x_0 + x_1) + \frac{1}{4}(1 - x_1)$. Note that the price asked by the seller lies above the midpoint of the seller's interval $[x_0, x_1]$, due to the possibility of the extremely optimistic buyer.

Proposition 4 stands in contrast to the no-trade theorem under common values as predicted by PBE (Milgrom and Stokey, 1982). We observe that trade occurs in our PCE whenever the median assessment $\frac{1}{2}(y_0 + y_1)$ of the buyer exceeds the price $p^*(x_0, x_1)$. The equilibrium price can be seen as an exaggeration of the seller's median assessment, because $p^*(x_0, x_1) > \frac{1}{2}(x_0 + x_1)$ unless $x_1 = 1$. The trade is possible because the traders cannot rule out the possibility of two opposing situations: winning and losing from trade. They do not want to miss a winning opportunity, but also they do not want to lose from trade. They compromise by choosing their decision thresholds so that they do not lose too much either way.

We hasten to point out that the PCE presented in Proposition 4 is not unique. For example, there is a no-trade PCE, where the seller always asks $p = 1$, and the buyer accepts to buy the good at a price p if and only if $p < y_0$. This equilibrium relies on a specific out-of-equilibrium speculated belief of the buyer that $v = x_0 = x_1 = y_0$ whenever p is different from 1. So, if the seller deviates to some price $p < 1$, either the buyer rejects it, or the seller makes a loss.

4. Conclusion

We introduce a formal methodology to better understand how players deal with uncertainty in dynamic strategic contexts. We are particularly interested in modeling players who have an intuitive understanding of uncertainty that can be expressed in terms of bounds. The general setting looks at players who have ambiguous preferences that are modeled as multiple priors. Learning occurs by updating prior by prior using Bayes' rule whenever possible. Decisions are made under ambiguity by finding best compromises.

Our objective is to present a solution concept that is as close as possible to Perfect Bayesian Equilibrium. The idea is to facilitate the understanding and acceptance of PCE and simplify the interpretation of new insights. This design objective also allows us to build on the discipline underlying the concept of a PBE.

We identify at least six reasons that motivated us to create this new solution concept, each of them motivated by contexts where PBE is not adequate. These reasons are robustness, ambiguity, non-probabilistic reasoning, parsimony, tractability, and accessibility. We explain each of these in more detail.

Robustness. The PCE can be used to investigate the robustness of a PBE to the priors of the players in a context where each of the players seeks a strategy that also performs well for very similar priors. Similarly it can be used to analyze how play changes for a given PBE when players only have an approximate understanding of the priors of others.

Ambiguity. Ambiguous preferences have become popular. Our concept allows to include players with such preferences. The formalism we introduce is not limited to the use of best compromises as the solution concept. We could have also inserted any alternative concept for decision making under ambiguity. The most prominent alternative is maximin utility preferences that leads to a pessimistic mindset. We prefer the flavor of finding compromises. Compromises seems necessary in a globalizing world where decision making is made in front of growing audiences and when there is less willingness to base decisions on specific distributional assumptions.

Non-probabilistic reasoning. Uncertainty per se seems to mean that details are hard to describe. And yet traditional models focus on two types of workers, high and low, or assume linear demand functions. Uncertainty seems to preclude that players agree on likelihoods of events, and yet this is done in PBE. We introduce PCE to open the door to understanding more realistic uncertainty.

Parsimony. The traditional PBE framework reveals a different solution for each prior. Such flexibility can be useful to fit data. But flexibility in terms of a multitude of different answers gives little guidance to those who need to make choices. One easily loses the big picture if there are many details that determine what happens. To achieve clear and transparent results, one often gives up realism and adapts simplistic uncertainty with only a few types for each player. In contrast, the PCE concept under genuine ambiguity is by design very parsimonious. Making best compromises across many different situations allows to abstract from many details.

Tractability. The usefulness of our solution concept is demonstrated in relevant economic examples where uncertainty is rich. This richness limits a tractable analysis of PBE. PCE yields tractable results with simple proofs as players focus on extreme situations, allowing them to ignore intermediate constellations.

Accessibility. The PCE concept under genuine ambiguity is undemanding and easy to teach. Uncertainty can be described with bounds. There is no need for probabilities, and Bayes' rule can be put back on the shelf.

The common acceptance of priors is dwindling. The literature on decision making and game playing under uncertainty has now developed alternative concepts. We hope to add to this literature. Numerous paths to future research open up in a search for new insights and for a clearer exposition of existing understanding of economic and strategic principles.

Appendix A. Proofs.

A.1. **Proof of Theorem 1.** Let $\mathcal{S}$ be the set of strategy profiles. Let $\bar{B}$ be a profile of speculated beliefs given as follows. For each $i \in N$ and each $\phi_i \in \Phi_i$, let $\bar{B}_i(\phi_i)$ contain all probability distributions over the decision nodes in ϕ_i, so $\bar{B}_i(\phi_i) = \Delta(\phi_i)$. As follows from Definition 1, $\bar{B}$ is consistent with every strategy profile in $\mathcal{S}$. This is because $\bar{B}_i(\phi_i)$ always contains every consistent belief at ϕ_i, for all priors and all strategies. We now argue that there exists $\bar{s} \in \mathcal{S}$ such that $(\bar{s}, \bar{B})$ is a PCE.

Consider an arbitrary $s \in \mathcal{S}$. Let $U_{\phi_i}(s)$ be the negative of player i's maximum loss at ϕ_i when player i follows her strategy s_i, so

$$\begin{aligned} U_{\phi_i}(s) &= -l(s_i(\phi_i)|s, B, \phi_i) \\ &= \inf_{\beta_i \in \bar{B}_i(\phi_i)} \left(\bar{u}_i(s_i(\phi_i)|s, \phi_i, \beta_i) - \sup_{a_i \in A(\phi_i)} \bar{u}_i(a_i|s, \phi_i, \beta_i) \right). \end{aligned} \tag{14}$$

We now construct an augmented game $(\Phi, \mathcal{G}, \Omega, \mu, U)$ as follows. Let Φ be the set of information sets excluding the initial node ϕ_0, so $\Phi = \bigcup_{i \in N} \Phi_i$. Let each information set $\phi \in \Phi$ be associated with a different player, so the set of players is the set of information sets Φ. The game tree $\mathcal{G}$ and the set of states Ω remain unchanged. Let μ be a common prior over the states, and assume that μ has full support over Ω. Nature moves first by choosing a state $\omega \in \Omega$ according to the prior μ. Each player $\phi \in \Phi$ moves only once, at her information set ϕ, by choosing an action from the set $A(\phi)$.

A strategy profile s describes a choice $s_\phi \in \Delta(A(\phi))$ of each player ϕ. The interim payoff of each player $\phi \in \Phi$ at the information set ϕ is given by $U_\phi(s)$. Let $U = (U_\phi)_{\phi \in \Phi}$.

The augmented game $(\Phi, \mathcal{G}, \Omega, \mu, U)$ can be seen as a game of incomplete information with a nonstandard specification of the players' payoffs. While in a standard game the payoffs are specified ex-post at each terminal node, in this augmented game the payoff U_ϕ of each player $\phi \in \Phi$ is specified in the interim, at the information set where the player makes a move. Because each player moves only once, the specification of the interim payoffs is sufficient to apply the concept of PBE or sequential equilibrium to the augmented game.

Another nonstandard feature of the augmented game is that each player's interim payoff $U_\phi(s)$ is independent of state ω. This is because by (14) the interim payoff $U_\phi(s)$ is defined as the infimum over all possible beliefs about the decision nodes in the information set ϕ. So, the prior μ does not affect the best-response actions by the players, it only affects the likelihood of reaching different information sets in the game tree.

Let $(s'_\phi, s_{-\phi})$ denote the strategy profile where $s'_\phi \in \Delta(A(\phi))$ is played by player ϕ and $s_{-\phi}$ is the profile of strategies at all other players. Observe that maximizing $U_\phi(s'_\phi, s_{-\phi})$ with respect to player ϕ's own decision $s'_\phi \in \Delta(A(\phi))$ is the same as minimizing the maximum loss at ϕ in the original game. Consequently, if $\bar{s}$ is a strategy profile in a sequential equilibrium of the augmented game, then $(\bar{s}, \bar{B})$ is a PCE of the original game. The existence of PCE follows from the existence of sequential equilibrium for finite games. We refer the reader to Chakrabarti and Topolyan (2016) for the backward-induction proof of existence of sequential equilibrium that uses interim payoffs at information sets to determine players' best-response correspondences. □

A.2. **Proof of Proposition 1.** To prove the existence of a unique PCE, we find a unique profile of best-compromise strategies and a unique profile of speculated beliefs that satisfy Definition 1.

First, we find the speculated beliefs. The firms have genuine ambiguity, so the set of priors Π_i of firm i is equal to the set of degenerate beliefs over $\mathcal{P}$. By Definition 1 and the consistency requirement in PCE, the set $B_i(\phi_i)$ of speculated beliefs of firm i at its unique information set ϕ_i must be equal to the set of priors, so $B_i(\phi_i) = \Pi_i$.

Next, we find each firm's equilibrium quantity. For derivations, we assume that the quantities and the price are always nonnegative, and then we verify that this is indeed the case in equilibrium.

Let $x_i^*(q_{-i}, P)$ be a best response strategy of player i given the knowledge of q_{-i} and the inverse demand function P. The loss of firm i from choosing quantity q_i, given q_{-i} and P, is denoted by $\Delta u_i(q_i, q_{-i}; P)$ and given by

$$\Delta u_i(q_i, q_{-i}; P) = P(x_i^*(q_{-i}, P) + q_{-i})x_i^*(q_{-i}, P) - P(q_i + q_{-i})q_i.$$

By (2), the marginal revenue of firm i satisfies

$$\underline{P}(q_i+q_{-i})+\underline{P}'(q_i+q_{-i})q_i \le P(q_i+q_{-i})+P'(q_i+q_{-i})q_i \le \bar{P}(q_i+q_{-i})+\bar{P}'(q_i+q_{-i})q_i.$$

Therefore, for given q_j and P, the best-response quantity $x_i^*(q_{-i}, P)$ of firm i always lies between $x_i^*(q_{-i}, \underline{P})$ and $x_i^*(q_{-i}, \bar{P})$. While the profit function need not be concave in general, it is concave when $P = \underline{P}$ or when $P = \bar{P}$. So the highest loss will always be attained in one of these two extreme cases:

$$l_i(q_i, q_{-i}) = \sup_P \Delta u_i(q_i, q_{-i}; P) = \max\{\Delta u_i(q_i, q_{-i}; \underline{P}), \Delta u_i(q_i, q_{-i}; \bar{P})\}.$$

It is easy to see that the maximum loss is minimized by balancing the two expressions under the maximum:

$$\Delta u_i(q_i, q_{-i}; \bar{P}) = \Delta u_i(q_i, q_{-i}; \underline{P}).$$

Substituting $\underline{P}$ and $\bar{P}$ and simplifying the expressions yields the equation

$$\frac{(\bar{a} - \bar{b}q_{-i})^2}{4\bar{b}} - (\bar{a} - \bar{b}(q_i + q_{-i}))q_i = \frac{(\underline{a} - \underline{b}q_{-i})^2}{4\underline{b}} - (\underline{a} - \underline{b}(q_i + q_{-i}))q_i. \tag{15}$$

Solving for q_i yields the unique best compromise quantity:

$$q_i^* = \frac{\underline{a}\sqrt{\bar{b}} + \bar{a}\sqrt{\underline{b}}}{2(\underline{b}\sqrt{\bar{b}} + \bar{b}\sqrt{\underline{b}})} - \frac{q_j}{2}, \quad i = 1, 2.$$

Solving this pair of equations for (q_1^*, q_2^*), we find (3). It is easy to verify that under our assumptions, $q_i^* > 0$, and moreover, $P(q_1^* + q_2^*) \geq \underline{P}(q_1^* + q_2^*) > 0$. Substituting the solution into (15) yields the maximum loss of each firm (4). □

A.3. **Proof of Proposition 2.** Similarly to the proof of Proposition A.2, to prove the existence of a unique PCE, we find a unique profile of best-compromise strategies and a unique profile of speculated beliefs that satisfy Definition 1.

First, we find the speculated beliefs. The firms have genuine ambiguity, so the set of priors Π_i of firm i is equal to the set of degenerate beliefs over $[\underline{c}, \bar{c}]^2$. By Definition 1 and the consistency requirement in PCE, firm i with cost c_i must have the set $B_i(c_i)$ of speculated beliefs equal to the set of priors, so $B_i(\phi_i) = \Pi_i$.

Next, we find each firm's equilibrium quantity. For derivations, we assume that each firm prices at or above marginal cost, and then we verify that this is indeed the case in equilibrium.

Consider firm i with type $c_i \in [\underline{c}, \bar{c}]$. Let $s^m(c_i)$ be the profit-maximizing pricing strategy if firm i were the monopoly, so $s^m(c_i) = (a + c_i)/2$. Since we have assumed that $\bar{c} \leq a/2$, this means that $s^m(c_i) \geq \bar{c}$ for all c_i. The monopoly profit is $(a - c_i)^2/(4b)$.

Fix the other firm's strategy $s_{-i}^*(c_{-i})$ and let $\bar{p}$ be the maximum price of the other firm, so $\bar{p} = \sup_{c_{-i} \in [\underline{c}, \bar{c}]} s_{-i}^*(c_{-i})$. Given the other firm's cost c_{-i}, and thus the price $p_{-i} = s^*_{-i}(c_{-i})$, firm i's maximum profit is

$$u_i^*(p_{-i};c_i) = \sup_{x_i \geq 0} u_i(x_i, p_{-i}; c_i) = \begin{cases} 0, & \text{if } p_{-i} \leq c_i, \\ (p_{-i} - c_i)\frac{a-p_{-i}}{b}, & \text{if } c_i < p_{-i} \leq s^m(c_i), \\ \frac{(a-c_i)^2}{4b}, & \text{if } p_{-i} > s^m(c_i) \end{cases}$$

$$= \max\left\{0, (p_{-i} - c_i)\frac{a-p_{-i}}{b}, \frac{(a-c_i)^2}{4b}\right\}.$$

Let p_i be a price of firm i. We now find the maximum loss of firm i from choosing p_i, given its marginal cost c_i and the strategy s^*_{-i} of the other firm. There are three cases.

First, suppose that $p_{-i} \leq c_i \leq p_i$. Then firm i cannot make positive profit, so p_i is a best response. Thus, firm i behaves optimally in this case, so the loss is zero.

Second, suppose that $c_i < p_{-i} \leq p_i$. Then firm i could have been better off by marginally undercutting p_{-i}. Maximizing the loss over $p_{-i} \in (c_i, p_i]$, we obtain

$$\sup_{p_{-i} \in (c_i, p_i)} (u_i^*(p_{-i}; c_i) - u_i(p_i, p_{-i}; c_i)) = \begin{cases} (p_i - c_i)\frac{a-p_i}{b}, & \text{if } p_i \leq s^m(c_i), \\ \frac{(a-c_i)^2}{4b}, & \text{if } p_i > s^m(c_i). \end{cases} \tag{16}$$

Third, suppose that $p_i < p_{-i}$. Then firm i could have made more profit by increasing its price, so its maximum loss is

$$\sup_{p_{-i} \in (p_i, \bar{p}]} (u_i^*(p_{-i}; c_i) - u_i(p_i, p_{-i}; c_i)) = u_i^*(\bar{p}; c_i) - u_i(p_i, \bar{p}; c_i)$$

$$= -(p_i - c_i)\frac{a-p_i}{b} + \begin{cases} (\bar{p} - c_i)\frac{a-\bar{p}}{b}, & \text{if } p_i \leq s^m(c_i), \\ \frac{(a-c_i)^2}{4b}, & \text{if } p_i > s^m(c_i). \end{cases} \tag{17}$$

To minimize the maximum loss, we need to minimize the greater of the expressions in (16) and (17). Observe that, by the definition of $s^m(c_i)$, the right-hand side in (16) is constant and the right-hand side in (17) is strictly increasing in p_i for $p_i > s^m(c_i)$. So we only need to consider $p_i \leq s^m(c_i)$. Under this assumption, the greater of the expressions in (16) and (17) can be simplified to

$$l_i(p_i, s^*_{-i}; c_i) = \max\left\{(p_i - c_i)\frac{a-p_i}{b}, (\bar{p} - c_i)\frac{a-\bar{p}}{b} - (p_i - c_i)\frac{a-p_i}{b}\right\}.$$

Because one expression is increasing and the other is decreasing in p_i for $p_i \leq s^m(c_i)$, the maximum loss is minimized at the solution of

$$(p_i - c_i)\frac{a-p_i}{b} = (\bar{p} - c_i)\frac{a-\bar{p}}{b} - (p_i - c_i)\frac{a-p_i}{b}. \tag{18}$$

Solving the above for p_i and assigning $s_i^*(c_i) = p_i$, we obtain (5).

To see that $s_i^*(c_i) \geq c_i$, observe that

$$s_i^*(c_i) - c_i = \frac{1}{2}\left(a - c_i - \sqrt{(a-\bar{c})^2 + (\bar{c} - c_i)^2}\right) \geq 0$$

by the triangle inequality and $a > \bar{c} \geq c_i$. Moreover, $s_i^*(c_i) > c_i$ when $c_i < \bar{c}$, and $s_i^*(\bar{c}) = \bar{c}$. Finally, substituting $s_i^*(c_i)$ into the maximum loss expression in (18) yields (6). □

A.4. **Proof of Proposition 3.** First we find the equilibrium wages w^H and w^L after the worker's level of education e_H and e_L. For each $j = L, H$, each firm i has the set of speculated states $S_i(e_j) \subset \Omega$. Let this set be the same for each firm. Denote this set by $S(e_j)$, so $S(e_j) = S_1(e_j) = S_2(e_j)$.

Let $\underline{\theta}_j$ and $\bar{\theta}_j$ be the lowest and highest productivity levels given e_j, so

$$\underline{\theta}_j = \inf\{\theta : (\theta, c) \in S(e_j)\} \quad \text{and} \quad \bar{\theta}_j = \sup\{\theta : (\theta, c) \in S(e_j)\}, \quad j = L, H. \quad (19)$$

Consider a firm i, some wages w_i and w_{-i}, and a state (θ, c). Firm i's maximum profit $u_i^*(w_{-i}; \theta)$ is obtained by marginally outbidding w_{-i} when it is below θ, and by choosing the wage below w_{-i} and thus giving up the worker if $\theta \leq w_{-i}$, so

$$u_i^*(w_{-i}; \theta) = \sup_{w_i \geq 0} u_i(w_i, w_{-i}; \theta) = \max\{\theta - w_{-i}, 0\}.$$

Observe that we only need to consider w_i and w_{-i} in $[\underline{\theta}_j, \bar{\theta}_j]$. A wage above $\bar{\theta}_j$ is dominated and cannot be a best compromise; a wage below $\underline{\theta}_j$ will always be overbid by the rival's wage, as there is common knowledge that $\theta \geq \underline{\theta}_j$.

Suppose that $w_i < w_{-i}$, so $u_i(w_i, w_{-i}; \theta) = 0$. Then the largest loss is obtained when θ is the greatest:

$$\sup_{\theta : (\theta, c) \in S(e_j)} (u_i^*(w_{-i}; \theta) - u_i(w_i, w_{-i}; \theta)) \leq \max\{\bar{\theta}_j - w_{-i}, 0\}.$$

Next, suppose that $w_i > w_{-i}$, so $u_i(w_i, w_{-i}; \theta) = \theta - w_i$. Then the largest loss is obtained when θ is the smallest:

$$\sup_{\theta : (\theta, c) \in S(e_j)} (u_i^*(w_{-i}; \theta) - u_i(w_i, w_{-i}; \theta)) = \max\{\theta - w_{-i}, 0\} - (\theta - w_i) \leq w_i - \underline{\theta}_j.$$

Finally, suppose that $w_i = w_{-i}$, so $u_i(w_i, w_{-i}; \theta) = (\theta - w_i)/2$. Then

$$\sup_{\theta : (\theta, c) \in S(e_j)} (u_i^*(w_{-i}; \theta) - u_i(w_i, w_{-i}; \theta)) = \max\{\theta - w_{-i}, 0\} - \frac{\theta - w_i}{2}$$
$$\leq \max\{0, \bar{\theta}_j - w_{-i}, (w_i - \underline{\theta}_j)/2\}.$$

The maximum loss $l_i(w_i, w_{-i})$ is given by the greatest of the three expressions, so

$$l_i(w_i, w_{-i}) = \max\{0, \bar{\theta}_j - w_{-i}, w_i - \underline{\theta}_j.\}.$$

The wages w_i that minimizes the maximum loss satisfies

$$w_i = \bar{\theta}_j + \underline{\theta}_j - w_{-i}, \quad i = 1, 2.$$

So, we have obtained two equations, one for each $i = 1, 2$. Solving this pair of equations for w_1 and w_2 yields the best compromise $w_i^*(e_j)$ for each firm i, where

$$w_i^*(e_j) = \frac{\bar{\theta}_j + \underline{\theta}_j}{2}, \quad i = 1, 2. \quad (20)$$

The associated maximum losses are

$$l_i(w_i^*(e_j), w_{-i}^*(e_j)) = w_i^*(e_j) - \underline{\theta}_j. \tag{21}$$

Next, observe that the worker operates under complete information. Given each choice of e_j, she anticipates the wages $w^j = w_1^*(e_j) = w_2^*(e_j)$, $j \in \{L, H\}$. So, given a state (θ, c), the worker chooses $e = e_H$ if and only if[3]

$$w^H - c(\theta) \geq w^L.$$

Recall that $c(\theta)$ is strictly decreasing, and denote by c^{-1} its inverse. Then, the worker chooses $e = e_H$ if and only if her type θ satisfies

$$\theta \geq c^{-1}(w^H - w^L).$$

Pooling PCE. If $w^H \leq w^L$, then every type chooses low level of education e_L, so the equilibrium is pooling. After observing $e = e_L$, the consistent set of speculated states $S(e_L)$ is thus the entire set of states, so $S(e_L) = \Omega$. By (7), the highest and lowest θ in $S(e_L)$ are $\bar{\theta}_L = 1$ and $\underline{\theta}_L = 0$. By (20), we obtain the equilibrium wages $w_i(e_L) = 1/2$. After observing an out-of-equilibrium education $e = e_H$, the set of speculated states $S(e_H)$ must induce the wage $w_i^*(e_H) \leq w_i^*(e_L)$. In particular, we can assume $S(e_H) = \Omega$, and thus $w_i^*(e_H) = 1/2$.

Substituting the wage of $w_i^*(e) = 1/2$ and the lower bound productivity $\underline{\theta}_L = 0$ into (21), we obtain the maximum loss for each firm i,

$$l_i(w_i^*(e_j), w_{-i}^*; e_j) = \frac{1}{2}, \quad i = 1, 2, \quad j = L, H.$$

Separating PCE. Consider now $w^H > w^L$, so that the worker with cost $c \leq w^H - w^L$ chooses high education. Let

$$S(e_L) = \{(\theta, c) \in \Omega : c > w^H - w^L\} \quad \text{and} \quad S(e_H) = \{(\theta, c) \in \Omega : c(\theta) \leq w^H - w^L\}$$

be the sets of speculated beliefs of each firm when the level of education is e_L and e_H, respectively. So, $S(e_L)$ and $S(e_H)$ contain all pairs (θ, c) such that low and high education is chosen, respectively. These sets thus satisfy the consistency requirement (Definition 1).

By (7) and (19), the highest and lowest θ in $S(e_H)$ are given by

$$\bar{\theta}_H = 1 \quad \text{and} \quad \underline{\theta}_H = \frac{1 - w^H + w^L}{b}. \tag{22}$$

Similarly, the highest and lowest θ in $S(e_L)$ are given by

$$\bar{\theta}_L = \frac{1 + \delta - w^H + w^L}{b} \quad \text{and} \quad \underline{\theta}_L = 0. \tag{23}$$

From (20), we have

$$w^H = \frac{\bar{\theta}_H + \underline{\theta}_H}{2} \quad \text{and} \quad w^L = \frac{\bar{\theta}_L + \underline{\theta}_L}{2}. \tag{24}$$

[3] The tie breaking is arbitrary, because the set of types is a continuum.

Solving the system of six equations in (22), (23), and (24), with six unknowns (w^H, w^L, $\bar{\theta}_H$, $\underline{\theta}_H$, $\bar{\theta}_L$, and $\underline{\theta}_L$), we obtain the equilibrium wages and the bounds on the productivity types as shown in (10) and (11).

Observe that the lowest possible cost of high education is $\inf\{c : (\theta, c) \in \Omega\} = 1 - b$. Therefore, there exist states (θ, c) where high education e_H is chosen if and only if $w^H - w^L > 1 - b$. Substituting our solution for w^H and w^L given by (10), we obtain that $w^H - w^L > 1 - b$ if and only if

$$\delta < 2b^2 - b.$$

This condition is thus necessary and sufficient for the existence of separating PCE.

Finally, substituting the wage w^H and the productivity lower bound $\underline{\theta}_H$ into (21), we obtain firm i's maximum loss when $e = e_H$,

$$l_i(w_i^*(e_H), w_{-i}^*(e_H); e_H) = w^H - \underline{\theta}_H = \frac{1}{2} - \frac{b+\delta}{4b^2}.$$

Substituting the wage w^L and the productivity lower bound $\underline{\theta}_L$ into (21), we obtain the maximum loss when $e = e_L$,

$$l_i(w_i^*(e_L), w_{-i}^*(e_L); e_L) = w^L - \underline{\theta}_L = \frac{\delta}{2b} + \frac{b+\delta}{4b^2}. \qquad \square$$

A.5. **Proof of Proposition 4.** Consider how a buyer who knows that v is in $[y_0, y_1]$ reacts when the seller asks p. Suppose that $p < 1/2$. Then the buyer speculates that v in $\{y_0\}$. This is consistent with the strategy of the seller as $p < 1/2$ is out of equilibrium. Given this speculation, accepting p if and only if $p \le y_0$ is a best compromise.

Now suppose that $p \ge 1/2$. The largest interval $[x_0, x_1] \subset [0, 1]$ that satisfies (13) is $[2p - 1, 1]$. So the buyer concludes that

$$v \in V_b(p, y_0, y_1) = [y_0, y_1] \cap [2p - 1, 1] = [\max\{y_0, 2p - 1\}, y_1].$$

Given this information about the set of possible values, the buyer now compares her maximum losses when accepting ($\alpha = 1$) and rejecting ($\alpha = 0$) the price p. The maximum loss from rejecting p is

$$l_b(0; p, y_0, y_1) = \sup_{v \in [\max\{y_0, 2p-1\}, y_1]} (v - p) = y_1 - p.$$

The maximum loss from accepting p is

$$l_b(1; p, y_0, y_1) = \sup_{v \in [\max\{y_0, 2p-1\}, y_1]} (p - v) = \min\{p - y_0, 1 - p\}.$$

Because $y_1 \le 1$, it is easy to verify that $l_b(0; p, y_0, y_1) \ge l_b(1; p, y_0, y_1)$ if and only if $p \le \frac{1}{2}(y_0 + y_1)$. Thus, it is a best compromise to buy the good when $p \le \frac{1}{2}(y_0 + y_1)$ and not to buy it otherwise.

Let us consider the first stage of the game. Anticipating the buyer's equilibrium behavior α^*, the seller chooses a price that minimizes his maximal loss. Observe that choosing a price $p < 1/2$ is dominated by $p = 1/2$. This is because when

$p < 1/2$, the buyer accepts p if and only if the value v is guaranteed to be at least as high as the price p. In this case, the seller's payoff cannot be positive.

Let $p \geq 1/2$. Suppose first that $p > \frac{1}{2}(y_0 + y_1) > v$. So p is rejected, but it would be optimal to reduce the price so that the buyer accepts it, specifically, to ask $p' = (y_0 + y_1)/2$, and thus gain $p' - v$. The supremum of this loss is given by

$$\sup_{\substack{(v,y_0,y_1):\, p>\frac{1}{2}(y_0+y_1)>v,\\ v\in[x_0,x_1]\cap[y_0,y_1]}} \left(\frac{y_0+y_1}{2} - v\right) = p - x_0.$$

Second, suppose that $p \leq \frac{1}{2}(y_0+y_1) < v$. So p is accepted, but it would be optimal not to sell, and thus gain $v - p$. The supremum of this loss is given by

$$\sup_{\substack{(v,y_0,y_1):\, p\leq\frac{1}{2}(y_0+y_1)<v,\\ v\in[x_0,x_1]\cap[y_0,y_1]}} (v - p) = x_1 - p.$$

Third, suppose that $p \leq \frac{1}{2}(y_0 + y_1)$ and $v \leq \frac{1}{2}(y_0 + y_1)$. So p is accepted, but it would be optimal to sell at a higher price, specifically, at $p' = \frac{1}{2}(y_0+y_1)$, and thus gain $p' - p$. The supremum of this loss is given by

$$\sup_{\substack{(v,y_0,y_1):\, p,v\leq\frac{1}{2}(y_0+y_1),\\ v\in[x_0,x_1]\cap[y_0,y_1]}} \left(\frac{y_0+y_1}{2} - p\right) = \frac{x_1+1}{2} - p.$$

Finally, suppose that $p > \frac{1}{2}(y_0 + y_1)$ and $v \geq \frac{1}{2}(y_0 + y_1)$. So, p is rejected, but any price $p' > v$ would have been rejected too, so the loss is zero in this case.

The maximum loss associated with the price $p \geq 1/2$ is the largest of the four losses computed above, so

$$l_s(p; x_0, x_1) = \max\left\{p - x_0, x_1 - p, \frac{x_1+1}{2} - p, 0\right\} = \max\left\{p - x_0, \frac{x_1+1}{2} - p\right\}.$$

A best compromise price minimizes the maximum loss $l_s(p; x_0, x_1)$ among all prices $p \geq 1/2$, leading to the seller's equilibrium strategy (13). □

Appendix B (**For Online Publication**)

In this appendix we analyze two additional examples: public good provision and forecasting.

B.1. **Public Good Provision.**

Here we investigate the provision of a public good under genuine ambiguity. We assume that each beneficiary knows her own value of the public good, but not those of the others.

There are n agents, each has a private value $v_i \in [0, \bar{v}]$ for the public good. Each agent i commits to contribute at most $x_i \in [0, \bar{v}]$ in case the public good is provided. Agents make their commitments simultaneously.

The cost of providing the public good is $c > 0$. We assume that this cost is relatively small, specifically,

$$c \leq \tfrac{1}{2}(n-1)\bar{v}. \tag{25}$$

This assumption simplifies the exposition. The complementary case can also be easily analysed.

The payoffs are as follows. If the sum of committed contributions does not cover the cost, so $\sum_{i=1}^n x_i < c$, then the public good is not provided, and each agent i obtains zero payoff. Otherwise, if $\sum_{i=1}^n x_i \geq c$, then the public good is provided, and each agent i obtains the payoff

$$v_i - t_i(x),$$

where $t_i(x)$ is the final transfer of agent i that depends on the profile of committed contributions $x = (x_1, ..., x_n)$. For all x such that $\sum_{i=1}^n x_i \geq c$, the transfers $t_i(x)$ must satisfy:

(a) $t_i(x) \leq x_i$ for each i, so no agent pays more than her committed contribution,
(b) $\sum_{i=1}^n t_i(x) \geq c$, so the payments cover the cost of the public good;
(c) $t = (t_1, ..., t_n)$ is symmetric, so the agents are treated ex-ante equally.

We compare three simple transfer rules that determine payments whenever the good is provided.

(i) *Pay-as-you-bid rule.* Each agent pays what she commits to contribute, so

$$t_i(x) = x_i. \tag{26}$$

(ii) *Proportional rule.* Each agent pays proportionally to her commitment, so

$$t_i(x) = \frac{cx_i}{\sum_{j=1}^n x_j}. \tag{27}$$

(iii) *Additive rule.* Each agent pays the equal share c/n plus the difference between her commitment and the average commitment, so

$$t_i(x) = \frac{c}{n} + x_i - \frac{1}{n}\sum_{j=1}^n x_j. \tag{28}$$

Let $s_i(v_i)$ be a strategy of agent i, so $x_i = s_i(v_i)$ specifies the committed contribution of agent i whose private value is v_i. We restrict attention to strategies that satisfy the following assumption.

$$\begin{aligned} &\text{Strategies } s_i \text{ are continuous, strictly increasing, and} \\ &s_i(v) = s_j(v) \ \text{ for all } i, j = 1, ..., n \text{ and all } v \in [0, \bar{v}]. \end{aligned} \tag{29}$$

We will measure the inefficiency of a strategy profile s by the maximum welfare loss as compared to to the complete information case. Our measure is denoted by $L(s)$ and is given by

$$L(s) = \sup_{(v_1,...,v_n)\in[0,\bar{v}]^n} \begin{cases} \max\{0, \sum_i v_i - c\}, & \text{if } \sum_i s_i(v_i) < c, \\ \sum_i t_i(s_i(v_i)) - c, & \text{if } \sum_i s_i(v_i) \geq c \text{ and } \sum_i v_i \geq c, \\ \sum_i t_i(s_i(v_i)) - \sum_i v_i, & \text{if } \sum_i s_i(v_i) \geq c \text{ and } \sum_i v_i < c. \end{cases}$$

The first case describes the loss when the public good is not provided. The second case describes the loss when the public good is provided with the total payment of $\sum_i t_i(s_i(v_i))$, but it would have been best to provide it with the total payment equal to c. The third case describes the loss when the public good is provided with the total payment of $\sum_i t_i(s_i(v_i))$, but it would have been best not to provide it. Note that the second case has zero loss under the proportional and additive rules. This is because under these rules the total payment is always equal to c whenever the good is provided. Also note that the third case has zero loss if the agents do not commit to contribute more than their values, so $s_i(v_i) \leq v_i$ for each i.

Proposition 5. *For each of the three transfer rules the is a unique PCE strategy profile $s^* = (s_1^*, ..., s_n^*)$ that satisfies Assumption (29). We present the strategies together with the associated welfare losses. For each $i = 1, ..., n$ and each $v_i \in [0, \bar{v}]$,*

(i) if $t_i(x)$ is the pay-as-you-bid rule, then

$$s_i^*(v_i) = \frac{v_i}{2} \quad \textit{and} \quad L(s^*) = \max\left\{c, \frac{n\bar{v}}{2} - c\right\};$$

(ii) if $t_i(x)$ is the proportional rule, then

$$s_i^*(v_i) = \frac{v_i}{2} - c + \frac{1}{2}\sqrt{v_i^2 + 4c^2} \quad \textit{and} \quad L(s^*) = \frac{n}{n+1}c;$$

(iii) if $t_i(x)$ is the additive rule, then

$$s_i^*(v_i) = \frac{n}{2n-1}v_i \quad \textit{and} \quad L(s^*) = \frac{n-1}{n}c.$$

Note that $\max\left\{c, \frac{n\bar{v}}{2} - c\right\} > \frac{n}{n+1}c > \frac{n-1}{n}c$. So, the additive rule is the most efficient among these three transfer rules.

Proof. We prove Proposition 5 for the proportional transfer rule given by (27), so we derive part (ii). The proof of parts (i) and (iii) for the other two rules is analogous but easier, and thus omitted.

Let us first derive an agent i's best compromise strategy s_i^*. Agent i who chooses x_i worries about two possible situations. It could be that the total contribution is marginally below c, so $\sum_j x_j = c - \varepsilon$ for a small $\varepsilon > 0$. The good is not provided, but had i contributed ε more it would have been provided. As $\varepsilon \to 0$, agent i's loss is $\max\{v_i - x_i, 0\}$.

Alternatively, it could be that all other agents contribute enough to cover c, so $\sum_{j\neq i} x_j \geq c$. Thus the agent could have contributed nothing and still received the good. In this case the loss is the amount of contribution, $t_i(x) = (cx_i)/\sum_j x_j$. This loss is maximized when the other agents' contributions exactly equal to the cost, so $\sum_{j\neq i} x_j = c$. This loss is given by

$$t_i^*(x) = \frac{cx_i}{x_i + \sum_{j\neq i} x_j} = \frac{cx_i}{x_i + c}.$$

The loss in the first case is weakly decreasing and the loss in the second case is strictly increasing in x_i. To find x_i that minimizes the maximum loss, we solve the equation

$$\max\{v_i - x_i, 0\} = \frac{cx_i}{x_i + c}$$

for x_i. Denote the solution by $s_i^*(v_i)$. It is easy to verify that it is as given in part (ii) of the statement of the proposition.

The above argument requires that there exist values $v_j \in [0, \bar{v}]$ such that $\sum_{j \neq i} s_j^*(v_j) = c$. Observe that $s_i^*(0) = 0$ and $s_i^*(v_i)$ is increasing in v_i. So, we only need to verify that $\sum_{j \neq i} s_j^*(\bar{v}) \geq c$. This inequality holds under condition (25).

It remains to determine the maximum welfare loss $L(s^*)$. Note that whenever the good is provided, $\sum_i t_i(s_i^*(v_i)) = c$ by construction. So, the maximum welfare loss is given by the maximum surplus lost when the good is not provided,

$$L(s) = \sup_{(v_1,\ldots,v_n)\in[0,\bar{v}]^n} \max\left\{0, \sum\nolimits_i v_i - c\right\} \quad \text{subject to} \sum\nolimits_i s_i^*(v_i) < c.$$

As $s_i^*(v_i)$ is increasing in v_i, the constraint must be binding. Moreover, it is easy to verify that $s_i^*(v_i)$ is convex in v_i. Thus, by Jensen's inequality we have

$$\sum\nolimits_i s_i^*(v_i) \geq n s_i^* \left(\frac{1}{n}\sum_i v_i\right).$$

So the maximum is attained for $v_1 = ... = v_n = z$ for $z \in [0, \bar{v}]$ such that $ns_i(z) = c$. Solving the equation

$$n\left(\frac{z}{2} - c + \frac{1}{2}\sqrt{z^2 + 4c^2}\right) = c$$

for z yields

$$z = \frac{2n+1}{n(n+1)}c.$$

So

$$L(s^*) = nz - c = \frac{2n+1}{n+1}c - c = \frac{n}{n+1}c.$$

□

B.2. **Forecasting.** Here we consider forecasting by a single agent of a random variable based on a noisy signal. In this example we illustrate how noise influences learning when the agent makes best compromise choices.

Formally, consider an agent who has to forecast a random variable θ that belongs to $[0, 1]$. The agent's payoff is the quadratic loss:

$$u(a, \theta) = -(a - \theta)^2.$$

Before making a forecast, the agent observes a noisy signal z drawn from some distribution conditional on θ.

We analyze two variations of this model. In one variation, the agent knows how the noisy signal z is generated but she is uncertain about the distribution of the fundamental variable θ. In the other variation, the agent knows the distribution of θ but she is uncertain about the signal generating process. In addition, at the end we deal with the case where the agent is uncertain about both aspects.

B.2.1. *Unknown Distribution of Variable θ.* Here we are interested in how to forecast a random variable with known mean based on a noisy signal with a known distribution.

Suppose that the agent does not know the distribution F of θ. She only knows the mean of this distribution, denoted by θ_0. We allow for any such distribution F that admits a density f such that $\delta \leq f(\theta) \leq 1/\delta$ for some $\delta \in (0,1)$. This assumption excludes holes in the support and point masses. The parameter δ can be interpreted as a lower bound on the degree of dispersion of θ. The set of such distributions is given by

$$\mathcal{F}_\delta = \{F \in \Delta([0,1]) : \mathbb{E}_F[\theta] = \theta_0 \text{ and } \delta \leq f(\theta) \leq 1/\delta \text{ for all } \theta \in [0,1]\}.$$

The agent can condition her forecast on a noisy signal z about θ. The signal generating process is known and given by the conditional probability distribution $G_\varepsilon(z|\theta)$ with a parameter $\varepsilon \in [0,1]$ specified as follows. Signal z reveals the true value θ with probability $1-\varepsilon$ and is drawn uniformly from $[0,1]$ with probability ε, so

$$G_\varepsilon(z|\theta) = \begin{cases} \varepsilon z, & \text{if } z < \theta, \\ 1-\varepsilon+\varepsilon z, & \text{if } z \geq \theta. \end{cases} \tag{30}$$

Had the agent known the distribution $F \in \mathcal{F}_\delta$, she could have formed a belief about θ conditional on the signal z. Let $\mathbb{E}_{F,G_\varepsilon}[\cdot|z]$ denote the conditional mean under this belief.

The maximum loss of a forecast $a \in [0,1]$ given a signal $z \in [0,1]$ is

$$l(a;z) = \sup_{F \in \mathcal{F}_\delta} \left(\sup_{a' \in [0,1]} \mathbb{E}_{F,G_\varepsilon}[-(a'-\theta)^2|z] - \mathbb{E}_{F,G_\varepsilon}[-(a-\theta)^2|z] \right).$$

A best compromise is a forecast $a^*(z)$ that achieves the smallest maximum loss, so

$$a^*(z) \in \arg\min_{a \in [0,1]} l(a;z).$$

This problem can be embedded in our formal setting as described in Section 2. A state is a pair $(\theta, z) \in [0,1]^2$, so the set of states is $\Omega = [0,1]^2$. The set of priors consists of all pairs (F,G) such that $F \in \mathcal{F}_\delta$ and G is given by (30). The set of speculated beliefs of the agent who observes z is the set of posteriors derived from the priors conditional on z.

Proposition 6. *The agent's best compromise is*

$$a^*(z) = (1-\lambda)z + \lambda\theta_0,$$

where

$$\lambda = \frac{\varepsilon}{2}\left(\frac{\delta}{1-\varepsilon(1-\delta)} + \frac{1}{\delta+\varepsilon(1-\delta)}\right).$$

The proof is presented at the end of this subsection.

Let us present some intuition behind Proposition 6. Due to the quadratic penalty of making inaccurate forecasts, the loss of a forecast is equal to its distance from the expected mean conditional on the signal. The forecaster is worried about two possible situations, namely, when this conditional mean is high and when it is low. Consequently, the best compromise involves a forecast at the midpoint of these two extreme conditional means. Solving for this midpoint yields the formulae given in the statement of the proposition. In particular, the best compromise forecast lies between the ex-ante mean θ_0 and the signal z.

Note that the agent's best compromise forecast depends on the precision ε of her signal and on the degree of the dispersion δ of the variable of interest. We show how each of these two parameters independently influences the best compromise forecast.

Fix the degree of dispersion δ. When the agent's signal is not very noisy, then her forecast is close to the signal. This is because a^* is continuous in ε and $\lim_{\varepsilon\to 0} a^*(z) = z$. When the signal is very noisy, then her prediction is close to the ex-ante mean, as $\lim_{\varepsilon\to 1} a^*(z) = \theta_0$.

Now we fix the precision ε of the noise and vary the bound δ on the degree of dispersion of θ. As we relax the constraints on F imposed by δ, we obtain that the forecast approximates the midpoint between θ_0 and z. Formally, $\lim_{\delta\to 0} a^*(z) = (\theta_0 + z)/2$. When δ is small, on the one hand, it could that F has very high dispersion, thus making the signal extremely valuable. On the other hand, it could that F has very low dispersion, in which case the signal has very little value. The forecast seeks a best compromise between these two situations and selects the midpoint.

Note that the above analysis and discussion reveals a discontinuity in the forecast a^* at $\varepsilon = \delta = 0$.

We now prove Proposition 6. To do this, we first present a simple lemma on how the loss of a forecast is computed.

Lemma 1. $l(a; z) = \sup_{F\in\mathcal{F}_\delta}(a - \mathbb{E}_{F,G_\varepsilon}[\theta|z])^2$.

The intuition is as follows. The variance of θ conditional on a signal z enters the payoffs additively, and thus cancels out when computing the loss. As a result, the maximum loss $l(a; z)$ is simply the maximum quadratic distance between a forecast a and the mean value of θ conditional on z.

Proof of Lemma 1. Fix G_ε. Let $\bar{a}_F(z) = \mathbb{E}_{F,G_\varepsilon}[\theta|z]$. Observe that

$$\bar{a}_F(z) \in \arg\max_{a'\in[0,1]} \mathbb{E}_{F,G_\varepsilon}[-(a' - \theta)^2|z]. \tag{31}$$

So, we have

$$\sup_{a'\in[0,1]} \mathbb{E}_{F,G_\varepsilon}[-(a'-\theta)^2|z] - \mathbb{E}_{F,G_\varepsilon}[-(a-\theta)^2|z] = \mathbb{E}_{F,G_\varepsilon}[-(\bar{a}_F(z)-\theta)^2 + (a-\theta)^2|z]$$

$$= \mathbb{E}_{F,G_\varepsilon}[(a - \bar{a}_F(z))(a + \bar{a}_F(z) - 2\theta)|z] = (a - \bar{a}_F(z))^2,$$

where the first equality is by (31) and the last equality is by $\mathbb{E}_{F,G_\varepsilon}[\theta|z] = \bar{a}_F(z)$. Thus,

$$l(a; z) = \sup_{F\in\mathcal{F}_\delta} (a - \bar{a}_F(z))^2 = \sup_{F\in\mathcal{F}_\delta} (a - \mathbb{E}_{F,G_\varepsilon}[\theta|z])^2. \qquad \square$$

Proof of Proposition 6. Different distributions $F \in \mathcal{F}_\delta$ induce different conditional means $\mathbb{E}_{F,G_\varepsilon}[\theta|z]$. Let $H(z)$ and $L(z)$ be the highest and lowest conditional means, respectively, so

$$H(z) = \sup_{F \in \mathcal{F}_\delta} \mathbb{E}_{F,G_\varepsilon}[\theta|z] \quad \text{and} \quad L(z) = \inf_{F \in \mathcal{F}_\delta} \mathbb{E}_{F,G_\varepsilon}[\theta|z]. \tag{32}$$

The loss of a forecast a given a signal z is

$$l(a; z) = \sup_{F \in \mathcal{F}_\delta} (a - \mathbb{E}_{F,G_\varepsilon}[\theta|z])^2 = \max\left\{(a - H(z))^2, (a - L(z))^2\right\}$$

where the first equality is by Lemma 1, and the last equality is by the convexity of the expression. Thus, the best compromise forecast is the midpoint between the highest and lowest conditional means, so

$$a^*(z) = \inf_{a \in [0,1]} l(a; z) = \frac{1}{2}\left(H(z) + L(z)\right).$$

It remains to find $H(z)$ and $L(z)$. Suppose that $z \geq \theta_0$. Observe that

$$\mathbb{E}_{F,G_\varepsilon}[\theta|z] = \frac{(1-\varepsilon)f(z)z + \varepsilon \int_0^1 \theta f(\theta)\mathrm{d}\theta}{(1-\varepsilon)f(z) + \varepsilon \int_0^1 f(\theta)\mathrm{d}\theta} = \frac{(1-\varepsilon)f(z)z + \varepsilon\theta_0}{(1-\varepsilon)f(z) + \varepsilon}$$

is increasing in $f(z)$. Using the assumption that $f(z) \leq 1/\delta$, we have

$$H(z) = \sup_{F \in \mathcal{F}_\delta} \frac{(1-\varepsilon)f(z)z + \varepsilon\theta_0}{(1-\varepsilon)f(z) + \varepsilon} = \left.\frac{(1-\varepsilon)f(z)z + \varepsilon\theta_0}{(1-\varepsilon)f(z) + \varepsilon}\right|_{f(z)=1/\delta} = \frac{(1-\varepsilon)z + \varepsilon\delta\theta_0}{1 - \varepsilon + \varepsilon\delta}.$$

Using the assumption that $f(z) \geq \delta$, we have

$$L(z) = \inf_{F \in \mathcal{F}_\delta} \frac{(1-\varepsilon)f(z)z + \varepsilon\theta_0}{(1-\varepsilon)f(z) + \varepsilon} = \left.\frac{(1-\varepsilon)f(z)z + \varepsilon\theta_0}{(1-\varepsilon)f(z) + \varepsilon}\right|_{f(z)=\delta} = \frac{(1-\varepsilon)\delta z + \varepsilon\theta_0}{(1-\varepsilon)\delta + \varepsilon}.$$

Analogously, for $z \leq \theta_0$ we obtain $H(z) = \frac{(1-\varepsilon)\delta z + \varepsilon\theta_0}{(1-\varepsilon)\delta + \varepsilon}$ and $L(z) = \frac{(1-\varepsilon)z + \varepsilon\delta\theta_0}{1-\varepsilon+\varepsilon\delta}$. Thus we obtain

$$a^*(z) = \frac{1}{2}\left(H(z) + L(z)\right) = \frac{1}{2}\left(\frac{(1-\varepsilon)z + \varepsilon\delta\theta_0}{1 - \varepsilon + \varepsilon\delta} + \frac{(1-\varepsilon)\delta z + \varepsilon\theta_0}{(1-\varepsilon)\delta + \varepsilon}\right). \qquad \square$$

B.2.2. *Unknown Distribution of Signal z.* Here we are interested in how to forecast a random variable with a known distribution after receiving a noisy signal that has an unknown distribution.

Suppose that the agent knows the distribution F of θ, but is uncertain about how the noisy signal z is generated. The following assumptions are made about this signal. The signal z is known to be not too far from the true value of θ, where a parameter $\delta > 0$ describes the maximal distance. So δ can also be interpreted as the precision of the signal. Let $y = z - \theta$ be called the noise. So it is known that $|y| \leq \delta$. The distribution of the noise y has a certain and an uncertain component. Let $\varepsilon \in [0, 1]$ be a known parameter. With probability $1 - \varepsilon$ the noise y is drawn from a known distribution G_0 and with probability ε it is drawn from an unknown distribution G_1. So ε measures how uncertain the agent is about how the noise is generated. Given the support restrictions on y, it follows that G_0 and G_1 both have support contained in $[-\delta, \delta]$. Let G_δ be the set of all distributions of y that satisfy the above description.

Let $\mathbb{E}_{F,G_\delta,\varepsilon}[\cdot|z]$ denote the conditional mean of θ given z for $G_\delta \in \mathcal{G}_\delta$. The maximum loss associated with a forecast $a \in [0,1]$ given a signal $z \in [0,1]$ is calculated as in Section B.2.1, so

$$l(a;z) = \sup_{G_\delta \in \mathcal{G}_\delta} \left(\sup_{a' \in [0,1]} \mathbb{E}_{F,G_\delta,\varepsilon}[-(a'-\theta)^2|z] - \mathbb{E}_{F,G_\delta,\varepsilon}[-(a-\theta)^2|z] \right).$$

Let $H(z)$ and $L(z)$ be the highest and lowest conditional means, so

$$H(z) = \sup_{G_\delta \in \mathcal{G}_\delta} \mathbb{E}_{F,G_\delta,\varepsilon}[\theta|z] \quad \text{and} \quad L(z) = \inf_{G_\delta \in \mathcal{G}_\delta} \mathbb{E}_{F,G_\delta,\varepsilon}[\theta|z].$$

It is straightforward to verify that

$$H(z) = \sup_{x \in [-\delta,\delta]} \frac{\varepsilon f(z-x)(z-x) + (1-\varepsilon)\int_{-\delta}^{\delta}(z-y)f(z-y)\mathrm{d}G_0(y)}{\varepsilon f(z-x) + (1-\varepsilon)\int_{-\delta}^{\delta} f(z-y)\mathrm{d}G_0(y)},$$

with an analogous expression for $L(z)$. We obtain the following result.

Proposition 7. *The agent's best compromise is*

$$a^*(z) = \frac{1}{2}\left(H(z) + L(z)\right).$$

The proof is analogous to that of Proposition 6 and thus omitted.

Just like in Section B.2.1, the best compromise is the midpoint between the highest and lowest conditional means. The agent's best compromise forecast depends on the precision δ of her signal, as well as on the degree ε of her uncertainty. We show how each of these two parameters independently influences the best compromise forecast.

Fix the degree of uncertainty ε. If the signal is very precise in the sense that δ is very small, then each of the two extreme conditional means are close to z. Hence, the best compromise forecast will also be close to z. Formally, $\lim_{\delta\to 0} a^*(z) = z$.

Fix the precision δ of the signal. As the degree of uncertainty ε vanishes, both extreme conditional means converge to the conditional mean under the benchmark distribution G_0. Formally, $\lim_{\varepsilon\to 0} a^*(z) = \mathbb{E}_{F,G_0,0}[\theta|z]$. For instance, if G_0 is the uniform distribution, then the best compromise forecast converges to the expected value of θ conditional on θ being within δ of the signal.

As the degree of uncertainty ε becomes large, the role of the benchmark G_0 diminishes and almost any noise within $[-\delta,\delta]$ becomes possible. When $\varepsilon = 1$, it could be that G_1 puts all mass on $-\delta$, in which case $\mathbb{E}_{F,G_\delta,\varepsilon}[\theta|z] = z+\delta$. This is the highest conditional mean given z, so $H(z) = z+\delta$. It could also be that G_1 puts all mass on δ, in which case $\mathbb{E}_{F,G_\delta,\varepsilon}[\theta|z] = z-\delta$. This is the lowest conditional mean given z, so $L(z) = z-\delta$. Consequently, the best compromise forecast is close to the signal z when the agent is very uncertain about how z is generated. Formally, $a^*(z) \to z$ as $\varepsilon \to 1$.

Note that the distribution F of the underlying variable of interest plays no role when the degree of uncertainty is extreme, so $\varepsilon = 1$. Consequently, we obtain that if the agent knows neither F nor the distribution of the noise, then the best compromise forecast is to choose the signal.

References

Azrieli, Yaron, and Roee Teper. 2011. "Uncertainty aversion and equilibrium existence in games with incomplete information." *Games and Economic Behavior*, 73: 310–317.

Chakrabarti, Subir K, and Iryna Topolyan. 2016. "An extensive form-based proof of the existence of sequential equilibrium." *Economic Theory Bulletin*, 4: 355–365.

Chiesa, Alessandro, Silvio Micali, and Zeyuan Allen Zhu. 2015. "Knightian Analysis of the Vickrey Mechanism." *Econometrica*, 83: 1727–1754.

Cremer, J, and R P McLean. 1985. "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent." *Econometrica*, 53: 345–361.

Epstein, Larry G, and Tan Wang. 1996. ""Beliefs About Beliefs" Without Probabilities." *Econometrica*, 64: 1343–1373.

Gilboa, I., and D. Schmeidler. 1989. "Maxmin Expected Utility with a Non-Unique Prior." *Journal of Mathematical Economics*, 18: 141–153.

Hanany, Eran, Peter Klibanoff, and Sujoy Mukerji. 2018. "Incomplete information games with ambiguity averse players." Mimeo.

Harsanyi, John C. 1967. "Games with incomplete information played by "Bayesian" players, I–III. Part I. The basic model." *Management Science*, 14: 159–182.

Huber, Peter J. 1965. "A Robust Version of the Probability Ratio Test." *The Annals of Mathematical Statistics*, 36: 1753–1758.

Kajii, Atsushi, and Stephen Morris. 1997. "The Robustness of Equilibria to Incomplete Information." *Econometrica*, 65: 1283–1309.

Kajii, Atsushi, and Takashi Ui. 2005. "Incomplete information games with multiple priors." *Japanese Economic Review*, 56: 332–351.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. "A smooth model of decision making under ambiguity." *Econometrica*, 73: 1849–1892.

Linhart, Peter B, and Roy Radner. 1989. "Minimax-regret strategies for bargaining over several variables." *Journal of Economic Theory*, 48: 152–178.

Milgrom, Paul, and Nancy Stokey. 1982. "Information, trade and common knowledge." *Journal of economic theory*, 26: 17–27.

Milnor, John. 1954. "Games Against Nature." In *Decision Processes.* , ed. R. M. Thrall, C. H. Coombs and R. L. Davis. Wiley, New York.

Radner, Roy. 1980. "Collusive behaviour in noncooperative epsilon-equilibria of oligopolies with long but finite lives." *Journal of Economic Theory*, 22: 136–154.

Renou, Ludovic, and Karl H Schlag. 2010. "Minimax regret and strategic uncertainty." *Journal of Economic Theory*, 145: 264–286.

Savage, L. J. 1951. "The Theory of Statistical Decision." *Journal of the American Statistical Association*, 46: 55–67.

Spence, Michael. 1973. "Job market signaling." *Quarterly Journal of Economics*, 87: 355–374.

Spulber, Daniel. 1995. "Bertrand Competition When Rivals' Costs Are Unknown." *Journal of Industrial Economics*, 43: 1–11.

Stauber, Ronald. 2011. "Knightian games and robustness to ambiguity." *Journal of Economic Theory*, 146: 248–274.

Stoye, Jörg. 2011. "Statistical Decisions Under Ambiguity." *Theory and Decision*, 70: 129–148.

Vives, Xavier. 1984. "Duopoly Information Equilibrium: Cournot and Bertrand." *Journal of Economic Theory*, 34: 71–94.

Wald, A. 1950. *Statistical Decision Functions.* Wiley, New York.