Compromise, Don't Optimize: Generalizing Perfect Bayesian Equilibrium to Allow for Ambiguity

Karl H. Schlag and Andriy Zapechelnyuk

Abstract. We introduce a solution concept for extensive-form games of incomplete information in which players can have multiple priors. Players' choices are based on the notions of complaints and compromises. Complaints come from hypothetical assessors who have different priors and evaluate the choices of the players. Compromises are choices that aim to make these complaints small. The resulting solution concept is called perfect compromise equilibrium and generalizes perfect Bayesian equilibrium. We use this concept to provide insights into how ambiguity influences Cournot and Bertrand markets, public good provision, markets for lemons, job market signaling, bilateral trade with common value, and forecasting.

JEL Classification: D81, D83

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

Date: June 28, 2023.

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, University of Edinburgh, 31 Buccleuch Place, Edinburgh, EH8 9JT, UK. *E-mail:* azapech@gmail.com.

The authors thank for helpful comments to Pierpaolo Battigalli, Jeffrey Ely, Simon Grant, Clara Ponsati, Roland Strausz, participants of various seminars where this paper has been presented, and anonymous referees.

1. Introduction

We wish to add to the research on economic models in which players do not necessarily have a unique prior and hence face ambiguity. We would like to capture settings where players can have different degrees of ambiguity. On the one hand, we want to allow players to have a single prior or a few different priors. On the other hand, we are particularly interested in modeling so-called *genuine ambiguity* where players have no probabilistic assessment of what they do not know. The objective under genuine ambiguity is to be able to formally model realistic agents who only focus on which states are possible, without assessing their likelihoods.

We present a new solution concept to capture strategic choices of ambiguous players in extensive-form games. We apply it in several prominent economic examples, the majority of which involve genuine ambiguity. Our approach leads to tractable analyses and parsimonious solutions. While the existing literature offers several solution concepts for extensive-form games with multiple priors, these concepts are not suitable to deal with genuine ambiguity. Moreover, economic examples in this literature have very simple parametric uncertainty as tractability issues arise quickly when uncertainty becomes richer.

The solution concept introduced in this paper is based on the everyday notion of complaints and compromises. It features a particular way of reasoning under uncertainty that simplifies tradeoffs and thereby can lead to tractable solutions. Imagine that each player has to justify each of her choices in front of a set of hypothetical assessors. Each assessor has a single prior. Typically, the player will not be able to find a choice that is best from the perspective of all of the assessors. Consequently, some of the assessors complain about the choice of the player. Confronted by these complaints we postulate that the player wishes to find a compromise. This compromise is an action under which none of the complaints of the assessors is too large. With this in mind, the player chooses an action that makes the largest complaint as small as possible. This methodology is applied to each decision separately, assuming common knowledge of the equilibrium profile of strategies. In particular, this means that the player anticipates the future choices of herself and others.

We hasten to point out three consequences of our approach. First of all, the way in which a compromise is found given the complaints is rooted in minmax regret (Savage, 1951). Second of all, each assessor updates his prior based on what has happened in the past. The player's aim to find a compromise among all the assessors leads to a solution concept that is based on prior-by-prior updating, also known as full Bayesian updating (Pires, 2002). Finally, we model the choices of players under common knowledge of the equilibrium profile. Consequently, each player takes strategies of all players including herself as given and anticipates future choices when making a decision. In the behavioural literature, this assumption is called *sophisticated behaviour* and is related to consistent planning (Strotz, 1955; O'Donoghue and Rabin, 1999; Siniscalchi, 2011). In summary, the need to find compromises leads to an equilibrium concept built on minmax regret, full Bayesian updating, and sophisticated behavior.

Our solution concept is called *perfect compromise equilibrium (PCE).* It generalizes perfect Bayesian equilibrium (PBE), and it exists in finite games. Formally, it specifies for each player a strategy and a belief mapping. The strategy identifies the action the player chooses at each of her information sets according to the compromise criterion. The belief mapping maps each prior of the player to a belief over decision nodes in each of her information sets by applying Bayes rule, prior by prior, whenever possible.

PCE relies on common knowledge of the equilibrium profile of strategies, just like PBE does. So there is strategic certainty. Players have a common belief of how others react to their information in equilibrium, they are only uncertain about what information others actually have. We hasten to point out that our concept can be used to incorporate strategic uncertainty, as outlined in Section 2.2.

PCE is a flexible concept as it can adapt to the mindset of each player by appropriately choosing her set of priors. When there is a single prior, then this player is Bayesian. When all priors are close to each other, then this player is concerned with robustness of her decisions to a slight misspecification of the prior. When the set of priors is large and dispersed, then the model captures the reasoning of a player who is very ambiguous. In the extreme case, when all priors are degenerate, we obtain the model of genuine ambiguity. As priors are degenerate, Bayes' rule no longer needs to be applied, one only checks if states are feasible or not.

Genuine ambiguity is a setting we are particularly interested in. Here we consider realistic people who have difficulty forming priors. Instead they only need to consider which states are possible, without assessing their likelihoods. For instance, it seems unlikely that firms conjecture a specific probability distribution when they think about what demand they will be facing. Yet it seems plausible that they put bounds on the uncertain demand. These bounds can come from the most optimistic and pessimistic scenarios provided by expertise. This way of modeling uncertainty without using priors comes with numerous advantages in comparison to PBE. It is easier to specify, justify, or estimate a set of possible states than to do this for a distribution (prior) over these states. This often makes solutions easier to obtain and more parsimonious. It also enables better understanding of how results depend on inputs. These advantages are demonstrated in our examples.

We investigate seven salient economic 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 marginal costs of their rivals. We consider public good provision where beneficiaries of a public good do not know each others' values and hypothesize an interval where these values can be. We consider Akerlof's market for "lemons" in which the buyer is ambiguous about the quality of the car. We consider Spence's job market where employers are uncertain about the cost of education and the productivity of workers, and conjecture bounds on these parameters. We consider bilateral trade with common value where each party knows an interval that contains the true value. Finally, we consider forecasting of a random variable with unknown distribution.

These examples highlight the value of the PCE concept. Novel realistic settings can be investigated, as we no longer have to confine ourselves to simplistic parametric models of uncertainty, such as when there are only two states (high and low). Realism enters as we can capture uncertainty without explicitly referring to distributions. Tractability is maintained by shifting the focus away from distributions to the worst-case analysis. New insights appear. We find that replacing priors by bounds on uncertain parameters has little impact on profits in Cournot and Bertrand competition settings where compromise values are small. 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 when playing PCE. Yet loosening these bounds causes firms to react differently. They become more competitive under Cournot competition and less competitive under Bertrand competition. In Akerlof's market for lemons, the buyer implicitly hesitates, by using a mixed strategy, when deciding whether or not to buy the car. In the public good game, we show the ease of comparing policies and the simplicity of the beneficiaries' contribution strategies. 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. The possibility that the trading partners have different valuations leads to trade with positive probability in a PCE, as ignoring this possibility generates losses that the traders want to minimize. This is true even though, unlike other papers that study trade with multiple priors (Billot et al., 2000; Kajii and Ui, 2006; Rigotti, Shannon and Strzalecki, 2008), we allow that the buyer makes inference from the price set by the seller. Finally, when forecasting a random variable with a known mean and unknown distribution based on a noisy signal, the best-compromise forecast is a weighted average of the mean and the signal.

Related Literature. Our paper contributes to the literature on robustness and ambiguity in games. The main contribution of this paper is that we offer an operationalizable solution concept for extensive-form games with ambiguity that can allow for simple analysis in problems with rich state spaces that often have tractability issues under PBE.

Two closely related papers are Hanany, Klibanoff and Mukerji (2020) and, concurrently with our paper, Pahlke (2022). Both papers consider extensive-form games with ambiguity. In Hanany, Klibanoff and Mukerji (2020), players have smooth ambiguity preferences (see also Klibanoff, Marinacci and Mukerji, 2005).[1] Specifically, a player aggregates expected utilities calculated under different priors using a distribution over these priors and a concave aggregator function. In Pahlke (2022), players have maxmin utility preferences and, like in our setting, update the priors individually, one-by-one. The central focus of both papers is sequential optimality, which means that a player's ex-ante optimal strategy remains optimal conditional on reaching every information set where that player moves. Hanany, Klibanoff and Mukerji (2020) show that their way of smooth aggregation of multiple priors is necessary and sufficient to have sequential optimality for general sets of priors, whereas Pahlke (2022) guarantees sequential optimality by restricting the sets of priors to have the rectangularity property similar to Epstein and Schneider (2003).

Our paper complements Hanany, Klibanoff and Mukerji (2020) and Pahlke (2022) in two respects, allowing for distinct results in applications, such as our examples. First, in our paper the players have minmax-regret-type preferences, which can be more suitable than maxmin utility and smooth ambiguity for some applications. Second, we do not bind ourselves by the constraint of sequential optimality. Instead, we apply a weaker requirement: our players make optimal choices when anticipating their own future choices, so they have *sophisticated behavior*. The difference of our approach from the above papers becomes apparent in the context of genuine ambiguity, where players only have degenerate priors over states. The set of degenerate priors generally fails the rectangularity property of Pahlke (2022). Smooth ambiguity requires to assign specific weights to priors, which collapses to a single prior when all these priors are degenerate. Thus, in the genuine ambiguity setting, the approach of Hanany, Klibanoff and Mukerji (2020) collapses to PBE with a given prior over states.

Let us compare the compromise (maxmin regret) approach used in our paper and a popular alternative approach, maxmin utility (e.g., Wald, 1950; Gilboa and Schmeidler, 1989; Epstein and Wang, 1996; Kajii and Ui, 2005; Azrieli and Teper, 2011). On an intuitive level, maxmin utility is applicable when players are

[1]See also Battigalli et al. (2019) who study players with smooth ambiguity preferences, but in the context of repeated population games.

pessimistic, while best compromises make more sense when players are interested in making decisions that are good in different contingencies. The underlying philosophy is very different. The maxmin choice is best when payoffs are the lowest, without taking into account the performance in other situations. The best compromise choice is the closest possible to the optimum in all situations.[2] In the salient examples investigated in this paper, the maxmin approach leads to unintuitive results. For instance, in Bertand duopoly with ambiguity about the rival's cost, maxmin utility leads firms to shut down. In contrast, our approach reveals economically relevant insights, and does this in a simple manner with minimal structural assumptions.

The concept of best compromise has origins in minmax regret (Savage, 1951) and connects to approximate optimality. Our optimization criterion differs from minmax regret as evaluation occurs at each information set, while minmax 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 minmax 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 in our paper.

We proceed as follows. In Section 2 we introduce our solution concept, prove existence, and discuss its properties. In Section 3 we illustrate PCE in seven self-contained economic examples. Section 4 concludes. All proofs are in Appendix A. An alternative forecasting example is in Appendix B.

[2]This difference can be illustrated in a laboratory game studying the Ellsberg paradox. A respondent who worries about the lowest payoff chooses the bet with a known probability. On the other hand, a respondent who is concerned that she could have done better may potentially choose the bet with an ambiguous probability. Different motives can lead to different behavior in this laboratory experiment.

2. Perfect Compromise Equilibrium

We introduce a solution concept called *perfect compromise equilibrium (PCE).* The concept is formally defined in Section 2.1 and discussed in Section 2.2. 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 economic examples.

2.1. **Formal Setting.** Consider a finite extensive-form game described by $(N, \mathcal{G}, \Omega, (\Pi_1, ..., \Pi_n), (u_1, ..., u_n))$, where $N = \{1, ..., n\}$ is a set of players, $\mathcal{G}$ is a finite game tree, Ω is a finite set of states, $\Pi_i \subset \Delta(\Omega)$ is a finite set of priors of player i, and u_i is a payoff function of player i. In particular, this embeds a nonprobabilistic view of uncertainty by letting Π_i contain only degenerate priors that put all weight on one of the states. We refer to this case as *genuine ambiguity.* Also note that we allow players to have different sets of priors.

The game tree $\mathcal{G}$ describes the order of players' moves, their information sets, and actions that are available at each information set. It is defined by a set of linked nodes that form a tree. The game starts with the initial node ϕ_0 assigned to nature, followed by decision nodes assigned to players and terminal nodes that describe payoffs. Each decision node is assigned three elements: a player i, an information set ϕ_i, and a set of actions $\mathscr{A}_{\phi_i}$ 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_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. In the canonical case, the set of actions $\mathscr{A}_{\phi_i}$ is a set of mixed actions $\Delta(A_{\phi_i})$ where A_{ϕ_i} is a finite set of pure actions. This corresponds to the typical model of a finite sequential game. Of interest for applications is also the case where mixed actions are ruled out. In this case, $\mathscr{A}_{\phi_i}$ contains only the set of pure actions available at ϕ_i.

Motivated by Harsanyi (1967), all incomplete information is captured by a move of nature at the beginning of the game without loss of generality. At the initial node ϕ_0, nature chooses a *state* ω from the set of states Ω. The set of priors Π_i describes alternative beliefs (theories) of player i for how the state has been determined. Note that if each player has a single prior, then this constitutes a standard Bayesian game with heterogeneous priors.

The game terminates after finitely many moves at some terminal node where 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 profile s describes the behavior of all players throughout the game. It prescribes to each player $i \in N$ in each of her information sets $\phi_i \in \Phi_i$ an action $s_{\phi_i} \in \mathscr{A}_{\phi_i}$.

Like in Bayesian games, we also specify posterior beliefs of the players in their information sets. We do this for each prior separately. We specify a posterior belief at each information set for each prior using Bayes' rule whenever possible. Thus, there are potentially as many posteriors at each information set of player i as there are priors in Π_i. This procedure can be found in the literature in a different context under the name of full Bayesian updating (Pires, 2002).

Formally, for each player i and each information set $\phi_i \in \Phi_i$, let $\beta_{\phi_i} : \Pi_i \to \Delta(\phi_i)$ be a belief mapping that associates each prior $\pi_i \in \Pi_i$ of player i with a posterior probability distribution β_{ϕ_i} over the decision nodes in ϕ_i. Thus, in the information set ϕ_i, player i faces a set $B_{\phi_i}(\beta)$ of posterior beliefs derived from the set of priors Π_i, where

$$B_{\phi_i}(\beta) = \{\beta_{\phi_i}(\pi_i) : \pi_i \in \Pi_i\}.$$

We will refer to $B_{\phi_i}(\beta)$ as the *set of beliefs* at ϕ_i, and to the profile $\beta = (\beta_{\phi_i})_{\phi_i \in \Phi_i, i \in N}$ as the *belief system.*

Like in PBE, we will require consistency of beliefs.

Definition 1. A belief mapping β_{ϕ_i} is called *consistent under a strategy profile* s if for each prior $\pi_i \in \Pi_i$ such that the information set ϕ_i is reached with a strictly positive probability under strategy profile s, the belief $\beta_{\phi_i}(\pi_i)$ is derived by Bayes rule from π_i.

A belief system β is *consistent under a strategy profile* s if for each $i \in N$ and each $\phi_i \in \Phi_i$ the belief mapping β_{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 under a given prior π_i and a given strategy profile s, then every belief $\beta_{\phi_i}(\pi_i) \in \Delta(\phi_i)$ is consistent under s. Of course, not all out-of-equilibrium beliefs are sensible in applications. For example, if needed, it is natural to refine the concept of PCE in the same way as sequential equilibrium (Kreps and Wilson, 1982) refines the concept of PBE. We do not provide more details on this refinement in order not to distract the reader from the main messages of the paper. Yet we hasten to point out that Theorem 1 also applies to this refinement, and that all PCE found in our examples below satisfy this refinement.

Next we define how players choose their strategies. When making a choice at a given information set, the choices at all other information sets are treated as given according to the players' strategies. The difficulty of making a decision at ϕ_i is that the player does not know which belief in the set of beliefs $B_{\phi_i}(\beta)$ should be used to evaluate the expected payoff. We resolve this issue by assuming the player chooses a best compromise. This is an action that is never too far from the best action under each belief in $B_{\phi_i}(\beta)$.

Formally, consider a pair (s,β). Denote by $\bar{u}_i(s_{\phi_i}|\phi_i,s,b_i)$ the expected payoff of player i from choosing an action $s_{\phi_i}\in\mathscr{A}_{\phi_i}$ in an information set ϕ_i under the belief b_i over the decision nodes in ϕ_i, assuming that the play is given by s elsewhere in the game. The payoff difference

$$\sup_{x_i\in\mathscr{A}_{\phi_i}} \bar{u}_i(x_i|\phi_i,s,b_i)-\bar{u}_i(s_{\phi_i}|\phi_i,s,b_i)$$

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

$$l(s_{\phi_i}|\phi_i,s,\beta)=\max_{b_i\in B_{\phi_i}(\beta)}\left(\sup_{x_i\in\mathscr{A}_{\phi_i}} \bar{u}_i(x_i|\phi_i,s,b_i)-\bar{u}_i(s_{\phi_i}|\phi_i,s,b_i)\right).$$

So the maximum is evaluated over all beliefs of player i at ϕ_i.

Player i makes a decision that minimizes the maximum loss. Such a choice is called a *best compromise*. Formally she chooses an element of

$$\arg\min_{s_{\phi_i}\in\mathscr{A}_{\phi_i}} l(s_{\phi_i}|\phi_i,s,\beta) \tag{1}$$

at each of her information sets ϕ_i. In equilibrium s^*, this means that she chooses $s^*_{\phi_i}\in\arg\min_{s_{\phi_i}\in\mathscr{A}_{\phi_i}} l(s_{\phi_i}|\phi_i,s^*,\beta)$. Hence, when computing the maximum loss and finding the best compromise, each player assumes that the behavior is given by s^* at all other information sets, including her own. Thus the players anticipate their own choices at subsequent information sets, which is known as *sophisticated behavior* and is closely related to *consistent planning* (Strotz, 1955; O'Donoghue and Rabin, 1999; Siniscalchi, 2011). Each player deals with her ambiguity about the true state by choosing best compromises. At the same time, a player acknowledges the fact that she will be facing ambiguity at later information sets, and hence anticipates her equilibrium choices there. This leads to our equilibrium concept that is motivated by complaints and compromises.

Definition 2. A pair (s^*,β^*) is called a *perfect compromise equilibrium (PCE)* if
(a) each player chooses a best compromise in each of her information sets;
(b) the belief system β^* is consistent under the strategy profile s^*.

We begin by establishing the existence of PCE in finite extensive-form games. A game $(N,\mathcal{G},\Omega,(\Pi_1,...,\Pi_n),(u_1,...,u_n))$ is called finite if N and Ω are finite, Π_i is finite for each $i\in N$, and every information set ϕ_i of the game tree $\mathcal{G}$ is associated with a finite set of pure actions A_{ϕ_i}. The set of actions available to player i in that information set comprises all the mixed actions, so $\mathscr{A}_{\phi_i}=\Delta(A_{\phi_i})$.

Theorem 1. *In a finite game, a perfect compromise equilibrium exists.*

The proof is in Appendix A.1.

2.2. **Discussion.** 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 minmax regret, thus having an axiomatic underpinning (Milnor, 1954; Hayashi, 2008; Puppe and Schlag, 2009; Stoye, 2011). Traditionally, minmax regret is evaluated ex-post after all uncertainty is resolved. In contrast, to model a compromise in the face of several beliefs, we consider the loss attained at the interim (at a given information set) for a given belief. Stoye's (2011) axioms continue to hold from this interim 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 treated as a move of 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 maxmin utility criterion can be used to model pessimism or cautiousness, a world in which the player always anticipates the worst outcome.

Planning and Updating. By nature of a sequential game, a player's perspective can change during the game. Future choices that look optimal today might not be optimal when the actual choice has to be made. To account for the changing perspectives, we assume that the players plan ahead what they and the others will do. The players' strategies are taken as given, and future choices are determined by these strategies.

An alternative approach would be to design a solution concept where a player's plan of actions does not depend on when this plan is made. This has been an objective in the related literature on maxmin utility and related ambiguity models (e.g., Epstein and Schneider, 2003; Wang, 2003; Hanany and Klibanoff, 2007; Hanany, Klibanoff and Mukerji, 2020; Pahlke, 2022). However, the insights of that literature are that this leads to substantial constraints to what priors are allowed, as in Pahlke (2022), or dictate a specific way of aggregating expected utilities under multiple priors, as in Hanany, Klibanoff and Mukerji (2020), that collapses back to PBE when the ambiguity is genuine (that is, when all the priors are degenerate).

Strategic Certainty. PCE assumes common knowledge of the equilibrium profile of strategies, just like PBE does. So there is strategic certainty. Players have a common belief of how others react in equilibrium to their information, they are only ambiguous about what information others actually have. However, PCE can be also used to incorporate strategic uncertainty as follows. The situation where a player, call her A, is ambiguous about the strategy of another player, call him B, is interpreted as ambiguity of A about some information that is private to B. That is, had A known everything about B, she would have had certainty about B's strategy. This way, any strategic uncertainty can be reinterpreted as informational uncertainty or ambiguity about the state of the world.

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 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 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 state of the world. 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 PCE is related to 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 $a_i \in \mathscr{A}_{\phi_i}$ at an information set ϕ_i is *strictly dominated* for player i if there exists another action $x_i \in \mathscr{A}_{\phi_i}$ 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 action that strictly dominates a_i will achieve a strictly lower loss for each belief, and hence its maximal loss will be strictly smaller. Thus, a strictly dominated action cannot be a part of a PCE. This argument can be iterated, so any iterated strictly dominated action cannot be a part of a PCE.

3. Examples

We illustrate our solution concept with a few economic examples that are prominent in the literature. We consider Cournot and Bertrand duopoly, public good provision, Akerlof's market for "lemons", Spence's job market signaling, bilateral trade with common value, and forecasting.[3] The examples presented in this section are self-contained as they do not require knowledge of the formalities presented in Section 2.

We are particularly interested in understanding strategic play under uncertainty when the players cannot or are unwilling to assess the likelihood of different states of the world at the beginning of the game. Formally, players can only have degenerate priors that put probability one on a single state of the world. We call this *genuine ambiguity.*

Apart from the market for "lemons" and forecasting, the examples presented below deal with genuine ambiguity. Therein, ambiguity is specified 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. In addition, we assume that 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.

3.1. **Cournot Duopoly with Unknown Demand.**

We investigate two firms that compete in quantities when neither firm knows the demand. We show that in a perfect compromise equilibrium the firms respond by slight increase of their quantities when they face such uncertainty. Each firms' potential loss is small relative to the case when it knows the demand exactly.

Consider 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

[3]An alternative forecasting model is presented in Appendix B.

number of units $q_i \geq 0$ to produce. Choices are made simultaneously. The firms face an inverse demand function $P(q_1 + q_2)$. 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 that the other firm produces 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* (PCE) 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.

Remark 1. 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).

Let us discuss the strategic concerns underlying the PCE in this game. Consider firm i who faces unknown demand and is deciding about how much to produce. This firm 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 best compromise q_i^* balances these two losses, assuming that the other firm follows its equilibrium strategy q_{-i}^*.

Our equilibrium analysis can shed light on how the firms respond 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,

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

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 quantities. There is a pressure to increase the quantity to account for the possibility of higher demand, and to decrease it to account for the possibility of lower demand. As a result, the quantity does not change very much. In fact, it increases slightly due to a larger pie size when the demand is high.

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. For example, 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 when allowing for a 10% error in the demand specification. Thus, uncertainty does not have a substantial impact on performance.

3.2. **Bertrand Duopoly with Private Costs.** We investigate two firms that compete in prices when the cost of the competitor is unknown. We show that in a perfect compromise equilibrium the firms charge prices above their marginal cost and hence make profits when they face such uncertainty. Moreover, the markup of each firm is decreasing in its own cost.

Consider two firms $i = 1, 2$ that produce a homogeneous good. They choose prices p_1 and p_2 simultaneously. The consumers only buy from the firm that offers a lower price. 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 knows its own marginal cost but not that of its competitor. It is common knowledge that the marginal costs belong to a given interval, so

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

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 the price $p_i^* = s_i^*(c_i)$ 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* (PCE) 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

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

The proof is in Appendix A.3.

Remark 2. 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).

Let us discuss the strategic concerns underlying the PCE in this game. For the sake of argument, suppose that the PCE price is strictly increasing in the cost. Each firm i that chooses a price above its marginal cost worries about two possible situations. It could be that the competitor has a weakly lower cost, and hence charges 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 competitor's price marginally undercuts p_i. Alternatively, it could be that the competitor has a higher cost and hence charges a higher price, $p_{-i} > p_i$. Thus, unless p_i is already profit maximizing, firm i is losing profit by charging too little. The greatest such loss occurs when the competitor's price is the highest possible (attained when $c_{-i} = \bar{c}$). The best compromise $p_i = s_i^*(c_i)$ balances these two losses, assuming that the competitor follows its equilibrium strategy.

Note that the firm's worry about losing profit when the competitor happens to have high cost leads to best compromise pricing above marginal cost. It is the upper bound on the competitor's possible cost that influences pricing. The lower bound plays no role, as the worst case for the firm is attained when the competitor's cost (and thus price) is only marginally lower. This leads to an upward pressure on pricing, the more so the higher the upper bound on the competitor's cost and the smaller the firm's cost. In particular, we obtain that the markup, $s_i^*(c_i) - c_i$, is decreasing in cost c_i.

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,

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

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,

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

Moreover, for small ε we obtain

$$\frac{\mathrm{d}s_i^\varepsilon(c_0)}{\mathrm{d}\varepsilon} = \frac{c_0}{4} + O(\varepsilon).$$

Thus, the firms optimally respond to the growing uncertainty about the demand by increasing their prices. The rate of increase is substantial as it does not vanish when ε tends to 0.

Next, consider the associated maximum losses as shown in (6). We find

$$l_i(s_i^\varepsilon(c_i), s_{-i}^\varepsilon, c_i) \leq \frac{2\varepsilon}{3} - \frac{\varepsilon^2}{9}, \quad i = 1, 2.$$

The maximum losses increase approximately linearly as the uncertainty increases. For example, if $\varepsilon = 0.1$, then the maximum losses are bounded by 0.07. So the firms lose no more than about 7% of the maximum profit when allowing for a 10% error about the rival's marginal cost.

3.3. **Public Good Provision.** Here we investigate how to provide a discrete public good when beneficiaries of the good are uncertain about the valuations of others. We assume that the beneficiaries fund the cost of provision with their own contributions. There are no external subsidies, and thus the Vickrey-Clarke-Groves (VCG) mechanism is not feasible in our setting (e.g., d'Aspremont and Gérard-Varet, 1979). Without making any distributional assumptions, we are able to analyze several focal mechanisms. Interestingly, our analysis of the perfect compromise equilibrium shows that these mechanisms can be applied under even less information than we assume. One does not depend on the number of players, one does not depend on the cost of the public good, and one does not depend on either of these parameters.

Consider n agents. Each agent has a private value $v_i \in [0, \bar{v}]$ for a public good. Agents know their own values of the good, but not those of the others. Each agent i chooses how much to contribute for the public good provision. Let $x_i \in [0, \bar{v}]$ be agent i's contribution. The agents make their choices simultaneously.

A commonly known cost of providing the public good is $c > 0$. To avoid considering multiple cases, we assume that this cost is not too high, specifically,

$$\frac{c}{n-1} \leq \frac{\bar{v}}{2}. \tag{7}$$

The payoffs are as follows. If the sum of the contributions does not cover the cost, so $\sum_{i=1}^{n} x_i < c$, then the public good is not provided, and the agents' contributions are returned to them. In this case 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 obtains the value of the good net of the contribution. In addition, the agents may be refunded the excess contribution, $\sum_{i=1}^{n} x_i - c$. The payoff of each agent i is

$$v_i - x_i + r_i(x),$$

where $r_i(x)$ is a refund to agent i that depends on the profile of contributions $x = (x_1, ..., x_n)$. We compare three simple refund rules.

(i) *No-refunds rule.* The excess contribution is not refunded to the agents, so

$$r_i(x) = 0, \quad i = 1, ..., n. \tag{8}$$

(ii) *Equal-split rule.* The excess contribution is divided equally among to the agents, so

$$r_i(x) = \frac{1}{n}\left(\sum\nolimits_{j=1}^{n} x_j - c\right), \quad i = 1, ..., n. \tag{9}$$

(iii) *Proportional rule.* The excess contribution is divided proportionally to the agents' individual contributions, so

$$r_i(x) = \left(1 - \frac{c}{\sum_{j=1}^{n} x_j}\right) x_i, \quad i = 1, ..., n. \tag{10}$$

Let $s_i(v_i)$ be a strategy of agent i, so $x_i = s_i(v_i)$ specifies the contribution of agent i whose private value is v_i. We restrict attention to strategies that are symmetric and undominated. Specifically, we assume that

$$s_i(v) = s_j(v) \text{ and } s_i(v) \leq v \text{ for all } v \in [0, \bar{v}] \text{ and all } i, j \in \{1, ..., n\}. \tag{11}$$

The assumption that the strategies are symmetric is substantive, as we rule out potential asymmetric equilibria. The assumption that the strategies are undominated is inconsequential for the results and introduced for notational convenience.

An agent i's *maximum loss* of choosing contribution x_i when the other agents choose a profile of contributions $s_{-i}(v_{-i})$ describes how much more payoff agent i could have obtained if she had known the true values of everybody else, anticipating that they follow their strategies. To determine the maximum loss, observe that agent i worries about two possible situations. It could be that the total contribution is marginally below c, so $x_i + \sum_{j \neq i} s_j(v_j) = c - \varepsilon$ 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 $v_i - x_i$. Alternatively, it could be that all other agents contribute enough to cover c, so $\sum_{j \neq i} s_j(v_j) \geq c$. Thus the agent could have contributed nothing and still received the good. In this case the loss is the amount of contribution net of the refund, $x_i - r_i(x_i, s_{-i}(v_{-i}))$. Agent i's maximum loss is thus given by

$$l_i(x_i, s_{-i}; v_i) = \sup_{v_{-i} \in [0,\bar{v}]^{n-1}} \max \{v_i - x_i, x_i - r_i(x_i, s_{-i}(v_{-i}))\}.$$

Agent i's *best compromise* given v_i is a strategy $s_i^*(v_i)$ that achieves the lowest maximum loss for a given strategy profile s_{-i} of the other agents:

$$s_i^*(v_i) \in \arg\min_{x_i \in [0, v_i]} l_i(x_i, s_{-i}; v_i).$$

A strategy profile $s^* = (s_1^*, ..., s_n^*)$ is a *perfect compromise equilibrium* (PCE) if each agent i chooses a best compromise given her value v_i when facing the strategy profile s_{-i}^* of the other agents.

In this application we are interested in how the agents' equilibrium behavior and total efficiency (welfare) changes in PCE induced by different refund rules. We measure the efficiency of a strategy profile s by the maximum welfare loss as compared to the complete information case. Because $s_i(v) \leq v$ by assumption (11), the welfare loss only emerges in the case of $\sum_i s_i(v_i) < c \leq \sum_i v_i$ where the good is not provided when it is efficient to do so. Our inefficiency measure is denoted by $L(s)$ and is given by

$$\begin{aligned} L(s) = & \sup_{(v_1,...,v_n) \in [0,\bar{v}]^n} \sum\nolimits_{i=1}^{n} v_i - c \\ & \text{subject to} \sum\nolimits_{i=1}^{n} s_i(v_i) < c \leq \sum\nolimits_{i=1}^{n} v_i. \end{aligned} \tag{12}$$

We now characterize the PCE and the associated welfare losses for each of the three refund rules.

Proposition 3. *For each of the three refund rules there is a unique PCE strategy profile $s^* = (s_1^*, ..., s_n^*)$ that satisfies assumption* (11)*. For each $i = 1, ..., n$ and each $v_i \in [0, \bar{v}]$,*

(i) if $r_i(x)$ is the no-refunds rule, then

$$s_i^*(v_i) = \frac{v_i}{2} \quad \textit{and} \quad L(s^*) = c;$$

(iii) if $r_i(x)$ is the equal-split rule, then

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

(iii) if $r_i(x)$ is the proportional rule, then

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

The proof is in Appendix A.4.

Note that under each rule the player contributes at least half of her value. Note also that

$$c > \frac{n}{n+1}c > \frac{n-1}{n}c.$$

So, the equal split rule confers a smaller welfare loss as compared to the other two rules. However, this difference between the mechanisms is small, and the welfare loss is close c when the number of beneficiaries is large. This is reminiscent of the VCG mechanism in this problem where the expected subsidy of the designer is close to c when n is large.

To summarize, we have investigated the behavior in a PCE when players have little information about the values of others. We find for each of the refund rules that even less information is needed. The equal-split rule leads to equilibrium behavior that requires no information about the cost c. Hence this rule can also be applied when the cost is unknown. The proportional rule leads to equilibrium behavior that does not depend on the number of players n and hence can also be applied when n is unknown. The no-refund rule leads to equilibrium behavior that does not depend on either of these two parameters. This reliance on minimal information makes these refund rules, and our solution concept in general, practically appealing.

3.4. **Market for "Lemons".** The following example presents a variation of the market for "lemons" (Akerlof, 1970) where a buyer is ambiguous about the quality of a car offered by a seller. We show that trade can only occur in a perfect compromise equilibrium if the car is offered at a (pooling) price that does not depend on the quality of the car. Whenever there is trade, the buyer randomizes between buying or not buying the car. This behavior reflects the buyer's desire to balance his losses from buying a low-quality car and not buying a high-quality car. This example highlights the difference between a PCE and a PBE. Namely, in a PBE the buyer's best response is a pure strategy, except when the buyer is exactly indifferent between buying and not buying.

Consider a seller (she) and a buyer (he). The seller has a car whose quality is either high (θ_H) or low (θ_L). She observes the quality of the car and decides at what price $p \in [0, 1]$ to offer it for sale. The buyer observes the price but not the quality, and decides whether or not to buy the car at this price. Let v_H and v_L be the buyer's value of high and low quality car, respectively, and let c_H and c_L be the seller's cost of high and low quality car, respectively. Assume that

$$v_L = 0 < 1 < v_H \quad \text{and} \quad 0 < c_L < c_H < 1. \tag{13}$$

In words, no matter which price in $[0, 1]$ is asked, the buyer always wants to buy a high quality car and never wants to buy a low quality car. In contrast, depending on the price, the seller might be willing to sell both types of car, only a low quality car, or neither type of car.

The buyer's ambiguity about the quality of the car is captured by a set of K priors $\Pi = [\pi_1, ..., \pi_K]$, where $0 < \pi_1 < ... < \pi_K < 1$. When $K = 1$, the buyer has a single prior, so this model becomes the classic market of lemons with standard uncertainty of the buyer.[4]

The seller's strategy $\sigma_S^* : \{\theta_H, \theta_L\} \to \Delta([0,1])$ specifies for each type $\theta \in \{\theta_L, \theta_H\}$ a discrete probability distribution $\sigma_S^*(\cdot|\theta)$ over prices in $[0,1]$.[5] The buyer's strategy $\sigma_B^* : [0,1] \to [0,1]$ specifies for each price $p \in [0,1]$ a probability $\sigma_S^*(p)$ that the buyer buys the car. The beliefs are as follows. The seller knows the type of the car, so her beliefs are trivial. The buyer does not know the type but observes the price. So, for each price $p \in [0,1]$, the buyer has a distinct information set, denoted by ϕ_p. In this information set, the buyer updates each of his priors to obtain a set of posterior beliefs $\{\beta_{\phi_p}^*(\pi)\}_{\pi\in\Pi}$, where $\beta_{\phi_p}^*(\pi)$ denotes the posterior belief in the information set ϕ_p given a prior $\pi \in \Pi$. The buyer's belief system is given by $\beta^* = (\beta_{\phi_p}^*)_{p\in[0,1]}$.

We now define the conditions that a profile $(\sigma_S^*, \sigma_B^*, \beta^*)$ of strategies and beliefs must satisfy to be called a *perfect compromise equilibrium.* Because the seller has no ambiguity, the seller's choice of price must be her best response. The seller chooses with a positive probability only those prices that maximize the expected payoff given the buyer's strategy σ_B^*. Formally,

$$Supp(\sigma_S^*(\cdot|\theta_j)) \subset \arg\max_{p\in[0,1]} (p - c_j)\sigma_B^*(p) \quad \text{for each } j = H, L. \tag{14}$$

Next, we define the buyer's *maximum loss* and the associated *best compromise strategy.* For each probability $b \in [0,1]$ that the car has high quality, the buyer obtains $bv_H - p$ if he buys the car and zero if he does not. Thus, given a belief b and a price p, the optimal choice yields the payoff of $\max\{bv_H - p, 0\}$. The buyer's loss from a strategy $\sigma_B^*(p)$ describes how much more payoff the buyer could have obtained if he made the optimal choice under b instead of $\sigma_B^*(p)$. This loss is given by

$$\begin{aligned}&\max\{bv_H - p, 0\} - (bv_H - p)\sigma_B^*(p)\\ &\qquad = \max\big\{(bv_H - p)(1 - \sigma_B^*(p)), (p - bv_H)\sigma_B^*(p)\big\}.\end{aligned}$$

[4] To streamline the exposition, we consider a continuum of prices $[0,1]$, rather than a finite set as assumed in Section 2.1. The insights of this illustrative example do not change if we discretize the set of prices.

[5] We restrict attention to discrete probability distributions to simplify the derivation of posterior beliefs.

Consequently, the buyer's *maximum loss* from the choice $\sigma_B^*(p)$ in the information set ϕ_p given the seller's strategy σ_S^* and the belief system β^* is

$$l_B(\sigma_B^*(p)|\phi_p, \sigma_S^*, \beta^*) = \max_{b \in \{\beta^*_{\phi_p}(\pi)\}_{\pi\in\Pi}} \max\left\{(bv_H - p)(1-\sigma_B^*(p)), (p - bv_H)\sigma_B^*(p)\right\}.$$

Observe that among the set of the buyer's beliefs, only two – the highest and the lowest – are relevant for the calculation of this maximum loss. Let

$$\underline{b}(p) = \min_{\pi\in\Pi} \beta^*_{\phi_p}(\pi) \ \text{ and } \ \bar{b}(p) = \max_{\pi\in\Pi} \beta^*_{\phi_p}(\pi).$$

On the one hand, the probability of high quality can be high, so the buyer makes a loss by rejecting the offer. The greatest such loss occurs when the posterior belief is the highest, namely, when the prior $\pi \in \Pi$ is such that $\beta^*_{\phi_p}(\pi) = \bar{b}(p)$. On the other hand, the probability of high quality can be low, so the buyer makes a loss by buying the car. The greatest such loss occurs when the belief is the lowest, namely, when the prior $\pi \in \Pi$ is such that $\beta^*_{\phi_p}(\pi) = \underline{b}(p)$. Consequently, the maximum loss can be summarized as

$$l_B(\sigma_B^*(p)|\phi_p, \sigma_S^*, \beta^*) = \max\left\{(\bar{b}(p)v_H - p)(1-\sigma_B^*(p)), (p - \underline{b}(p)v_H)\sigma_B^*(p)\right\}.$$

The buyer's best-compromise strategy σ_B^* minimizes the maximum loss, so for each $p \in [0,1]$ the probability of buying the car satisfies

$$\sigma_B^*(p) \in \arg\min_{q\in[0,1]} \left(\max\left\{(\bar{b}(p)v_H - p)(1-q), (p - \underline{b}(p)v_H)q\right\}\right). \tag{15}$$

Next, consider the buyer's beliefs. For any price $p \in [0,1]$, the buyer has, potentially, multiple beliefs in his information set ϕ_p. In order to be a part of a PCE, the buyer's belief mapping must be consistent with the strategy σ_S^* of the seller. Specifically, each prior $\pi \in \Pi$ is transformed into a posterior belief $\beta^*_{\phi_p}(\pi)$ using Bayes' rule whenever possible, so

$$\beta^*_{\phi_p}(\pi) = \frac{\pi\sigma_S^*(p|\theta_H)}{(1-\pi)\sigma_S^*(p|\theta_L) + \pi\sigma_S^*(p|\theta_H)} \ \text{ if } (1-\pi)\sigma_S^*(p|\theta_L) + \pi\sigma_S^*(p|\theta_H) > 0, \tag{16}$$

and otherwise $\beta^*_{\phi_p}(\pi)$ can have any value in $[0,1]$.

In summary, a profile $(\sigma_S^*, \sigma_B^*, \beta^*)$ of strategies and beliefs is a *perfect compromise equilibrium* if it satisfies conditions (14), (15), and (16).

Before characterizing the PCE, we introduce the following notation. We say that a PCE *involves no-trade* if the buyer does not buy the car at any price that can be offered in equilibrium. Specifically, for each $p \in [0,1]$,

$$\text{if } \pi\sigma_S^*(p|\theta_H) + (1-\pi)\sigma_S^*(p|\theta_L) > 0 \text{ then } \sigma_B^*(p) = 0.$$

We say that a PCE is pooling on a single price if the seller offers the car at a fixed, nonrandom price irrespective of the car quality. Formally, a profile $(\sigma_S^*, \sigma_B^*, \beta^*)$ is *a PCE that is pooling on single price* p^* if

$$\sigma_S^*(p^*|\theta_H) = \sigma_S^*(p^*|\theta_L) = 1.$$

Observe that in such a PCE the buyer's beliefs must satisfy the following. When observing p^*, the buyer's set of posterior beliefs is equal to the set of priors, so $\beta^*_{\phi_{p^*}}(\pi) = \pi$ for each $\pi \in \Pi$. When observing $p \neq p^*$, the buyer's beliefs and behavior are such that no type of seller has incentive to deviate to p. For example, the buyer is certain that the car has low quality and does not buy it.

In addition, given such beliefs, the buyer's probability to buy the car at the equilibrium price p^* can be easily derived from (15). When there is a single prior, so $\Pi = \{\bar{\pi}\}$, it satisfies

$$\sigma_B^*(p^*) \in \begin{cases} \{1\} & \text{if } p^* < \bar{\pi} v_H, \\ [0,1] & \text{if } p^* = \bar{\pi} v_H, \\ \{0\} & \text{if } p^* > \bar{\pi} v_H. \end{cases} \tag{17}$$

When there are multiple priors, so $\Pi = \{\pi_1, ..., \pi_K\}$ with $K > 1$, it is given by

$$\sigma_B^*(p^*) = \begin{cases} 1 & \text{if } p^* < \pi_1 v_H, \\ \frac{\pi_K v_H - p^*}{(\pi_K - \pi_1) v_H} & \text{if } \pi_1 v_H \leq p^* \leq \pi_K v_H, \\ 0 & \text{if } p^* > \pi_1 v_H. \end{cases} \tag{18}$$

Note that the probability that the buyer buys the car at price p^* is also equal to the probability that trade takes place. In particular, such a PCE need not involve trade.

Proposition 4. *Every PCE either involves no trade, or is pooling on a single price, or has both of these properties.*

The proof is in Appendix A.5. Intuitively, suppose by contradiction there is a PCE that has trade and is not pooling on a single price. Specifically, suppose that there are at least two prices that can be offered in equilibrium by the seller such that the buyer buys with positive probabilities at these prices. Because $c_L < c_H$, so the low-type car is cheaper than the high-type car, it follows that whenever the seller of one type is indifferent among several prices and plays a mixed strategy, the seller of the other type strictly prefers a single price in this set. This means that all but one equilibrium prices reveal the car type to the buyer. If the revealed type is low, then the buyer does not buy. If the revealed type is high, then the buyer buys with probability one, but then the low-type seller would want to deviate and choose that price to pretend to be the high type.

Proposition 4 has established that every PCE that has trade must be pooling on a single price. We now characterize the set of equilibrium prices for such PCE, both for the case of a single prior where it coincides with a PBE and for the more general case of multiple priors.

Proposition 5.
(a) Suppose that there is a single prior, so $\Pi = \{\bar{\pi}\}$*. There exists a PBE that has trade and is pooling on price* p^* *if and only if* $p^* \in [c_H, \bar{\pi} v_H]$*.*
(b) Suppose that there are multiple priors, so $\Pi = \{\pi_1, ..., \pi_K\}$ *with* $K > 1$*. There exists a PCE that has trade and is pooling on price* p^* *if and only if* $p^* \in [c_H, \pi_K v_H)$*.*

The proof is in Appendix A.6.

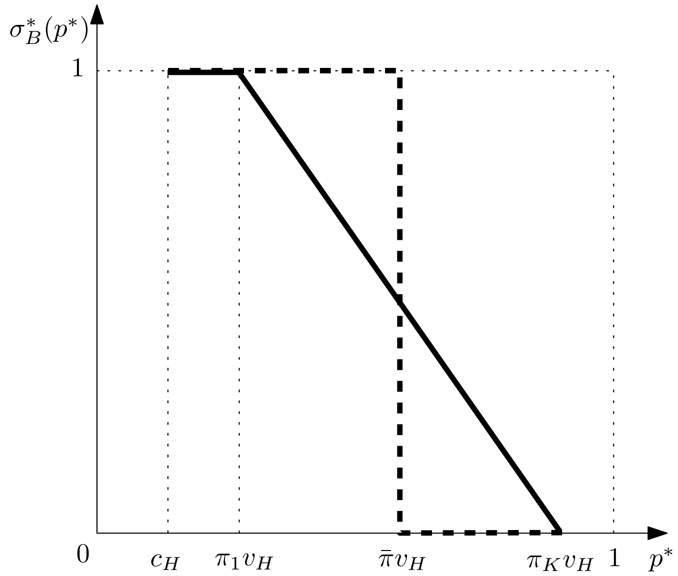

FIGURE 1. Trade under different pooling equilibrium prices in the market of "lemons". Dashed line shows the probability of trade under PBE with a single prior $\bar{\pi}$. Solid line shows the probability of trade under PCE with a set of priors $\Pi = \{\pi_1, ..., \pi_K\}$.

The crucial difference between the PCE under multiple priors and the PBE (that has a single prior) is in the equilibrium behavior of the buyer. Figure 1 illustrates how the buyer's equilibrium behavior $\sigma_B^*(p^*)$, which is equal to the probability of trade, depends on the equilibrium price p^* in the continuum of the pooling equilibria as the price increases from c_H to $\pi_K v_H$. The dashed line, which corresponds to equation (17), shows what happens in the PBE under a single prior $\bar{\pi}$. The buyer compares the price p^* with the expected value $\bar{\pi} v_H$, and then buys the car when $p^* < \bar{\pi} v_H$, is indifferent when $p^* = \bar{\pi} v_H$, and does not buy the car when $p^* > \bar{\pi} v_H$. In contrast, the probability of trade is different in the PCE under a set of priors Π. This is described by equation (18) and depicted by the solid line in Figure 1. When the price is below the most pessimistic expected value, so $p^* < \pi_1 v_H$, then the buyer buys the car. Otherwise, as long as the price does not exceed the most optimisitic expected value, $\pi_1 v_H < p^* < \pi_K v_H$, the buyer smoothes out his response by randomizing his choice, where the probability of buying decreases as p^* goes up. This shows the feature of PCE that captures the buyer's hesitancy or the desire to balance losses under different contingencies when facing multiple priors.

3.5. **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. We find that, unlike in the classic setting with a single prior, there is no longer a clear separation between workers with different levels of productivity.

Consider a single worker and two firms. The worker has productivity $\theta \in [0,1]$. 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

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

Each firm i's payoff is given by

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

where γ_i is the probability that the worker chooses firm i when she is indifferent between the offers of two firms. We refer to $\gamma = (\gamma_1, \gamma_2)$ as *tie-breaking type*, where

$$\gamma \in \Delta_2 = \{(\gamma_1, \gamma_2) \in [0,1]^2 : \gamma_1 + \gamma_2 = 1\}.$$

The worker knows her productivity type θ, her cost of high education c, and her tie-breaking type γ. The firms know none of these. 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 $a-\theta-\varepsilon/2$ and $a - \theta + \varepsilon/2$, where a and ε are commonly known parameters. Parameter a is interpreted as the benchmark cost of education of the lowest productivity type $\theta = 0$, and ε is related to the amount of uncertainty about the cost of education for a given productivity. Formally, the firms know that (θ, c, γ) belongs to the set Ω given by

$$\Omega = \left\{(\theta, c, \gamma) \in [0,1] \times \mathbb{R}_+ \times \Delta_2 : a - \theta - \frac{\varepsilon}{2} \leq c \leq a - \theta + \frac{\varepsilon}{2}\right\}. \tag{19}$$

We assume that

$$a \in [1, 3/2],\ \varepsilon \in [0, 1/3],\ \text{and } 0 \le a - 1 - \frac{\varepsilon}{2} \le \frac{1}{2} - \frac{3\varepsilon}{2}. \tag{20}$$

The last condition implies that the lower bound on the cost of education for the highest productivity type $\theta = 1$, which is given by $a - \theta - \varepsilon/2 = a - 1 - \varepsilon/2$, is nonnegative but not too large. We impose this condition to reduce the number of cases to consider, thus simplifying the exposition.

The worker's strategy $e^*(\theta, c, \gamma)$ describes her choice of the education level for each profile $(\theta, c, \gamma) \in \Omega$. 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*. Formally, these are the firms' degenerate beliefs that put probability one on specific states. Speculated states are the profiles (θ, c, γ) 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 $(\theta, c, \gamma) \in S_i(e)$ if $e^*(\theta, c, \gamma) = e$.

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

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

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:

$$w_i^*(e) \in \arg\min_{w_i \ge 0} l_i(w_i, w_{-i}^*; e). \tag{21}$$

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

$$e^*(\theta, c, \gamma) \in \arg\max_{e \in \{e_L, e_H\}} v(w_1^*(e), w_2^*(e), e; \theta, c). \tag{22}$$

A profile $(e^*, w_1^*, w_2^*, S_1, S_2)$ of strategies and speculated states is a *perfect compromise equilibrium* (PCE) if two conditions hold. First, the strategies satisfy (21) and (22), 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 $(\theta, c, \gamma) \in \Omega$. 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 6.
(i) There exists a pooling PCE in which the worker chooses low education, so

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

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) There exists a separating PCE in which the worker chooses high education if and only if her cost c is at most $\frac{1}{2}$, so for all $(\theta, c, \gamma) \in \Omega$

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

and the firms' wages are given by

$$w_i^*(e_H) = \frac{2a+1}{4} \quad \textit{and} \quad w_i^*(e_L) = \frac{2a-1+\varepsilon}{4}, \quad i = 1, 2. \tag{23}$$

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

$$\theta \in \left[0, a - \frac{1}{2} + \varepsilon\right] \quad \textit{if } e = e_L, \quad \textit{and} \quad \theta \in \left[a - \frac{1}{2}, 1\right] \quad \textit{if } e = e_H. \tag{24}$$

The proof is in Appendix A.7.

Let us discuss the strategic concerns underlying these PCE. Each firm i, when facing unknown productivity of the worker and deciding about the wage offer w_i, 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 prevent employing a worker whose productivity is below the wage. The greatest such loss occurs when the productivity is the lowest possible. The firm thus offers the best compromise wage that balances these the losses of not hiring a productive worker and hiring an unproductive worker, assuming that the other firm follows its equilibrium strategy.

The particular wage offer depends on the greatest and smallest productivities that are inferred from the level of education e that the worker chooses. In the pooling equilibrium, $e = e_L$ or $e = e_H$ do not provide any useful information, so all productivity types are possible. However, in the separating equilibrium, the firms believe that the productivity belongs to different intervals when observing different levels of education. For example, if $a = 7/6$ and $\varepsilon = 1/3$, then the firms believe that $\theta \in [0, 1]$ if the education is low, and that $\theta \in [2/3, 1]$ if the education is high. This leads to different wages charged in the separating equilibrium. In anticipation of the difference between wages associated with high and low education, the workers base their choice of education on their own cost, independent of their own productivity. Due to the heterogeneity in costs and productivity, a worker who chooses high education might be less productive than another worker who chooses low education, just because the cost of the former worker is lower. Without knowing the cost, one cannot predict which education level a worker with productivity between $a - 1/2$ and $a - 1/2 + \varepsilon$ will choose.

Parameter ε captures the degree of uncertainty about the productivity conditional on the cost of education, and thus, conditional on the chosen education level that reveals information about the cost. The higher the uncertainty, the less reliably the education level signals about productivity. This results in a higher wage to low educated workers and a smaller wage gap, making the differentiation of higher and lower productivity types less effective. To see why this occurs, observe that an increase in ε expands the sets of speculated states conditional on each education level. The threshold $\bar{\theta}_L$ increases and $\underline{\theta}_H$ decreases, so both $[0, \bar{\theta}_L]$ and $[\underline{\theta}_H, 1]$ expand. This results in a higher wage to low educated workers and a lower wage to higher educated workers. In turn, this reduction in the wage gap makes more types of workers prefer low education, thus shifting both thresholds $\bar{\theta}_L$ and $\underline{\theta}_H$ upward. This leads to even higher wages for low educated workers, although it partially offsets the initial negative effect on the wage for high educated workers.

In summary, our analysis shows that Spence's insights carry over to this novel setting that does not rely on probabilities and distributions. It also shows how easy the analysis captures richer uncertainty about workers than that in the classic Spence's framework. A particular consequence of richer uncertainty is that, unlike in the traditional model, the strict separation between the productivity of the workers choosing different education levels in the separating equilibrium no longer holds in our setting.

3.6. **Bilateral Trade with Common Value.** In this example we consider how prices emerge in bilateral trade when two traders value the good the same and have uncertainty about this value. Bilateral trade is modeled by assuming that the seller sets the price and then the buyer decides whether to buy at this price. This yields a signaling game where the price set by the seller may reveal information about the value to the buyer. We show that in a perfect compromise equilibrium trade can occur, even though the buyer makes inference about the seller's information from observing the price.

The fact that trade can occur stands in stark contrast to the no-trade theorem under common values as predicted by PBE (Milgrom and Stokey, 1982). This extends the insights on the possibility of trade with common value under multiple priors found by Billot et al. (2000), Kajii and Ui (2006), and Rigotti, Shannon and Strzalecki (2008). Our novelty is that we have an explicit sequential move game that describes how trade takes place. Thus the occurrence of trade has to incorporate the information revealed when the traders take actions. In contrast, the existing literature on trade under multiple priors has identified whether there are allocations of goods that benefit both traders. However, it remains unclear whether this trade would take place in a market game where offers are being made. The worry is that information revealed by the choice of an offer might eliminate the incentive to trade.

Consider the following model of bilateral trade. A seller wants to sell an indivisible good to a buyer. The value of the good is the same for each of the traders, it is denoted by v. 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.[6]

The traders commonly know an upper and a lower bound on the possible values. These will be normalized to be 0 and 1, so $v \in [0, 1]$. Each trader also has private information about this value. Specifically, the seller knows the value belongs to $[x_0, x_1]$, and the buyer knows it belongs to $[y_0, y_1]$. As both have correct information, it follows that

$$v \in [x_0, x_1] \cap [y_0, y_1]. \tag{25}$$

An interpretation is that each trader privately consults an independent expert. The expert privately informs the trader about the most pessimistic and the most optimistic assessments of the true value.

The traders have no prior beliefs about each other's information. Instead, all they know is the setting. Specifically, the seller with private information $[x_0, x_1]$ knows that the buyer's private information $[y_0, y_1]$ has a nonempty intersection with $[x_0, x_1]$, and the buyer holds symmetric knowledge.

Trade occurs according to the following take-it-or-leave-it protocol. First, the seller chooses a price $p \in [0, 1]$. Then, the buyer decides whether or not to trade at this price, and the game is over.

Let us describe the traders' strategies. The seller's strategy p^*, referred to as pricing rule, specifies a price $p^*(x_0, x_1) \in [0, 1]$ given the seller's information

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

$[x_0, x_1]$. The buyer's strategy α^*, referred to as acceptance rule, specifies a decision $\alpha^*(p, y_0, y_1) \in \{0, 1\}$ whether or not to trade at price p given the buyer's 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 chosen by the seller. This is formalized with the concept of speculated values. These are the values of v that the buyer thinks are possible after he observes the price chosen by the seller. Formally, a set of speculated values, denoted by $V_b(p, y_0, y_1)$, is a nonempty subset of $[y_0, y_1]$ that depends on the price p.

In equilibrium, the set of speculated values comprises the values that can emerge under a given pricing rule p^* of the seller. Formally, for given $[y_0, y_1] \subset [0, 1]$ and $p \in [0, 1]$, we say that $V_b(p, y_0, y_1)$ is *consistent with pricing rule* p^* if the following conditions hold. If price p can occur under the pricing rule p^*, that is, if there exists $[x_0, x_1] \subset [0, 1]$ whose intersection with $[y_0, y_1]$ is nonempty such that $p^*(y_0, y_1) = p$, then

$$V_b(p, y_0, y_1) = [y_0, y_1] \cap \left(\bigcup_{0 \le x_0 \le x_1 \le 1} \left\{ [x_0, x_1] : p^*(x_0, x_1) = p \text{ and } [x_0, x_1] \cap [y_0, y_1] \neq \varnothing \right\} \right). \quad (26)$$

Otherwise, if p cannot occur under the pricing rule p^*, then $V_b(p, y_0, y_1)$ can be an arbitrary nonempty subset of $[y_0, y_1]$.

The buyer's maximum loss from his choice $\alpha \in \{0, 1\}$ given price p and 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 rule 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 the acceptance rule α^*. Each trader's *best compromise* is a choice that achieves the lowest maximum loss for a given strategy of the other trader. A profile (p^*, α^*, V_b) is a *perfect compromise equilibrium* (PCE) if each trader chooses a best compromise given the strategy of the other trader, and the buyer's set of speculated values $V_b(p, y_0, y_1)$ is consistent with the seller's pricing rule p^*.

Proposition 7. *For each lower bound on the price $p_0 \in [1/2, 1]$, a perfect compromise equilibrium is given as follows. The seller asks*

$$p^*(x_0, x_1) = \max\left\{\frac{x_0 + x_1}{2} + \frac{1 - x_1}{4}, p_0\right\}. \tag{27}$$

If the seller asks $p \geq p_0$, then the buyer speculates that

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

and accepts this price if and only if

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

If the seller asks $p < p_0$, then the buyer speculates that $v \in V_b(p, y_0, y_1) = \{y_0\}$ and accepts this price if and only if $p \leq y_0$.

The formal proof is in Appendix A.8.

Let us discuss the strategic concerns underlying this PCE. Note that in absence of any strategic interaction, the buyer would buy if and only if the price is below his midpoint $(y_0 + y_1)/2$. This is because the buyer is balancing two worst cases where the true value is at the extreme points of the interval $[y_0, y_1]$. Similarly, the seller would sell if and only if the price is above her midpoint $(x_0 + x_1)/2$.

In the strategic trade setting, the seller would want to sell the good at least at her midpoint. However, she is worried that the buyer might be extremely optimistic and willing to accept even higher price. Hence, as a compromise, she optimally sets the price above her midpoint.

In equilibrium, the buyer agrees to trade when the price is below his midpoint. He does not take into account the information contained in the price as this information is coarser that what he already knows when the price is above a given lower bound $p_0 \geq 1/2$. This leads to a range of equilibrium prices, $[p_0, 1]$, which is narrow enough so that the seller who anticipates the behavior of the buyer never reveals information that the buyer would want to use. Note that when $p_0 = 1$, we obtain the no-trade equilibirum where the seller always sets $p = 1$ and the buyer rejects it (except when he is certain that the value is 1).

We obtain the possibility of trade under common values when $p_0 < 1$. The trade is possible because the traders do not want to miss out on a good trade opportunity, but also they do not want to make an unprofitable deal. They make compromise decisions so that they do not lose too much either way.

3.7. **Forecasting.** In this final example we are interested in how to forecast a random variable with known mean but unknown dispersion based on a signal with known distribution. This is a classical updating problem under multiple priors. We find that the best compromise forecast is given by a convex combination between the signal and the known mean. The weights depend on the precision of the signal and on the bound on the dispersion of the random variable. For instance, it is close to the midpoint between the signal and the known mean when the random variable is very dispersed.

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

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

where $a \in [0,1]$ denotes a forecast.

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

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

The agent bases her forecast on a noisy signal z of the parameter of interest θ. She knows how this signal is generated. Specifically we assume that signal z reveals the true value θ with probability $1-\varepsilon$ and is drawn uniformly from $[0,1]$ with probability ε. So, the conditional distribution of z given θ is

$$G_\varepsilon(z|\theta) = \begin{cases} \varepsilon z, & \text{if } z < \theta, \\ 1-\varepsilon+\varepsilon z, & \text{if } z \geq \theta. \end{cases} \tag{28}$$

Let $\mathbb{E}_{F,G_\varepsilon}[\cdot|z]$ denote the conditional mean of θ when the agent speculates that θ is distributed according to F. The *maximum loss* of a forecast $a \in [0,1]$ given a signal $z \in [0,1]$ is

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

It describes how much greater payoff the agent could have obtained if she knew the distribution F. A forecast $a^*(z)$ is a *best compromise* if it achieves the smallest maximum loss,

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

Proposition 8. *The best compromise forecast is given by*

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

where

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

The proof is in Appendix A.9.

Let us present some intuition behind Proposition 8. 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 best compromise forecast depends on the signal's precision ε and on the degree of dispersion δ of the variable of interest. We now show how each of these two parameters influences the best compromise forecast.

Fix the degree of dispersion δ. On the one hand, when the signal is very precise, then the best compromise forecast is close to the signal. This is because a^* is continuous in ε and $\lim_{\varepsilon \to 0} a^*(z) = z$. On the other hand, when the signal is very noisy, then the best compromise forecast 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, $\lim_{\delta \to 0} a^*(z) = (\theta_0 + z)/2$. This is because the best compromise balances two extreme situations. It could be 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 agent seeks the 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 $\varepsilon = \delta = 0$.

In summary, when using a signal to update information about a random variable with known mean but unknown dispersion, the best compromise forecast has a simple form. It is a convex combination of the known mean and the observed signal. Under extreme uncertainty where the imposed bounds on the density vanish, the forecast is particularly simple, namely, it is equal to the midpoint between the mean and the signal.

In Appendix B we consider an alternative setting, where the agent knows the distribution of the random variable but she does not know the conditional distribution of the noisy signal. We also deal with the case where the agent is ambiguous about both distributions (Remark 3).

4. Conclusion

This paper adds to the literature aimed at a better understanding of how players can deal with uncertainty in dynamic contexts without necessitating a single prior. We are particularly interested in allowing players to have only an intuitive understanding of uncertainty that can be expressed in terms of bounds. The uncertainty of a player is modeled by confronting the player with multiple "assessors", each of whom holding a different prior. The assessors process information and compute posterior beliefs independently, leading to updating prior by prior using Bayes' rule whenever possible. The player searches for a compromise among all the assessors anticipating future moves of her own and the other players. This leads to the best compromise choices and sophisticated behavior.

Our objective is to present a solution concept that is as close as possible to perfect Bayesian Equilibrium while allowing for multiple priors. The proximity to PBE should facilitate the understanding and acceptance of the new concept and simplify the interpretation of new insights. This design objective also allows us to build on the discipline underlying the concept of PBE.

We identify six reasons that motivated us to create this new solution concept, each of them is associated with contexts where PBE has deficiencies. These reasons are robustness, ambiguity, non-probabilistic reasoning, parsimony, tractability, and accessibility. We explain each of these in more detail.

Robustness. The PCE concept can be used to investigate the robustness of insights gained by PBE analyses when players are not willing to commit to a specific prior. Similarly it can be used to understand how predictions depend on the degree of understanding of the different players.

Ambiguity. Preferences that allow for decision makers to care about ambiguity have become popular. Our concept allows us to include players with such preferences and to estimate the degree of ambiguity of players in the data. 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 maxmin 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 often focus on two types of workers, high and low, or capture the uncertainty by a small number of parameters. 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 can prevent a tractable analysis of PBE. In our examples, PCE is shown to yield 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.

Of course, there are several alternative approaches to learning under ambiguity. We hope to add to this literature, and to see more future work on economics applications, and empirical testing and comparison of different theories.

Appendix A. Proofs.

A.1. **Proof of Theorem 1.** Consider a game $\Gamma = (N, \mathcal{G}, \Omega, (\Pi_1, ..., \Pi_n), (u_1, ..., u_n))$. Let Φ be the set of information sets excluding the initial node ϕ_0, so $\Phi = \bigcup_{i\in N} \Phi_i$. Recall that A_ϕ is the set of pure actions of the player who moves at information set $\phi \in \Phi$. A strategy profile s associates with each information set ϕ a mixed action $s_\phi \in \mathscr{A}_\phi = \Delta(A_\phi)$ at ϕ.

We now define an ε-perturbed game. Let ε be a small enough positive number. Let $\Delta_\varepsilon(A(\phi))$ be the set of mixed actions at information set ϕ such that each pure action in $A(\phi)$ is played with probability at least ε. Let $\mathcal{S}_\varepsilon$ be the set of strategy profiles such that $s_\phi \in \Delta_\varepsilon(A(\phi))$ for each $\phi \in \Phi$. So the strategies in $\mathcal{S}_\varepsilon$ are completely mixed. An *ε-perturbed game* Γ_ε is the original game Γ where the players' strategies are confined to $\mathcal{S}_\varepsilon$.

Consider a strategy profile $s \in \mathcal{S}_\varepsilon$. Because s is fully mixed, the belief system that is consistent with s is uniquely defined by Bayes' rule. Denote this belief system by $\beta(s)$, and let $\beta_\phi(\pi; s)$ is the posterior probability distribution over the decision nodes in the information set ϕ derived from a prior π. Let

$$B_{\phi_i}(\beta, s) = \{\beta_{\phi_i}(\pi_i; s) : \pi_i \in \Pi_i\}$$

be the set of beliefs at each $\phi_i \in \Phi_i$ for each player $i \in N$. Let $U_{\phi_i}(s)$ be the negative of player i's maximum loss at $\phi_i \in \Phi_i$ when player i follows her strategy s_{ϕ_i}, so

$$\begin{aligned} U_{\phi_i}(s) &= -l(s_{\phi_i}|s, \beta(s), \phi) \\ &= \inf_{b_i \in B_{\phi_i}(\beta,s)} \left(\bar{u}_i(s_{\phi_i}|s, \phi_i, b_i) - \sup_{a_i \in A(\phi_i)} \bar{u}_i(a_i|s, \phi_i, b_i) \right). \end{aligned} \tag{29}$$

Two observations are in order. First, $U_{\phi_i}(s) = U_{\phi_i}(s_{\phi_i}, s_{-\phi_i})$ is continuous in s_{ϕ_i}. This is because $\bar{u}$ is continuous, and the set $B_{\phi_i}(\beta, s)$ of beliefs at ϕ_i is independent of s_{ϕ_i} (it only depends on the choices in the information sets preceding ϕ_i). Second, $U_{\phi_i}(s_{\phi_i}, s_{-\phi_i})$ is also continuous in $s_{-\phi_i}$ when $s \in S_\varepsilon$, so the strategies are fully mixed. This is because $B_{\phi_i}(\beta, s)$ is a continuous correspondence w.r.t. $s \in S_\varepsilon$, as it is derived by Bayes' rule from the set of priors pointwise, and Bayes' rule is a well defined and continuous operator for $s \in S_\varepsilon$. In addition, both $B_{\phi_i}(\beta, s)$ and $A(\phi_i)$ are compact. The continuity of $U_{\phi_i}(s_{\phi_i}, s_{-\phi_i})$ in $s_{-\phi_i}$ then follows from the Maximum Theorem (Berge, 1963).

We now construct an augmented game $(\Phi, \mathcal{G}, \Omega, \pi_0, U)$ as follows. 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 π_0 be a common prior over the states, and assume that π_0 has full support over Ω. Nature moves first by choosing a state $\omega \in \Omega$ according to the prior π_0. Each player $\phi \in \Phi$ moves only once, at her information set ϕ, by choosing a mixed action from the set $\Delta_\varepsilon(A(\phi))$. 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 $(\Phi, \mathcal{G}, \Omega, \pi_0, U)$ 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 $U_\phi(s)$ depends on the set of beliefs $B_\phi(s)$ at ϕ, but it is independent of the state ω itself. So, the prior π_0 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 $(s'_\phi, s_{-\phi}) \in \mathcal{S}_\varepsilon$ denote the strategy profile where s'_ϕ 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 $s'_\phi \in \Delta_\varepsilon(A(\phi))$ is the same as minimizing the maximum loss at ϕ in the perturbed game Γ_ε. Consequently, if $\bar{s}$ is a strategy profile in a sequential equilibrium of the augmented game, then $(\bar{s}, \beta(\bar{s}))$ is a PCE of Γ_ε. 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.

Thus we have shown the existence of a PCE in every perturbed game Γ_ε. It remains to show the existence of a PCE in the original, unperturbed game Γ. Consider a sequence $(\varepsilon_k)_{k=1}^\infty$ such that $\lim_{k\to\infty} \varepsilon_k = 0$. Let (s^k, β^k) be a PCE for the perturbed game Γ_{ε_k}. By Bolzano-Weierstrass theorem there exists a subsequence $(k_t)_{t=1}^\infty$ such that (s^{k_t}, β^{k_t}) converges to some (s^*, β^*) as $t \to \infty$. Observe that the belief system β^* is consistent with s^*. This is because for each player i, each information set $\phi_i \in \Phi_i$, and each prior $\pi_i \in \Pi_i$, either $\beta^*_{\phi_i}(\pi_i)$ is derived by Bayes rule that is continuous as (s^{k_t}, β^{k_t}) approaches (s^*, β^*), or Bayes rule is undefined in the limit, in which case $\beta^*_{\phi_i}(\pi_i)$ is also consistent by definition. Next, for all $\varepsilon > 0$, all t such that $\varepsilon \geq \varepsilon_{k_t}$, and all $s'_\phi \in \Delta_\varepsilon(A_\phi)$ we have

$$\begin{aligned} 0 \leq & U_\phi(s^{k_t}_\phi, s^{k_t}_{-\phi}) - U_\phi(s'_\phi, s^{k_t}_{-\phi}) = -l(s^{k_t}_\phi | s^{k_t}, \beta^{k_t}, \phi) + l(s'_\phi | s^{k_t}, \beta^{k_t}, \phi) \xrightarrow{t\to\infty} \\ & - l(s^*_\phi | s^*, \beta^*, \phi) + l(s'_\phi | s^*, \beta^*, \phi) = U_\phi(s^*_\phi, s^*_{-\phi}) - U_\phi(s'_\phi, s^*_{-\phi}), \end{aligned}$$

where the inequality is by $s^{k_t}_\phi$ being a best response in the augmented game, the first equality is by (29), the limit is by the continuity of $l(s_\phi | s, \beta, \phi)$ in s and β, and the second equality is because the set $B_\phi(s^{k_t})$ of beliefs at ϕ is independent of the mixed action $s^{k_t}_\phi$ at ϕ. It follows that s^*_ϕ is a best response to $s^*_{-\phi}$. So s^* is a best compromise strategy profile in the unperturbed game Γ. We thus have shown that (s^*, β^*) is a PCE of Γ. □

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 beliefs that satisfy Definition 1.

First, we find the 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 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

$$\frac{(\bar{a} - \bar{b}q_{-i})^2}{4\bar{b}} - (\bar{a} - \bar{b}(q_i + q_{-i}))q_i = \frac{(\underline{a} - \underline{b}q_{-i})^2}{4\underline{b}} - (\underline{a} - \underline{b}(q_i + q_{-i}))q_i. \tag{30}$$

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, $P(q_1^* + q_2^*) \ge \underline{P}(q_1^* + q_2^*) > 0$. Substituting the solution into (30) yields the maximum loss of each firm (4). □

A.3. **Proof of Proposition 2.** Similarly to the 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 beliefs that satisfy Definition 1.

First, we determine the 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 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

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

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

$$\sup_{p_{-i} \in (c_i, p_i)} (u^*_i(p_{-i}; c_i) - u_i(p_i, p_{-i}; c_i)) = \begin{cases} (p_i - c_i)\frac{a - p_i}{b}, & \text{if } p_i \leq s^m(c_i), \\ \frac{(a-c_i)^2}{4b}, & \text{if } p_i > s^m(c_i). \end{cases} \tag{31}$$

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

$$\begin{aligned} \sup_{p_{-i} \in (p_i, \bar{p}]} (u^*_i(p_{-i}; c_i) - u_i(p_i, p_{-i}; c_i)) &= u^*_i(\bar{p}; c_i) - u_i(p_i, \bar{p}; c_i) \\ &= -(p_i - c_i)\frac{a - p_i}{b} + \begin{cases} (\bar{p} - c_i)\frac{a - \bar{p}}{b}, & \text{if } p_i \leq s^m(c_i), \\ \frac{(a-c_i)^2}{4b}, & \text{if } p_i > s^m(c_i). \end{cases} \end{aligned} \tag{32}$$

To minimize the maximum loss, we need to minimize the greater of the expressions in (31) and (32). Observe that, by the definition of $s^m(c_i)$, the right-hand side in (31) is constant and the right-hand side in (32) 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 (31) and (32) 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

$$(p_i - c_i)\frac{a - p_i}{b} = (\bar{p} - c_i)\frac{a - \bar{p}}{b} - (p_i - c_i)\frac{a - p_i}{b}. \tag{33}$$

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 (33) yields (6). □

A.4. **Proof of Proposition 3.** We prove only part (iii) of Proposition 3 for the proportional rule given by (10). The proof of parts (i) and (ii) for the other two rules is analogous but easier, and thus omitted.

Let the refunds r_i be given by the proportional rule (10). First we 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 $x_i + \sum_{j \neq i} s_j(v_j) = c - \varepsilon$ 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 $v_i - x_i$. Alternatively, it could be that all other agents contribute enough to cover c, so $\sum_{j \neq i} s_j(v_j) \geq c$. Thus the agent could have contributed nothing and still received the good. In this case the loss is the amount of contribution net of the refund, $x_i - r_i(x)$. This loss is maximized when the other agents' contributions exactly equal to the cost, so $\sum_{j \neq i} s_j(v_j) = c$, so by (10) we have

$$x_i - r_i(x) = \frac{cx_i}{x_i + \sum_{j \neq i} s_j(v_j)} \leq \frac{cx_i}{x_i + c}.$$

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

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

for x_i. Denote the solution by $s^*(v_i)$. It is easy to verify that it is as given in part (iii) of the statement of Proposition 3. Note that it is symmetric across the players, so we drop the subscript i.

The above argument requires that there exist values $v_j \in [0, \bar{v}]$ such that $\sum_{j \neq i} s^*(v_j) = c$. Observe that $s^*(0) = 0$ and $s^*(v_i)$ is increasing in v_i. So, we only need to verify that $\sum_{j \neq i} s^*(\bar{v}) \geq c$, which holds under condition (7).

It remains to determine the maximum welfare loss $L(s^*)$ as defined in (12). As $s^*(v_i)$ is increasing in v_i, the constraint $\sum_{i=1}^n s^*(v_i) < c$ must be binding. Moreover, it is easy to verify that $s^*(v_i)$ is convex in v_i. Thus, by Jensen's inequality we have

$$\sum\nolimits_{j=1}^{n} s^*(v_j) \geq ns^*\left(\frac{1}{n}\sum\nolimits_{j=1}^{n} v_j\right).$$

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

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

for z yields

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

We thus obtain

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

A.5. **Proof of Proposition 4.** Let $(\sigma_S^*, \sigma_B^*, \beta^*)$ be a PCE, and let P^* be the set of equilibrium prices, so

$$P^* = \mathrm{Supp}(\sigma_S^*(\cdot|\theta_H)) \cup \mathrm{Supp}(\sigma_S^*(\cdot|\theta_L)).$$

Recall that $\sigma_S^*(\cdot|\cdot)$ is discrete, so P^* is countable. By contradiction, suppose that P^* contains at least two prices and that there is a positive probability of trade, so

$$|P^*| \geq 2 \;\text{ and }\; \sum_{p \in p^*} \sigma_B^*(p) > 0.$$

Consider a price $p \in P^*$ that might be offered by the low-type seller, so $\sigma_S^*(p|\theta_L) > 0$. Then p must also be offered with a positive probability by the high-type seller, and the trade must be possible at p, so

$$\sigma_S^*(p|\theta_L) > 0 \text{ implies } \sigma_S^*(p|\theta_H) > 0 \text{ and } \sigma_B^*(p) > 0. \tag{34}$$

Otherwise, if $\sigma_S^*(p|\theta_H) = 0$, then the buyer would have inferred that the car had low quality with certainty, and therefore would not buy it at that price, so $\sigma_B^*(p) = 0$. But if $\sigma_B^*(p) = 0$, then the low-type seller's payoff from choosing p would be zero, and thus she would have had a profitable deviation to another price p' such that $\sigma_B^*(p') > 0$.

By (34) and the definition of P^*, it follows that $\mathrm{Supp}(\sigma_S^*(\cdot|\theta_H)) = P^*$, so the high-type seller randomizes over all prices in P^*. Hence she must be indifferent among them all, so

$$(p' - c_H)\sigma_B^*(p') = (p'' - c_H)\sigma_B^*(p'') \;\text{ for all } p', p'' \in P^*. \tag{35}$$

However, because $c_L < c_H$, the low-type seller cannot be indifferent between the prices in P^*, in fact, she strictly prefers the smallest of these prices, so

$$(p' - c_L)\sigma_B^*(p') > (p'' - c_L)\sigma_B^*(p'') \;\text{ for all } p', p'' \in P^*,\ p' < p''.$$

Hence the low-type seller assigns probability one on the smallest price in P^*, so

$$\mathrm{Supp}(\sigma_S^*(\cdot|\theta_L)) = \{p_*\}, \quad \text{where } p_* = \min\{p : p \in P^*\}.$$

Consequently, for each price $p' \in P^*$ such that $p' > p_*$, the buyer infers that that the type is high with certainty, irrespective of the prior. That is, $\beta^*_{\phi_{p'}}(\pi) = 1$ for all $\pi \in \Pi$. In this case, by (15), the buyer buys with certainty, so $\sigma_B^*(p') = 1$. But then we have

$$(p' - c_H)\sigma_B^*(p') = p' - c_H > p_* - c_H \geq (p_* - c_H)\sigma_B^*(p_*),$$

which contradicts (35). Thus we have reached a contradiction. $\square$

A.6. **Proof of Proposition 5.** Here we only prove claims for PCE, as this is equivalent to a PBE when there is a single prior. Consider a PCE $(\sigma_S^*, \sigma_B^*, \beta^*)$ that involves trade and thus, by Proposition 4, is pooling on a single price p^*. So,

$$\sigma_S^*(p^*|\theta_H) = \sigma_S^*(p^*|\theta_L) = 1, \quad \text{and} \quad \sigma_B^*(p^*) > 0. \tag{36}$$

Moreover, when observing p^*, the buyer's set of posterior beliefs is equal to the set of priors, so

$$\beta^*_{\phi_{p^*}}(\pi) = \pi \text{ for each } \pi \in \Pi. \tag{37}$$

We show that p^* necessarily satisfies the conditions on its range that are specified in Proposition 5.

First, in case of single prior, so $\Pi = \{\bar\pi\}$, by (15) and (37), $\sigma_B^*(p^*)$ must satisfy (17) when $p^* \leq \bar\pi v_H$, and $\sigma_B^*(p^*) = 0$ when $p^* > \bar\pi v_H$. Thus, if there exists a buyer's strategy σ_B^* that satisfies (15) and (36), then $p^* \leq \bar\pi v_H$. In case of multiple priors, so $\Pi = \{\pi_1, ..., \pi_K\}$ with $K \geq 2$, by (15) and (37), $\sigma_B^*(p^*)$ must satisfy (18) when $p^* < \pi_K v_H$, and $\sigma_B^*(p^*) = 0$ when $p^* \geq \pi_K v_H$. Thus, if there exists a buyer's strategy σ_B^* that satisfies (15) and (36), then $p^* < \pi_K v_H$.

Second, by (14), when $p^* < c_H$, the seller would not want to sell a high-quality car. So, as $\sigma_B^*(p^*) > 0$, this seller would want to deviate to another p' such that $\sigma_B^*(p') = 0$, for instance, $p' = 1$. Thus, if the seller's strategy σ_S^* satisfies (14) and (36), then $p^* \geq c_H$.

It remains to show that for every $p^* \in [c_H, \bar\pi v_H]$ in the case of $\Pi = \{\bar\pi\}$ and for every $p^* \in [c_H, \pi_K v_H)$ in the case of $\Pi = \{\pi_1, ..., \pi_K\}$ with $K \geq 2$, we can find the buyer's out-of-equilibrium beliefs and behavior to support the PCE that is pooling on p^*. Indeed, when observing any $p \neq p^*$, let the buyer be certain that the car has low quality, and let $\sigma_B^*(p) = 0$. Given that $p^* \geq c_H$ and $\sigma_B^*(p^*) > 0$, by (14), the seller will prefer p^* to any price $p \neq p^*$ irrespective of the car type. $\square$

A.7. **Proof of Proposition 6.** 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

$$\underline{\theta}_j = \inf\{\theta : (\theta, c, \gamma) \in S(e_j)\}, \quad \bar{\theta}_j = \sup\{\theta : (\theta, c, \gamma) \in S(e_j)\}, \quad j = L, H. \tag{38}$$

Consider a firm i, some wages w_i and w_{-i}, and a state (θ, c, γ). 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 not hiring the worker if $\theta \leq w_{-i}$, so

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

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 $u_i(w_i, w_{-i}; \theta, \gamma) = 0$. Then the largest loss is obtained when θ is the greatest conditional on e_j, so

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

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

$$\begin{aligned}\sup_{\theta:(\theta,c,\gamma)\in S(e_j)} (u_i^*(w_{-i};\theta) - u_i(w_i, w_{-i}; \theta, \gamma)) &= \sup_{\theta \in [\underline{\theta}_j, \bar{\theta}_j]} (\max\{\theta - w_{-i}, 0\} - (\theta - w_i)) \\ &= w_i - \underline{\theta}_j.\end{aligned}$$

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

$$\begin{aligned}\sup_{\theta:(\theta,c,\gamma)\in S(e_j)} (u_i^*(w_{-i};\theta) - u_i(w_i, w_{-i}; \theta, \gamma)) &= \sup_{\theta \in [\underline{\theta}_j, \bar{\theta}_j], \gamma \in \Delta_2} (\max\{\theta - w_{-i}, 0\} - (\theta - w_i)\gamma_i) \\ &= \max\left\{\bar{\theta}_j - w_{-i}, w_i - \underline{\theta}_j\right\},\end{aligned}$$

Consequently, the maximum loss $l_i(w_i, w_{-i})$ is given by

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

Clearly, in the best compromise equilibrium, $w_i = w_{-i}$, and no one can reduce the loss by deviating to w_i above or below w_{-i}, so the best compromise $w_i^*(e_j)$ is the solution of

$$w_i^*(e_j) - \underline{\theta}_j = \bar{\theta}_j - w_i^*(e_j).$$

Solving the above equation yields

$$w_i^*(e_j) = \frac{\bar{\theta}_j + \underline{\theta}_j}{2}, \quad i = 1, 2. \tag{39}$$

The associated maximum losses are

$$l_i(w_i^*(e_j), w_{-i}^*(e_j)) = w_i^*(e_j) - \underline{\theta}_j = \frac{\bar{\theta}_j - \underline{\theta}_j}{2}. \tag{40}$$

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, γ), the worker chooses $e = e_H$ if and only if[7]

$$c \leq w^H - w^L.$$

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 (19), the highest and lowest θ in $S(e_L)$ are $\bar{\theta}_L = 1$ and $\underline{\theta}_L = 0$. By (39), 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 upper and lower productivity bounds $\underline{\theta}_j = 1$ and $\underline{\theta}_j = 0$ into (40), 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

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

be the sets of 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 profiles (θ, c, γ) such that low and high education is chosen, respectively. These sets thus satisfy the consistency requirement (Definition 1).

By (19) and (38), the highest and lowest θ in $S(e_H)$ are given by

$$\bar{\theta}_H = 1 \quad \text{and} \quad \underline{\theta}_H = a - \frac{\varepsilon}{2} - w^H + w^L. \tag{41}$$

So, $S(e_H)$ is nonempty when $\underline{\theta}_H = a - \frac{\varepsilon}{2} - w^H + w^L \leq 1$.

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

$$\bar{\theta}_L = \max\left\{a + \frac{\varepsilon}{2} - w^H + w^L, 1\right\} \quad \text{and} \quad \underline{\theta}_L = 0. \tag{42}$$

Because $w^H - w^L \leq 1$ and by assumption $a + \frac{\varepsilon}{2} \geq 1$, we have $\bar{\theta}_L \geq 0$. So $S(e_L)$ is always nonempty.

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

From (39), we have

$$w^H = \frac{\bar{\theta}_H + \underline{\theta}_H}{2} \quad \text{and} \quad w^L = \frac{\bar{\theta}_L + \underline{\theta}_L}{2}. \tag{43}$$

Solving the system of six equations in (41), (42), and (43) with six unknowns (w^H, w^L, $\bar{\theta}_H$, $\underline{\theta}_H$, $\bar{\theta}_L$, and $\underline{\theta}_L$) when taking into account the condition (20), we obtain the unique equilibrium wages and the bounds on the productivity types as shown in (23) and (24).

Finally, substituting the upper and lower productivity bounds $\underline{\theta}_j$ and $\underline{\theta}_j$ into into (40), we obtain firm i's maximum loss when $e = e_H$,

$$l_i(w_i^*(e_H), w_{-i}^*(e_H); e_H) = \frac{\bar{\theta}_H - \underline{\theta}_H}{2} = \frac{3}{4} - \frac{a}{2},$$

and the maximum loss when $e = e_L$,

$$l_i(w_i^*(e_L), w_{-i}^*(e_L); e_L) = \frac{\bar{\theta}_L - \underline{\theta}_L}{2} = \frac{a}{2} - \frac{1}{4} + \frac{\varepsilon}{2}. \qquad \square$$

A.8. **Proof of Proposition 7.** Consider how a buyer who knows that v is in $[y_0, y_1]$ reacts when the seller asks p. Let $p_0 \in [1/2, 1]$. Suppose that $p < p_0$. Assume that the buyer speculates that v in $\{y_0\}$. This is consistent with the strategy of the seller as $p < p_0$ is out of equilibrium. Given this speculation, accepting p if and only if $p \leq y_0$ is a best compromise.

Now suppose that $p \geq p_0$. The largest interval $[x_0, x_1] \subset [0, 1]$ that satisfies (27) 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 $y_1 \leq 1$, it is easy to verify that $l_b(0; p, y_0, y_1) \geq l_b(1; p, y_0, y_1)$ if and only if $p \leq \frac{1}{2}(y_0 + y_1)$. Thus, it is the best compromise to buy the good when $p \leq \frac{1}{2}(y_0 + y_1)$ 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 $p < p_0$ is dominated by $p = p_0$. This is because when $p < p_0$, 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 $p \geq p_0$. 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 $p \geq p_0$ is the largest of the four losses computed above, so

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

The best compromise price minimizes the maximum loss $l_s(p; x_0, x_1)$ among all prices $p \geq p_0$, leading to the seller's equilibrium strategy (27). $\square$

A.9. **Proof of Proposition 8.** Before proving Proposition 8, we 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

$$\bar{a}_F(z) \in \arg\max_{a'\in[0,1]} \mathbb{E}_{F,G_\varepsilon}[-(a' - \theta)^2|z]. \tag{44}$$

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

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

We now prove Proposition 8. 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

$$H(z) = \sup_{F\in\mathcal{F}_\delta} \mathbb{E}_{F,G_\varepsilon}[\theta|z] \quad \text{and} \quad L(z) = \inf_{F\in\mathcal{F}_\delta} \mathbb{E}_{F,G_\varepsilon}[\theta|z]. \tag{45}$$

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

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

Appendix B. Alternative Model of Forecasting

This section considers an alternative an alternative forecasting model to the one presented in Section 3.7. 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

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

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 9. *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 8 and thus omitted.

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 $\varepsilon \to 1$.

Remark 3. 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

Akerlof, George A. 1970. "The Market for "Lemons": Quality Uncertainty and the Market Mechanism." *The Quarterly Journal of Economics*, 84: 488–500.

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

Battigalli, Pierpaolo, Emiliano Catonini, Giacomo Lanzani, and Massimo Marinacci. 2019. "Ambiguity attitudes and self-confirming equilibrium in sequential games." *Games and Economic Behavior*, 115: 1–29.

Berge, Claude. 1963. *Topological Spaces.* Oliver and Boyd.

Billot, Antoine, Alain Chateauneuf, Itzhak Gilboa, and Jean-Marc Tallon. 2000. "Sharing beliefs: Between agreeing and disagreeing." *Econometrica*, 68: 685–694.

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

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

d'Aspremont, Claude, and Louis-André Gérard-Varet. 1979. “Incentives and incomplete information.” *Journal of Public Economics*, 11: 25–45.

Epstein, Larry G, and Martin Schneider. 2003. “Recursive multiple-priors.” *Journal of Economic Theory*, 113: 1–31.

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, and Peter Klibanoff. 2007. “Updating preferences with multiple priors.” *Theoretical Economics*, 2: 261–298.

Hanany, Eran, Peter Klibanoff, and Sujoy Mukerji. 2020. “Incomplete information games with ambiguity averse players.” *American Economic Journal: Microeconomics*, 12: 135–87.

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

Hayashi, Takashi. 2008. “Regret aversion and opportunity dependence.” *Journal of Economic Theory*, 139: 242–268.

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

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

Kajii, Atsushi, and Takashi Ui. 2006. “Agreeable bets with multiple priors.” *Journal of Economic Theory*, 128: 299–305.

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

Kreps, David M, and Robert Wilson. 1982. “Sequential Equilibria.” *Econometrica*, 50: 863–894.

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.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing it now or later." *American Economic Review*, 89: 103–124.

Pahlke, Marieke. 2022. "Dynamic consistency in incomplete information games with multiple priors." *Games and Economic Behavior*, 133: 85–108.

Pires, Cesaltina Pacheco. 2002. "A Rule for Updating Ambiguous Beliefs." *Theory and Decision*, 53: 137–152.

Puppe, Clemens, and Karl H Schlag. 2009. "Choice under complete uncertainty when outcome spaces are state dependent." *Theory and Decision*, 66: 1–16.

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.

Rigotti, Luca, Chris Shannon, and Tomasz Strzalecki. 2008. "Subjective beliefs and ex ante trade." *Econometrica*, 76: 1167–1190.

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

Siniscalchi, Marciano. 2011. "Dynamic Choice Under Ambiguity." *Theoretical Economics*, 6: 379–421.

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.

Strotz, R. H. 1955. "Myopia and inconsistency in dynamic utility maximization." *Review of Economic Studies*, 23: 165–180.

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.

Wang, Tan. 2003. "Conditional preferences and updating." *Journal of Economic Theory*, 108: 286–321.