

Hanany, Eran; Klibanoff, Peter; Mukerji, Sujoy

Working Paper

Incomplete information games with ambiguity averse players

Working Paper, No. 868

Provided in Cooperation with:

School of Economics and Finance, Queen Mary University of London

Suggested Citation: Hanany, Eran; Klibanoff, Peter; Mukerji, Sujoy (2018) : Incomplete information games with ambiguity averse players, Working Paper, No. 868, Queen Mary University of London, School of Economics and Finance, London

This Version is available at:

<https://hdl.handle.net/10419/210425>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Incomplete Information Games with Ambiguity Averse Players

Eran Hanany, Peter Klibanoff and Sujoy Mukerji

Working paper No. 868

September 2018

ISSN1473-0278

School of Economics and Finance



Queen Mary
University of London

Incomplete Information Games with Ambiguity Averse Players*

Eran Hanany[†] Peter Klibanoff[‡] Sujoy Mukerji[§]

This version: September 13, 2018

Abstract

We study incomplete information games with ambiguity averse players. Our focus is on equilibrium concepts satisfying *sequential optimality* – each player’s strategy is optimal at each information set given opponents’ strategies. We show sequential optimality, which does not make any explicit assumption on updating, is equivalent to sequential optimality with respect to beliefs updated using a particular generalization of Bayesian updating. Ambiguity aversion expands the set of equilibria compatible with players sharing common ambiguous beliefs. We connect ambiguity aversion with belief robustness. Examples illustrate new strategic behavior, including strategic use of ambiguity, under ambiguity aversion.

1 Introduction

Dynamic games of incomplete information are the subject of a large literature, both theory and application, with diverse fields including models of firm competition, agency theory, auctions, search, insurance and many others. In such games, how players perceive and react to uncertainty, and the way it evolves over the course of the game, is of central importance. In

*We thank Itzhak Gilboa, Frederic Koessler, Fabio Maccheroni, Ludovic Renou, Jean-Marc Tallon and participants at various conferences and seminars. Klibanoff and Mukerji additionally thank the Robust Finance research group at the ZiF, Bielefeld University for helpful comments and hospitality during their time as visiting fellows in Spring/Summer 2015. This research was partially supported by grant No. 2014350 from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel to Hanany and Klibanoff.

[†]Faculty of Engineering, Tel Aviv University, Tel Aviv 69978, Israel. Email: hananye@tau.ac.il

[‡]Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA. E-mail: peterk@kellogg.northwestern.edu

[§]School of Economics and Finance, Queen Mary University of London, London E1 4NS, United Kingdom. Email: s.mukerji@qmul.ac.uk

the theory of decision making under uncertainty, preferences that allow for decision makers to care about ambiguity¹ have drawn increasing interest (Gilboa and Marinacci, 2013). That ambiguity may remain relevant in a steady-state has been demonstrated in e.g., Epstein and Schneider (2003b), Maccheroni and Marinacci (2005) and Klibanoff, Marinacci and Mukerji (2009). We propose equilibrium notions for incomplete information games involving ambiguity about parameters (which could be, for example, privately known types of players). The parameter space is a modelling device to allow for players to be uncertain (which in our setting may include both ambiguity and risk) about the payoffs they and the others face and/or the strategies played by the other players. This allows us to examine effects of introducing ambiguity aversion in strategic settings, static and dynamic. The definition and analysis of solution concepts capturing dynamic considerations, such as optimality of continuation strategies at each information set, are the main contributions of the paper. Such optimality is absent from almost all existing literature on games with ambiguity averse players.

In our analysis, players have smooth ambiguity preferences (Klibanoff, Marinacci and Mukerji, 2005) and may be ambiguity averse. Such preferences for a player i evaluate a behavior strategy profile σ by

$$\sum_{\pi \in \Delta(\Theta)} \phi_i \left(\sum_{\theta \in \Theta} U_i(\sigma, \theta) \pi(\theta) \right) \mu_i(\pi),$$

where Θ is the parameter space modeling the incomplete information, μ_i is a subjective probability over $\Delta(\Theta)$ (i.e., a second-order probability over Θ), $U_i(\sigma, \theta)$ is i 's expected payoff from σ given θ , and ϕ_i is an increasing function, the concavity of which reflects ambiguity aversion. All else equal, as ϕ_i becomes more concave, player i becomes more ambiguity averse (see e.g., Theorem 2 in Klibanoff, Marinacci, Mukerji 2005). The presence of ambiguity is captured by non-degeneracy of μ_i . In the smooth ambiguity model it is possible to hold the players' information fixed (by fixing μ_i) while varying their ambiguity attitude from aversion to neutrality (i.e., replacing a more concave ϕ_i with an affine one, which reduces preferences to expected utility). This facilitates a natural way to understand the effect of introducing ambiguity aversion into a strategic environment. Our focus is on extensive form games, specifically multistage games with perfect recall, and on equilibrium notions capturing perfection analogous to those in standard theories for ambiguity neutral players, such as subgame perfect equilibrium (Selten, 1965), sequential equilibrium (Kreps and Wilson, 1982) and perfect Bayesian equilibrium (PBE) (e.g., Fudenberg and Tirole,

¹In this literature, ambiguity refers to subjective uncertainty about probabilities (see e.g., Ghirardato, 2004).

1991a,b).

We use as a running example a variation on the peace negotiation game proposed by Greenberg (2000) motivated by the common practice of governments to manipulate expectations of rewards and punishments so as to encourage negotiating parties to cooperate. The game involves two small countries in peace negotiations but who, left on their own, would not agree to peace, and a large country that has the power to affect the small countries and desires peace between them. The large country hopes to induce cooperation by the small countries by leading each to believe that it will likely be punished by the large country if negotiations break down (or favored if negotiations succeed). However, uncertainty created by any single mixture over which country to punish is inadequate to induce both small countries to agree to peace. More precisely, for the payoffs in the example, any (mixture over) choice of whom to punish/favor that the large country might make contingent on the success/failure of the negotiation is insufficient to convince both small countries to reach agreement – any mixture that punishes one of them often enough is inadequate to incentivize the other. This is true irrespective of ambiguity aversion. However, if the small countries are ambiguity averse, the large country can, by taking steps to obscure the likelihood of who will be punished/favored, create ambiguity in the minds of the small countries and push both of them towards peace. Given sufficient ambiguity aversion of the small countries, it is thus in the strategic interest of the large country’s government to try to behave in a way that makes it difficult for both negotiating parties to be confident that it will not be punishing that party with high probability if negotiations break down.

We would like to model such behavior as an equilibrium (i.e., all parties best responding to the strategies of the others), capturing, through the use of ambiguous incomplete information about parameters and the possibility of conditioning strategies on such parameters, the idea that, in the presence of ambiguity aversion, some players may choose to play in a manner that is perceived as ambiguous. In the context of the example, one can think of this use of ambiguous incomplete information about parameters as a reduced form way of accounting for the fact that opportunities for fully learning which (mixture over) actions the large country will play after the negotiation are restricted. Motivation for such a reduced form may be, for instance, that political parties hold government office only temporarily and, over their time in office, mediate negotiations that are not identical, plausibly giving its strategies an unpredictably changing nature that Bewley (1988, p.35) identifies as essential when discussing how “Knightian” uncertainty may exist in a steady state. Or it may be that players see only coarse descriptions of the parameter profiles realized in previous plays of the game, implying that they are not able to pin down a unique empirical frequency for the parameter profile (coarse observability motivates, e.g., Lehrer (2012)’s notion of equilibrium

with partially specified probabilities).

The uncertainty in a game could also be about payoffs (e.g., an entry game where the entrant is uncertain about the incumbent's cost as in Section 5.2) and, because of the relative novelty of the situation to one or more players, it is plausible that they (e.g., the entrant) view the uncertainty at least partially as ambiguity. In that example, the entrant is expert enough to figure out what (possibly mixed) action the incumbent would deploy were they to know the realized cost. However, unlike the incumbent, since the entrant has not actively participated in the market/industry they are entering, the entrant is not privy to the (possibly stochastic) law governing the realized cost, and also, unlike the incumbent, does not directly see the realized cost before having to make the entry decision. Thus we model the entrant as having ambiguous uncertainty about the parameter (e.g., cost) while being knowledgeable enough (e.g., about optimal pricing practices given costs and demand) to correctly anticipate the incumbent's best response (e.g., pricing strategy) contingent on the parameter. More generally, except for the ambiguous nature of the parameter uncertainty, this combination of uncertainty about parameters and correct anticipation of strategic behavior given the parameter is central to the standard notion of Bayesian (Nash) equilibrium (BNE). This motivates the approach of this paper, which starts by generalizing BNE with regard to the uncertainty about parameters and how players react to that uncertainty.

We first define an ex-ante (Nash) equilibrium concept allowing for aversion to ambiguity about parameters. When there is no parameter uncertainty, this is simply Nash equilibrium under complete information. When there are common beliefs and ambiguity neutrality, it becomes Bayesian Nash equilibrium. Next, we refine ex-ante equilibrium by imposing perfection in the form of a *sequential optimality* requirement – each player i 's strategy must be optimal at each information set given the strategies of the other players and i 's beliefs at that information set. Sequential optimality does not make any explicit assumption on updating. When all players are ambiguity neutral, the definition of sequential optimality reduces to the definition of ex-ante equilibrium plus Kreps and Wilson (1982)'s sequential rationality. In this ambiguity neutral case, our results show that sequential optimality is equivalent to the version of Perfect Bayesian Equilibrium (PBE) described in Gibbons (1992, Chapter 4.1) despite the fact that Bayesian updating is assumed in the latter. As with PBE, a main motivation for sequential optimality is ruling out ex-ante equilibria that depend crucially on non-credible off-path behavior and doing so in a way that strengthens subgame perfection. Sequential optimality and our subsequent analysis and extensions of it are the main contributions of the paper.

We show that sequential optimality is equivalent to sequential optimality with respect to beliefs updated using the smooth rule (Hanany and Klibanoff 2009), a generalization of

Bayesian updating for smooth ambiguity preferences, which coincides with Bayes' rule under ambiguity neutrality.² Thus, we show that analysis of sequential optima of a game may be undertaken under the *as if* assumption that all players use smooth rule updating, and in this sense, that sequential optimality, which does not assume particular updating, nonetheless provides a way of cutting through the vexing issue of what update rule to impose in dynamic games with ambiguity aversion. Moreover, we show that under a slight strengthening of the smooth rule, the absence of profitable one-stage deviations can be used for verifying sequential optimality.

Our characterization of sequential optimality implies that it does not effectively restrict player i 's beliefs at information sets immediately following a deviation (though it does effectively restrict beliefs at off-path information sets that are not immediately off-path). We propose a refinement of sequential optimality restricting such beliefs: sequential equilibrium with ambiguity (SEA). In addition to sequential optimality, SEA imposes a generalization of Kreps and Wilson (1982)'s consistency condition from their definition of sequential equilibrium. Our 'as if' result on updating under sequential optimality motivates the use of the smooth rule in this generalization. We show that in the definition of SEA, sequential optimality may be equivalently replaced by the absence of profitable one-stage deviations. This implies that checking only one-stage deviations with respect to smooth rule consistent beliefs is sufficient to establish that a strategy profile is an SEA. Under ambiguity neutrality, SEA and sequential equilibrium are equivalent. Finally, we establish that SEA exist for any finite multistage game with perfect recall and incomplete information, and for any specification of players' ambiguity aversions and ex-ante beliefs.

Section 4.1 provides results on comparative statics of the equilibrium set in ambiguity aversion that apply to any of the above three notions of equilibrium. First, for fixed beliefs, ambiguity aversion may change the equilibrium set in a variety of ways – it can expand, shrink or simply change the set of equilibria. Second, we take the point of view of an outside observer who is not willing to assume particular beliefs when describing the equilibrium predictions of the theory. Holding payoffs and the structure of the game fixed, ambiguity aversion expands the set of equilibria compatible with players sharing a common belief (i.e., $\mu_i = \mu$ for all players i , running over all possible μ). Common beliefs are essential to this result. With unrestricted heterogeneous beliefs (i.e., running over all possible μ_i), ambiguity aversion does not affect the set of equilibria. If, as in some existing literature, we were instead to limit attention to pure strategies (both in terms of the equilibrium profile and,

²Under ambiguity aversion, the smooth rule may be thought of as applying Bayes' rule to the measure in the local linear approximation of preferences at the given strategy profile. Such local measures have previously proved useful in economics and decision theory. See e.g., Rigotti, Shannon and Strzalecki (2008), Hanany and Klbanoff (2009), Ghirardato and Siniscalchi (2012).

crucially, in terms of the deviations against which optimality is checked), ambiguity aversion expands the set of such equilibria even when we run over unrestricted heterogeneous beliefs.

Ambiguity averse behavior is often viewed as a robust response to doubts about beliefs (e.g., Hansen (2007)). We describe a sense in which this robustness extends to properties of equilibria. Section 4.2 defines robustness of an equilibrium to increases in ambiguity aversion and shows that this is related to a type of belief robustness (Theorem 4.5).

Section 5 contains an example of a Milgrom and Roberts (1982)-style limit pricing entry game with ambiguity about the incumbent’s cost. We show that limit pricing arises in an SEA. The limit pricing in our example is part of a semi-pooling equilibrium and succeeds in deterring some entry. Pooling equilibria are often sensitive to beliefs. We provide conditions under which this limit pricing is robust to increased ambiguity aversion on the part of the entrant and then apply Theorem 4.5 to conclude that under these same conditions ambiguity aversion can make the set of beliefs supporting limit pricing as large as desired.

Another example in Section 5 is a principal, multi-agent communication game. The principal is shown to strictly benefit from conditioning his cheap talk message to the agents on a payoff irrelevant ambiguous event. The recent literature (see Bose and Renou (2014), Beauchêne, Li and Li (2017), Kellner and Le Quement (2018)) on the value of ambiguous communication in dynamic settings uses violations of sequential optimality in an essential way. Left open was the question of whether such value existed when respecting sequential optimality. Our analysis of this example establishes that valuable strategic use of ambiguity can be sequentially optimal, and, moreover, can occur as part of an SEA. In the context of communication games and mechanism design, sequential optimality may be viewed as ensuring that players both react optimally to any information they receive and that participation or design are taken optimally from an ex-ante perspective.

Section 6 discusses some possible extensions of our approach, including to other models of ambiguity averse players’ preferences. Finally, Section 7 discusses closely related literature (especially Battigalli et al. (2017) and Pahlke (2018)) and compares it to our theory. A supplementary appendix contains all proofs and some further analysis.

2 Model

We begin by defining the central domain of the paper, finite multistage games with incomplete information and perfect recall where players have (weakly) ambiguity averse smooth ambiguity preferences. It is on this domain that we develop and apply our equilibrium concepts. Such games allow for both imperfectly observed actions and private observations as the game proceeds. Other than perfect recall and finiteness, the multistage structure (i.e.,

the assumption that all players move simultaneously at each point) is the additional potential limitation on the game forms we consider. While not entirely without loss of generality, if one doesn't object to giving a player singleton action sets at stages where this player has no "real" move, the multistage assumption is not restrictive. Note that (finite) normal form games with incomplete information and (weakly) ambiguity averse smooth ambiguity preferences are the special case where there is a single stage.

Formally, a *finite extensive-form multistage game with incomplete information and perfect recall and (weakly) ambiguity averse smooth ambiguity preferences*, Γ , is a tuple $(N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$ where:

- N is the finite set of players.
- H is the finite set of (terminal) histories, each of which is a finite sequence of length $T + 2$ of the form $h = (h_{-1}, (h_{0,i})_{i \in N}, \dots, (h_{T,i})_{i \in N})$.
For $0 \leq t \leq T + 1$, let $H^t \equiv \{h^t \equiv (h_{-1}, (h_{0,i})_{i \in N}, \dots, (h_{t-1,i})_{i \in N}) \mid h \in H\}$ be the set of partial histories up to (but not including) stage t . The set of all partial histories is $\mathcal{H} \equiv \{\emptyset\} \cup \bigcup_{0 \leq t \leq T+1} H^t$. For each $i \in N$, $0 \leq t \leq T$ and $h^t \in H^t$, $A_i(h^t) \equiv \{\hat{h}_{t,i} \mid \hat{h} \in H, \hat{h}^t = h^t\}$ is the set of actions available to player i at h^t . The set of initial partial histories, $\Theta \equiv H^0$, is called the set of "parameters" or "types".
- $\mathcal{I}_i \equiv \bigcup_{0 \leq t \leq T} \mathcal{I}_i^t$ are the information sets for player i , where each \mathcal{I}_i^t is a partition of H^t such that, for all $h^t, \hat{h}^t \in H^t$, $\hat{h}^t \in I_i(h^t)$ implies $A_i(h^t) = A_i(\hat{h}^t)$ (where $I_i(h^t)$ is the unique element of \mathcal{I}_i^t containing h^t).
Perfect recall means: for each player i , stage $0 \leq t \leq T$ and partial histories $h^t, \hat{h}^t \in H^t$, $I_i(h^t) = I_i(\hat{h}^t)$ implies $R_i(h^t) = R_i(\hat{h}^t)$, where, for each partial history $\bar{h}^t \in H^t$, $R_i(\bar{h}^t)$ is the ordered list of information sets i encounters and the actions i takes under \bar{h}^t .³
- $u_i : H \rightarrow \mathbb{R}$ is the (utility) payoff of player i given the history.⁴
- μ_i is a probability over $\Delta(\Theta)$ having finite support such that $\sum_{\pi \in \Delta(\Theta)} \mu_i(\pi) \pi(\theta) > 0$ for all $i \in N$ and $\theta \in \Theta$, where $\Delta(\Theta)$ is the set of all probability measures over Θ .⁵
- $\phi_i : \text{co}(u_i(H)) \rightarrow \mathbb{R}$ is a continuously differentiable, (weakly) concave and strictly increasing function.

³Formally, $R_i(\bar{h}^t) \equiv ((I_i(\bar{h}^s), \bar{h}_{s,i})_{0 \leq s < t}, I_i(\bar{h}^t))$. For future reference, note that we extend both A_i and R_i to information sets in the natural way.

⁴As is usual for preferences in games, we assume that u_i may be extended to a larger domain such that $u_i(H)$ is interior to the convex hull of the image of u_i on the larger domain, and that ϕ_i may be similarly extended. This ensures the validity of the interior optimality characterizations we use throughout.

⁵All of our results (except for Theorem 4.3) also hold if the class of games is restricted to those with a common μ such that $\mu_i = \mu$ for all players i . None of our examples rely on differences in the μ_i .




Figure 2.1: A peace negotiation game

The vectors give utility payoffs for players 1, 2 and 3, in that order, for each path.

The first three bullet points above describe the *game form*, while the rest describe *preferences*. Observe that at each partial history each player (not just those with non-trivial moves at this point) has an information set. The non-standard preference parts of this definition are ϕ_i and μ_i which are part of the specification of smooth ambiguity preferences, with the degree of concavity of ϕ_i reflecting ambiguity aversion and μ_i indicating the presence of ambiguity when $\text{supp } \mu_i$ contains multiple probability measures.⁶

We remark that the parameter space Θ may possibly include payoff relevant and/or payoff irrelevant components. The role of payoff irrelevant components is to facilitate our modeling of the strategic use of ambiguity via conditioning actions on these components.

Though largely standard, as formal objects such games might seem complex. To aid understanding, we next introduce a concrete example to which we will return at several points.

Running Example: The game in Figure 2.1 is a variation on the main example in Greenberg (2000). There are three players. Players 1 and 2 are small countries involved in peace negotiations. Player 3 is a large country that has the power to affect the small countries. The parameter space is $\Theta = \{I, II\}$ and represents payoff irrelevant private information of player 3. Specifically, one should think of this parameter as the outcome of an ambiguous device that player 3 has access to and *may* choose to condition its action on. At $t = 0$, only player 1 has a non-trivial move, which is either to c (operate) or to d (effect)

⁶ supp denotes the support of a measure.

in the negotiations. If player 1 plays c , then player 2 observes this and may choose either to c (ooperate) or to d (efect) in the negotiations. If both 1 and 2 cooperate, the negotiation leads to peace, and then player 3 chooses between favoring 1 under the peace ($f1$) and favoring 2 under the peace ($f2$). The favored country receives a payoff (u_i) of 5, while the other players receive 4. If either player defects, then negotiations break down. At that point, player 3, knowing that negotiations have failed, but not able to observe which country defected, must choose between an action that punishes player 1 ($p1$), an action that punishes player 2 ($p2$), and a neutral action (n). A player who is punished gets payoff 0, while the unpunished country is favored and gets a payoff of 10. If player 3 punishes, it prefers to punish the defector: player 3's payoff is 1 when punishing the defector, and 0 when punishing the other player. Player 3 staying neutral leads to a payoff of 6 for players 1 and 2, and a payoff of x for player 3. The payoff of 6 means that both small countries would prefer negotiations to fail if they knew there would be no punishment. We leave x as a variable in order to capture different relative merits for player 3 of staying neutral compared to punishing. A (terminal) history h consists of complete path through the game tree. For example, one history is $(I, c, c, f2)$ and $(I), (I, c)$ and (I, c, c) are partial histories.⁷

The beliefs μ for all players are $\frac{1}{2}-\frac{1}{2}$ over distributions π_1 and π_2 with $\pi_1(I) = 1$ and $\pi_2(I) = 0$, reflecting the common perception of the ambiguity concerning Θ . The starkness of the π s is purely for convenience – any beliefs reflecting ambiguity about whether I or II is more likely would suffice for our analysis. ■

A strategy for player i specifies the distribution over i 's actions conditional on each information set of player i . Formally:

Definition 2.1 (Behavior Strategy) A (behavior) strategy for player i in a game Γ is a function σ_i such that $\sigma_i(I_i) \in \Delta(A_i(I_i))$ for each $I_i \in \mathcal{I}_i$.

Let Σ_i denote the set of all strategies for player i . A strategy profile, $\sigma \equiv (\sigma_i)_{i \in N}$, is a strategy for each player.

Given a strategy profile σ , history h and $0 \leq r \leq t \leq T + 1$, the probability of reaching h^t starting from h^r is $p_\sigma(h^t|h^r) \equiv \prod_{j \in N} \prod_{r \leq s < t} \sigma_j(I_j(h^s))(h_{s,j})$.⁸ It is useful to separate this probability into a part affected only by σ_i and a part affected only by σ_{-i} . These are

⁷To formally write some histories, one may need to include some of the "dummy" actions chosen from singleton action sets needed to create the multi-stage structure of the game. These actions are not depicted in the game tree. For example, if player 1 plays d , to keep the multi-stage structure, this must be followed by a node where all three players have singleton action sets, before arriving at the stage where player 3 has a non-trivial move. So an example of such a history could be written as $(I, (d, \emptyset, \emptyset), (\emptyset, \emptyset, \emptyset), (\emptyset, \emptyset, n))$ where \emptyset is used to denote such "dummy" actions and the triples represent the actions of players 1, 2 and 3 respectively.

⁸If $r = t$, so that the product is taken over an empty set, invoke the convention that a product over an empty set is 1.

$p_{i,\sigma_i}(h^t|h^r) \equiv \prod_{r \leq s < t} \sigma_i(I_i(h^s))(h_{s,i})$ and $p_{-i,\sigma_{-i}}(h^t|h^r) = \prod_{j \neq i} \prod_{r \leq s < t} \sigma_j(I_j(h^s))(h_{s,j})$ respectively, with $p_{i,\sigma_i}(h^t|h^r)p_{-i,\sigma_{-i}}(h^t|h^r) = p_\sigma(h^t|h^r)$. With this notation, we can now state formally the assumption that players' ex-ante preferences over strategies are smooth ambiguity preferences (Klibanoff, Marinacci and Mukerji 2005) with the u_i, ϕ_i and μ_i as specified by the game.

Assumption 2.1 (Ex-ante Preferences) *Fix a game Γ . Ex-ante (i.e., given the empty partial history), given σ_{-i} , each player i ranks strategies σ'_i according to*

$$V_i(\sigma'_i, \sigma_{-i}) \equiv \sum_{\pi \in \Delta(\Theta)} \phi_i \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi(h^0) \right) \mu_i(\pi). \quad (2.1)$$

If there were only a single distribution π over parameters, and so no ambiguity, μ_i would be degenerate and ϕ_i irrelevant and (2.1) specializes to the usual ex-ante expected payoffs (i.e., the expression in the argument of ϕ_i) in the context of a game. In the presence of ambiguity about the parameters, player i aggregates such expected payoffs for each $\pi \in \text{supp } \mu_i$ using ϕ_i and beliefs μ_i . Despite this non-linear aggregation, it proves helpful, in understanding both optimality and the effect of ambiguity aversion on this aggregation, to examine modified expected payoffs using local linear approximations.

As is true for any preference represented by a smooth, increasing and concave objective function, σ_i is optimal if and only if it is optimal according to the local linear approximation of the objective function around σ_i . In the context of (2.1), σ_i an ex-ante best response to σ_{-i} for player i given ϕ_i and μ_i is equivalent to σ_i maximizing, among all σ'_i , the following modified expected payoff,

$$\sum_{h \in H} u_i(h) p_{i,\sigma'_i}(h|h^0) q^{(\sigma, \mu_i), i}(h) \quad (2.2)$$

where $q^{(\sigma, \mu_i), i}$ is i 's ex-ante (σ, μ_i) -local measure over histories, defined for each $h \in H$ by,

$$p_{-i,\sigma_{-i}}(h|h^0) \sum_{\pi \in \Delta(\Theta)} \pi(h^0) \phi'_i \left(\sum_{\hat{h} \in H} u_i(\hat{h}) p_\sigma(\hat{h}|\hat{h}^0) \pi(\hat{h}^0) \right) \mu_i(\pi). \quad (2.3)$$

Notice that i 's ambiguity aversion leads the marginal of this local measure on parameters, $h^0 \in \Theta$, to tilt, via the ϕ'_i term, towards parameters given more weight by π 's for which i expects to fare less well under σ . This tilting is not present under ambiguity neutrality (since ϕ'_i is constant). As ambiguity aversion increases, this tilting becomes more severe.

Now turn to defining preferences beyond the ex-ante stage. To do so, we need to define the notion of a belief at an information set. The only property we require of such a belief is that it puts weight only on distributions over the partial histories belonging to that information

set. Formally:

Definition 2.2 (Belief at an Information Set) *A belief for player i in a game Γ given information set I_i is a finite support probability measure ν_{i,I_i} over $\Delta(I_i)$. An interim belief system $\nu \equiv (\nu_{i,I_i})_{i \in N, I_i \in \mathcal{I}_i}$ is a belief for each player at each of that player's information sets.*

Since at any partial history each player has an information set, an interim belief system specifies a belief for each player at each partial history. Given these beliefs, the following defines a player's preferences at an information set:

Assumption 2.2 (Preferences at an Information Set) *Fix a game Γ and a strategy profile σ . Any player i at information set I_i ranks strategies σ'_i according to*

$$V_{i,I_i}(\sigma'_i, \sigma_{-i}) \equiv \sum_{\pi \in \Delta(I_i)} \phi_i \left(\sum_{h|h^t \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^t) \pi(h^t) \right) \nu_{i,I_i}(\pi), \quad (2.4)$$

where t is the stage at which the information set I_i occurs.

Compared to the ex-ante preferences given in (2.1), the preferences (2.4) at I_i differ only in that (1) the beliefs may have changed in light of I_i and σ (i.e., μ_i is replaced by some belief ν_{i,I_i} concentrated on I_i), and (2) the probabilities of reaching various histories according to the strategy profile are now calculated starting from I_i rather than from the beginning of the game. Given a strategy σ_i for player i , the *continuation strategy at information set I_i* , $\sigma_i^{I_i}$, is the restriction of σ_i to the information sets \hat{I}_i such that $I_i \in R_i(\hat{I}_i)$. Preferences at I_i may be equivalently thought of as ranking continuation strategies (i.e., any two strategies with identical continuations from I_i are treated identically by preferences at I_i).

Just as for the ex-ante preferences, it is useful to observe that σ_i being a best response to σ_{-i} for player i at information set I_i given ϕ_i and ν_{i,I_i} is equivalent to σ_i maximizing, among all σ'_i , the following modified expected payoff,

$$\sum_{h|h^t \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^t) q^{(\sigma, \nu), i, I_i}(h) \quad (2.5)$$

where $q^{(\sigma, \nu), i, I_i}$ is i 's (σ, ν) -local measure over histories given I_i , defined for each $h \in H$ such that $h^t \in I_i$ by,

$$p_{-i, \sigma_{-i}}(h|h^t) \sum_{\pi \in \Delta(I_i)} \phi'_i \left(\sum_{\hat{h}|\hat{h}^t \in I_i} u_i(\hat{h}) p_{\sigma}(\hat{h}|\hat{h}^t) \pi(\hat{h}^t) \right) \pi(h^t) \nu_{i,I_i}(\pi). \quad (2.6)$$

Now, i 's ambiguity aversion leads the marginal of this local measure on I_i to tilt towards partial histories given more weight by π 's for which i expects to fare less well *under the continuation* $\sigma_i^{I_i}$.

Using the above preferences, Section 3 turns to equilibrium analysis. The main focus is describing and analyzing sequential optimality – the assumption that each player is best responding to σ_{-i} according to their preferences at each information set – including its relation with updating of beliefs. Observe that Definition 2.2 assumes no particular connection between beliefs at different information sets, the ex-ante beliefs, and the strategy profile. Our Theorem 3.1 will show that strategic behavior under sequential optimality is as if beliefs are related by a particular update rule. In the case of expected utility preferences, this update rule is exactly Bayes' rule given the strategy profile. More generally, this update rule ensures that, given the strategy profile, i 's local measure at I_i , $q^{(\sigma, \nu), i, I_i}$, is the Bayesian update of i 's local measure at the previous information set (or of i 's ex-ante local measure if I_i is an initial information set).

3 Equilibrium

3.1 Ex-ante Equilibrium

As a step toward defining sequential optimality, we use the ex-ante preferences to define ex-ante (Nash) equilibrium:

Definition 3.1 (Ex-ante Equilibrium) *Fix a game Γ . A strategy profile σ is an ex-ante (Nash) equilibrium if, for all players i ,*

$$V_i(\sigma) \geq V_i(\sigma'_i, \sigma_{-i})$$

for all $\sigma'_i \in \Sigma_i$.

An equilibrium requires each player i , given ex-ante beliefs μ_i , to best respond to the (parameter-contingent) strategies of the other players. To the extent that the play that strategies specify varies with the parameter, ambiguity about parameters (as reflected in beliefs) translates into ambiguity about play in equilibrium. In the case of ambiguity neutrality, where the ϕ_i are linear (i.e., subjective expected utility), and homogeneous ex-ante beliefs, $\mu_i = \mu$ for all players i , the definition reduces to the usual (ex-ante) Bayesian Nash Equilibrium definition. More generally, ex-ante equilibrium is the same as the ex-ante equilibrium defined in the context of strategic form games with ambiguity averse players by Azrieli and

Teper (2011) applied to our setting of extensive form games and smooth ambiguity preferences. Similar ex-ante concepts also appear in Kajii and Ui (2005), Bade (2011), Riedel and Sass (2013) and Kellner (2015).

When there is ambiguity about parameters, the motive for conditioning play on them can come from a mix of direct payoff concerns and/or indirect strategic advantage from having others be ambiguous about one's play. When, however, some or all of the parameter space is payoff irrelevant, the only motive for conditioning play on these payoff irrelevant aspects is indirect strategic advantage. This approach to modeling "pure" strategic ambiguity in equilibrium via strategies conditioning play on payoff irrelevant parameters follows Bade (2011), who extends to ambiguity Aumann (1974)'s modeling of equilibrium strategic uncertainty through conditioning on such parameters.

An alternative approach is to model ambiguity directly over opponents' strategies without requiring a parameter space. In any such approach, unlike in ours, the Nash assumption that each player is best responding to the strategies of the others is problematic as it leaves no possibility of strategic ambiguity. Thus, such approaches, while maintaining that each player's strategy is a best response to their beliefs about strategies, necessarily relax the assumption that these beliefs are correct. Examples of such approaches in the literature include Dow and Werlang (1992), Klibanoff (1996), Lo (1996, 1999), Eichberger and Kelsey (2000), Marinacci (2000), Lehrer (2012), Battigalli et al. (2015) and Battigalli et al. (2017).

When there is some ambiguity aversion (one or more ϕ_i concave) and ambiguity (μ non-degenerate), ex-ante equilibrium behavior can differ from that in Bayesian Nash Equilibrium. Intuition for this is that at an ex-ante equilibrium, each player is behaving as if they maximize expected payoff with respect to their local measure (2.3) at that strategy profile. Even with common beliefs μ , ambiguity averse players will generally have different local measures, in particular, different local marginals on parameters. Thus, ambiguity aversion leads players with common beliefs to act, locally, as if they were standard ambiguity neutral players with heterogeneous beliefs.

Running Example continued: Returning to the peace negotiation example, we analyze the ex-ante equilibria of the game which helps illustrate some of the above. First, suppose that both small countries are ambiguity neutral (i.e., ϕ_1 and ϕ_2 are affine). In this case, no ex-ante equilibrium results in a positive probability of peace. In contrast, whenever players 1 and/or 2 are sufficiently ambiguity averse (e.g., $\phi_1 = \phi_2 = -e^{-ax}$ with $a \geq \ln(\frac{5}{4})$), peace with probability 1 is an ex-ante equilibrium outcome. Formally:

Proposition 3.1 *(i) If players 1 and 2 are ambiguity neutral (i.e., ϕ_1 and ϕ_2 are affine), no ex-ante equilibrium results in a positive probability of peace. This also holds for any other specification of common belief μ .*

(ii) If players 1 and/or 2 are sufficiently ambiguity averse, there is an ex-ante equilibrium yielding peace with probability 1 (i.e., with $\sigma_1(c) = \sigma_2(c) = 1$).

The proof of this and all subsequent results in the paper may be found in the Appendices. Intuitively, players 1 and 2 are willing to cooperate only if each is sufficiently worried that they will be punished if negotiations break down. Given the specified payoffs and beliefs in combination with ambiguity neutrality, this is impossible because any strategy of 3 that incentivizes one of them to cooperate will lead the other to defect. Peace does not occur in a Bayesian Nash equilibrium.

How do ambiguity and ambiguity aversion change the story? Consider the strategy profile where 1 and 2 cooperate and player 3, when the payoff irrelevant parameter is I , punishes 1 if there is defection and favors 2 if there is peace, and does the opposite when the parameter is II . Since the beliefs μ reflect ambiguity about the parameters, 3's strategy creates ambiguity in the minds of players 1 and 2 about who will be punished or favored. Observe that player 1 does worse under parameter I than under parameter II and the reverse is true for player 2. Recall from the discussion following equation (2.3) that ambiguity aversion would then, at this strategy profile, tilt the marginals of their local measures on parameters in different directions – 1's towards I and 2's towards II . Given sufficient ambiguity aversion (to generate enough tilting), cooperation will be a best response for both because deviating is unattractive under parameter I for player 1 and unattractive under parameter II for player 2. Finally, given cooperation by 1 and 2, any strategy of player 3 is an ex-ante best response because 3 is indifferent among all actions that follow cooperation, and actions following defection are off the equilibrium path. ■

Observe that, as in the example, equilibrium disciplines the strategic ambiguity that appears. Given that a player knows that by conditioning his play on ambiguous parameters he can induce strategic ambiguity in the minds of others, the decision to condition in this manner is a decision about whether it is advantageous to play ambiguously. Equilibrium requires that such ambiguity inducing play is permitted only when it is a best response.

3.2 Sequential Optimality

Fundamental to our theory will be sequential optimality. It requires that, given the strategies of the others, each player's strategy is optimal at each of that player's information sets. Since at each partial history each player has an information set containing that partial history, sequential optimality imposes an analogue of ex ante equilibrium at each partial history in the game. Under ambiguity neutrality, the definition of sequential optimality specializes to ex-ante equilibrium plus Kreps and Wilson (1982)'s *sequential rationality*. Formally:

Definition 3.2 Fix a game Γ . A pair (σ, ν) consisting of a strategy profile and interim belief system is sequentially optimal if, for all players i and all information sets I_i ,

$$V_i(\sigma) \geq V_i(\sigma'_i, \sigma_{-i}) \quad (3.1)$$

and

$$V_{i,I_i}(\sigma) \geq V_{i,I_i}(\sigma'_i, \sigma_{-i}) \quad (3.2)$$

for all $\sigma'_i \in \Sigma_i$, where the V_i and V_{i,I_i} are as specified in (2.1) and (2.4).⁹

A strategy profile σ is said to be sequentially optimal whenever there exists an interim belief system ν such that (σ, ν) is sequentially optimal.

As is standard for perfection-like equilibrium concepts, a major motivation for sequential optimality is to rule out non-credible off-path behavior, and off-path restrictions are the only means through which it refines ex-ante equilibrium. Any ex-ante equilibrium for which all information sets are on-path is also sequentially optimal (see Theorem A.1 in the Appendix). How does sequential optimality relate to familiar concepts? Sequential optimality with a common μ implies (and is stronger than) subgame perfection adapted to allow for smooth ambiguity preferences. To see this, recall that a proper subgame follows a partial history at which all information sets are singletons. Since in our games all players have an information set at each partial history, for any proper subgame (3.2) ensures that the continuation strategy profile derived from σ forms an ex-ante equilibrium of the subgame with respect to degenerate beliefs. For the overall game, (3.1) ensures σ is an ex-ante equilibrium. When preferences are expected utility and there is a common μ , we show (see Corollary 3.1) that sequential optimality identifies the same strategy profiles as the version of Perfect Bayesian Equilibrium (PBE) defined in Gibbons (1992, pp. 177-180) (sometimes called weak PBE).¹⁰

Both PBE and sequential optimality go well beyond subgame perfection in several respects. First, they rule out off-path play that is not optimal against any belief given the strategy profile at partial histories that generate non-singleton information sets (where subgame perfection has no bite). Second, if, for example, a simultaneous move game is triggered immediately following a deviation, they require play of an equilibrium of that incomplete information simultaneous move game given beliefs and the strategy profile. Note that this

⁹Note that since $V_{i,I_i}(\tilde{\sigma}_i, \sigma_{-i}) = V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i})$ if $\tilde{\sigma}_i^{I_i} = \hat{\sigma}_i^{I_i}$, requiring the inequalities for the V_{i,I_i} to hold as i changes only her continuation strategy given I_i would result in an equivalent definition.

¹⁰In games where at some partial histories information sets are specified for only a subset of players (for example, as is frequently done, only for the player(s) having a non-trivial move at that point), ex-ante equilibrium plus sequential rationality need not imply subgame perfection (see e.g., Mas-Colell et al. (1995, p. 290)'s Example 9.C.5) and thus the same is true of sequential optimality. However, once the “missing” information sets are added (in the case of Example 9.C.5, specify an information set for the incumbent containing the partial history “In”) the implication of subgame perfection is restored.

second implication is generally more restrictive than simply ruling out strategies that are (iteratively) strictly dominated given the information set.

Running Example continued: Earlier we showed that only when players 1 and/or 2 are sufficiently ambiguity averse is there an ex-ante equilibrium resulting in peace with probability 1. This equilibrium is also subgame perfect, since there are no off-path proper subgames. Is such behavior sequentially optimal? The answer depends on x , the payoff player 3 receives when staying neutral after negotiations break down. Only when x is sufficiently low (specifically $x \leq 1$) will there be a sequentially optimal strategy profile leading to peace for sure (or even to a positive probability of peace). Why? When $x > 1$, staying neutral is more attractive for player 3 than punishing, and thus punishment, at least some of which is necessary to incentivize successful peace negotiations, becomes a non-credible threat and cannot be sustained as an optimal response to any beliefs after negotiations break down. When $x \leq 1$, some ambiguous punishment strategies by 3 (e.g., playing $p1$ if the parameter is I and $p2$ if the parameter is II) are best responses to some beliefs of 3 about who deviated from playing c . Given sufficient ambiguity aversion of players 1 and 2, such punishment strategies are enough to ensure cooperation. Formally:

Proposition 3.2 *If $x > 1$, then in all sequential optima players 1 and 2 play d with probability 1. If $x \leq 1$ and players 1 and 2 are sufficiently ambiguity averse, then there is a sequential optimum yielding peace with probability 1.*

Remark 3.1 *When $0.5 < x \leq 1$, to have a sequential optimum yielding peace with probability 1 may require strictly more ambiguity aversion than would be necessary to have an ex-ante equilibrium doing so. For example, if $\phi_1(x) \equiv -e^{-x \ln(5/3)} = -(\frac{5}{3})^{-x}$ (so that $\frac{\phi'_1(x)}{\phi'_1(y)} = (\frac{5}{3})^{y-x}$) and player 2 is ambiguity neutral, then, as is shown in the proof of Proposition 3.1, there is an ex-ante equilibrium yielding peace with probability 1. However, there is no sequential optimum yielding peace with probability 1. The intuition is that when $0.5 < x \leq 1$, optimality for player 3 when given the move limits the ability to punish the less ambiguity averse player 2 by ruling out mixtures that include both $p1$ and $p2$. In contrast, when $x \leq 0.5$ enough punishment strategies are credible that there is a sequential optimum yielding peace whenever there is an ex-ante equilibrium doing so.*

■

The definition of sequential optimality does not assume particular updating. Nevertheless, we show that analysis of sequential optima of a game may be undertaken under the *as if* assumption that all players use an update rule, proposed by Hanany and Klibanoff (2009), called the *smooth rule*. Under ambiguity neutrality, the smooth rule reduces to Bayes' rule.

More generally, the smooth rule applied to beliefs implies the local measures, $q^{(\sigma, \nu), i, I_i}(h)$, as defined in (2.6), are related by Bayes' rule.

We now define the smooth rule in the game-theoretic context. Though it is notationally complex, at this point all that is important to take from this definition is that i 's updated beliefs are proportional to i 's beliefs at the previous information set times i 's subjective likelihood of reaching from the previous information set to the current one given σ_{-i} , weighted by a term (involving a ratio of ϕ'_i expressions) that can be non-constant only when the player is ambiguity averse. Thus, smooth rule updating is a re-weighted version of Bayesian updating. When this weighting term is constant, as is the case under ambiguity neutrality, the smooth rule reduces to Bayes' rule.¹¹ For clarity, the smooth rule formula is stated in two pieces, one for updating to an initial information set from the ex-ante stage, and the other for updating to a current information set from the previous one. The smooth rule pins down i 's updating at all information sets (both on and off path) except those that i does not view as reachable from the immediately preceding information set given σ . We also define a strong version of the smooth rule which additionally pins down updating at information sets that i does not view as reachable given σ only because of i 's own strategy σ_i (under ambiguity neutrality, this will be referred to as strong Bayes' rule¹²). Notice, as with Bayes' rule in PBE, that the updating formula applies "wherever possible" including at all on-path and some off-path information sets. We defer further discussion of the rule, and suggest that the reader may also wish to defer parsing its formal definition.

Notation 3.1 For information set I_i , define $s(I_i)$ to be such that $I_i \in \mathcal{I}_i^{s(I_i)}$, i.e., at what stage of the game is I_i . Given a partial history $h^t \in \mathcal{H}$ and $-1 \leq s \leq t-1$, h^s is the partial history formed by truncating h^t just before stage s .

Notation 3.2 For information set $I_i \not\subseteq \Theta$, define I_i^{-1} to be the information set immediately preceding I_i in $R_i(I_i)$.

Definition 3.3 An interim belief system ν satisfies the smooth rule using strategy profile σ if the following holds for each player i and information set I_i , letting t denote the stage at which the information set I_i occurs (i.e., $t = s(I_i)$):

¹¹More generally, these weighting terms reflect differences in the motive to hedge against ambiguity at different information sets (see Hanany and Klibanoff 2009 and Baliga, Hanany and Klibanoff 2013).

¹²Such a version of Bayes' rule is discussed, e.g., in Hendon, Jacobsen and Sloth (1996) and Perea (2002).

If $I_i \subseteq \Theta$, then for all $\pi \in \Delta(I_i)$,

$$\nu_{i,I_i}(\pi) \propto \sum_{\hat{\pi} \in \Delta(\Theta) | \hat{\pi}_{I_i} = \pi} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \hat{\pi}(h^0) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi(h^t) \right)} \hat{\pi}(I_i) \mu_i(\hat{\pi}), \quad (3.3)$$

where $\hat{\pi}_{I_i} \in \Delta(I_i)$ is given by $\hat{\pi}_{I_i}(\theta) = \frac{\hat{\pi}(\theta)}{\sum_{\hat{\theta} \in I_i} \hat{\pi}(\hat{\theta})}$; and

if $I_i \not\subseteq \Theta$ and $\sum_{h^t \in I_i} \sum_{\hat{\pi} \in \Delta(I_i^{-1})} p_\sigma(h^t|h^{t-1}) \hat{\pi}(h^{t-1}) \nu_{i,I_i^{-1}}(\hat{\pi}) > 0$, then for all $\pi \in \Delta(I_i)$,

$$\begin{aligned} \nu_{i,I_i}(\pi) \propto \sum_{\hat{\pi} \in \Delta(I_i^{-1}) | \hat{\pi}_{I_i} = \pi} & \frac{\phi'_i \left(\sum_{h|h^{t-1} \in I_i^{-1}} u_i(h) p_\sigma(h|h^{t-1}) \hat{\pi}(h^{t-1}) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi(h^t) \right)} \\ & \cdot \left(\sum_{h^t \in I_i} p_{-i,\sigma_{-i}}(h^t|h^{t-1}) \hat{\pi}(h^{t-1}) \right) \nu_{i,I_i^{-1}}(\hat{\pi}), \end{aligned} \quad (3.4)$$

where $\hat{\pi}_{I_i} \in \Delta(I_i)$ is given by $\hat{\pi}_{I_i}(h^t) = \frac{p_{-i,\sigma_{-i}}(h^t|h^{t-1}) \hat{\pi}(h^{t-1})}{\sum_{\hat{h}^t \in I_i} p_{-i,\sigma_{-i}}(\hat{h}^t|\hat{h}^{t-1}) \hat{\pi}(\hat{h}^{t-1})}$ for all $h^t \in I_i$.

If (3.4) additionally holds for all $I_i \not\subseteq \Theta$ for which $\sum_{h^t \in I_i} \sum_{\hat{\pi} \in \Delta(I_i^{-1})} p_{-i,\sigma_{-i}}(h^t|h^{t-1}) \hat{\pi}(h^{t-1}) \nu_{i,I_i^{-1}}(\hat{\pi}) > 0$, then we say ν satisfies the strong smooth rule using strategy profile σ .

We are now ready to state our result that analysis of sequential optima of a game may be undertaken under the “as if” assumption that all players use smooth rule updating.

Theorem 3.1 *Fix a game Γ and a strategy profile σ . Then σ is sequentially optimal if and only if there exists an interim belief system $\hat{\nu}$ satisfying the smooth rule using σ such that $(\sigma, \hat{\nu})$ is sequentially optimal.*

Outline of the proof: The if direction follows by definition. The only if direction proceeds by considering each player i and information set I_i separately (and has analogues for the ex-ante stage). Pair each I_i with its collection of immediate successors J_i . The argument makes use of the following construction of the interim belief system $\hat{\nu}$: let ν be an interim belief system such that (σ, ν) is sequentially optimal; at information sets where the

smooth rule is unrestrictive let $\hat{\nu}_{i,I_i} = \nu_{i,I_i}$; given this and the ex-ante beliefs, the smooth rule formula then pins down $\hat{\nu}$ at all other information sets.

Three steps are key to showing that $(\sigma, \hat{\nu})$ is sequentially optimal. First, as is true for any preference represented by a smooth, increasing and concave objective function, σ_i is optimal at I_i if and only if it is optimal according to the local linear approximation around σ_i given in (2.5). Second, perfect recall and the linearity of this objective allows us to conclude that σ_i maximizes (2.5) implies that, for each J_i , σ_i maximizes the part of the summation in (2.5) taken only over J_i . Finally, the smooth rule formula in (3.4) implies that the $(\sigma, \hat{\nu})$ - local measure given each J_i is the Bayesian update, given J_i and σ , of (and therefore proportional to) the $(\sigma, \hat{\nu})$ - local measure given I_i . Thus, applying this and again the first step, for each J_i , σ_i is optimal at J_i given $\hat{\nu}$. Therefore $(\sigma, \hat{\nu})$ is sequentially optimal. ■

Note that Theorem 3.1 would be false if we were to replace the smooth rule with Bayes' rule – restricting attention to interim belief systems satisfying Bayesian updating generally rules out some (or all) sequentially optimal strategies. This is so because applying Bayes' rule to beliefs does *not* generate Bayesian updating of the local measures, q . The latter was the implication of smooth rule updating of beliefs essential to proving the theorem.

One characteristic of smooth rule updating that will be unfamiliar to many readers is that its formula includes expected payoffs under σ at all (terminal) histories reachable from the previous information set, as well as ambiguity aversion via ϕ_i . Recall from Definition 2.4, however, that preferences at any information set I_i rank strategies only through their continuation from I_i since beliefs, however determined, are concentrated on measures over I_i . The reader might nevertheless be worried by the fact that the evaluation of a continuation strategy at an information set depends on updated beliefs, but updated beliefs under the smooth rule depend on the continuation strategy at the information set. This should not be a concern, because the only comparisons that are meaningful when checking if σ_i is a best response to σ_{-i} at an information set I_i are those evaluating σ_i and each alternative σ'_i according to the same belief, ν_{i,I_i} , whatever it is. Theorem 3.1 says that this belief *may be determined by updating according to the smooth rule using $\sigma = (\sigma_i, \sigma_{-i})$* .

Recall that, restricting attention to expected utility preferences, the smooth rule specializes to Bayes' rule. In this case, we have the following corollary:

Corollary 3.1 *Fix a game Γ with all players ambiguity neutral and a strategy profile σ . Then σ is sequentially optimal if and only if there exists an interim belief system $\hat{\nu}$ satisfying Bayes' rule such that $(\sigma, \hat{\nu})$ is sequentially rational.*

Thus, under ambiguity neutrality, sequential optimality identifies the same set of strategy profiles as Kreps and Wilson (1982)'s sequential rationality plus the assumption of Bayesian

updating given σ , which are, in turn, the same as perfect Bayesian equilibrium (PBE) as defined in, for example, Gibbons (1992).¹³

When applying concepts like subgame perfection, PBE and sequential equilibrium, it is extremely useful, when verifying optimality, to only need to check “one-stage” deviations (as opposed to general deviations), i.e., check that for each player i and information set I_i , there are no profitable deviations by i at I_i alone. These one-stage deviations are typically a small fraction of the deviations available to players. Formally, the absence of these profitable one-stage deviations is the following:

Definition 3.4 *The pair (σ, ν) has no profitable one-stage deviations if for each player i and each information set I_i , $V_{i,I_i}(\sigma) \geq V_{i,I_i}(\sigma'_i, \sigma_{-i})$ for all σ'_i agreeing with σ_i everywhere except possibly at I_i .*

For finite horizon games, in the standard ambiguity neutral case and under the assumption that beliefs are related by strong Bayes’ rule given the strategy profile σ , having no profitable one-stage deviations is sufficient for sequential optimality of σ (see Hendon, Jacobsen and Sloth 1996). Is there an analogous statement that applies under ambiguity aversion? As we show next, the answer is yes. In this analogue, the role of strong Bayesian updating given σ is played by the strong smooth rule.

Theorem 3.2 *Fix a game Γ and a pair (σ, ν) such that ν satisfies the strong smooth rule using σ . Then (σ, ν) is sequentially optimal if and only if (σ, ν) has no profitable one-stage deviations.*

The small body of existing literature on dynamic games with ambiguity aversion that also attempts to deal with non-credible threats does this by requiring no profitable one-stage deviations (or the slightly stronger Strotzian consistent planning, see Appendix B) with respect to exogenously imposed particular update rules.^{14,15} We noted above that requiring no profitable one-stage deviations with respect to beliefs updated by Bayes’ rule may, under ambiguity aversion, rule out some (or all) sequential optima. It may even rule out those for which all information sets are on-path. Thus, differently from sequential optimality, it may refine ex-ante equilibria by restricting on-path as well as off-path behavior. In fact, it is not a refinement of ex-ante equilibrium – it may also admit strategy profiles that are not ex-ante equilibria. This is illustrated by the following example.

¹³Shimoi and Watson (1998) prove a related result in the context of defining extensive form rationalizability – the set of such rationalizable strategies when defined using best responses given any conjectures about others’ play remains the same when limiting attention to conjectures consistent with Bayes’ rule.

¹⁴Examples of such literature include Bose and Renou (2014), Battigalli et al. (2017), Beauchêne, Li and Li (2017) and Kellner and Le Quement (2018).

¹⁵A second approach appearing in the literature is based on recursive preferences. We relate sequential optimality to this approach in Section 7 when discussing Ellis (2018) and Pahlke (2018).




Figure 3.1: Game illustrating violation of sequential optimality. The vectors give utility payoffs for players 1 and 2, in that order, for each path.

Example: Consider the game in Figure 3.1. Player 2 is privately informed of $\theta \in \{I, II, III\}$ at the beginning of the game. Observe that for each θ , player 2 has a strictly dominant strategy if given the move: types *I* and *II* play *U*, and type *III* plays *D*. Given this strategy for player 2, observe that for player 1, the payoff to playing *i* followed by *d* if *U* is, θ -by- θ , strictly higher than the payoff to playing *o* followed by anything. Thus, no strategy involving playing *o* with positive probability can be a best reply to 2's optimal strategy no matter how player 1 perceives and treats the uncertainty about θ . This implies that *o* is not part of any ex-ante (Nash) equilibrium, let alone a sequential optimum.

In contrast, it is easy to specify ambiguity averse preferences under which player 1 playing *o* with positive probability and playing *u* with probability $1/2$ if given the move together with player 2 playing her strictly dominant strategy if given the move has no profitable one-stage deviations under Bayesian updating of beliefs. Notice that from the perspective of player 1, taking as given playing *u* with probability $1/2$, *i* does worse than *o* under *II* and better than *o* otherwise. If there is ambiguity concerning the probability of types *II* versus *III* and player 1 is sufficiently ambiguity averse, when considering playing *i* with probability 1 and

evaluating moving toward o , that change will be evaluated using weights tilted sufficiently toward II to make that movement attractive. The details may be found in Appendix B.¹⁶

■

From a decision theoretic perspective, the vast majority of atemporal models of ambiguity averse preferences adopt state-by-state monotonicity as a fundamental building block. Thus, another reason for interest in sequential optimality as a solution concept for dynamic games is that it rules out the play of type-by-type (iteratively) strictly dominated strategies in dynamic settings. Sequential optimality is a stronger requirement than consistent planning (without exogenous assumptions on updating), which is, in turn, stronger than requiring no profitable one-stage deviations. It then follows from Theorem 3.2 that, when restricted to beliefs satisfying the strong smooth rule, all three concepts are equivalent. This finding generalizes the fact that under expected utility, when restricted to beliefs satisfying strong Bayes' rule, the three are equivalent.

Do sequential optima always exist? In the next section we explore a refinement of sequential optimality. We show existence for this refinement, thus implying existence of sequential optima.

3.3 Sequential Equilibrium with Ambiguity

To define sequential equilibrium with ambiguity (SEA), we consider a condition, *smooth rule consistency*, that imposes requirements on beliefs that may have bite even at information sets that i does not view as reachable from their immediate predecessor given σ .¹⁷ Our condition extends Kreps and Wilson (1982)'s *consistency* condition used for the same purpose in defining sequential equilibrium. We extend consistency in order to accommodate ambiguity aversion by replacing Bayes' rule in their definition with the smooth rule. This replacement is justified by Theorem 3.1, since, by that result, sequential optima are as if players are responding to beliefs generated using smooth rule updating. Just as consistency uses limits of Bayesian updates to deliver beliefs consistent with small trembles converging to sequentially optimal strategies under ambiguity neutrality, limits of smooth rule updates deliver this under ambiguity aversion. Recall that if we simply limited attention to Bayes' rule, then sequentially optimal strategies might fail to exist under ambiguity aversion. Smooth rule consistency is defined as follows:

¹⁶There we also show that strengthening the no profitable one-stage deviation criterion to a Strotzian consistent planning requirement does not eliminate the play of o with positive probability in the example.

¹⁷If there are no such information sets, any sequentially optimal strategy profile is also part of an SEA (see Theorem A.3).

Definition 3.5 (Smooth Rule Consistency) *Fix a game Γ . A pair (σ, ν) consisting of a strategy profile and interim belief system satisfies smooth rule consistency if there exists a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$, with $\lim_{k \rightarrow \infty} \sigma^k = \sigma$, such that $\nu = \lim_{k \rightarrow \infty} \nu^k$, where each ν^k is the interim belief system satisfying the smooth rule using σ^k .*

Observe that smooth rule consistency is a true extension of Kreps and Wilson's consistency because Bayes' rule and the smooth rule coincide under ambiguity neutrality. SEA strengthens sequential optimality exactly by adding the requirement of smooth rule consistency:

Definition 3.6 (SEA) *A sequential equilibrium with ambiguity (SEA) of a game Γ is a pair (σ, ν) consisting of a strategy profile and interim belief system such that (σ, ν) is sequentially optimal and satisfies smooth rule consistency.*

We may use Theorem 3.2 and that smooth rule consistency delivers a ν satisfying the strong smooth rule using σ to conclude that replacing sequential optimality in the definition of SEA by having no profitable one-stage deviations would not change the set of equilibrium strategies.

Theorem 3.3 *A pair (σ, ν) is an SEA if and only if (σ, ν) has no profitable one-stage deviations and satisfies smooth rule consistency.*

An implication of this last result together with Corollary 3.1 is that SEA and sequential equilibrium are equivalent under the assumption of ambiguity neutrality.

Running Example continued: Let's apply SEA to our running example. We show that if peace is an ex-ante equilibrium outcome, then it is also an SEA outcome if and only if the large country's payoff x to staying neutral after a breakdown is at most 0.5. Notice that this is strictly lower than the bound of 1 we saw when applying sequential optimality.

Proposition 3.3 *Suppose there is an ex-ante equilibrium yielding peace with probability 1. Then there exists an SEA yielding peace with probability 1 if and only if $x \leq 0.5$.*

The key is that smooth rule consistency forces beliefs of player 3 about which player defected to be the same at both information sets that follow defection. This is not implied by sequential optimality alone, and when $0.5 < x \leq 1$, attaining peace requires player 3 to hold different beliefs across these information sets. Once these beliefs are forced to be the same, it is not in 3's interest to play in a way such that both players 1 and 2 think there

is a chance they will be punished following defection. Therefore, at least one of the small countries will find it in their interest to defect.

In contrast, when $x \leq 0.5$, at each of 3's information sets following defection, any mixture over punishing 1 and punishing 2 is a best response to beliefs such that the two nodes in the information set are unambiguously equally likely. Thus, given such beliefs, 3 is willing to punish in a way sufficient to ensure peace. Such beliefs satisfying smooth rule consistency can be generated, for example, by a completely mixed sequence of strategies converging to always cooperating that give, along the sequence, the same probability to 1 defecting as to 2 defecting after 1 cooperates. ■

We next show that every game Γ has at least one SEA (and thus also at least one sequential optimum and ex-ante equilibrium). Since the functions ϕ_i describing players' ambiguity attitudes are part of the description of Γ , this result goes beyond the observation that an SEA would exist if players were ambiguity neutral, and ensures existence given any specified ambiguity aversion and ex-ante beliefs. The manner in which the smooth rule generalizes Bayes' rule allows us to prove this result by adapting known techniques from existence proofs for sequential equilibrium.

Theorem 3.4 *An SEA exists for any game Γ .*

In Appendix A (see Theorem A.2), we show that SEA implies that beliefs are uniquely defined at all information sets according to a version of the smooth rule formula using limits of likelihoods. This provides a method for constructing beliefs satisfying smooth rule consistency.

4 Effects of ambiguity aversion on equilibria

4.1 Comparative statics in ambiguity aversion

In this section, we explore the extent to which changes in ambiguity aversion affect equilibrium play. When we say equilibrium in this section, it will not matter whether we refer to ex-ante equilibria, sequential optima or SEAs, as the comparative statics in ambiguity aversion will be the same for all of these. We start with the simplest and most direct comparative statics question: Holding all else fixed about a game (which, recall, includes specifying beliefs), do changes in ambiguity aversion affect the set of equilibrium strategy profiles (and play paths)? The answer is yes they can, as was true in our running example. In fact, the set of equilibria can change entirely, as the following result shows:

Theorem 4.1 *For some game form, payoffs and beliefs, the set of equilibrium strategy profiles under ambiguity neutrality is disjoint from that under some ambiguity aversions.*

What if, as an outside observer, one is not willing to fix particular beliefs when describing the equilibrium predictions of the theory, but is willing to assume that all players share the same belief? How do such predictions change when ambiguity aversion is introduced?

Theorem 4.2 *Fix any game form and payoffs. Taking the union, over beliefs μ , of the set of equilibria generated if the common belief were μ , ambiguity aversion makes this union weakly larger (in the superset sense) compared to ambiguity neutrality, and, for some games and ambiguity aversions, this expansion is strict.*

Thus, under an assumption of common beliefs, ambiguity aversion may not only generate new equilibrium behavior (and new paths of play), but also does not eliminate equilibria possible under ambiguity neutrality. For instance, strict expansion occurs in the running example with sufficient ambiguity aversion and $x \leq 0.5$.

Does dropping the restriction to common beliefs change the answer to the question in the previous paragraph? It does, and quite dramatically so – in this case, we show that the predictions of the theory do not change with ambiguity aversion:

Theorem 4.3 *Fix any game form and payoffs. Ambiguity aversion never affects the union, over beliefs $(\mu_i)_{i \in N}$, of the set of equilibria generated if beliefs were $(\mu_i)_{i \in N}$.*

To gain intuition for the previous two results, it is useful to rewrite the linear approximation in (2.2) as follows:

$$\sum_{\pi \in \Delta(\Theta)} \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi(h^0) \right) \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi(h^0) \right) \mu_i(\pi) \quad (4.1)$$

Equation (4.1) is a linear aggregation of the expected payoffs for each π with weights on the π given by the ϕ'_i term times μ_i . We will refer to the normalized-to-one version of these weights on the π as i 's *effective beliefs* (at σ), since they are the local analogue of beliefs μ_i in the ex-ante preferences. We see from (4.1) that the impact of changing ambiguity aversion on equilibria comes from the effect that ambiguity aversion has in generating different effective beliefs for different players. In Theorem 4.2, this is what may generate new equilibria. In Theorem 4.3, observe that any change in effective beliefs coming from changes in ϕ'_i can be offset by corresponding changes in μ_i . This offsetting is not possible under the restriction to a common μ because the offsetting required for each player may differ.

Theorem 4.3 has analogues in the literature. For a result that in individual decision problems, under standard assumptions (including reduction, broad framing, statewise dominance and expected utility evaluation of objective lotteries), all observed behavior optimal according to ambiguity averse preferences is also optimal for *some* subjective expected utility preferences, see e.g., Kuzmics (2015). Bade (2016) independently shows that without restrictions on beliefs, predictions using ex-ante equilibria do not change with ambiguity aversion. Considering a type of self-confirming equilibria, Battigalli et al. (2015, p. 667) show that the set of these equilibria does not change as ambiguity aversion changes.

Battigalli et al. (2015) have as their main finding a result (their Theorem 1 together with an example of strict inclusion), which Battigalli et al. (2017) focuses on partially extending to dynamic games, in which they show that the set of their self-confirming equilibria increases as ambiguity aversion increases and that this increase can be strict. This result relies crucially on limiting attention to pure strategies (both in terms of the equilibrium profile and in terms of the deviations against which optimality is checked). If we were also to limit attention to pure strategies in both these respects, an analogous result would apply to our equilibria:

Theorem 4.4 *Fix any game form and payoffs. Taking the union, over beliefs $(\mu_i)_{i \in N}$, of the set of pure equilibria with respect to pure strategy deviations generated if the beliefs were $(\mu_i)_{i \in N}$, increasing ambiguity aversion of one or more players weakly expands (in the superset sense) this union. For some games and increases in ambiguity aversion, this expansion is strict.*

4.2 Robustness

To introduce our discussion of robustness, we begin with an example. Consider a two player one-stage game, where each player has a choice between two actions A and B . The parameter space about which there is ambiguity is $\Theta = \{\theta_1, \theta_2\}$. Both players have beliefs μ such that $\mu(\pi_1) = \mu(\pi_2) = \frac{1}{2}$, where $\pi_1(\theta_1) = \frac{2}{3}$ and $\pi_2(\theta_1) = \frac{1}{2}$, and do not learn anything about θ before choosing their action. They both share the same ϕ . Payoffs as a function of the actions and θ are as follows:

θ_1	A	B	θ_2	A	B
A	0, 0	1, -8	A	6, 6	1, 16
B	-8, 1	-6, -6	B	16, 1	12, 12

(4.2)

Observe that given θ_1 , A is strictly dominant for each player, while given θ_2 , B is strictly dominant. Under ambiguity neutrality, i.e. ϕ affine, both (A, A) and (B, B) are equilibrium strategy profiles. We claim that (A, A) is robust to increased ambiguity aversion (i.e., remains

an equilibrium when ϕ becomes more concave), but (B, B) is not. To see that (A, A) is robust, note that, assuming her opponent plays A , a player evaluates the mixed strategy $\lambda A + (1 - \lambda)B$ according to $\frac{1}{2}\phi(2\lambda) + \frac{1}{2}\phi(4 - \lambda)$, which is maximized at $\lambda = 1$ for any concave ϕ . To see that (B, B) is not robust, note that for example, if $\phi(x) = -e^{-\alpha x}$ with $\alpha > \ln(\frac{1+\sqrt{5}}{2}) \approx 0.48$, it is profitable to deviate to A .

Another sense of robustness is that an equilibrium supported for a wider range of beliefs is more robust. Consider the set of weights μ on π_1 and π_2 that support (A, A) as an equilibrium. Such weights are those satisfying $\mu(\pi_1) \geq \frac{\phi'(3)}{2\phi'(2)+\phi'(3)}$. Notice that as ϕ becomes more concave, $\frac{\phi'(3)}{2\phi'(2)+\phi'(3)}$ decreases, approaching 0 in the limit as $\frac{\phi'(3)}{\phi'(2)}$ approaches 0, and thus sufficient ambiguity aversion results in a large set of weights μ supporting (A, A) . The fact that ambiguity aversion leads to such a large set of beliefs supporting (A, A) is not special to this example. We show, under some conditions, that equilibria that are robust to increased ambiguity aversion *must* be supported by a large set of beliefs for sufficient ambiguity aversion, and furthermore, this supporting set of beliefs may be made as large as desired (see Theorem 4.5 and Remark 4.1). We refer to this as ambiguity aversion making an equilibrium belief robust.

One use of our robustness result is as follows: Consider a population having heterogeneous beliefs. Equilibria that, under ambiguity neutrality, are not supported by many beliefs might not be expected to occur often. Our result offers ambiguity aversion as a possible explanation for unexpected prevalence of such equilibria. Specifically, if such an equilibrium is, like (A, A) , robust to increased ambiguity aversion, ambiguity aversion can make it an equilibrium for more of the population (i.e., for more beliefs).

We turn to formal definitions of these two robustness notions and our result relating them. An equilibrium strategy profile is robust to increased ambiguity aversion if it remains an equilibrium whenever one or more of the ϕ_i becomes more concave:

Definition 4.1 *For a game Γ , an equilibrium σ is robust to increased ambiguity aversion if it remains an equilibrium whenever, for each i , ϕ_i is replaced by an at least as concave $\hat{\phi}_i$.*

Ambiguity aversion makes an equilibrium strategy profile σ belief robust if sufficient increases in players' ambiguity aversion, holding the π 's in the supports of players' beliefs $(\mu_i)_{i \in N}$ fixed, make all beliefs placing sufficient weight on each such π support σ as an equilibrium:

Definition 4.2 *For a game Γ , consider an equilibrium σ . Ambiguity aversion makes σ belief robust if, for each i and $\varepsilon_i \in (0, \frac{1}{|\text{supp } \mu_i|})$, there exists $\bar{\phi}_i^{\varepsilon_i}$ at least as concave as ϕ_i so that:*

σ is an equilibrium of this game with $(\hat{\mu}_i)_{i \in N}$ and $(\hat{\phi}_i)_{i \in N}$ whenever the $(\hat{\mu}_i)_{i \in N}$ have the same

supports as the $(\mu_i)_{i \in N}$ and, for each i , $\min_{\pi \in \text{supp } \mu_i} \hat{\mu}_i(\pi) > \varepsilon_i$ and $\hat{\phi}_i$ at least as concave as $\bar{\phi}_i^{\varepsilon_i}$.

The next result shows, under some conditions on how expected payoffs vary with π , the tight connection between robustness to increased ambiguity aversion and belief robustness.

Theorem 4.5 *Fix a game Γ . The following is true when either ex-ante equilibrium or sequential optimality are used as the notion of equilibrium:*

If an equilibrium σ is robust to increased ambiguity aversion and, for each player i , the expected payoff $\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0)$ can be strictly ordered across the π in the support of μ_i , then ambiguity aversion makes σ belief robust.

Theorem 4.5 also holds when SEA is used as the equilibrium concept under the following modification: In addition to SEA, require that there is a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$, such that $\lim_{k \rightarrow \infty} \sigma^k = \sigma$, with respect to which smooth rule consistency simultaneously holds for any interim belief systems used to support σ as an SEA in the theorem or associated definitions.

Remark 4.1 *Suppose ϕ_i is twice continuously differentiable with $\phi_i' > 0$ everywhere. Then Theorem 4.5 remains true when, in the definition of belief robust, $\bar{\phi}_i^{\varepsilon_i}$ is restricted to be of the constant absolute ambiguity aversion form, $\bar{\phi}_i^{\varepsilon_i}(x) \equiv -e^{-\alpha(\varepsilon_i)x}$ where $\alpha(\varepsilon_i) > 0$.*

Intuition for the theorem in the case of ex-ante equilibrium is as follows: Robustness to increased ambiguity aversion implies that σ_i must be a best response given the minimizing π , since, as long as there is a unique expected payoff minimizing π , one can always find increases in ambiguity aversion that generate effective beliefs that are *any* interior convex combination of the effective beliefs under the original ϕ_i and degenerate beliefs on the minimizing π . Given this, if one were to go to the limit (i.e., all effective weight placed on the minimizing π), then the beliefs over the π cease to matter and all beliefs with the same support make σ_i a best response. The proof of the theorem reveals that the arguments required along the way toward the limit are more subtle, making use of concave transformations tailored to generate specific shifts in effective beliefs when defining threshold $\bar{\phi}_i^{\varepsilon_i}$ that do the job and relaxing all beliefs to all beliefs up to the ε_i constraints. As Remark 4.1 indicates, if one doesn't mind bounds that may be much less tight, comparison to the threshold $\bar{\phi}_i^{\varepsilon_i}$ may be simplified by taking the threshold to have a constant coefficient of ambiguity aversion $\alpha(\varepsilon_i)$ (see Klibanoff, Marinacci and Mukerji (2005), pp. 1865-6).

Without the assumption on expected payoffs, the theorem would be false. The role of this assumption is to ensure enough flexibility in the manner in which more ambiguity aversion can shift the effective weight placed on expected payoffs for the various π .

In Section 5.2, we apply Theorem 4.5 in analyzing the effect of the entrant's ambiguity aversion on the robustness of limit-pricing equilibria.

5 Examples

5.1 Example: Ambiguous Cheap Talk

The example is a game in which deliberately introducing ambiguity about actions *without* payoff consequences (“ambiguous cheap talk”) proves valuable in equilibrium (ex-ante, sequentially optimal or SEA). The equilibrium we identify would not be a solution under no profitable one-stage deviations (or Strotzian consistent planning) together with Bayesian updating.

There are three players, a principal, P , and two agents, r (ow) and c (olumn). P wishes to induce the agents to w (ork) for him rather than start their own b (usinesses). There is uncertainty about the value to an agent of b relative to w . P has private information about these values. By making an informative public announcement, P hopes to improve the chances that the agents find w attractive. We show an ambiguous communication strategy is optimal for P and is part of an SEA.

The parameter space has two components, a payoff relevant component, which can take the value I or II , related to market-relevant characteristics of a technology, and a payoff irrelevant component, which can take the value U or D , related to the findings of a laboratory experiment. Thus the parameter space is $\Theta = \{IU, ID, IIU, IID\} = \{I, II\} \times \{U, D\}$. At stage $t = 0$, only P has a non-trivial move, which is to send a message α or β . At $t = 1$, only the agents have non-trivial moves, and each chooses b or w . P is privately informed of $\theta \in \Theta$ before sending his message. The message is publicly observed by both agents before they choose their actions. Payoffs are given in the following matrices, where each cell lists the payoff to P , r , and c in that order:

I	b	w	II	b	w
b	0, 0, 5	0, 0, 1	b	0, 5, 0	0, 5, 1
w	2, 1, 5	2, 2, 2	w	0, 1, 0	2, 2, 2

Notice that P 's message is cheap talk. To understand the above payoffs, begin with P . He has an idea concerning the use of the technology and the skills and labor of the agents to

make a product. Full success of the product occurs under technology I if r works for P (no matter what c does), but under technology II requires both agents to work for P as both of their skills are crucial in this case. Partial success occurs under I if only c works for P , and under technology II if either of the agents works alone for P . If neither agent works for P , nothing is accomplished with regard to the product and no payment is made by P . Now turn to the agents. If an agent works for P , she gets some benefit, but she also incurs an effort cost that is higher than when both agents work for P (thus her payoff of 1 from working alone increases to 2 when working together). If an agent does not work for P , she starts an independent business based on her own idea for using the technology. Agent r 's business idea will be a huge success under technology II but amount to nothing under technology I , while the reverse is true of c 's business idea. Agent r is ambiguity averse with $\phi_r(x) = -e^{-11x}$. The exact specification of ϕ_P and ϕ_c will not be important for our analysis of the game.

The beliefs μ for all players are $\frac{1}{2}-\frac{1}{2}$ over distributions π_1 and π_2 given by (where the numbers on the top and left edges are the corresponding marginals):

		3/4	1/4			1/5	4/5
	π_1	I	II		π_2	I	II
1/3	U	3/12	1/12	1/4	U	1/20	4/20
2/3	D	6/12	2/12	3/4	D	3/20	12/20

Notice that there is ambiguity about the payoff relevant component of θ and, fixing that component, ambiguity about the payoff irrelevant component of θ . This belief structure is, for example, consistent with there being an underlying factor which affects (the likelihood of) *both* components. The factor might be some scientific principle that is not well understood, which influences both the functioning of the technology (I vs. II) as well as the findings of the laboratory experiment (U vs. D) not affecting any of the players' business ventures.

First consider the case where all players are ambiguity neutral. If P plays an uninformative strategy (e.g., sends the same message for all values of the parameter), then calculation shows that both agents will respond by playing b for sure, and P would get a payoff of 0. However, P can do better. The following strategy profile is an ex-ante equilibrium under ambiguity neutrality: P fully reveals the payoff relevant component of the parameter, and the agents play their dominant strategies in response. That is: If the payoff relevant component of the parameter is I , P sends message α , otherwise P sends message β ; after message α , r plays w and c plays b , while after message β , r plays b and c plays w . Under this strategy, P gets his maximal payoff of 2 when I occurs, but gets 0 when II occurs.

Why isn't there an equilibrium where P does better than this? Any possible improvement must involve incentivizing both agents to play w with positive probability when II occurs. However, since the only way to convince r to play w is to have her put sufficient weight on I occurring while c is convinced to play w only if she puts sufficient weight on II occurring, it is impossible under ambiguity neutrality for P to have it both ways.

Next reintroduce r 's ambiguity aversion ($\phi_r(x) = -e^{-11x}$), and consider the following strategy profile, σ^* : If the parameter is IU , P sends message α , otherwise P sends message β ; after either message, r plays w ; after message α , c plays b , and after β , c plays w . Observe that P is making use of the payoff irrelevant component of the parameter. We show that this strategy profile is an equilibrium (Proposition 5.1), and, that the principal does strictly better than if he were not allowed to use the payoff irrelevant component (Proposition 5.2). Formally:

Proposition 5.1 *The strategy profile σ^* is an SEA. In this equilibrium, P attains his maximum possible payoff for each parameter.*

Remark 5.1 The strategy profile σ^* remains an equilibrium for any ϕ_r more concave than the one in the example.

Proposition 5.2 *If P were not allowed to make his strategy depend on the payoff irrelevant component of the parameter (i.e., U or D), there would be no ex-ante equilibrium yielding P the maximum possible payoff for each parameter.*

How does playing the ambiguous communication strategy σ_P^* help P do better in the example? It allows P to expose r to more ambiguity in equilibrium than r would be exposed to under the optimal communication strategy that does not make use of U vs. D . This additional ambiguity helps P provide incentives to both agents simultaneously, by moving their effective beliefs further apart in the desired directions.

Notice that the only interesting updated beliefs are those of the agents after having observed the message β (as following α the agents know the payoff relevant component of the parameter is I). By Theorem 3.1, it is sufficient to consider smooth rule updated beliefs. Recall that π_1 puts more weight on I than does π_2 . Since σ^* performs worse for r under I than under II , r does worse under π_1 than under π_2 in equilibrium. Therefore, ambiguity aversion leads r 's smooth rule updated belief to place more weight on π_1 than Bayesian updated beliefs would. This is crucial in ensuring sequential optimality of σ^* following β , as r placing more weight on I pushes r towards playing w . For c it is the reverse, i.e., since σ^* performs better for c under I than II , c does better under π_1 than under π_2 in equilibrium, and therefore c 's smooth rule updated belief places weight on π_1 that is (weakly) less than

the Bayesian updated belief. To ensure that both players coordinate on playing w after β , differing updated beliefs are crucial:¹⁸ if they shared a common updated belief, at least one agent would deviate.

Remark 5.2 If agent r becomes sufficiently more ambiguity averse, Proposition 5.2 no longer holds: in addition to the equilibrium in Proposition 5.1, there will be an equilibrium where P conditions his play only on I vs. II and also obtains his maximum possible payoff for each parameter. Intuitively, with enough ambiguity aversion on the part of r , the additional ambiguity generated by conditioning on U vs. D is no longer needed.

5.2 Example: Limit Pricing under Ambiguity

In this section, we use a parametric class of games based on Milgrom and Roberts (1982)'s limit pricing entry model with the twist that the entrant has ambiguity about the incumbent's cost and is ambiguity averse. In this application, SEA refines ex-ante equilibrium and we find conditions under which the entrant's ambiguity aversion makes limit pricing behavior more robust compared to ambiguity neutrality.

The game is as follows: An incumbent monopolist has private information concerning its per-unit production costs c_I (which is one of $c_L < c_M < c_H$). Thus the parameter space is $\Theta = \{L, M, H\}$.¹⁹ In the first stage, the incumbent chooses a quantity that, together with inverse market demand, $P(Q) = a - bQ$, $a, b > 0$, and c_I determines its first period profit. A potential entrant with known per-unit production costs c_E observes this quantity and decides whether or not to enter at the second stage. If no entry is chosen, in the final stage the incumbent remains a monopolist and again chooses a quantity while facing the same market demand and costs as in the first stage, and the entrant gets a payoff of zero. If entry is chosen, the entrant pays a fixed cost $K \geq 0$, the incumbent's cost is learned by the entrant, and in the final stage the two firms compete in a complete information Cournot duopoly with the same market demand. To make this a finite game, suppose a finite set of feasible quantities \mathcal{Q} (including at least the monopoly quantities for each possible production cost and the complete information Cournot quantities).²⁰ Denote the entrant's beliefs and ambiguity aversion by μ and ϕ respectively. The incumbent's beliefs and ambiguity aversion play no role in our analysis.

¹⁸In the context of individual decisions, such belief polarization under ambiguity is explored in Baliga, Hanany and Klibanoff (2013).

¹⁹The use of at least three costs is necessary to have non-trivial updating on the equilibrium path under pure strategy limit pricing. With only two possible costs, pure limit pricing strategies involve full pooling.

²⁰The strategies we construct remain SEA strategies no matter what finite set of feasible quantities is assumed as long as the monopoly and Cournot quantities for each cost are included.

We construct an SEA strategy profile σ^{LP} where in the first stage, incumbent types M and L pool at the monopoly quantity for L , and type H plays the monopoly quantity for H . Then the entrant enters after observing any quantity strictly below the monopoly quantity for L and does not enter otherwise, and in the final stage they play the monopoly or duopoly quantities accordingly. These strategies involve limit pricing by incumbent type M – it raises its quantity (thus lowering price) in the first stage in order to successfully deter entry.

For later reference, we collect here conditions assumed explicitly or implicitly already plus restrictions equivalent to all monopoly and duopoly quantities being positive:

Assumption 5.1 $a, b > 0$, $K \geq 0$, $c_H > c_M > c_L \geq 0$, $c_E \geq 0$, $a > c_H$, $a + c_E - 2c_H > 0$ and $a + c_L - 2c_E > 0$.

The following proposition provides sufficient conditions for σ^{LP} to be not only part of an SEA, but also robust to increased ambiguity aversion and more belief robust. One way in which SEA refines ex-ante equilibrium in this example is by requiring that the Cournot quantities in the complete information duopoly game following entry are played (there are ex-ante equilibria violating sequential optimality that involve the incumbent deterring all entry by threatening to flood the market if entry occurs). The robustness results tell us that ambiguity aversion can enlarge the circumstances under which limit pricing can be equilibrium behavior.

What is the role of the conditions in the proposition? The first three conditions correspond to the following incentives in the game: ICH for I ensures that a type H incumbent does not want to pool with the other types to deter entry, ICM for I ensures that a type M incumbent does not want to separate from type L and stop deterring entry, and ICH for E ensures that the entrant strictly wants to enter when it is sure the incumbent is type H . The combination of the two subsequent conditions on the beliefs and the assumption of sufficient ambiguity aversion of the entrant ensure that it does not want to enter after observing the limit price (i.e., the monopoly quantity for type L).

Proposition 5.3 *Under Assumption 5.1, the limit pricing strategy profile σ^{LP} is part of an SEA if*

$$\left(\frac{a + c_E - 2c_H}{3}\right)^2 \geq \frac{a - c_L}{2} \left(a - \frac{a - c_L}{2} - c_H\right), \quad (\text{ICH for I})$$

$$\frac{a - c_L}{2} \left(a - \frac{a - c_L}{2} - c_M\right) \geq \left(\frac{a + c_E - 2c_M}{3}\right)^2, \quad (\text{ICM for I})$$

$$b \left(\frac{a + c_H - 2c_E}{3b}\right)^2 > K, \quad (\text{ICH for E})$$

some $\pi \in \text{supp } \mu$ makes entry conditional on $\{L, M\}$ strictly unprofitable, all $\pi \in \text{supp } \mu$ can be ordered in the likelihood-ratio ordering, and the entrant is sufficiently ambiguity averse.

Moreover, under the same conditions, σ^{LP} is SEA robust to increased ambiguity aversion, and ambiguity aversion makes it SEA belief robust.

The proof uses the formula for an interim belief system ν satisfying smooth rule consistency provided by Theorem A.2 to establish that (σ^{LP}, ν) is an SEA and uses Theorem 4.5 to establish belief robustness. It follows from the above result that for any beliefs such that some $\pi \in \text{supp } \mu$ makes entry conditional on $\{L, M\}$ strictly unprofitable and all $\pi \in \text{supp } \mu$ can be ordered in the likelihood-ratio ordering and that lead an ambiguity neutral entrant to want to enter even after observing the limit price, there exists a large enough increase in the entrant's ambiguity aversion such that the entrant would be deterred by the limit price. In this way, increasing ambiguity aversion leads to expansion in the set of beliefs μ that can support such a limit pricing SEA.

6 Extensions

6.1 Other models of ambiguity averse players

We have assumed players have smooth ambiguity preferences (Klibanoff, Marinacci and Mukerji, 2005), which proved very convenient in many respects. Can our approach be applied to players with other kinds of ambiguity averse preferences? We suggest how to do so for any preferences that can be represented by $W_i(U_i(\sigma, \theta)_{\theta \in \Theta})$, where W_i is a continuous, monotonic and quasi-concave aggregator (across the parameters $\theta \in \Theta$) of the vector $U_i(\sigma, \theta)_{\theta \in \Theta}$ of i 's expected utilities of σ . Quasi-concavity of W_i reflects ambiguity aversion. This includes smooth ambiguity preferences along with many other models in the literature, and is essentially what Cerreia-Vioglio et al. (2011) call Uncertainty Averse preferences. Note that Maxmin Expected Utility (Gilboa and Schmeidler, 1989) is in this class, and if the set of probability measures in the Maxmin EU representation is taken to be the (convex hull of) the support of μ_i , then these preferences can be interpreted as a model of an infinitely ambiguity averse player with beliefs given by the support of μ_i .

By modifying our framework to specify W_i rather than μ_i and ϕ_i , the definition of ex-ante preferences and equilibrium are easily adapted. However, since such preferences do not necessarily have separately specified beliefs and ambiguity aversion, the notion of interim belief system would need to be replaced by an interim preference system (i.e., an interim preference for each player and information set). Given that change, sequential optimality could be defined. Based on our proof of Theorem 3.1, we conjecture the following would be true: σ is sequentially optimal if and only if there exists an interim preference system derived by updating preferences so that the local measure in some local linear approximation of the

updated preferences at σ is the Bayesian update of the local measure in some local linear approximation of the preferences from the previous information set at σ with respect to which σ is sequentially optimal. An analogously modified version of Theorem 3.2 is conjectured to hold as well. Observe that there are two key differences from our current results: first, the reference to *some* local linear approximation is needed to reflect the possibility of non-smooth preferences, and, second and more importantly, only updating of local approximations is specified in the new result, and not updating of beliefs themselves or even of the preference representation as a whole. Specifying an update rule for the preferences themselves requires more structure. While smooth rule updating of beliefs generates such updating for smooth ambiguity preferences, updates generating the local approximations property for, among others, Maxmin EU and Variational preferences (Maccheroni, Marinacci and Rustichini, 2006) are described in Hanany and Klibanoff (2007, 2009). In defining SEA, replacement of smooth rule consistency with a consistency condition based on preference updates satisfying a similar local approximations updating property would be needed.

Providing results and examples involving comparative statics in ambiguity aversion and robustness to increased ambiguity aversion and belief robustness would, even to pose the relevant questions properly, require some kind of separate specification and manipulation of ambiguity aversion and of beliefs. Here the smooth ambiguity model, with ambiguity aversion (via ϕ_i) and beliefs (via μ_i) separately and conveniently specified, was especially helpful. We conjecture that if one had some other class of Uncertainty Averse preferences where these components could be sensibly specified then one could investigate these issues.

6.2 Implementation of mixed actions

Recall that the objects of choice of a player are behavior strategies, which, for each type of the player, specify a mixture over the available actions at each point in the game where the player has an opportunity to move. Suppose at some point a player's strategy specifies a non-degenerate mixture, and, as can happen under ambiguity aversion, this strategy is strictly better than any specifying a pure action. If such a mixture is to be implemented by means of playing pure actions contingent on the outcome of a (possibly existing in the player's mind only) randomization device, then an additional sequential optimality concern beyond that formally reflected in Definition 3.2 may be relevant. Specifically, after the realization of the randomization device is observed, will it be optimal for the player to play the corresponding pure action? A way to ensure this is true is to consider behavior strategies that, instead of specifying mixed actions, specify pure actions contingent on randomization devices, and extend the specification of beliefs and preferences of a player to include points

after realization of her randomization device but before she has taken action contingent on the device, and add to Definition 3.2 the requirement of optimality also at these points. The properties of sequential optimality shown and used in this paper would remain true under these modifications.

6.3 Other extensions

We briefly discuss a final pair of possible extensions. First, in the running example, there is only one payoff irrelevant component of the parameter. Suppose one wanted to allow the players to condition on any payoff irrelevant ambiguity they wish. This may be (approximately) achieved by enriching the parameter space to include many such components, and specify μ so that these reflect a rich (but finite) collection of ambiguous devices. Such enrichment would allow, for instance, explicitly modeling the large country in the running example as choosing to condition on a payoff irrelevant component having the “optimal” ambiguity about it. Furthermore, our point that the large country will strictly want to condition its play on some such ambiguous component is robust to any enrichment of this form. The same applies to the ambiguous cheap talk example.

Second, as written, our theory does not allow a player to be uncertain about the ambiguity aversion, i.e. ϕ_i , of other players. This might be done as follows: Introduce a ϕ_i -type component of the parameter space that ϕ_i is allowed to depend on, and assume that the first thing that happens in the game is that each player learns their own ϕ_i -type. The point immediately after this occurs would be treated as the ex-ante stage of the game, and the analysis would then proceed exactly as in the paper.

7 Closely related literature

To the best of our knowledge, we are the first to propose an equilibrium notion for dynamic games with incomplete information that requires sequential optimality while allowing for ambiguity averse preferences. In this section we relate our approach to the few papers which investigate general dynamic games with incomplete information and ambiguity aversion.²¹

²¹In addition to the papers we mention when discussing ex-ante equilibrium in Section 3.2, a number of previous papers have analyzed incomplete information games with ambiguity sensitive preferences in settings without dynamics, including Salo and Weber (1995), Levin and Ozdenoren (2004), Bose, Ozdenoren and Pape (2006), Chen, Katuscak and Ozdenoren (2007), Lopomo, Rigotti and Shannon (2010), Bodoh-Creed (2012), Wolitzky (2013, 2016), Ayouni and Koessler (2017), di Tillio, Kos and Messner (2017) and Auster (2018). Additional papers on dynamic games with ambiguity not discussed here include Eichberger and Kelsey (1999, 2004), Dominiak and Lee (2017) and Muraviev, Riedel and Sass (2017).

There have been only a very few papers investigating general dynamic games with incomplete information and ambiguity aversion. The two most closely related to ours are Battigalli et al. (2017) and Pahlke (2018). Battigalli et al. (2017) explores a notion of self-confirming equilibrium in dynamic games where players are ambiguity averse with smooth ambiguity preferences (building on Battigalli et al. (2015), which did the same for games in strategic form, and so took a purely ex-ante perspective). There are a number of key differences from our approach. First, self-confirming equilibrium places essentially no restriction on off-path beliefs or behavior and so is not designed to address strategic concerns such as perfection and credibility of off-path threats that are central to our analysis. Second, they limit attention to Bayesian updating of beliefs μ_i , leading to violations of sequential optimality even at on-path information sets. Instead of sequential optimality, they require “unimprovability” which can be thought of as roughly no profitable one-stage deviations at on-path information sets.²² Third, they assume players choose only pure actions at each information set, with any randomization being modeled by explicitly included moves of an artificial separate player who is assumed to randomize over actions using commonly known probabilities. A practical consequence of this is that their no profitable one-stage deviations requirement is generally strictly more permissive under ambiguity aversion than ours at a given information set because they are only checking against a (at most) finite selection of mixed actions (which, recall, under ambiguity aversion may be strictly better than any pure action). Though mixed strategy profiles appear in their definition of equilibrium, those are mixtures only in the population sense of distributions over a population of players in the same role who may have some heterogeneity in the pure strategy they play. The main result of Battigalli et al. (2015) was a comparative static: the set of self-confirming equilibria was shown to expand as players became more ambiguity averse. A main focus in Battigalli et al. (2017) is investigating the extent to which this result carries over to dynamic games. They find it does not carry over in general due to possible on-path dynamic inconsistency, but does extend under conditions where this is not an issue (see their Section 6). This finding complements our Theorem 4.4 showing that, when similarly limiting attention to pure strategies and pure-strategy deviations, the sets of ex-ante, sequentially optimal and SEA profiles expand as players become more ambiguity averse (see Section 4.1 and Appendix A.4.2). As Battigalli et al. (2015) showed for self-confirming equilibria, this expansion depends crucially on the restriction on mixed strategies and becomes equality once the full set of mixtures is considered (our Theorem 4.3).

²²In an extension, they explore a rationalizable version of their self-confirming equilibrium, where rationality is defined in terms of unimprovability at all information sets. This yields some restrictions on off-path behavior.

Subsequent to our paper, Pahlke (2018) explores a notion of sequential equilibrium in dynamic games where players are ambiguity averse with Recursive Maxmin Expected Utility preferences (preferences shown by Epstein and Schneider (2003a) to be equivalent to Maxmin Expected Utility (MEU) preferences (Gilboa and Schmeidler, 1989) together with prior-by-prior Bayesian updating, where the set of priors satisfies a condition known as rectangularity). In Pahlke’s definition of a game, a notable non-standard aspect is that the specific preferences, even ex-ante, are not a primitive. Rather, these ex-ante preferences for each player i are derived according to a procedure that generates a rectangular set $\Pi_i \subseteq \Delta(\Theta)$ for each player depending on the (other players’ part of the) strategy profile σ under consideration, a set $\mathcal{P} \subseteq \Delta(\Theta)$ which is a primitive of the game, and the filtration defined by player i ’s information sets. That the Π_i may vary across players is important in ensuring existence of Pahlke’s equilibria when the Π_i are non-singleton, i.e., when there is non-trivial ambiguity. This heterogeneity explains how Pahlke’s approach overcomes the specialization of the results of Ellis (2018) to MEU preferences, which implies that when Θ consists of privately known types for each player and there is a common rectangular set of priors across players, ex-ante preferences must be ambiguity neutral. We make four observations comparing Pahlke (2018) to our framework and approach when adapted to MEU preferences (see Section 6.1). Fix a game form and payoffs. First, for each (\mathcal{P}, σ) , Pahlke’s analysis of whether σ is an equilibrium corresponds to checking whether it is so according to our approach in the game in which each player i has ex-ante MEU preferences with set of priors Π_i . Second, as generally $\Pi_i \neq \mathcal{P}$, equilibria Pahlke identifies for a given \mathcal{P} will differ from the ones we would identify for that \mathcal{P} . Third, since the Π_i vary with σ , Pahlke’s determination of the set of equilibria given a particular \mathcal{P} involves potentially analyzing a different game in our framework for each candidate σ . Fourth, since each generated Π_i must be rectangular, the strategy profiles that are an equilibrium for at least one \mathcal{P} according to Pahlke form a subset of those that are an equilibrium according to our approach for at least one assignment of sets of priors to each player i .

The literature also includes papers presenting analysis restricted to specific applications of dynamic games of incomplete information with ambiguity. These include Bose and Daripa (2009), Bose and Renou (2014), Kellner and Le Quement (2017, 2018), Beauchêne, Li and Li (2017), Auster and Kellner (2018), and all focus on behavior that cannot occur under ambiguity neutrality. Their approach is to assume MEU preferences and prior-by-prior Bayesian updating and use as an equilibrium concept optimality under consistent planning in the spirit of Strotz (1955-56) (see Appendix B for a definition). Thus, as was discussed and illustrated in Section 3.2, the set of equilibria identified by this approach include strategy profiles that are not sequentially optimal, and exclude some or all of the sequentially optimal

strategy profiles. In fact, all of the behavior they emphasize violates sequential optimality.

References

- [1] R. AUMANN, Subjectivity and correlation in randomized strategies, *Journal of Mathematical Economics*, **1** (1974), 67-96.
- [2] S. AUSTER, Robust contracting under common value uncertainty, *Theoretical Economics*, **13** (2018), 175-204.
- [3] S. AUSTER, C. KELLNER, Robust Bidding and Revenue in Descending Price Auctions, mimeo. (2018).
- [4] M. AYOUNI, F. KOESSLER, Hard Evidence and Ambiguity Aversion, *Theory and Decision*, **82** (2017), 327-339.
- [5] Y. AZRIELI, R. TEPER, Uncertainty aversion and equilibrium existence in games with incomplete information, *Games and Economic Behavior* **73(2)** (2011), 310-317.
- [6] S. BADE, Ambiguous act equilibria, *Games and Economic Behavior* **71(2)** (2011), 246-260.
- [7] S. BADE, Dynamic Semi-Consistency, mimeo. (2016).
- [8] S. BALIGA, E. HANANY, P. KLIBANOFF, Polarization and Ambiguity, *American Economic Review* **103(7)** (2013), 3071-3083.
- [9] P. BATTIGALLI, S. CERREIA-VIOGLIO, F. MACCHERONI, M. MARINACCI, Self-confirming Equilibrium and Model Uncertainty, *American Economic Review* **105(2)** (2015), 646-77.
- [10] P. BATTIGALLI, E. CATONINI, G. LANZANI, M. MARINACCI, Ambiguity Attitudes and Self-Confirming Equilibrium in Sequential Games, Working Paper 607, IGIER, Bocconi. (2017).
- [11] D. BEAUCHÊNE, J. LI, M. LI, Ambiguous Persuasion, mimeo. (2017).
- [12] T. BEWLEY, Knightian Decision Theory and Econometric Inference, Cowles Foundation Discussion Paper no. 868, (1988).
- [13] A.L. BODOH-CREED, Ambiguous beliefs and mechanism design, *Games and Economic Behavior* **75(2)** (2012), 518-537.

- [14] S. BOSE, A. DARIPA, A dynamic mechanism and surplus extraction under ambiguity, *Journal of Economic Theory* **144**(5) (2009), 2084-2114.
- [15] S. BOSE, E. OZDENOREN, A. PAPE, Optimal Auctions with Ambiguity, *Theoretical Economics* **1**(4) (2006), 411-438.
- [16] S. BOSE, L. RENOU, Mechanism Design With Ambiguous Communication Devices, *Econometrica* **82** (2014), 1853-1872.
- [17] S. CERREIA-VIOGLIO, F. MACCHERONI, M. MARINACCI, L. MONTRUCCHIO, Uncertainty Averse Preferences, *Journal of Economic Theory* **146**(4) (2011), 1275-1330.
- [18] Y. CHEN, P. KATUSCAK, E. OZDENOREN, Sealed Bid Auctions with Ambiguity: Theory and Experiments, *Journal of Economic Theory*, **136**(1) (2007), 513-535.
- [19] A. DOMINIAK, M. S. LEE, Coherent Dempster-Shafer Equilibrium and Ambiguous Signals, *Journal of Mathematical Economics*, **68** (2017), 42-54.
- [20] J. EICHBERGER, D. KELSEY, Education-Signalling and Uncertainty, in Machina, M.J., and Munier, B.R. (eds.), *Beliefs, Interactions and Preferences in Decision Making* (1999), 135-157, Kluwer.
- [21] J. EICHBERGER, D. KELSEY, Non-Additive Beliefs and Strategic Equilibria, *Games and Economic Behavior* 30, (2000), 183-215.
- [22] J. EICHBERGER, D. KELSEY, Sequential Two-Player Games with Ambiguity, *International Economic Review* **45** (2004), 1229-1261.
- [23] A. ELLIS, On Dynamic Consistency in Ambiguous Games, *Games and Economic Behavior*, (2018), forthcoming.
- [24] L. G. EPSTEIN, M. SCHNEIDER, Recursive multiple-priors, *Journal of Economic Theory*, **113**(1) (2003a), 1-31.
- [25] L. G. EPSTEIN, M. SCHNEIDER, IID: independently and indistinguishably distributed, *Journal of Economic Theory*, **113**(1) (2003b), 32-50.
- [26] D. FUDENBERG, J. TIROLE, *Game Theory* (1991a), MIT Press.
- [27] D. FUDENBERG, J. TIROLE, Perfect Bayesian Equilibrium and Sequential Equilibrium, *Journal of Economic Theory*, **53** (1991b), 236-260.

- [28] P. GHIRARDATO, Defining Ambiguity and Ambiguity Aversion, *Uncertainty in Economic Theory: A Collection of Essays in Honor of David Schmeidler's 65th Birthday*, I. Gilboa. et al. (Eds.), (2004), London: Routledge.
- [29] P. GHIRARDATO, M. SINISCALCHI, Ambiguity in the small and in the large, *Econometrica* **80** (2012), 2827-2847.
- [30] R. GIBBONS, *Game Theory for Applied Economists* (1992), Princeton University Press.
- [31] I. GILBOA, M. MARINACCI, Ambiguity and the Bayesian Paradigm, *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society*, D. Acemoglu, M. Arellano, and E. Dekel (Eds.), (2013), New York: Cambridge University Press.
- [32] I. GILBOA, D. SCHMEIDLER, Maxmin Expected Utility with a Non-Unique Prior, *Journal of Mathematical Economics*, **18** (1989), 141-153.
- [33] I.L. GLICKSBERG, A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points, *Proceedings of the American Mathematical Society* **3(1)** (1952), 170-174.
- [34] J. GREENBERG, The Right to Remain Silent, *Theory and Decision*, **48** (2000), 193-204.
- [35] E. HANANY, P. KLIBANOFF, Updating Preferences with Multiple Priors, *Theoretical Economics* **2** (2007), 261-298.
- [36] E. HANANY, P. KLIBANOFF, Updating Ambiguity Averse Preferences, *The BE Journal of Theoretical Economics (advances)* **9(1)** (2009), article 37.
- [37] L. P. HANSEN, Beliefs, Doubts and Learning: Valuing Macroeconomic Risk, *American Economic Review* **97(2)** (2007), 1-30.
- [38] E. HENDON, J. JACOBSEN, B. SLOTH, The one-shot-deviation principle for sequential rationality, *Games and Economic Behavior*, **12** (1996), 274-282.
- [39] A. KAJII, T. UI, Incomplete Information Games with Multiple Priors, *Japanese Economic Review* **56(3)** (2005), 332-351.
- [40] C. KELLNER, Tournaments as a Response to Ambiguity Aversion in Incentive Contracts, *Journal of Economic Theory*, **159(A)** (2015), 627-655.
- [41] C. KELLNER, M. LE QUEMENT, Endogenous ambiguity in cheap talk, *Journal of Economic Theory* **173** (2018), 1-17.

- [42] C. KELLNER, M. LE QUEMENT, Modes of ambiguous communication, *Games and Economic Behavior* **104** (2017), 271-292.
- [43] P. KLIBANOFF, M. MARINACCI, S. MUKERJI, A Smooth Model of Decision Making under Ambiguity, *Econometrica* **73** (2005), 1849-1892.
- [44] P. KLIBANOFF, M. MARINACCI, S. MUKERJI, Recursive Smooth Ambiguity Preferences, *Journal of Economic Theory* **144** (2009), 930-976.
- [45] D.M. KREPS, R. WILSON, Sequential Equilibria, *Econometrica* **50(4)** (1982), 863-894.
- [46] C. KUZMICS, Abraham Wald's complete class theorem and Knightian uncertainty, *Games and Economic Behavior*, **104** (2017), 666-673.
- [47] E. LEHRER, Partially Specified Probabilities: Decisions and Games, *American Economic Journal: Microeconomics*, **4(1)** (2012), 70-100.
- [48] D. LEVIN, E. OZDENOREN, Auctions with Uncertain Numbers of Bidders, *Journal of Economic Theory*, **118(2)**, (2004) 229-251.
- [49] K.C. LO, Extensive Form Games with Uncertainty Averse Players, *Games and Economic Behavior* **28** (1999), 256-270.
- [50] G. LOPOMO, L. RIGOTTI, C. SHANNON, Uncertainty in Mechanism Design, mimeo. (2014).
- [51] F. MACCHERONI, M. MARINACCI, A Strong Law of Large Numbers for Capacities, *Annals of Probability* **33** (2005), 1171-1178.
- [52] F. MACCHERONI, M. MARINACCI, A. RUSTICHINI, Ambiguity aversion, robustness, and the variational representation of preferences, *Econometrica* **74(6)** (2006), 1447-1498.
- [53] M. MARINACCI, Ambiguous Games, *Games and Economic Behavior* **31** (2000), 191-219.
- [54] P. MILGROM, J. ROBERTS, Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis, *Econometrica* **50(2)** (1982), 443-459.
- [55] I. MURAVIEV, F. RIEDEL, L. SASS, Kuhn's Theorem for Extensive Form Ellsberg Games, *Journal of Mathematical Economics*, **68** (2017), 26-41.

- [56] S. MUKERJI, J-M. TALLON, An overview of economic applications of David Schmeidler's models of decision making under uncertainty. In Gilboa, I. (ed.), *Uncertainty in Economic Theory* (2004), 283-302, Routledge, London.
- [57] R.B. MYERSON, *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA (1991).
- [58] M. PAHLKE, Dynamic Consistency in Incomplete Information Games with Multiple Priors, mimeo. (2018).
- [59] A. PEREA, A note on the one-deviation property in extensive form games, *Games and Economic Behavior* **40** (2002), 322-338.
- [60] F. RIEDEL, L. SASS, Ellsberg games, *Theory and Decision* **76** (2013), 469-509.
- [61] L. RIGOTTI, C. SHANNON, T. STRZALECKI, Subjective Beliefs and Ex Ante Trade, *Econometrica* **76** (2008), 1167-1190.
- [62] A. SALO, M. WEBER, Ambiguity Aversion in First Price Sealed Bid Auctions, *Journal of Risk and Uncertainty* **11** (1995), 123-137.
- [63] R. SELTEN, Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit, *Zeitschrift für die gesamte Staatswissenschaft* **121**, (1965), 301-324 and 667-689.
- [64] M. SHIMOJI, J. WATSON, Conditional Dominance, Rationalizability, and Game Forms, *Journal of Economic Theory* **83** (1998), 161-195.
- [65] M. SINISCALCHI, Dynamic choice under ambiguity, *Theoretical Economics* **6** (2011), 379-421.
- [66] R.H. STROTZ, Myopia and Inconsistency in Dynamic Utility Maximization, *The Review of Economic Studies* **23(3)** (1955-56), 165-180.
- [67] A. DI TILLIO, N. KOS, M. MESSNER, The Design of Ambiguous Mechanisms, *The Review of Economic Studies* **84(1)** (2017), 237-276.
- [68] A. WOLITZKY, Bilateral Trading with Maxmin Agents, mimeo. Stanford University (2013).
- [69] A. WOLITZKY, Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade, *Theoretical Economics* **11** (2016), 971-1004.

A Appendix: Proofs

A.1 Proofs related to the running example

Proof of Proposition 3.1. To prove (i), suppose ϕ_1 and ϕ_2 are affine. Fix any σ such that peace (i.e., 1 and 2 play c) occurs with positive probability. We will show that at least one player has a profitable deviation. Denoting the total probability $(\sum_{\pi \in \Delta(\Theta)} \mu(\pi) \pi(\theta) \sigma_3(\theta)(a))$ that player 3 assigns to each action a by $\rho(a)$, player 1 and 2's ex-ante preferences are, respectively, given by:

$$V_1(\sigma) = \sigma_1(c) \sigma_2(c) (4\rho(f2) + 5\rho(f1)) + (1 - \sigma_1(c) \sigma_2(c)) (6\rho(n) + 10\rho(p2)),$$

and

$$V_2(\sigma) = \sigma_1(c) \sigma_2(c) (5\rho(f2) + 4\rho(f1)) + (1 - \sigma_1(c) \sigma_2(c)) (6\rho(n) + 10(1 - \rho(n) - \rho(p2))).$$

If it is not profitable to deviate to $\sigma_1(c) = 0$ for player 1, then

$$6\rho(n) + 10\rho(p2) \leq 4\rho(f2) + 5(1 - \rho(f2)) = 5 - \rho(f2),$$

while if it is not profitable to deviate $\sigma_2(c) = 0$ for player 2, then

$$4\rho(n) + 10\rho(p2) \geq -5\rho(f2) - 4(1 - \rho(f2)) + 10 = 6 - \rho(f2),$$

a contradiction. Thus at least one player wants to deviate to playing d . Observe that this argument holds for any common μ not just the one specified in the example.

Turning to (ii), fix a strategy profile σ defined by $\sigma_1(c) = \sigma_2(c) = 1$ and $\sigma_3(I)(p1) = \sigma_3(II)(p2) = \sigma_3(I)(f2) = \sigma_3(II)(f1) = 1$. First, observe that any strategy by player 3 is an ex-ante best response, since on the equilibrium path 3 receives a payoff of 4 no matter what strategy 3 plays. Second, player $i \in \{1, 2\}$, when choosing the probability λ_i with which to play c , is ex-ante best responding if and only if

$$1 \in \arg \max_{\lambda_i \in [0,1]} \frac{1}{2} \phi_i(4\lambda_i) + \frac{1}{2} \phi_i(5\lambda_i + 10(1 - \lambda_i)).$$

This is equivalent to

$$4\phi'_i(4) - 5\phi'_i(5) \geq 0$$

which is

$$\frac{\phi'_i(4)}{\phi'_i(5)} \geq \frac{5}{4}. \quad (\text{A.1})$$

Observe that by increasing the concavity of ϕ_i one can increase $\frac{\phi'_i(4)}{\phi'_i(5)}$ as much as desired.

Suppose one player is sufficiently ambiguity averse that (A.1) is satisfied, but the other is not. When is there an ex-ante equilibrium yielding peace with probability 1? It is necessary to consider more general strategies of player 3 than specified above. Again let σ denote the strategy profile and again require $\sigma_1(c) = \sigma_2(c) = 1$. We leave player 3's strategy flexible. Player 1 is ex-ante best responding if and only if

$$1 \in \arg \max_{\lambda_1 \in [0,1]} \frac{1}{2} \phi_1 \left(\begin{array}{l} (4\sigma_3(I)(f2) + 5(1 - \sigma_3(I)(f2))\lambda_1 \\ + (10\sigma_3(I)(p2) + 6\sigma_3(I)(n))(1 - \lambda_1) \end{array} \right) \\ + \frac{1}{2} \phi_1 \left(\begin{array}{l} (4\sigma_3(II)(f2) + 5(1 - \sigma_3(II)(f2))\lambda_1 \\ + (10\sigma_3(II)(p2) + 6\sigma_3(II)(n))(1 - \lambda_1) \end{array} \right).$$

This is equivalent to

$$\begin{aligned} & (4\sigma_3(I)(f2) + 5(1 - \sigma_3(I)(f2)) - 10\sigma_3(I)(p2) - 6\sigma_3(I)(n)) \\ & \cdot \phi'_1(4\sigma_3(I)(f2) + 5(1 - \sigma_3(I)(f2))) \\ & + (4\sigma_3(II)(f2) + 5(1 - \sigma_3(II)(f2)) - 10\sigma_3(II)(p2) - 6\sigma_3(II)(n)) \\ & \cdot \phi'_1(4\sigma_3(II)(f2) + 5(1 - \sigma_3(II)(f2))) \\ & \geq 0. \end{aligned} \quad (\text{A.2})$$

Player 2 is ex-ante best responding if and only if

$$1 \in \arg \max_{\lambda_2 \in [0,1]} \frac{1}{2} \phi_2 \left(\begin{array}{l} (5\sigma_3(I)(f2) + 4(1 - \sigma_3(I)(f2))\lambda_2 \\ + (10(1 - \sigma_3(I)(n) - \sigma_3(I)(p2)) + 6\sigma_3(I)(n))(1 - \lambda_2) \end{array} \right) \\ + \frac{1}{2} \phi_2 \left(\begin{array}{l} (5\sigma_3(II)(f2) + 4(1 - \sigma_3(II)(f2))\lambda_2 \\ + (10(1 - \sigma_3(II)(n) - \sigma_3(II)(p2)) + 6\sigma_3(II)(n))(1 - \lambda_2) \end{array} \right).$$

This is equivalent to

$$\begin{aligned}
& (5\sigma_3(I)(f2) + 4(1 - \sigma_3(I)(f2)) - 10(1 - \sigma_3(I)(n) - \sigma_3(I)(p2)) - 6\sigma_3(I)(n)) \\
& \cdot \phi'_2(5\sigma_3(I)(f2) + 4(1 - \sigma_3(I)(f2))) \\
& + (5\sigma_3(II)(f2) + 4(1 - \sigma_3(II)(f2)) - 10(1 - \sigma_3(II)(n) - \sigma_3(II)(p2)) - 6\sigma_3(II)(n)) \\
& \cdot \phi'_2(5\sigma_3(II)(f2) + 4(1 - \sigma_3(II)(f2))) \\
& \geq 0.
\end{aligned} \tag{A.3}$$

Suppose, for example, that player 2 is not ambiguity averse enough to satisfy (A.1). If we set $\sigma_3(II)(p2) = \sigma_3(I)(f2) = \sigma_3(II)(f1) = 1$ and $\sigma_3(I)(p2) = 1 - \sigma_3(I)(p1)$, then (A.2) and (A.3) become

$$(4 - 10\sigma_3(I)(p2))\phi'_1(4) - 5\phi'_1(5) \geq 0$$

and

$$(5 - 10(1 - \sigma_3(I)(p2)))\phi'_2(5) + 4\phi'_2(4) \geq 0.$$

If we set $\sigma_3(I)(p2) = 0.1$, then even an ambiguity neutral player 2 will have c as a best response. As long as player 1 has

$$\frac{\phi'_1(4)}{\phi'_1(5)} \geq \frac{5}{3}$$

then this is an ex-ante equilibrium.

Similarly, if player 2 has

$$\frac{\phi'_2(4)}{\phi'_2(5)} \geq \frac{5}{3}$$

then setting $\sigma_3(I)(p1) = \sigma_3(I)(f2) = \sigma_3(II)(f1) = 1$ and $\sigma_3(II)(p2) = 0.9$ gives an ex-ante equilibrium without conditions on 1's ambiguity aversion. ■

Proof of Proposition 3.2. Suppose $x > 1$. Following the play of d by either 1 or 2, the only best response to any updated belief of player 3 is to play n with probability 1 as this yields 3 a payoff of $x > 1 > 0$. Thus $\sigma_3(I, d)(n) = \sigma_3(II, d)(n) = 1$ in any sequential optimum. Given that, both players 1 and 2 expect to get a payoff of 6 if they deviate to d , which is higher than any payoff to playing c (i.e., 4 or 5).

Suppose $x \leq 0.5$ and fix an ex-ante equilibrium σ yielding peace with probability 1. We will construct a sequential optimum $(\hat{\sigma}, \nu)$ yielding peace with probability 1. Set $\hat{\sigma}_3(I, d)(p1) = \sigma_3(I, d)(p1) + \frac{\sigma_3(I, d)(n)}{2}$, $\hat{\sigma}_3(I, d)(p2) = \sigma_3(I, d)(p2) + \frac{\sigma_3(I, d)(n)}{2}$, and $\hat{\sigma}_3(I, d)(n) = 0$, and use the analogous construction to determine $\hat{\sigma}_3(II, d)$. At all other information sets, let $\hat{\sigma} = \sigma$. Observe that $\hat{\sigma}$ yields peace with probability 1. We next verify that $\hat{\sigma}$ is an ex-ante equilibrium. Player 3 is trivially ex-ante best responding, so it suffices to show that 1 and

2 are as well. Since $\frac{10}{2} < 6$, for both parameters I and II , player 1 and 2's expected payoff to deviating towards d is less under $\hat{\sigma}_3$ than under σ_3 . Therefore c remains an ex-ante best response for both players and $\hat{\sigma}$ is an ex-ante equilibrium. It remains to specify beliefs and check optimality for all players at all information sets. Given beliefs satisfying the strong smooth rule using $\hat{\sigma}$, by Theorem 3.2, it is sufficient to check optimality against one-stage deviations, and therefore only at information sets where the player has a non-trivial move. Begin with player 3. Following the play of c, c , 3 is indifferent among any mixture over $f1$ and $f2$ and is thus best responding. Following the breakdown of negotiations, observe that at either of player 3's two non-singleton information sets the strong smooth rule does not restrict 3's beliefs, and any mixture over actions $p1$ and $p2$ is a best response to *some* beliefs. Let 3's respective beliefs at these information sets be such that $\hat{\sigma}_3(I, d)$ and $\hat{\sigma}_3(II, d)$ are, respectively, best responses. Next, specify beliefs for players 1 and 2 that place probability $\frac{1}{2}$ on each of the two degenerate π on the corresponding information set. Since these beliefs maintain the ex-ante μ -weights and there is no change in payoffs compared to the ex-ante evaluations, these beliefs satisfy the strong smooth rule using $\hat{\sigma}$ and playing c is a best response for 1 and 2 at the information sets where they move because it was an ex-ante best response.

Suppose $0.5 < x \leq 1$. Then no non-degenerate mixture over $p1$ and $p2$ can be a best response at either of player 3's non-singleton information sets because it would be dominated by replacing one of them in the mixture by n . If players 1 and 2 are both sufficiently ambiguity averse so that (A.1) is satisfied, the proof of part (ii) of Proposition 3.1 shows that there is an ex-ante equilibrium yielding peace with probability 1 in which player 3 uses only degenerate mixtures of $p1$ and $p2$. The arguments in the $x \leq 0.5$ case applied to this ex-ante equilibrium show that it is also sequentially optimal. ■

Proof of Remark 3.1. Observe that under $0.5 < x \leq 1$, player 3 is limited to mixtures of $p1$ and n or $p2$ and n in order to be best responding. Inequalities (A.2) and (A.3) are necessary for players 1 and 2 to be best responding. We now show that, under (A.3), the maximal value of the left-hand side of (A.2) is negative, contradicting (A.2). Since the payoffs do not depend on the parameter (I or II), we may assume without loss of generality that $\sigma_3(I)(f2) \geq \sigma_3(II)(f2)$. Fixing $\sigma_3(I)(f2)$ and $\sigma_3(II)(f2)$, since the coefficients of $\sigma_3(I)(n)$, $\sigma_3(I)(p2)$, $\sigma_3(II)(n)$ and $\sigma_3(II)(p2)$ are negative in (A.2) and positive in (A.3), the maximal value of the left-hand side of (A.2) is obtained when (A.3) is binding. This equality under ambiguity neutrality of player 2 is equivalent to

$$\sigma_3(II)(p2) = \frac{12 - \sigma_3(I)(f2) - \sigma_3(II)(f2) - 4\sigma_3(I)(n) - 10\sigma_3(I)(p2) - 4\sigma_3(II)(n)}{10}.$$

Substituting using this and $\phi_1(x) = -e^{-x \ln(5/3)}$, the left-hand side of (A.2) becomes (up to multiplication by a positive constant)

$$\begin{aligned} & (5 - \sigma_3(I)(f2))\left(\frac{3}{5}\right)^{5-\sigma_3(I)(f2)} + (-7 + \sigma_3(I)(f2))\left(\frac{3}{5}\right)^{5-\sigma_3(II)(f2)} \\ & - \sigma_3(I)(n)[6\left(\frac{3}{5}\right)^{5-\sigma_3(I)(f2)} - 4\left(\frac{3}{5}\right)^{5-\sigma_3(II)(f2)}] \\ & - \sigma_3(I)(p2)[10\left(\frac{3}{5}\right)^{5-\sigma_3(I)(f2)} - 10\left(\frac{3}{5}\right)^{5-\sigma_3(II)(f2)}] - 2\sigma_3(II)(n)\left(\frac{3}{5}\right)^{5-\sigma_3(II)(f2)}. \end{aligned}$$

To maximize this expression, which is linear in $\sigma_3(I)(n)$, $\sigma_3(I)(p2)$ and $\sigma_3(II)(n)$ with negative coefficients, the three variables must be as low as possible. Therefore $\sigma_3(II)(n) = 0$ and $\sigma_3(II)(p2) = 1$, implying $\sigma_3(I)(n) = \frac{1}{2} - \frac{\sigma_3(I)(f2) + \sigma_3(II)(f2) + 10\sigma_3(I)(p2)}{4}$ and $\sigma_3(I)(p1) = \frac{1}{2} + \frac{\sigma_3(I)(f2) + \sigma_3(II)(f2) + 6\sigma_3(I)(p2)}{4}$. Since $\sigma_3(I)(p2) > 0$ implies $\sigma_3(I)(p1) > 0$, the restriction to mixtures of $p1$ and n or $p2$ and n implies $\sigma_3(I)(p2) = 0$. Simplifying using these values yields

$$\frac{4 + \sigma_3(I)(f2) + 3\sigma_3(II)(f2)}{2} \left(\frac{3}{5}\right)^{5-\sigma_3(I)(f2)} - (5 + \sigma_3(II)(f2)) \left(\frac{3}{5}\right)^{5-\sigma_3(II)(f2)}.$$

This expression is increasing in $\sigma_3(I)(f2)$, and under $\sigma_3(I)(f2) = 1$, is decreasing in $\sigma_3(II)(f2)$, thus it is maximized when $\sigma_3(I)(f2) = 1$ and $\sigma_3(II)(f2) = 0$ at the value of $-\frac{81}{1250}$. Therefore (A.2) and (A.3) cannot be simultaneously satisfied.

The $x \leq 0.5$ statement follows from the proof of Proposition 3.2. ■

Proof of Proposition 3.3. Fix an ex-ante equilibrium σ yielding peace with probability 1. Suppose $x > 0.5$, and that an SEA yielding peace with probability 1 exists. Theorem A.2 implies that formula (A.12) is necessary for smooth rule consistency. For player 3, since any σ yielding peace with probability 1 has $\bar{p}_{-i, \sigma_{-i}}(I, d|I) = \bar{p}_{-i, \sigma_{-i}}(II, d|II)$ and $\bar{p}_{-i, \sigma_{-i}}(I, c, d|I) = \bar{p}_{-i, \sigma_{-i}}(II, c, d|II)$ irrespective of the sequence σ^k of completely mixed strategy profiles chosen to converge to σ , formula (A.12) implies that beliefs are the same at both information sets where 3 has a non-trivial move. Furthermore, $x > 0.5$ implies that given such beliefs, either all of 3's best responses never involve $p1$ or all of them never involve $p2$. In the former case, player 1 will play d with probability 1, and in the latter case, player 2 will play d with probability 1. This contradicts the existence of such an SEA.

Suppose $x \leq 0.5$. We will construct an SEA $(\hat{\sigma}, \nu)$ yielding peace with probability 1. Construct $\hat{\sigma}$ as in the part of the proof of Proposition 3.2 that assumes $x \leq 0.5$. By the argument there, $\hat{\sigma}$ is an ex-ante equilibrium. It remains to specify an interim belief system satisfying smooth rule consistency and check optimality for all players and information sets. Consider any sequence $\hat{\sigma}^k$ of completely mixed strategies converging to $\hat{\sigma}$ such that, for

players 1 and 2, $\hat{\sigma}_1^k(c) = 1 - \frac{1}{k+1}$, $\hat{\sigma}_1^k(d) = \frac{1}{k+1}$ and $\hat{\sigma}_2^k(c) = 1 - \frac{1}{k+1}$, $\hat{\sigma}_2^k(d) = \frac{1}{k+1}$. By Lemma A.4 and Theorem A.2, formula (A.12) identifies an interim belief system $\hat{\nu}$ satisfying smooth rule consistency. By Theorem 3.3, for optimality, it is sufficient to check against one-stage deviations, and therefore only at information sets where the player has a non-trivial move. Begin with player 3. Following the play of c, c , 3 is indifferent among any mixture over $f1$ and $f2$ and is thus best responding. Off the equilibrium path, observe that at either of player 3's two information sets where 3 has a non-trivial move, since $x \leq 0.5$, any mixture over actions $p1$ and $p2$ is a best response to beliefs placing all weight on a half-half measure over the two elements of the information set. At either of these information sets, since $\bar{p}_{-i, \sigma_{-i}}(h^t | h^0) = \lim_{k \rightarrow \infty} \frac{\frac{1}{k+1}}{\frac{1}{k+1} + \frac{k}{(k+1)^2}} = \frac{1}{2}$, $\hat{\nu}_{3, I_i}$ is degenerate on a half-half measure over the two elements of the information set. Therefore $\hat{\sigma}_3(I, d)$ and $\hat{\sigma}_3(II, d)$ are, respectively, best responses given $\hat{\nu}_{3, I_i}$ at those information sets. Next turn to players 1 and 2. Since no uncertainty resolves for either player before they make their respective non-trivial move, their beliefs $\hat{\nu}_{i, I_i}$ place probability $\frac{1}{2}$ on each of the two degenerate π on the corresponding information set, where the $\frac{1}{2}$ is inherited from μ . Therefore playing c is a best response given $\hat{\nu}_{i, I_i}$ at those information sets for 1 and 2 because it was an ex-ante best response. ■

A.2 Proofs of results in Section 3.2

We next state and prove a key lemma on the preservation of optimality under smooth rule updating:

Definition A.1 (Reachability) *Player i views information set $I_i \not\subseteq \Theta$ as reachable from information set I_i^{-1} given σ and ν if $\sum_{h^{s(I_i)} \in I_i} \sum_{\hat{\pi} \in \Delta(I_i^{-1})} p_\sigma(h^{s(I_i)} | h^{s(I_i)-1}) \hat{\pi}(h^{s(I_i)-1}) \nu_{i, I_i^{-1}}(\hat{\pi}) > 0$.*

Lemma A.1 *Fix a game Γ , a (σ, ν) such that σ is an ex-ante equilibrium, a player i and an information set I_i such that either $I_i \subseteq \Theta$ or i views I_i as reachable from I_i^{-1} given σ and ν . If ν_{i, I_i} is derived from $\nu_{i, I_i^{-1}}$ (or, if $I_i \subseteq \Theta$, from μ_i) via the smooth rule using σ and, for all $\sigma'_i \in \Sigma_i$,*

$$V_{i, I_i^{-1}}(\sigma) \geq V_{i, I_i^{-1}}(\sigma'_i, \sigma_{-i}),$$

(or, if $I_i \subseteq \Theta$, given ex-ante optimality), then, for all $\sigma'_i \in \Sigma_i$,

$$V_{i, I_i}(\sigma) \geq V_{i, I_i}(\sigma'_i, \sigma_{-i}).$$

Proof of Lemma A.1. The inequalities $V_{i, I_i}(\sigma) \geq V_{i, I_i}(\sigma'_i, \sigma_{-i})$ (respectively, $V_i(\sigma) \geq V_i(\sigma'_i, \sigma_{-i})$) for all σ'_i are equivalent to the condition that $\sigma'_i = \sigma_i$ maximizes (2.5) (respectively, $\sigma'_i = \sigma_i$ maximizes (2.2)).

We want to show that $V_{i,I_i}(\sigma) \geq V_{i,I_i}(\sigma'_i, \sigma_{-i})$ for all σ'_i . By the above, it is sufficient to show that $\sigma'_i = \sigma_i$ maximizes (2.5).

Let $t = s(I_i)$. Consider the case where $t > 0$ (the case where $t = 0$ is similar, using (2.2) instead of (A.4), and is omitted). By assumption in the statement of the lemma, $V_{i,I_i^{-1}}(\sigma) \geq V_{i,I_i^{-1}}(\sigma'_i, \sigma_{-i})$ for all $\sigma'_i \in \Sigma_i$. As in (2.5), this is equivalent to the condition that $\sigma'_i = \sigma_i$ maximizes

$$\sum_{h|h^{t-1} \in I_i^{-1}} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i^{-1}}(h). \quad (\text{A.4})$$

Notice that, since i 's strategy is a function only of i 's information sets and, by perfect recall, $R_i(h^t) = R_i(I_i)$ for any h such that $h^t \in I_i$, $p_{i,\sigma'_i}(h^t|h^{t-1})$ is the same for any such h . Thus, the objective function in (A.4) can be equivalently written as

$$\begin{aligned} & \sum_{\substack{h|h^{t-1} \in I_i^{-1} \\ \text{and } h^t \notin I_i}} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i^{-1}}(h) \\ & + p_{i,\sigma'_i}(\bar{h}^t|\bar{h}^{t-1}) \sum_{h|h^t \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^t) q^{(\sigma,\nu),i,I_i^{-1}}(h) \end{aligned}$$

for any \bar{h} such that $\bar{h}^t \in I_i$. The advantage of doing so is making clear that only the term

$$\sum_{h|h^t \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^t) q^{(\sigma,\nu),i,I_i^{-1}}(h) \quad (\text{A.5})$$

is affected by the specification of σ'_i from I_i onward and no other part of σ'_i affects (A.5). Therefore, since reachability implies $p_{i,\sigma'_i}(\bar{h}^t|\bar{h}^{t-1}) > 0$, σ_i maximizes (A.4) implies that σ_i maximizes (A.5). For that to imply σ_i maximizes (2.5), it is sufficient to show that $q^{(\sigma,\nu),i,I_i}(h) \propto q^{(\sigma,\nu),i,I_i^{-1}}(h)$ holds for $\{h \mid h^t \in I_i\}$. This proportionality may be shown by using the local measure definition (2.6), applying the smooth rule to substitute for $\nu_{i,I_i}(\pi)$ for all $\pi \in \Delta(I_i)$ and then using the expression for $\hat{\pi}_{I_i}$ and cancelling terms. ■

Theorem A.1 *Fix a game Γ . Suppose σ is an ex-ante equilibrium and, for each player i and each information set $I_i \not\subseteq \Theta$, $\sum_{h^{s(I_i)} \in I_i} p_\sigma(h^{s(I_i)}|h^{s(I_i)-1}) > 0$. Then, σ is sequentially optimal.*

Proof of Theorem A.1. By ex-ante optimality of σ , (3.1) in the definition of sequential optimality is satisfied. Since $\sum_{h^{s(I_i)} \in I_i} p_\sigma(h^{s(I_i)}|h^{s(I_i)-1}) > 0$ for all $I_i \not\subseteq \Theta$, recursive substitution in the smooth rule formula using σ starting from $I_i \subseteq \Theta$ implies that i views I_i as reachable from I_i^{-1} given σ and any ν satisfying the smooth rule using σ . Therefore, by Lemma A.1, (3.2) in the definition of sequential optimality is satisfied for all i and I_i . ■

Proof of Theorem 3.1. The if direction follows by definition. We show the only if direction. Suppose (σ, ν) is sequentially optimal. We show that $(\sigma, \hat{\nu})$, where, for all i , I_i , if $I_i \subseteq \Theta$ or if i views I_i as reachable from I_i^{-1} given σ and ν , $\hat{\nu}_{i,I_i}$ is derived via the smooth rule, and $\hat{\nu}_{i,I_i} = \nu_{i,I_i}$ everywhere else, is sequentially optimal. By construction, $\hat{\nu}$ satisfies the smooth rule using σ except, possibly, for I_i not viewed as reachable from I_i^{-1} given σ and ν . However, from the definition of the smooth rule (Definition 3.3), observe that it is exactly in such cases where the smooth rule allows any updated beliefs. Thus $\hat{\nu}$ satisfies the smooth rule using σ . Since $\hat{\nu}$ does not enter into the ex-ante function V_i , the fact that (σ, ν) is sequentially optimal directly implies that $V_i(\sigma) \geq V_i(\sigma'_i, \sigma_{-i})$ for all $\sigma'_i \in \Sigma_i$. To see that $(\sigma, \hat{\nu})$ satisfies the optimality conditions $V_{i,I_i}(\sigma) \geq V_{i,I_i}(\sigma'_i, \sigma_{-i})$ for all $\sigma'_i \in \Sigma_i$, observe that (a) where $\hat{\nu}_{i,I_i} = \nu_{i,I_i}$, it directly inherits this from (σ, ν) and (b) everywhere else, Lemma A.1 shows that smooth rule updating ensures the required optimality. ■

Proof of Theorem 3.2. The only if direction follows by definition. For the if direction, suppose ν is an interim belief system satisfying the strong smooth rule using σ such that (σ, ν) has no profitable one-stage deviations. First, for each player i , observe that having no profitable one-stage deviations implies optimality of σ_i according to V_{i,I_i} for all $I_i \in \mathcal{I}_i^T$. Proceed by induction on the stage t . Fix any t such that $0 < t \leq T$, and suppose that, for each player i , σ_i is optimal according to V_{i,I_i} for all $I_i \in \mathcal{I}_i^t$. We claim that, for each player i , σ_i is optimal according to V_{i,I_i} for all $I_i \in \mathcal{I}_i^{t-1}$. The argument for this is as follows. Fix a player i and $I_i \in \mathcal{I}_i^{t-1}$. Consider any strategy σ'_i for player i . For any $J_i \in \mathcal{I}_i^t$, the optimality of σ_i according to V_{i,J_i} implies (see (2.5))

$$\begin{aligned} & \sum_{h|h^t \in J_i} u_i(h) p_{i,\sigma_i}(h|h^t) q^{(\sigma,\nu),i,J_i}(h) \\ & \geq \sum_{h|h^t \in J_i} u_i(h) p_{i,\sigma'_i}(h|h^t) q^{(\sigma,\nu),i,J_i}(h). \end{aligned} \tag{A.6}$$

Since ν satisfies strong smooth rule updating using σ , for all such J_i for which $I_i = J_i^{-1}$ and $\sum_{h^t \in J_i} \sum_{\hat{\pi} \in \Delta(I_i)} p_{-i,\sigma_{-i}}(h^t|h^{t-1}) \hat{\pi}(h^{t-1}) \nu_{i,I_i}(\hat{\pi}) > 0$, $q^{(\sigma,\nu),i,J_i}(h) \propto q^{(\sigma,\nu),i,I_i}(h)$ holds for $\{h \mid h^t \in J_i\}$. This proportionality follows from using the local measure definition (2.6), applying the strong smooth rule iteratively to substitute for ν_{i,I_i} and simplifying. After substituting in (A.6) for $q^{(\sigma,\nu),i,J_i}$, cancelling the constant of proportionality and multiplying by $p_{i,\sigma'_i}(h^t|h^{t-1})$, which is constant for any h such that $h^t \in J_i$ because i 's strategy is a function only of i 's information sets and, due to perfect recall, $R_i(h^t) = R_i(J_i)$ for any h such that $h^t \in J_i$, (A.6)

becomes

$$\begin{aligned}
& \sum_{h|h^t \in J_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h) \\
& \geq \sum_{h|h^t \in J_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma'_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h) \\
& = \sum_{h|h^t \in J_i} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h).
\end{aligned} \tag{A.7}$$

If $I_i = J_i^{-1}$ but $\sum_{h^t \in J_i} \sum_{\hat{\pi} \in \Delta(I_i)} p_{-i,\sigma_{-i}}(h^t|h^{t-1}) \hat{\pi}(h^{t-1}) \nu_{i,I_i}(\hat{\pi}) = 0$, then $q^{(\sigma,\nu),i,I_i}(h) = 0$ for all h with $h^t \in J_i$. Thus, summing (A.7) for all $J_i \in \mathcal{I}_i^t$ for which $I_i = J_i^{-1}$ and $\sum_{h^t \in J_i} \sum_{\hat{\pi} \in \Delta(I_i)} p_{-i,\sigma_{-i}}(h^t|h^{t-1}) \hat{\pi}(h^{t-1}) \nu_{i,I_i}(\hat{\pi}) > 0$ is the same as summing for all $J_i \in \mathcal{I}_i^t$ such that $J_i^{-1} = I_i$, yielding:

$$\begin{aligned}
& \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h) \\
& \geq \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h).
\end{aligned} \tag{A.8}$$

The absence of profitable one-stage deviations implies σ_i is optimal according to V_{i,I_i} among all strategies deviating only at I_i . By (A.6) applied to I_i and restricted to such deviations,

$$\begin{aligned}
& \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h) \\
& \geq \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h).
\end{aligned} \tag{A.9}$$

Combining (A.9) and (A.8) implies

$$\begin{aligned}
& \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h) \\
& \geq \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h).
\end{aligned} \tag{A.10}$$

Since (A.10) holds for any σ'_i , it is the same as (A.6) with $t-1$ in the role of t and I_i in the role of J_i . Therefore σ_i is optimal according to V_{i,I_i} . Since this conclusion holds for any $I_i \in \mathcal{I}_i^{t-1}$, the induction step is completed. It follows that (σ, ν) satisfies the optimality conditions (3.2) in the definition of sequentially optimal.

It remains to show that σ also satisfies the ex-ante optimality conditions (3.1). Since ν satisfies smooth rule updating using σ , for all $I_i \subseteq \Theta$, $q^{(\sigma, \nu), i, I_i}(h) \propto q^{(\sigma, \mu_i), i}(h)$ holds for $\{h \mid h^0 \in I_i\}$. Using this to substitute for $q^{(\sigma, \nu), i, I_i}$ in (A.10) with $t = 1$, cancelling the constant of proportionality and summing for all I_i , yields:

$$\begin{aligned} & \sum_h u_i(h) p_{i, \sigma_i}(h | h^0) q^{(\sigma, \mu_i), i}(h) \\ & \geq \sum_h u_i(h) p_{i, \sigma'_i}(h | h^0) q^{(\sigma, \mu_i), i}(h). \end{aligned} \tag{A.11}$$

Since (A.11) holds for any σ'_i , σ maximizes (2.2) which is equivalent to the ex-ante optimality condition (3.1). ■

Proof of Corollary 3.1. Assume ambiguity neutrality. From Theorem 3.1 and the fact that under ambiguity neutrality the smooth rule specializes to Bayes' rule, σ is sequentially optimal if and only if there exists an interim belief system $\hat{\nu}$ satisfying Bayes' rule such that $(\sigma, \hat{\nu})$ is sequentially optimal. Sequential optimality is sequential rationality plus ex-ante equilibrium, immediately implying the only if direction of the corollary. To show the if direction, repeat the argument in the last paragraph of the proof of Theorem 3.2 showing that, for each i , optimality of σ_i at all $I_i \subseteq \Theta$ using smooth rule updated beliefs implies ex-ante optimality of σ_i . ■

A.3 Proofs of results in Section 3.3

Lemma A.2 *If (σ, ν) satisfies smooth rule consistency, then ν satisfies the strong smooth rule using σ .*

Proof of Lemma A.2. By definition of smooth rule consistency, there exists a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$, with $\lim_{k \rightarrow \infty} \sigma^k = \sigma$, such that $\nu = \lim_{k \rightarrow \infty} \nu^k$, where each ν^k is the interim belief system satisfying the smooth rule using σ^k . Since $\sum_{h^s(I_i) \in I_i} p_{\sigma^k}(h^{s(I_i)} | h^{s(I_i)-1}) > 0$ for all $I_i \not\subseteq \Theta$, recursive substitution in the smooth rule formula using σ^k starting from $I_i \subseteq \Theta$ implies that, for any k , i views I_i as reachable from I_i^{-1} given σ^k and ν^k and so the formulas (3.3) and (3.4) hold for all I_i . For any player i and I_i , if either $I_i \subseteq \Theta$ or $\sum_{h^t \in I_i} \sum_{\hat{\pi} \in \Delta(I_i^{-1})} p_{-i, \sigma_{-i}}(h^t | h^{t-1}) \hat{\pi}(h^{t-1}) \nu_{i, I_i^{-1}}(\hat{\pi}) > 0$, then, by continuity in the strategy profile of the formulas (3.3) and (3.4), ν satisfies the strong smooth rule using σ at such information sets. Finally, notice that at all remaining information sets, ν trivially satisfies the strong smooth rule given σ since this rule does not restrict beliefs there. ■

Proof of Theorem 3.3. One direction is by definition. For the other direction it is enough to show that if (σ, ν) has no profitable one-stage deviations and satisfies smooth rule consistency then (σ, ν) is sequentially optimal. This follows directly from Lemma A.2 and Theorem 3.2. ■

Proof of Theorem 3.4. Fix a sequence $\varepsilon^k = (\varepsilon_I^k)_{I \in \cup_{i \in N} \mathcal{I}_i}$ of strictly positive vectors of dimension $|\cup_{i \in N} \mathcal{I}_i|$, converging in the sup-norm to 0 and such that $\varepsilon_{I_i}^k \leq \frac{1}{|A_i(I_i)|}$ for all players i and information sets I_i . For any k , let Γ^k be the restriction of the game Γ defined such that the set of feasible strategy profiles is the set of all completely mixed σ^k satisfying $\sigma_i^k(I_i)(a_i) \geq \varepsilon_{I_i}^k$ for all i , I_i and actions $a_i \in A_i(I_i)$. Consider the agent normal form G^k of the game Γ^k (see e.g., Myerson, 1991, p.61). Since the payoff functions are concave and the set of strategies of each player in G^k is non-empty, compact and convex, G^k has an ex-ante equilibrium by Glicksberg (1952). Let $\hat{\sigma}^k$ be the strategy profile in the game Γ^k corresponding to this equilibrium. Then $\hat{\sigma}^k$ is an ex-ante equilibrium of Γ^k . By Theorem A.1, since all information sets are on the equilibrium path, there exists an interim belief system ν^k such that $(\hat{\sigma}^k, \nu^k)$ is sequentially optimal. By Theorem 3.1, there exists an interim belief system $\hat{\nu}^k$ satisfying the smooth rule using $\hat{\sigma}^k$ such that $(\hat{\sigma}^k, \hat{\nu}^k)$ is a sequential optimum of Γ^k . By compactness of the set of strategy profiles, the sequence $\hat{\sigma}^k$ has a convergent sub-sequence, the limit of which is denoted by $\hat{\sigma}$. By continuity in the strategy profile of the smooth rule formula and compactness of the set of interim belief systems, an associated sub-sequence of $\hat{\nu}^k$ converges to a limit interim belief system which we denote by $\hat{\nu}$. By continuity of the payoff functions, $\hat{\sigma}$ is an ex-ante equilibrium of Γ . Given any information set I_i and continuation strategy $\tilde{\sigma}_i^{I_i}$ of player i in Γ , let $\tilde{\sigma}_i^{k, I_i}$ be a feasible strategy in Γ^k for this player that is closest (in the sup-norm) to $\tilde{\sigma}_i^{I_i}$. Since, by sequential optimality of $(\hat{\sigma}^k, \hat{\nu}^k)$ for each k , $\hat{\sigma}_i^{k, I_i}$ is weakly better than $\tilde{\sigma}_i^{k, I_i}$ for player i given belief $\hat{\nu}_{i, I_i}^k$, and since, along the sub-sequence, $\tilde{\sigma}_i^{k, I_i}$ converges to $\tilde{\sigma}_i^{I_i}$ and $\hat{\nu}_{i, I_i}^k$ converges to $\hat{\nu}_{i, I_i}$, continuity of the payoff functions implies that $\hat{\sigma}_i^{I_i}$ is weakly better than $\tilde{\sigma}_i^{I_i}$ for this player given belief $\hat{\nu}_{i, I_i}$. Therefore $(\hat{\sigma}, \hat{\nu})$ satisfies sequential optimality. Finally, observe that $(\hat{\sigma}, \hat{\nu})$ satisfies smooth rule consistency (since it is explicitly constructed as the limit of an appropriate sequence). Therefore $(\hat{\sigma}, \hat{\nu})$ is an SEA of Γ . ■

Our next result provides an explicit formula for interim belief systems satisfying smooth rule consistency. This smooth rule-like formula, which will be generally useful when working with SEA, uses a limit of likelihoods of the partial histories in an information set given θ and σ_{-i}^k . Before stating the result, we need some notation and a lemma:

Notation A.1 Let $\bar{p}_{-i, \sigma_{-i}}(h^t | h^0)$ denote $\lim_{k \rightarrow \infty} \frac{p_{-i, \sigma_{-i}^k}(h^t | h^0)}{\sum_{\hat{h}^t \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^t | \hat{h}^0)}$, where $t = s(I_i)$. For each

information set I_i , consider the smallest $0 \leq r \leq t$ such that there exists $h \in H$ and \hat{I}_i for which $h^t \in I_i$, $h^r \in \hat{I}_i$ and $p_\sigma(h^t|h^r) > 0$. This generates an \hat{I}_i for each I_i . Let $\mathcal{F}_{i,\sigma}$ denote the set of all these \hat{I}_i that are non-singleton and for which $\hat{I}_i \not\subseteq \Theta$.

The next lemma shows that existence of $\bar{p}_{-i,\sigma_{-i}}(h^t|h^0)$ need be checked only at information sets in $\mathcal{F}_{i,\sigma}$.

Lemma A.3 Fix a game Γ , a strategy profile σ and a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$ such that $\lim_{k \rightarrow \infty} \sigma^k = \sigma$. If $\bar{p}_{-i,\sigma_{-i}}(h^t|h^0)$ exists for each player i and each $h^t \in I_i \in \mathcal{F}_{i,\sigma}$, then $\bar{p}_{-i,\sigma_{-i}}(h^t|h^0)$ exists for each i and $h^t \in I_i \in \mathcal{I}_i$.

Proof of Lemma A.3. If I_i is a singleton, $\bar{p}_{-i,\sigma_{-i}}(h^t|h^0) = 1$. If $I_i \not\subseteq \Theta$, consider the $\hat{I}_i \in \mathcal{F}_{i,\sigma}$ corresponding to I_i , and if $I_i \subseteq \Theta$, set $\hat{I}_i = I_i$. Let $r = s(\hat{I}_i)$. Observe that for all

$$h^t \in I_i, \bar{p}_{-i,\sigma_{-i}}(h^t|h^0) = \lim_{k \rightarrow \infty} \frac{\frac{p_{-i,\sigma_{-i}}^k(h^t|h^r)}{\sum_{\hat{h}^r \in \hat{I}_i} p_{-i,\sigma_{-i}}^k(\hat{h}^r|\hat{h}^0)}}{\sum_{\hat{h}^t \in I_i} p_{-i,\sigma_{-i}}^k(\hat{h}^t|\hat{h}^r)} \frac{p_{-i,\sigma_{-i}}^k(\hat{h}^r|\hat{h}^0)}{\sum_{\hat{h}^r \in \hat{I}_i} p_{-i,\sigma_{-i}}^k(\hat{h}^r|\hat{h}^0)} = \frac{p_{-i,\sigma_{-i}}(h^t|h^r) \bar{p}_{-i,\sigma_{-i}}(h^r|h^0)}{\sum_{\hat{h}^t \in I_i} p_{-i,\sigma_{-i}}(\hat{h}^t|\hat{h}^r) \bar{p}_{-i,\sigma_{-i}}(\hat{h}^r|\hat{h}^0)},$$

where the last equality follows since $\sum_{\hat{h}^t \in I_i} p_{-i,\sigma_{-i}}(\hat{h}^t|\hat{h}^r) > 0$ and $\bar{p}_{-i,\sigma_{-i}}(h^r|h^0)$ exists for all $h^r \in \hat{I}_i$ (either by assumption since $\hat{I}_i \in \mathcal{F}_{i,\sigma}$, or, if $\hat{I}_i \subseteq \Theta$, because it is constant in k). Thus $\bar{p}_{-i,\sigma_{-i}}(h^t|h^0)$ exists. ■

Theorem A.2 Fix a game Γ , a strategy profile σ and a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$ such that $\lim_{k \rightarrow \infty} \sigma^k = \sigma$. Then (σ, ν) satisfies smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ if and only if, for $\{\sigma^k\}_{k=1}^\infty$, $\bar{p}_{-i,\sigma_{-i}}(h^t|h^0)$ exists for each player i and each $h^t \in I_i \in \mathcal{F}_{i,\sigma}$, and ν satisfies the formula

$$\begin{aligned} \nu_{i,I_i}(\pi) \propto & \sum_{\hat{\pi} \in \Delta(\Theta) | \bar{\pi}_{I_i} = \pi} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \hat{\pi}(h^0) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi(h^t) \right)} \\ & \cdot \left(\sum_{h^t \in I_i} \bar{p}_{-i,\sigma_{-i}}(h^t|h^0) \hat{\pi}(h^0) \right) \mu_i(\hat{\pi}) \end{aligned} \quad (\text{A.12})$$

for each i and I_i , where $\bar{\pi}_{I_i} \in \Delta(I_i)$ is given by $\bar{\pi}_{I_i}(h^t) = \frac{\bar{p}_{-i,\sigma_{-i}}(h^t|h^0) \hat{\pi}(h^0)}{\sum_{\hat{h}^t \in I_i} \bar{p}_{-i,\sigma_{-i}}(\hat{h}^t|\hat{h}^0) \hat{\pi}(\hat{h}^0)}$ for all $h^t \in I_i$.

Proof of Theorem A.2. We first establish that the formula and the assumed existence of $\bar{p}_{-i, \sigma_{-i}}(h^t|h^0)$ imply that (σ, ν) satisfies smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$. Fix i, I_i . Since each σ^k is completely mixed, the smooth rule using $\{\sigma^k\}_{k=1}^\infty$ has bite at each information set and so, applying the formulas in the smooth rule iteratively starting from μ_i , for each k , the belief ν_{i, I_i}^k determined by smooth rule updating using $\{\sigma^k\}_{k=1}^\infty$ satisfies: for $\pi \in \Delta(I_i)$,

$$\nu_{i, I_i}^k(\pi) \propto \sum_{\hat{\pi} \in \Delta(\Theta) | \hat{\pi}_{I_i} = \pi} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma^k}(h|h^0) \hat{\pi}(h^0) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_{\sigma^k}(h|h^t) \pi(h^t) \right)} \cdot \left(\sum_{h^t \in I_i} p_{-i, \sigma_{-i}^k}(h^t|h^0) \hat{\pi}(h^0) \right) \mu_i(\hat{\pi}), \quad (\text{A.13})$$

where

$$\hat{\pi}_{I_i}(h^t) = \frac{p_{-i, \sigma_{-i}^k}(h^t|h^0) \hat{\pi}(h^0)}{\sum_{\hat{h}^t \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^t|\hat{h}^0) \hat{\pi}(\hat{h}^0)} \text{ for all } h^t \in I_i.$$

It remains to show that the limit of ν_{i, I_i}^k equals ν_{i, I_i} (as defined in (A.12)). Divide (A.13) by $\sum_{\hat{h}^t \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^t|\hat{h}^0) > 0$, which is constant with respect to π . Then the limit of $\nu_{i, I_i}^k(\pi)$ is proportional to:

$$\lim_{k \rightarrow \infty} \sum_{\hat{\pi} \in \Delta(\Theta) | \hat{\pi}_{I_i} = \pi} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma^k}(h|h^0) \hat{\pi}(h^0) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_{\sigma^k}(h|h^t) \pi(h^t) \right)} \cdot \left(\sum_{h^t \in I_i} \frac{p_{-i, \sigma_{-i}^k}(h^t|h^0)}{\sum_{\hat{h}^t \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^t|\hat{h}^0)} \hat{\pi}(h^0) \right) \mu_i(\hat{\pi}). \quad (\text{A.14})$$

By Lemma A.3, $\bar{p}_{-i, \sigma_{-i}}(h^t|h^0)$ exists. Then (A.14) is proportional (in $\pi \in \Delta(I_i)$) to the right-

$$\text{hand side of (A.12) since, whenever it exists, } \lim_{k \rightarrow \infty} \hat{\pi}_{I_i}(h^t) = \lim_{k \rightarrow \infty} \frac{\frac{p_{-i, \sigma_{-i}^k}(h^t|h^0) \hat{\pi}(h^0)}{\sum_{\hat{h}^t \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^t|\hat{h}^0)}}{\sum_{\hat{h}^t \in I_i} \frac{p_{-i, \sigma_{-i}^k}(\hat{h}^t|\hat{h}^0) \hat{\pi}(\hat{h}^0)}{\sum_{\hat{h}^t \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^t|\hat{h}^0)}} =$$

$\frac{\bar{p}_{-i,\sigma_{-i}}(h^t|h^0)\hat{\pi}(h^0)}{\sum_{\hat{h}^t \in I_i} \bar{p}_{-i,\sigma_{-i}}(\hat{h}^t|\hat{h}^0)\hat{\pi}(\hat{h}^0)} = \bar{\pi}_{I_i}(h^t)$. Therefore (σ, ν) satisfies smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$.

Next assume that (σ, ν) satisfies smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$. Therefore $\nu = \lim_{k \rightarrow \infty} \nu^k$, where each ν^k is the interim belief system determined by smooth rule updating using σ^k (see (A.13) for the formula). Fix any i, I_i . The existence of $\lim_{k \rightarrow \infty} \nu^k$, as described in (A.14), requires the existence of $\lim_{k \rightarrow \infty} \frac{p_{-i,\sigma_{-i}^k}(h^{s(I_i)}|h^0)}{\sum_{\hat{h}^{s(I_i)} \in I_i} p_{-i,\sigma_{-i}^k}(\hat{h}^{s(I_i)}|\hat{h}^0)}$ for each $h^{s(I_i)} \in I_i$.

Therefore, using $\{\sigma^k\}_{k=1}^\infty$, $\bar{p}_{-i,\sigma_{-i}}(h^{s(I_i)}|h^0)$ exists for all information sets. Then (A.12) yields a well defined interim belief system $\hat{\nu}$, and, by the argument in the earlier direction of this proof, $(\sigma, \hat{\nu})$ satisfies smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$. Thus $\hat{\nu} = \lim_{k \rightarrow \infty} \nu^k = \nu$. Therefore ν satisfies (A.12) using $\{\sigma^k\}_{k=1}^\infty$. ■

Theorem A.3 Fix a game Γ . Suppose (σ, ν) is sequentially optimal and for each player i , and each information set $I_i \not\subseteq \Theta$, $\sum_{h^{s(I_i)} \in I_i} p_\sigma(h^{s(I_i)}|h^0) > 0$. Then there exists $\hat{\nu}$ satisfying (A.12) such that $(\sigma, \hat{\nu})$ is an SEA.

Proof of Theorem A.3. Consider a sequence of a completely mixed strategy profiles converging to σ . Since the limits $\bar{p}_{-i,\sigma_{-i}}(h^{s(I_i)}|h^0)$ in Theorem A.2 exist if $\sum_{h^{s(I_i)} \in I_i} p_\sigma(h^{s(I_i)}|h^0) > 0$ for all $I_i \not\subseteq \Theta$, there exists $\hat{\nu}$ satisfying (A.12). Theorem A.2 then implies that $(\sigma, \hat{\nu})$ satisfies smooth rule consistency, and is thus an SEA. ■

Finally, the next lemma shows that for many common specifications of completely mixed sequences $\{\sigma^k\}_{k=1}^\infty$, $\bar{p}_{-i,\sigma_{-i}^k}(h^t|h^0)$ exists everywhere and has a simple formula.

Lemma A.4 Suppose σ^k is such that, for each i and I_i , for all $a_i \in A_i(I_i)$, $\sigma_i^k(I_i)(a_i) \propto c_{a_i} k^{d_{a_i}}$ with $c_{a_i} > 0$ and $d_{a_i} \geq 0$. Then $\bar{p}_{-i,\sigma_{-i}^k}(h^t|h^0)$ exists for all i and I_i .

Proof of Lemma A.4. Under the assumption on σ^k , for each i, I_i and $h^t \in I_i$,

$$\bar{p}_{-i,\sigma_{-i}^k}(h^t|h^0) = \lim_{k \rightarrow \infty} \frac{\prod_{j \neq i} \prod_{0 \leq s < t} \sigma_j^k(I_j(h^s))(h_{s,j})}{\sum_{\hat{h}^t \in I_i} \prod_{j \neq i} \prod_{0 \leq s < t} \sigma_j^k(I_j(\hat{h}^s))(\hat{h}_{s,j})} = \lim_{k \rightarrow \infty} \frac{C_{h^t} k^{D_{h^t}}}{\sum_{\hat{h}^t \in I_i} C_{\hat{h}^t} k^{D_{\hat{h}^t}}}$$

for $C_{h^t} = \prod_{j \neq i} \prod_{0 \leq s < t} c_{h_{s,j}}$ and $D_{h^t} = \sum_{j \neq i} \sum_{0 \leq s < t} d_{h_{s,j}}$. This limit exists and is equal to

$$\frac{\sum_{\hat{h}^t \in \{h^t\} | D_{\hat{h}^t} = \max_{\bar{h}^t \in I_i} D_{\bar{h}^t}} \sum_{\hat{h}^t \in I_i | D_{\hat{h}^t} = \max_{\bar{h}^t \in I_i} D_{\bar{h}^t}} C_{\hat{h}^t}.$$

■

A.4 Proofs of results in Section 4.1

This subsection contains an example and formal results on comparative statics in ambiguity aversion.

A.4.1 Example: New Strategic Behavior in Equilibrium

We present a 3-player game, with incomplete information about player 1, in which a path of play can occur as part of an SEA when players 2 and 3 are sufficiently ambiguity averse, but never occurs as part of even an ex-ante equilibrium if we modify the game by making players 2 and 3 ambiguity neutral (expected utility). Furthermore, under the SEA we construct, player 1 achieves a higher expected payoff than under any ex-ante equilibrium of the game with ambiguity neutral players, and even outside the convex hull of such ex-ante equilibrium payoffs. The game is depicted in Figure A.1.

There are three players: 1, 2 and 3. First, it is determined whether player 1 is of type I or type II and 1 observes her own type. Players 2 and 3 have only one type, so there is complete information about them. The payoff triples in Figure A.1 describe vNM utility payoffs given players' actions and player 1's type (i.e., (u_1, u_2, u_3) means that player i receives u_i). Players 2 and 3 have ambiguity about player 1's type and have smooth ambiguity preferences with an associated $\phi_2 = \phi_3 = \phi$ and $\mu_2 = \mu_3 = \mu$. Player 1 also has smooth ambiguity preferences, but nothing in what follows depends on either ϕ_1 or μ_1 . Player 1's first and only move in the game is to choose between action $P(\text{lay})$ which leads to players 2 and 3 playing a simultaneous move game in which their payoffs depend on 1's type, and action $Q(\text{uit})$, which ends the game.²³

²³Note that to eliminate any possible effects of varying players' risk aversion, think of the payoffs being generated using lotteries over two "physical" outcomes, the better of which has utility u normalized to 5/2 and the worse of which has u normalized to 0. So, for example, the payoff 1 can be thought of as generated by the lottery giving the better outcome with probability 2/5 and the worse outcome with probability 3/5.




Figure A.1: New equilibrium behavior with ambiguity aversion

Proposition A.1 *Suppose players 2 and 3 are ambiguity neutral and have a common belief μ . There is no ex-ante equilibrium such that player 1 plays P with positive probability.*

Proof of Proposition A.1. Observe that player 1 is willing ex-ante to play P with positive probability if and only if, after the play of P , (U, R) will be played with probability at least $\frac{1}{2}$. Suppose there is an ex-ante equilibrium, σ , in which P is played with positive probability. Let p_I and p_{II} denote the probabilities according to σ that types I and II , respectively, of player 1 play P . Then player 2 finds it optimal to play U with positive probability if and only if

$$p_I \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(\pi(I)) + p_{II} \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(1 - \pi(I)) \geq \frac{5}{2} p_{II} \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(1 - \pi(I))$$

which is equivalent to

$$p_I \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(\pi(I)) \geq \frac{3}{2} p_{II} \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(1 - \pi(I)). \quad (\text{A.15})$$

Similarly, player 3 finds it optimal to play R with positive probability if and only if

$$p_I \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(\pi(I)) + p_{II} \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(1 - \pi(I)) \geq \frac{5}{2} p_I \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(\pi(I))$$

which is equivalent to

$$p_I \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(\pi(I)) \leq \frac{2}{3} p_{II} \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(1 - \pi(I)). \quad (\text{A.16})$$

Since (A.15) and (A.16) cannot both be satisfied when $p_I + p_{II} > 0$ (i.e., P is played with positive probability), σ must specify that (P, U, R) is never realized as part of a history. This implies that player 1 has an ex-ante profitable deviation to the strategy of always playing Q , contradicting the assumption that σ is an ex-ante equilibrium. ■

Since σ being part of a sequentially optimal (σ, ν) implies σ is an ex-ante equilibrium, Proposition A.1 immediately implies that none of the stronger concepts such as SEA, PBE or sequential equilibrium can admit the play of P with positive probability under ambiguity neutrality. The next result shows that the situation changes dramatically under sufficient ambiguity aversion.

Proposition A.2 *There exist ϕ and μ (e.g., $\phi(x) \equiv -e^{-x}$ and $\mu(\pi_0) = \mu(\pi_1) = \frac{1}{2}$, where $\pi_0(I) = 1$ and $\pi_1(I) = 0$) such that in an SEA both types of player 1 play P with probability*

1, and (U, R) is played with probability greater than $\frac{1}{2}$.

Proof of Proposition A.2. Let μ put probability $\frac{1}{2}$ on π_0 and $\frac{1}{2}$ on π_1 , where $\pi_0(I) = 1$ and $\pi_1(I) = 0$.²⁴ Let $\phi(x) \equiv -e^{-x}$.²⁵ Let σ be a strategy profile specifying that both types of player 1 play P with probability 1, player 2 plays U with probability λ^* if given the move and player 3 plays R with probability λ^* if given the move, where $\lambda^* = 1 - \frac{2}{5} \ln(3/2)$. Notice that according to σ , (P, U, R) occurs with probability $(1 - \frac{2}{5} \ln(3/2))^2 > \frac{7}{10}$. Observe that player 1 strictly prefers ex-ante to play P with probability 1 for both types if and only if, after the play of P , (U, R) will be played with probability greater than $\frac{1}{2}$. The same is true for each type of player 1 after her type is realized as well. Player 2 ex-ante chooses the probability, $\lambda \in [0, 1]$, with which to play U if given the move to maximize

$$-\frac{1}{2}e^{-\lambda} - \frac{1}{2}e^{-(\lambda + \frac{5}{2}(1-\lambda))}.$$

One can verify that the maximum is reached at $\lambda = \lambda^*$. Similarly, player 3 ex-ante chooses the probability, $\lambda \in [0, 1]$, with which to play R if given the move to maximize

$$-\frac{1}{2}e^{-(\lambda + \frac{5}{2}(1-\lambda))} - \frac{1}{2}e^{-\lambda}$$

which is again maximized at $\lambda = \lambda^*$.

Now consider the following sequence of completely mixed strategies with limit σ : σ^k has each type of player 1 play P with probability $1 - \frac{1}{k+1}$, and leaves the strategies otherwise the same as in σ . By Lemma A.4, Theorem A.2 provides a formula (A.12) for an interim belief system ν satisfying smooth rule consistency. Recall that player 1 learns the parameter at the beginning of the game. Thus we need only be concerned with the beliefs of players 2 and 3. Therefore (σ, ν) satisfies smooth rule consistency. It remains to show (σ, ν) is sequentially optimal. Since

$$\frac{\nu_{2, \{I, II\} \times \{P\}}(\delta_{(I, P)})}{\nu_{2, \{I, II\} \times \{P\}}(\delta_{(II, P)})} = \frac{\nu_{3, \{I, II\} \times \{P\} \times \{U, D\}}(\delta_{(II, P)})}{\nu_{3, \{I, II\} \times \{P\} \times \{U, D\}}(\delta_{(I, P)})} = \frac{\frac{\phi'_i(\lambda^*)}{\phi'_i(\lambda^*)} \frac{1}{2}}{\frac{\phi'_i(\lambda^* + \frac{5}{2}(1-\lambda^*))}{\phi'_i(\lambda^* + \frac{5}{2}(1-\lambda^*))} \frac{1}{2}} = 1,$$

σ remains optimal for players 2 and 3 following the play of P given ν . Thus, (σ, ν) is sequentially optimal. It is therefore an SEA. ■

As the proof of Proposition A.2 mentions, the example μ is chosen for simplicity, and

²⁴The degeneracy of the π in the support of μ is not necessary for the argument to go through – it merely shortens some calculations and reduces the ambiguity aversion required.

²⁵Any more concave ϕ will also work, as will any ϕ more concave than $-e^{-\alpha x}$ for $\alpha = \frac{-4(\ln(2/3))}{5(2-\sqrt{2})} \approx 0.554$.

degeneracy of the measures in its support is not necessary for the result.

A.4.2 Formal Comparative Statics in Ambiguity Aversion

Notation A.2 For a game $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$, let $E_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ denote the set of all ex-ante equilibria of the game $\hat{\Gamma} = (N, H, (\mathcal{I}_i)_{i \in N}, (\hat{\mu}_i)_{i \in N}, (u_i, \hat{\phi}_i)_{i \in N})$ differing from Γ only in ambiguity aversions and beliefs. Let $Q_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ denote the analogous set of sequentially optimal strategy profiles and $S_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ denote the analogous set of SEA strategy profiles.

Notation A.3 Denote the identity function by ι .

Proof of Theorem 4.1. We show that there exists a game Γ and $(\hat{\phi}_i)_{i \in N}$ such that

$$E_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) \cap E_\Gamma((\mu_i)_{i \in N}, (\iota)_{i \in N}) = \emptyset.$$

Modify Example A.4.1 by removing the action Q for Player 1. For each player i , let $\mu_i = \mu$ where μ puts probability $\frac{1}{2}$ on π_0 and $\frac{1}{2}$ on π_1 , where $\pi_0(I) = 1$ and $\pi_1(I) = 0$, and let $\hat{\phi}_i(x) = -e^{-x}$ for all i . With these preferences, the unique ex-ante equilibrium has player 2 play U with probability λ^* and player 3 play R with probability λ^* , where $\lambda^* = 1 - \frac{2}{5} \ln(3/2)$. In contrast, when $\phi_i = \iota$ for all i , using the same μ , then the unique ex-ante equilibrium has player 2 playing D with probability 1 and player 3 play L with probability 1. ■

Examination of the proof shows that, fixing beliefs, not only are the equilibrium strategies distinct under ambiguity aversion compared to ambiguity neutrality, but it can also be that the strategies under ambiguity aversion generate paths of play that do not occur in equilibrium under ambiguity neutrality. An analogue of Theorem 4.1 is true for sequential optima, SEA and any other refinement of ex-ante equilibria as well, as they are all ex-ante equilibria. Thus, with fixed beliefs, change in ambiguity aversion can impact the set of equilibrium strategies and realized play.

Further examination of the proof shows that ambiguity aversion continues to affect the equilibrium set even if we impose *common beliefs* (i.e., $\mu_i = \mu$ for all players i). The next result addresses the question of whether ambiguity aversion plus the assumption of common beliefs has equilibrium implications that are different from ambiguity neutrality plus the assumption of common beliefs. It shows that, in this case, ambiguity aversion always weakly expands the set of equilibria compared to ambiguity neutrality and may do so strictly:

Proof of Theorem 4.2. We show: For all games Γ and $(\hat{\phi}_i)_{i \in N}$, $\bigcup_{\hat{\mu}} E_\Gamma((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} E_\Gamma((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$, and the same holds when Q or S replaces E ; moreover, there exists a game Γ and $(\hat{\phi}_i)_{i \in N}$ such that all these inclusions are strict and some of the new equilibrium

strategies induce new paths of play. That $\bigcup_{\hat{\mu}} E_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} E_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ follows by considering only degenerate beliefs on the left-hand side and choosing them to have the same reduced measure as the right-hand side beliefs. $\bigcup_{\hat{\mu}} Q_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} Q_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ follows using the same construction and additionally taking the left-hand side updated beliefs at each information set to be degenerate with the same reduced measure as the right-hand side updated beliefs at the corresponding information set and noting that this preserves optimality at each information set. $\bigcup_{\hat{\mu}} S_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} S_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ follows using the same construction as for sequential optima, observing that the left-hand side degenerate beliefs satisfy smooth rule consistency since the right-hand side beliefs do so. As shown by Propositions 3.1 and 3.3, in the running example with sufficient ambiguity aversion and $x \leq 0.5$ the inclusion is strict and the new strategies generate new paths of play. ■

Next, in the constructive proof of Theorem 4.3, we show that beliefs $\hat{\mu}_i$ and $\hat{\nu}_{i, I_i}$ that support a given equilibrium profile σ are related to the beliefs μ_i and ν_{i, I_i} in the game with the original ambiguity aversion(s) by the formulae in (A.17) and (A.18) where the ϕ_i are the original and $\hat{\phi}_i$ the new specifications of ambiguity aversions.

Proof of Theorem 4.3. Fix a game Γ . We show: For all $(\hat{\phi}_i)_{i \in N}$,

$$\bigcup_{(\hat{\mu}_i)_{i \in N}} E_{\Gamma}((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) = \bigcup_{(\hat{\mu}_i)_{i \in N}} E_{\Gamma}((\hat{\mu}_i)_{i \in N}, (\phi_i)_{i \in N}),$$

and the same holds when Q or S replaces E .

Let $\sigma \in E_{\Gamma}((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$. Ex-ante equilibrium is equivalent to ex-ante optimality for all players i of σ_i according to i 's preferences given σ_{-i} . This ex-ante optimality is equivalent to $\sigma'_i = \sigma_i$ maximizing (4.1) with respect to σ'_i . Let $\hat{\mu}_i$ be the probability measure such that

$$\hat{\mu}_i(\pi) \propto \frac{\phi'_i(\sum_h u_i(h) p_{\sigma}(h|h^0) \pi(h^0))}{\hat{\phi}'_i(\sum_h u_i(h) p_{\sigma}(h|h^0) \pi(h^0))} \mu_i(\pi). \quad (\text{A.17})$$

Observe that replacing ϕ_i and μ_i with $\hat{\phi}_i$ and $\hat{\mu}_i$ leaves the effective beliefs at σ , and so also (4.1), unchanged up to proportionality. Thus σ_i is ex-ante optimal for player i given $\hat{\phi}_i$, $\hat{\mu}_i$ and σ_{-i} . As this is true for each player i , $\sigma \in E_{\Gamma}((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

Turn now to sequentially optimal strategy profiles. Suppose $\sigma \in Q_{\Gamma}((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (σ, ν) is sequentially optimal for Γ . Let $\hat{\mu}_i$ be defined as in (A.17) and for each I_i , $\hat{\nu}_{i, I_i}$ be the probability measure such that

$$\hat{\nu}_{i, I_i}(\pi) \propto \frac{\phi'_i\left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{\sigma}(h|h^{s(I_i)}) \pi(h^{s(I_i)})\right)}{\hat{\phi}'_i\left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{\sigma}(h|h^{s(I_i)}) \pi(h^{s(I_i)})\right)} \nu_{i, I_i}(\pi). \quad (\text{A.18})$$

By the argument in the ex-ante equilibrium part of this proof, $\sigma \in E_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. Optimality for player i at I_i as a function of σ'_i is equivalent (see (2.5) and (2.6)) to $\sigma'_i = \sigma_i$ maximizing

$$\sum_{\pi \in \Delta(I_i)} \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^s(I_i)) \pi(h^s(I_i)) \right) \phi'_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_\sigma(h|h^s(I_i)) \pi(h^s(I_i)) \right) \nu_{i, I_i}(\pi) \quad (\text{A.19})$$

with respect to σ'_i . Observe that replacing ϕ_i and ν_{i, I_i} with $\hat{\phi}_i$ and $\hat{\nu}_{i, I_i}$ leaves (A.19) unchanged up to proportionality. This is true for each player i and I_i . Thus, $(\sigma, \hat{\nu})$ is sequentially optimal in $\hat{\Gamma}$.

We now extend the argument to SEA. Suppose $\sigma \in S_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (σ, ν) is an SEA for Γ (with corresponding sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$). By Theorem A.2, ν satisfies (A.12) using $\{\sigma^k\}_{k=1}^\infty$. As above, let $\hat{\mu}_i$ be as in (A.17) and for each I_i , $\hat{\nu}_{i, I_i}$ be defined as in (A.18). By our previous arguments, $(\sigma, \hat{\nu})$ is sequentially optimal in $\hat{\Gamma}$. It remains to show that $(\sigma, \hat{\nu})$ satisfies smooth rule consistency in $\hat{\Gamma}$. Observe that replacing ϕ_i , μ_i and ν_{i, I_i} with $\hat{\phi}_i$, $\hat{\mu}_i$ and $\hat{\nu}_{i, I_i}$ in (A.12) preserves its validity. Thus, by Theorem A.2, $(\sigma, \hat{\nu})$ satisfies smooth rule consistency in $\hat{\Gamma}$. Therefore $(\sigma, \hat{\nu})$ is an SEA of $\hat{\Gamma}$.

The above arguments have shown $E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} E_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$, $Q_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} Q_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ and $S_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} S_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. Applying these arguments twice (the second time with the roles of ϕ_i and $\hat{\phi}_i$ interchanged), we obtain that, for any game, the union over all beliefs of the set of equilibrium strategy profiles is independent of ambiguity aversion. ■

Finally, turn to the case of pure strategies and only pure strategy deviations as in Battigalli et al. (2015). Modify the equilibrium set notation to restrict attention to pure strategies:

Definition A.2 For a game $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$, let $\tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ be the set of all ex-ante equilibria with respect to pure strategies of a game $\hat{\Gamma} = (N, H, (\mathcal{I}_i)_{i \in N}, (\hat{\mu}_i)_{i \in N}, (u_i, \hat{\phi}_i)_{i \in N})$ differing from Γ only in ambiguity aversions and beliefs. Let $\tilde{Q}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ and $\tilde{S}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ be the analogous respective sets of sequentially optimal and SEA strategy profiles with respect to pure strategies.

Proof of Theorem 4.4. Fix a game Γ . We show that: For all $(\hat{\phi}_i)_{i \in N}$ such that, for each i , $\hat{\phi}_i$ is at least as concave as ϕ_i , $\tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$, and the same holds when \tilde{Q} or \tilde{S} replaces \tilde{E} . There exists a game Γ and $(\hat{\phi}_i)_{i \in N}$ such that for each i , $\hat{\phi}_i$ is at least as concave as ϕ_i , all these inclusions are strict and some of the new equilibrium strategies induce new paths of play.

Suppose $\varsigma \in \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$, for each i , $\hat{\phi}_i = \chi_i(\phi_i)$ for some increasing, differentiable and concave χ_i (note that differentiability of χ_i is implied by the continuous differentiability of $\hat{\phi}_i$ in the class of games considered in this paper) and $\hat{\mu}_i$ is the probability measure such that

$$\hat{\mu}_i(\pi) \propto \frac{\mu_i(\pi)}{\chi'_i(\phi_i(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0)))}. \quad (\text{A.20})$$

By definition of $\tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$, for each i and each ς'_i ,

$$\sum_\pi \phi_i \left(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0) \right) \mu_i(\pi) \geq \sum_\pi \phi_i \left(\sum_h u_i(h)p_{(\varsigma'_i, \varsigma_{-i})}(h|h^0)\pi(h^0) \right) \mu_i(\pi). \quad (\text{A.21})$$

Since χ_i is increasing, differentiable and concave, for each π ,

$$\begin{aligned} & \chi_i \left(\phi_i \left(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0) \right) \right) - \chi_i \left(\phi_i \left(\sum_h u_i(h)p_{(\varsigma'_i, \varsigma_{-i})}(h|h^0)\pi(h^0) \right) \right) \\ & \geq \chi'_i \left(\phi_i \left(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0) \right) \right) \\ & \quad \cdot \left[\phi_i \left(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0) \right) - \phi_i \left(\sum_h u_i(h)p_{(\varsigma'_i, \varsigma_{-i})}(h|h^0)\pi(h^0) \right) \right]. \end{aligned}$$

Thus, dividing both sides by $\chi'_i(\phi_i(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0)))$ and taking the expectation with respect to μ_i yields

$$\begin{aligned} & \sum_\pi \chi_i \left(\phi_i \left(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0) \right) \right) - \chi_i \left(\phi_i \left(\sum_h u_i(h)p_{(\varsigma'_i, \varsigma_{-i})}(h|h^0)\pi(h^0) \right) \right) \hat{\mu}_i(\pi) \\ & \geq \sum_\pi \left[\phi_i \left(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0) \right) - \phi_i \left(\sum_h u_i(h)p_{(\varsigma'_i, \varsigma_{-i})}(h|h^0)\pi(h^0) \right) \right] \mu_i \geq 0, \end{aligned}$$

where the last inequality follows from A.21. Since this is true for each i and each ς'_i , $\varsigma \in \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. This shows $\tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

Turn now to the part of the theorem about sequentially optimal strategy profiles. Suppose $\varsigma \in \tilde{Q}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (ς, ν) is sequentially optimal for Γ with respect to pure strategies. Further suppose that for each i , $\hat{\phi}_i = \chi_i(\phi_i)$ for some increasing, differentiable and concave χ_i , $\hat{\mu}_i$ is defined as in (A.20), and for each I_i ,

$\hat{\nu}_{i,I_i}$ is the probability measure such that

$$\hat{\nu}_{i,I_i}(\pi) \propto \frac{\nu_{i,I_i}(\pi)}{\chi'_i \left(\phi_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_\varsigma(h|h^s(I_i)) \pi(h^s(I_i)) \right) \right)}.$$

By the argument in the ex-ante equilibrium part of this proof, $\varsigma \in \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. By definition of $\tilde{Q}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$, for each i , each I_i and each ς'_i ,

$$\begin{aligned} & \sum_{\pi} \phi_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_\varsigma(h|h^s(I_i)) \pi(h^s(I_i)) \right) \nu_{i,I_i}(\pi) \\ & \geq \sum_{\pi} \phi_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_{(\varsigma'_i, \varsigma_{-i})}(h|h^s(I_i)) \pi(h^s(I_i)) \right) \nu_{i,I_i}(\pi). \end{aligned} \quad (\text{A.22})$$

Since χ_i is increasing, differentiable and concave, for each π we repeat the argument in the ex-ante equilibrium part of this proof to conclude that $\varsigma \in \tilde{Q}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. This shows $\tilde{Q}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{Q}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

Finally, turn to the part of the theorem about SEA strategy profiles. Suppose $\varsigma \in \tilde{S}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (ς, ν) is an SEA for Γ with respect to pure strategies, where the sequence used in satisfying smooth rule consistency is $\{\sigma^k\}_{k=1}^\infty$. By Theorem A.2, ν satisfies (A.12) using $\{\sigma^k\}_{k=1}^\infty$. Further suppose that for each i , $\hat{\phi}_i = \chi_i(\phi_i)$ for some increasing, differentiable and concave χ_i , $\hat{\mu}_i$ is defined as in (A.20), and, for each I_i , $\hat{\nu}_{i,I_i}$ is defined as in (A.22). By the arguments in the sequentially optimal part of the proof, $\varsigma \in \tilde{Q}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. Observe that replacing ϕ_i , μ_i and ν_{i,I_i} with $\hat{\phi}_i$, $\hat{\mu}_i$ and $\hat{\nu}_{i,I_i}$ in (A.12) preserves its validity. Thus, by Theorem A.2, $(\sigma, \hat{\nu})$ satisfies smooth rule consistency in $\hat{\Gamma}$. Thus $\varsigma \in \tilde{S}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. This shows $\tilde{S}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{S}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

To prove that strict inclusions may happen, consider the game depicted in Figure A.2. There are two players, 1 and 2. First, it is determined whether player 2 is of type I or type II and 2 observes the type. Player 1 does not observe the type. The payoff pairs in Figure A.2 describe vNM utility payoffs given players' actions and type (i.e., (u_1, u_2) means that player i receives u_i). Player 1's first move in the game is to choose between action T (wo) which gives the move to player 2 and action B (et) (i.e., betting that player 2 is of type II) which reveals the type and ends the game. If T , then player 2's move is a choice between C (ontinue) which leads to player 1 again having a non-trivial move, and S (top) which reveals the type and ends the game. If C , then player 1 has a choice between G (amble) and H (edge) after which the game ends.




Figure A.2: Ambiguity aversion generates new equilibria with respect to pure strategies

Under ambiguity neutrality for both players, $\bigcup_{(\mu_i)_{i \in N}} \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\iota)_{i \in N}) = \{(B, (S, S), H), (B, (S, S), G), (B, (C, S), H), (B, (C, S), G), (B, (S, C), H), (B, (S, C), G), (B, (C, C), H), (B, (C, C), G)\}$. To see this, first note that if $\sum_\pi \pi(I) \mu_1(\pi) \in (0, \frac{2}{5})$, then all the pure profiles where 1 plays B are ex-ante equilibria under ambiguity neutrality. Second, any pure profile where 1 plays T cannot be an ex-ante equilibrium under ambiguity neutrality. Observe that 2 plays C following T (under either type) only if 1 plays H , 1 can play H rather than G on-path if and only if $2 \geq 6 \sum_\pi \pi(I) \mu_1(\pi)$, and 1 can play T followed by H rather than B if and only if $p(C)2 \geq 4(1 - \sum_\pi \pi(I) \mu_1(\pi))$ where $0 \leq p(C) \leq 1$ is 1's reduced probability that the type is such that 2 plays C . Since $\sum_\pi \pi(I) \mu_1(\pi)$ cannot be simultaneously $\leq \frac{1}{3}$ and $\geq \frac{1}{2}$, 1 cannot play T in pure strategy equilibrium under ambiguity neutrality.

By the weak inclusions already shown, and since SEA implies sequentially optimal, which in turn implies ex-ante equilibrium, it is enough to show that for some strictly concave $\hat{\phi}_1$ there is an SEA strategy profile with respect to pure strategies not contained in $\bigcup_{(\mu_i)_{i \in N}} \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\iota)_{i \in N})$. To this end, suppose $\hat{\phi}_1(x) \equiv -e^{-2x}$, $\hat{\phi}_2 \equiv \iota$ and $\mu_1(\pi_1) = \mu_1(\pi_2) = \frac{1}{2}$, where $\pi_1(I) = \frac{1}{4}$ and $\pi_2(I) = \frac{3}{4}$. Consider the pure strategy profile $(T, (C, C), H)$ and a sequence of completely mixed strategy profiles approaching it where the k^{th} element of the sequence has player 1 and each type of player 2 playing the action not assigned by $(T, (C, C), H)$ with probability $\frac{1}{k+1}$ at any point they are given the move. By Lemma A.4 and Theorem A.2, ν calculated using (A.12) with $\hat{\phi}_1$ and μ_1 satisfies smooth rule consistency. By Theorem 3.3, for sequential optimality, it is sufficient to check against one-stage deviations, and therefore only at information sets where the player has a non-trivial move.

For $I_1 = \{I, II\}$ and $I_1 = \{I, II\} \times \{T\} \times \{C\}$, $\nu_{1,I_1}((\pi_1)_{I_1}) = \nu_{1,I_1}((\pi_2)_{I_1}) = \frac{1}{2}$. Since $\hat{\phi}_1(2) > \frac{1}{2}(\hat{\phi}_1(3) + \hat{\phi}_1(1))$ and $\hat{\phi}_1(2) > \frac{1}{2}(\hat{\phi}_1(\frac{3}{2}) + \hat{\phi}_1(\frac{9}{2}))$, 1 is best responding, and since C is a best response for player 2 given any beliefs, $(T, (C, C), H)$ is sequentially optimal with respect to pure strategies given ν . Therefore $(T, (C, C), H) \in \tilde{S}_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ and the proof is complete. ■

A.5 Proofs of results in Section 4.2

Parts of the next proof (of Theorem 4.5) make use of the following particularly convenient set of $\hat{\phi}_i$ at least as concave as ϕ_i , parametrized by $l \geq 1$ and $b \geq 1$, with $\hat{\phi}_i$ strictly more concave than ϕ_i when $b > 1$ and equal to ϕ_i when $b = 1$:

Let e_i^l denote the l^{th} lowest distinct value of $\sum_{h \in H} u_i(h)p_\sigma(h|h^0)\pi(h^0)$ generated by π in the support, Π_i , of μ_i .

Definition A.3 For any $l \geq 1$ such that e_i^{l+1} exists, for $b \geq 1$ let $\hat{\phi}_i^l \equiv \psi_i^l \circ \phi_i$, where ψ_i^l is defined by

$$\psi_i^l(y) = \begin{cases} y + \frac{1}{2}(b-1)[\phi_i(e_i^l) + \phi_i(e_i^{l+1})] & , \quad y \geq \phi_i(e_i^{l+1}) \\ \frac{-(b-1)y^2 + 2[b\phi_i(e_i^{l+1}) - \phi_i(e_i^l)]y - (b-1)[\phi_i(e_i^l)]^2}{2[\phi_i(e_i^{l+1}) - \phi_i(e_i^l)]} & , \quad \phi_i(e_i^l) < y < \phi_i(e_i^{l+1}) \\ b \cdot y & , \quad y \leq \phi_i(e_i^l) . \end{cases}$$

When $b > 1$, it may be verified that any ψ_i^l is continuously differentiable, concave, strictly increasing and not affine. Notice that for all $x \leq e_i^l$, $\hat{\phi}_i^l(x) = b\phi_i'(x)$ and for all $x \geq e_i^{l+1}$, $\hat{\phi}_i^l(x) = \phi_i'(x)$.

Proof of Theorem 4.5. This proof makes use of the E_Γ , Q_Γ and S_Γ notations for sets of equilibria given in Notation A.2 (see p. 62). Fixing an ex-ante equilibrium $\sigma \in E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ (resp. Q_Γ with associated interim belief system ν or S_Γ with associated ν and sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$) and a player i , say that ambiguity aversion makes σ_i *ex-ante* (resp. *sequentially optimal or SEA*) *belief robust* if for each $\varepsilon_i \in (0, \frac{1}{|\text{supp } \mu_i|})$, there exists $\bar{\phi}_i^{\varepsilon_i}$ at least as concave as ϕ_i so that σ_i is an ex-ante best response to σ_{-i} given each $\hat{\mu}_i$ and $\hat{\phi}_i$ such that $\min_{\pi \in \text{supp } \mu_i} \hat{\mu}_i(\pi) > \varepsilon_i$ and $\hat{\mu}_i$ has the same support, Π_i , as μ_i , and such that $\hat{\phi}_i$ at least as concave as $\bar{\phi}_i^{\varepsilon_i}$ (resp. that plus also a best response to σ_{-i} at each information set I_i given $\hat{\phi}_i$ and ν_{i,I_i} or (for SEA) all of the previous plus satisfying the part for player i of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$). To prove that ambiguity aversion makes σ ex-ante (resp. sequentially optimal or SEA) belief robust, it is sufficient to show, for each player i , that ambiguity aversion makes σ_i ex-ante (resp. sequentially optimal or SEA) belief robust. The argument is the same for each player, so

for the remainder of the proof fix a player i . Also assume for the remainder of the argument that $|\Pi_i| > 1$, as otherwise the result follows immediately because there is only one possible belief with that support.

We begin by proving that ambiguity aversion makes σ_i ex-ante belief robust. Recall that σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\mu}_i$ and $\hat{\phi}_i$ if and only if σ_i maximizes, among all σ'_i , (4.1) with $\hat{\mu}_i$ replacing μ_i and $\hat{\phi}_i$ replacing ϕ_i . Observe that any strategies σ'_i that are weakly worse than σ_i (in terms of ex-ante expected payoff, $\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi(h^0)$) for all $\pi \in \Pi_i$ can never interfere with optimality of σ_i and will thus, without loss of generality, be ignored whenever making statements about strategies other than σ_i in what follows. For each l , denote by π_i^l the unique $\pi \in \Pi_i$ under which σ_i gives e_i^l , the l^{th} lowest distinct ex-ante expected payoff generated by Π_i . For each strategy σ'_i and $1 \leq l \leq |\Pi_i|$, denote

$$d_i^l(\sigma'_i) \equiv \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^l(h^0) - \sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi_i^l(h^0) \right) \cdot \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^l(h^0) \right).$$

The conclusion of the theorem in the ex-ante case is immediate when all strategies σ'_i are weakly worse than σ_i for all π_i^l (i.e., $d_i^l(\sigma'_i) \geq 0$ for all l), so assume that there exists a strategy σ'_i with $d_i^l(\sigma'_i) < 0$ for some $1 \leq l \leq |\Pi_i|$.

We next show that all strategies σ'_i must have $d_i^1(\sigma'_i) \geq 0$. To see this, suppose, to the contrary, there exists a strategy $\hat{\sigma}_i$ with $d_i^1(\hat{\sigma}_i) < 0$. Since σ is ex-ante robust to increased ambiguity aversion, σ_i is an ex-ante best response to σ_{-i} for player i given μ_i and $\hat{\phi}_i^1$ (from Definition A.3), and, in particular, is at least as good as $\hat{\sigma}_i$. Using (4.1) with $\hat{\phi}_i^1$ replacing ϕ_i , this implies

$$\sum_{l=1}^{|\Pi_i|} d_i^l(\hat{\sigma}_i) \mu_i(\pi) + (b-1) d_i^1(\hat{\sigma}_i) \geq 0.$$

Since the value of the first term is bounded and $d_i^1(\hat{\sigma}_i) < 0$, taking b large enough generates a contradiction.

For each pure strategy ς'_i , let $m(\varsigma'_i) < |\Pi_i|$ be the smallest number l for which $d_i^{l+1}(\varsigma'_i) < 0$. By the previous paragraph, $m(\varsigma'_i) \geq 1$. By the definition of $m(\varsigma'_i)$, $d_i^l(\varsigma'_i) \geq 0$ for all $1 \leq l \leq m(\varsigma'_i)$. Furthermore, $d_i^l(\varsigma'_i) > 0$ for some $1 \leq l \leq m(\varsigma'_i)$, because otherwise ς'_i could be used together with μ_i and $\hat{\phi}_i^{m(\varsigma'_i)+1}$ to generate a contradiction to σ_i being ex-ante robust to increased ambiguity aversion. Thus $\sum_{l=1}^{m(\varsigma'_i)} d_i^l(\varsigma'_i) > 0$ and $\min_{m(\varsigma'_i)+1 \leq l \leq |\Pi_i|} d_i^l(\varsigma'_i) < 0$. For each $1 \leq m < |\Pi_i|$, if there exists no pure strategy ς'_i with $m(\varsigma'_i) = m$, then let $B(m) = 1$,

otherwise let

$$B(m) \equiv \max\{1, \max_{\zeta'_i | m(\zeta'_i)=m} \frac{-\min_{m+1 \leq l \leq |\Pi_i|} d_i^l(\zeta'_i)}{\varepsilon_i \sum_{l=1}^m d_i^l(\zeta'_i)}\},$$

which is well defined because the set of pure strategies is finite. Define $\bar{\phi}_i^{\varepsilon_i} = \psi_i^1 \circ \dots \circ \psi_i^{|\Pi_i|-1} \circ \phi_i$, for ψ_i^m with $b = B(m)$ for each $1 \leq m < |\Pi_i|$. Consider $\hat{\phi}_i$ at least as concave as $\bar{\phi}_i^{\varepsilon_i}$, i.e., $\hat{\phi}_i = \psi_i \circ \bar{\phi}_i^{\varepsilon_i}$ for some ψ_i continuously differentiable, concave and strictly increasing. For any $\hat{\mu}_i$ such that $\min_{\pi \in \text{supp } \mu_i} \hat{\mu}_i(\pi) > \varepsilon_i$ and $\hat{\mu}_i$ has the same support as μ_i , and any pure strategy ζ'_i ,

$$\begin{aligned} & \sum_{\pi \in \Pi_i} \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi(h^0) - \sum_{h \in H} u_i(h) p_{(\zeta'_i, \sigma_{-i})}(h|h^0) \pi(h^0) \right) \\ & \cdot \hat{\phi}'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi(h^0) \right) \hat{\mu}_i(\pi) \\ & = \sum_{l=1}^{|\Pi_i|} d_i^l(\zeta'_i) \psi'_i \left(\bar{\phi}_i^{\varepsilon_i} \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^l(h^0) \right) \right) \left(\prod_{m=l}^{|\Pi_i|-1} B(m) \right) \hat{\mu}_i(\pi_i^l) \\ & \geq \psi'_i \left(\bar{\phi}_i^{\varepsilon_i} \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^{m(\zeta'_i)}(h^0) \right) \right) \left(\prod_{m=m(\zeta'_i)+1}^{|\Pi_i|-1} B(m) \right) \\ & \quad \cdot \left(B[m(\zeta'_i)] \varepsilon_i \sum_{l=1}^{m(\zeta'_i)} d_i^l(\zeta'_i) + \min_{m(\zeta'_i)+1 \leq l \leq |\Pi_i|} d_i^l(\zeta'_i) \right) \\ & \geq 0, \end{aligned}$$

where the last inequality follows by applying the definition of $B[m(\zeta'_i)]$. Therefore, σ_i does at least as well as any pure strategy ζ'_i given σ_{-i} according to (4.1) with $\hat{\mu}_i$ and $\hat{\phi}_i$ replacing μ_i and ϕ_i . Since (4.1) is linear in the mixing weights in σ'_i , this is sufficient to conclude that σ_i is a best response to σ_{-i} given $\hat{\mu}_i$ and $\hat{\phi}_i$. Therefore ambiguity aversion makes σ_i ex-ante belief robust.

Consider now sequential optimality. Consider $\bar{\phi}_i^{\varepsilon_i}$ and $\hat{\phi}_i$ as defined above. Since σ is sequentially optimal robust to increased ambiguity aversion, σ_i is an ex-ante best response to σ_{-i} given μ_i and $\hat{\phi}_i$, and for each information set I_i there exists a belief ν_{i,I_i} such that σ_i is a best response at I_i to σ_{-i} given ν_{i,I_i} and $\hat{\phi}_i$. Consider any $\hat{\mu}_i$ such that $\min_{\pi \in \text{supp } \mu_i} \hat{\mu}_i(\pi) > \varepsilon_i$ and $\hat{\mu}_i$ has the same support as μ_i . By the ex-ante equilibrium argument above, σ_i is an ex-ante best response to σ_{-i} given $\hat{\mu}_i$ and $\hat{\phi}_i$. Given $\hat{\mu}_i$ and $\hat{\phi}_i$, derive $\hat{\nu}_{i,I_i}$ from the smooth rule using σ for those information sets I_i for which that rule implies that $\hat{\nu}_{i,I_i}$ must vary with ex-ante beliefs. By Lemma A.1, σ_i an ex-ante best response to σ_{-i} given $\hat{\mu}_i$ and $\hat{\phi}_i$ implies

σ_i is a best response to σ_{-i} at these information sets given $\hat{\phi}_i$ and $\hat{\nu}_{i,I_i}$. Extend $\hat{\nu}$ by setting $\hat{\nu}_{i,I_i} = \nu_{i,I_i}$ elsewhere. Thus σ_i is a best response to σ_{-i} also at these remaining information sets given $\hat{\nu}_{i,I_i}$ and $\hat{\phi}_i$, as this fact is not affected by the shift from μ_i to $\hat{\mu}_i$. This shows that ambiguity aversion makes σ_i sequentially optimal belief robust.

Finally turn to SEA. We establish the existence of beliefs at each I_i so that i 's part of both sequential optimality and smooth rule consistency are satisfied. Since (σ, ν) satisfies smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$, Theorem A.2 yields that $\bar{p}_{-i,\sigma_{-i}}(h^t|h^0)$ exists for each player i and each $h^t \in I_i \in \mathcal{F}_{i,\sigma}$. Given any $\hat{\phi}_i$, for each I_i , construct a belief $\hat{\nu}_{i,I_i}$ as defined in (A.12) using $\{\sigma^k\}_{k=1}^\infty$ with $\hat{\mu}_i$ and $\hat{\phi}_i$ replacing μ_i and ϕ_i . Theorem A.2 applied with $\hat{\mu}_i$ and $\hat{\phi}_i$ replacing μ_i and ϕ_i (and noting that $\bar{p}_{-i,\sigma_{-i}}(h^t|h^0)$ is independent of the choice of $\hat{\mu}_i$ and $\hat{\phi}_i$) implies that σ together with $\hat{\nu}$ satisfies player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ given $\hat{\mu}_i$ and $\hat{\phi}_i$.

Showing that σ_i is a best response to σ_{-i} for player i at each I_i given $\hat{\nu}_{i,I_i}$ and $\hat{\phi}_i$ is equivalent to showing that

$$\begin{aligned} \sigma_i \in \arg \max_{\sigma'_i} \sum_{\hat{\pi} \in \Pi_i} \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^{s(I_i)}) \bar{p}_{-i,\sigma_{-i}}(h^{s(I_i)}|h^0) \hat{\pi}(h^0) \right) \quad (\text{A.23}) \\ \cdot \hat{\phi}'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \hat{\pi}(h^0) \right) \hat{\mu}_i(\hat{\pi}), \end{aligned}$$

as can be seen by considering (A.19) with $\hat{\nu}_{i,I_i}$ replacing ν_{i,I_i} and $\hat{\phi}_i$ replacing ϕ_i , substituting for $\hat{\nu}_{i,I_i}$ and $\bar{\pi}_{I_i}$ using (A.12), replacing the summation over $\pi \in \Delta(I_i)$ and $\hat{\pi} \in \Delta(\Theta)$ such that $\bar{\pi}_{I_i} = \pi$ with summation over $\hat{\pi} \in \Pi_i$ since each element in the support of $\hat{\nu}_{i,I_i}$ is $\bar{\pi}_{I_i}$ for some $\hat{\pi}$ in the support of $\hat{\mu}_i$, and simplifying, including, since π no longer appears in the expression, replacing the notation $\hat{\pi} \in \Delta(\Theta)$ with $\pi \in \Delta(\Theta)$.

Since σ is SEA robust to increased ambiguity aversion, σ_i is an ex-ante best response to σ_{-i} given μ_i and $\hat{\phi}_i^1$, and for any I_i there exists a belief ν_{i,I_i} , constructed as was $\hat{\nu}_{i,I_i}$ at the beginning of the SEA part of the proof except now using μ_i and $\hat{\phi}_i^1$, such that σ_i is a best response to σ_{-i} given ν_{i,I_i} and $\hat{\phi}_i^1$. From the definition of $\hat{\phi}_i^1$, the assumption that π_i^1 is well-defined and the assumption that the same sequence $\{\sigma^k\}_{k=1}^\infty$ can be used in smooth rule consistency for each $\hat{\phi}_i$ (which ensures use of the same $\bar{p}_{-i,\sigma_{-i}}(h^{s(I_i)}|h^0)$), (A.23) with $\hat{\mu}_i = \mu_i$ and $\hat{\phi}_i = \hat{\phi}_i^1$ with b large enough (i.e., $\hat{\phi}_i^1$ sufficiently concave) implies that, for each I_i , $\sigma'_i = \sigma_i$ must maximize the following expected payoff under π_i^1 ,

$$\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^{s(I_i)}) \bar{p}_{-i,\sigma_{-i}}(h^{s(I_i)}|h^0) \pi_i^1(h^0),$$

and, by the corresponding argument for ex-ante equilibrium, also maximizes the ex-ante expected payoff under π_i^1 .

For each I_i , σ'_i and $1 \leq l \leq |\Pi_i|$, denote

$$\begin{aligned} d_{i,I_i}^l(\sigma'_i) &\equiv \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_\sigma(h|h^{s(I_i)}) \bar{p}_{-i,\sigma_{-i}}(h^{s(I_i)}|h^0) \pi_i^l(h^0) \right. \\ &\quad \left. - \sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i,\sigma_{-i})}(h|h^{s(I_i)}) \bar{p}_{-i,\sigma_{-i}}(h^{s(I_i)}|h^0) \pi_i^l(h^0) \right) \cdot \phi'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi_i^l(h^0) \right). \end{aligned}$$

That σ_i is a best response to σ_{-i} for player i at I_i is immediate when $d_{i,I_i}^l(\zeta'_i) \geq 0$ for all pure strategies ζ'_i and l , so assume that $d_{i,I_i}^l(\zeta'_i) < 0$ for some ζ'_i and $1 \leq l \leq |\Pi_i|$. For each I_i and any such ζ'_i (as any other strategy can never interfere with optimality of σ_i at I_i and thus, without loss of generality, may be ignored), let $m_{i,I_i}(\zeta'_i) < |\Pi_i|$ be the smallest number l for which $d_{i,I_i}^{l+1}(\zeta'_i) < 0$. By the previous paragraph, $m_{i,I_i}(\zeta'_i) \geq 1$. By the definition of $m_{i,I_i}(\zeta'_i)$, $d_{i,I_i}^l(\zeta'_i) \geq 0$ for all $1 \leq l \leq m(\zeta'_i)$. Furthermore, $d_{i,I_i}^l(\zeta'_i) > 0$ for some $1 \leq l \leq m_{i,I_i}(\zeta'_i)$, because otherwise ζ'_i could be used together with μ_i and $\hat{\phi}_i^{m_{i,I_i}(\zeta'_i)+1}$ to generate a contradiction to σ_i being SEA robust to increased ambiguity aversion. Thus $\sum_{l=1}^{m_{i,I_i}(\zeta'_i)} d_{i,I_i}^l(\zeta'_i) > 0$ and $\min_{m_{i,I_i}(\zeta'_i)+1 \leq l \leq |\Pi_i|} d_{i,I_i}^l(\zeta'_i) < 0$. For each $1 \leq m < |\Pi_i|$, if there exists no pure strategy ζ'_i with $m_{i,I_i}(\zeta'_i) = m$, then let $B_{i,I_i}(m) = 1$, otherwise let

$$B_{i,I_i}(m) \equiv \max\left\{1, \max_{\zeta'_i | m_{i,I_i}(\zeta'_i) = m} \frac{-\min_{m+1 \leq l \leq |\Pi_i|} d_{i,I_i}^l(\zeta'_i)}{\varepsilon_i \sum_{l=1}^m d_{i,I_i}^l(\zeta'_i)}\right\},$$

which is well defined because the set of pure strategies is finite. Define $\bar{\phi}_i^{\varepsilon_i} = \psi_i^1 \circ \dots \circ \psi_i^{|\Pi_i|-1} \circ \phi_i$, for ψ_i^m with $b = \bar{B}(m) \equiv \max\{B(m), \max_{I_i \in \mathcal{I}_i} B_{i,I_i}(m)\}$ for each $1 \leq m < |\Pi_i|$, where $B(m)$ is as defined in the ex-ante part of the proof. Consider $\hat{\phi}_i$ at least as concave as $\bar{\phi}_i^{\varepsilon_i}$, i.e., $\hat{\phi}_i = \psi_i \circ \bar{\phi}_i^{\varepsilon_i}$ for some ψ_i continuously differentiable, concave and strictly increasing. Consider any $\hat{\mu}_i$ such that $\min_{\pi \in \text{supp } \mu_i} \hat{\mu}_i(\pi) > \varepsilon_i$ and $\hat{\mu}_i$ has the same support as μ_i . By the argument above for ex-ante equilibrium, σ_i is an ex-ante best response to σ_{-i} given $\hat{\mu}_i$ and $\hat{\phi}_i$. For any

I_i and pure strategy ς'_i ,

$$\begin{aligned}
& \sum_{\pi \in \Pi_i} \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_\sigma(h|h^{s(I_i)}) \bar{p}_{-i, \sigma_{-i}}(h^{s(I_i)}|h^0) \pi(h^0) \right. \\
& \quad \left. - \sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^{s(I_i)}) \bar{p}_{-i, \sigma_{-i}}(h^{s(I_i)}|h^0) \pi(h^0) \right) \cdot \hat{\phi}'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0) \right) \hat{\mu}_i(\pi) \\
&= \sum_{l=1}^{|\Pi_i|} d_{i, I_i}^l(\varsigma'_i) \psi'_i \left(\bar{\phi}_i^{\varepsilon_i} \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi_i^l(h^0) \right) \right) \left(\prod_{m=l}^{|\Pi_i|-1} \bar{B}(m) \right) \hat{\mu}_i(\pi_i^l) \\
&\geq \psi'_i \left(\bar{\phi}_i^{\varepsilon_i} \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi_i^{m_{i, I_i}(\varsigma'_i)}(h^0) \right) \right) \left(\prod_{m=m_{i, I_i}(\varsigma'_i)+1}^{|\Pi_i|-1} \bar{B}(m) \right) \\
&\quad \cdot \left(\bar{B}[m_{i, I_i}(\varsigma'_i)] \varepsilon_i \sum_{l=1}^{m_{i, I_i}(\varsigma'_i)} d_i^l(\varsigma'_i) + \min_{m_{i, I_i}(\varsigma'_i)+1 \leq l \leq |\Pi_i|} d_i^l(\varsigma'_i) \right) \\
&\geq 0,
\end{aligned}$$

where the last inequality follows by applying the definition of $\bar{B}[m(\varsigma'_i)]$. Therefore, σ_i does at least as well as any pure strategy ς'_i given σ_{-i} according to (A.23) with $\hat{\mu}_i$ and $\hat{\phi}_i$ replacing μ_i and ϕ_i . Since (A.23) is linear in the mixing weights in σ'_i , this is sufficient to conclude that σ_i is a best response to σ_{-i} for player i at each I_i given $\hat{\nu}_{i, I_i}$ and $\hat{\phi}_i$. Furthermore, by construction, σ together with beliefs $\hat{\nu}_{i, I_i}$ for player i satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ given $\hat{\mu}_i$ and $\hat{\phi}_i$. Therefore ambiguity aversion makes σ_i SEA belief robust. ■

Proof of Remark 4.1. Assume ϕ_i is twice continuously differentiable with strictly positive first derivative and recall that, all along, it was assumed to be strictly increasing and concave. In the proof of Theorem 4.5, $\bar{\phi}_i^{\varepsilon_i}$ was taken to be $\psi_i^1 \circ \dots \circ \psi_i^{|\Pi_i|-1} \circ \phi_i$. From the definition of the ψ_i^m , it follows that $\bar{\phi}_i^{\varepsilon_i}$ is twice differentiable, strictly increasing, and concave and has bounded second derivative. Take this $\bar{\phi}_i^{\varepsilon_i}$ and let $(\bar{\phi}_i^{\varepsilon_i})^{-1}$ be its inverse. We want to show that there exists an $\alpha(\varepsilon_i) > 0$ and an increasing, concave transformation ζ such that $-e^{-\alpha(\varepsilon_i)x} = \zeta[\bar{\phi}_i^{\varepsilon_i}(x)]$ for all $x \in \text{co}(u_i(H))$. For any $y \in \bar{\phi}_i^{\varepsilon_i}[\text{co}(u_i(H))]$, $-e^{-\alpha(\varepsilon_i)x} = \zeta[\bar{\phi}_i^{\varepsilon_i}(x)]$ implies $\zeta(y) = -e^{-\alpha(\varepsilon_i)(\bar{\phi}_i^{\varepsilon_i})^{-1}(y)}$, which is increasing. Thus $\zeta'(y) = \frac{\alpha(\varepsilon_i)e^{-\alpha(\varepsilon_i)(\bar{\phi}_i^{\varepsilon_i})^{-1}(y)}}{(\bar{\phi}_i^{\varepsilon_i})'[(\bar{\phi}_i^{\varepsilon_i})^{-1}(y)]}$, and the sign of $\zeta''(y)$ is the sign of $-\alpha(\varepsilon_i)e^{-\alpha(\varepsilon_i)(\bar{\phi}_i^{\varepsilon_i})^{-1}(y)} \left(\alpha(\varepsilon_i) - \left(-\frac{(\bar{\phi}_i^{\varepsilon_i})''[(\bar{\phi}_i^{\varepsilon_i})^{-1}(y)]}{(\bar{\phi}_i^{\varepsilon_i})'[(\bar{\phi}_i^{\varepsilon_i})^{-1}(y)]} \right) \right)$, so ζ is concave for all sufficiently large $\alpha(\varepsilon_i)$, since $-\frac{(\bar{\phi}_i^{\varepsilon_i})''(x)}{(\bar{\phi}_i^{\varepsilon_i})'(x)}$, the coefficient of ambiguity aversion at $x \in \text{co}(u_i(H))$ is non-negative and bounded above. Note that $(\bar{\phi}_i^{\varepsilon_i})''(x)$ is bounded because the composition of any functions f and g that have bounded second derivatives and

continuous, strictly positive first derivatives has bounded second derivatives and continuous, strictly positive first derivatives as follows from the formula $f[g(x)]'' = [f'[g(x)]g'(x)]' = f''[g(x)][g'(x)]^2 + f'[g(x)]g''(x)$ and the fact that since ϕ_i has bounded derivatives, as do the ψ_i^m . Observe that the $\alpha(\varepsilon_i)$ may need to be much higher than some $-\frac{(\bar{\phi}_i^{\varepsilon_i})''(x)}{(\bar{\phi}_i^{\varepsilon_i})'(x)}$ since it must be at least the supremum of this over x . ■

A.6 Proofs of results in Section 5

The next result relates to analysis of the ambiguous cheap talk example.

Proof of Proposition 5.1. Since all information sets are on-path under the given strategies, by Theorems A.1 and A.3 it is sufficient to establish that the given strategies form an ex-ante equilibrium. P 's strategy is an ex-ante best response because it leads to payoff 2 for all parameters, which is the highest feasible payoff for this player. Let γ_m be the probability with which agent r plays w after message $m \in \{\alpha, \beta\}$, and similarly let δ_m be the corresponding probabilities for agent c . The proposed strategies correspond to $\gamma_\alpha = \gamma_\beta = \delta_\beta = 1$ and $\delta_\alpha = 0$. We now verify that these are ex-ante best responses. Denoting $\pi_k(IIU) + \pi_k(IID)$ by $\pi_k(II)$, given the strategies of the others, r maximizes

$$\frac{1}{2} \sum_{k=1}^2 \phi_r \left(\pi_k(IU)\gamma_\alpha + 2\pi_k(ID)\gamma_\beta + \pi_k(II)[2\gamma_\beta + 5(1 - \gamma_\beta)] \right).$$

Since this function is strictly increasing in γ_α , it is clearly maximized at $\gamma_\alpha = 1$. The first derivative with respect to γ_β evaluated at $\gamma_\alpha = \gamma_\beta = 1$ is

$$\begin{aligned} & \frac{1}{2} \sum_{k=1}^2 [2\pi_k(ID) - 3\pi_k(II)] \phi_r'(2 - \pi_k(IU)) \\ &= \frac{11}{8} e^{-11 \cdot \frac{39}{20}} \left(e^{-11(\frac{7}{4} - \frac{39}{20})} - \frac{42}{5} \right) > 0, \end{aligned}$$

where the last line uses $\phi_r(x) = -e^{-11x}$ and the values of the π_k . Thus, by concavity in γ_β , the maximum is attained at $\gamma_\alpha = \gamma_\beta = 1$. Similarly, given the strategies of the others, c maximizes

$$\frac{1}{2} \sum_{k=1}^2 \phi_c \left(\pi_k(IU)[2\delta_\alpha + 5(1 - \delta_\alpha)] + \pi_k(ID)[2\delta_\beta + 5(1 - \delta_\beta)] + 2\pi_k(II)\delta_\beta \right).$$

Since this function is strictly decreasing in δ_α , it is clearly maximized at $\delta_\alpha = 0$. The first

derivative with respect to δ_β evaluated at $\delta_\alpha = 0$ and $\delta_\beta = 1$ is

$$\begin{aligned} & \frac{1}{2} \sum_{k=1}^2 [-3\pi_k(ID) + 2\pi_k(II)] \phi'_c(3\pi_k(IU) + 2) \\ &= -\frac{1}{2} \phi'_c\left(\frac{11}{4}\right) + \frac{23}{40} \phi'_c\left(\frac{43}{20}\right) \geq \frac{3}{40} \phi'_c\left(\frac{11}{4}\right) > 0, \end{aligned}$$

where the last line uses the values of the π_k . Since ϕ_c is weakly concave, the problem is weakly concave in δ_β , thus the maximum is attained at $\delta_\alpha = 0$ and $\delta_\beta = 1$. ■

Proof of Proposition 5.2.

Limit attention to strategies for P conditioning only on the payoff relevant component of the parameter, I and II . Denote P 's probability of playing α conditional on the payoff relevant component by ρ_I and ρ_{II} , respectively. Let γ_m be the probability with which r plays w after message $m \in \{\alpha, \beta\}$, and similarly let δ_m be the corresponding probabilities for c . Given ρ_I and ρ_{II} , r chooses $\gamma_\alpha, \gamma_\beta$ to maximize

$$\frac{1}{2} \sum_{k=1}^2 \phi_r \left(\begin{array}{l} \pi_k(I)[\rho_I(1 + \delta_\alpha)\gamma_\alpha + (1 - \rho_I)(1 + \delta_\beta)\gamma_\beta] \\ + \pi_k(II)[\rho_{II}((1 + \delta_\alpha)\gamma_\alpha + 5(1 - \gamma_\alpha)) \\ + (1 - \rho_{II})((1 + \delta_\beta)\gamma_\beta + 5(1 - \gamma_\beta))] \end{array} \right) \quad (\text{A.24})$$

and c chooses $\delta_\alpha, \delta_\beta$ to maximize

$$\frac{1}{2} \sum_{k=1}^2 \phi_c \left(\begin{array}{l} \pi_k(I)[\rho_I((1 + \gamma_\alpha)\delta_\alpha + 5(1 - \delta_\alpha)) \\ + (1 - \rho_I)((1 + \gamma_\beta)\delta_\beta + 5(1 - \delta_\beta))] \\ + \pi_k(II)[\rho_{II}(1 + \gamma_\alpha)\delta_\alpha + (1 - \rho_{II})(1 + \gamma_\beta)\delta_\beta] \end{array} \right). \quad (\text{A.25})$$

The proof proceeds by considering four cases, which together are exhaustive:

Case 1: When $\rho_I = \rho_{II} = 1$ (resp. $\rho_I = \rho_{II} = 0$) so that only one message is sent, for P to always receive the maximal payoff of 2 it is necessary that the agents play w, w with probability 1 after this message, i.e. $\gamma_\alpha = \delta_\alpha = 1$ (resp. $\gamma_\beta = \delta_\beta = 1$). But w is not a best response for c , as can be seen by the fact that the partial derivative of (A.25) with respect to δ_α (resp. δ_β) evaluated at those strategies is

$$\frac{1}{2} (4 - 5 \sum_{k=1}^2 \pi_k(I)) \phi'_c(2) = -\frac{3}{8} \phi'_c(2) < 0.$$

Similarly, one can show that w is not a best response for r .

Case 2: When $0 < \rho_{II} < 1$, since under II , P sends both messages with positive

probability, it is necessary that w, w is played with probability 1 after both messages in order that the principal always receive the maximal payoff of 2. A necessary condition for this to be a best response for c is that the partial derivatives of (A.25) with respect to $\delta_\alpha, \delta_\beta$ are non-negative at $\gamma_\alpha = \gamma_\beta = \delta_\alpha = \delta_\beta = 1$. This is, respectively, equivalent to $14\rho_{II} \geq 19\rho_I$ and $14(1 - \rho_{II}) \geq 19(1 - \rho_I)$, which implies $14 \geq 19$, a contradiction.

Case 3: When $\rho_{II} = 0$ and $0 < \rho_I \leq 1$, (A.25) is strictly decreasing in δ_α , thus the maximum is attained at $\delta_\alpha = 0$. For P to always receive the maximal payoff of 2, it is necessary that $\gamma_\alpha = \gamma_\beta = \delta_\beta = 1$. However, this is not a best response for r because the partial derivative of (A.24) with respect to γ_β evaluated at these strategies using the values for the π_k is,

$$\frac{3}{4}\left(\frac{1}{2} - \rho_I\right)\phi'_r\left(2 - \frac{3}{4}\rho_I\right) + \left(-\frac{1}{5}\rho_I - 1\right)\phi'_r\left(2 - \frac{\rho_I}{5}\right) < 0.$$

To see this, note that the second term is always negative, the first term is non-positive for $\frac{1}{2} \leq \rho_I \leq 1$, and, when $0 < \rho_I < \frac{1}{2}$, substituting $\phi_r(x) = -e^{-11x}$ yields that the left-hand side is negative.

Case 4: When $\rho_{II} = 1$ and $0 \leq \rho_I < 1$, the argument is identical to Case 3 except the roles of the messages α and β are swapped. ■

The next result relates to analysis of the limit pricing example. Denote the entrant's Cournot profit net of entry costs when facing an incumbent of type θ by $w_\theta \equiv b\left(\frac{a+c_\theta-2c_E}{3b}\right)^2 - K$.

Lemma A.5 *Under Assumption 5.1, σ^{LP} is an ex-ante equilibrium if and only if (ICH for I), (ICM for I), $w_H \geq 0$ and*

$$\sum_{\pi} \mu(\pi)(\pi(L)w_L + \pi(M)w_M)\phi'(\pi(H)w_H) \leq 0. \quad (\text{ICL for E})$$

The conditions above correspond to the following incentives in the game: (ICH for I), (ICM for I) were described in the main text, $w_H \geq 0$ ensures that the entrant is willing to enter when it is sure the incumbent is type H , and ICL for E ensures the entrant does not want to enter after observing the monopoly quantity for type L .

Proof of Lemma A.5. Since there is complete information in the final stage, the Cournot or monopoly quantities respectively are ex-ante optimal there. Taking the incumbent's point of view, consider its action in the first stage. Since the incumbent learns its cost before taking any action and there is no other uncertainty, checking ex-ante optimality for the incumbent is equivalent to checking optimality for each incumbent type separately given the entrant's strategy. This is true no matter what the incumbent's ambiguity aversion or beliefs.

When does type H not prefer to pool with M,L at the monopoly quantity for L and thereby deter entry? Profits for H in the conjectured equilibrium are $b(\frac{a-c_H}{2b})^2 + b(\frac{a+c_E-2c_H}{3b})^2$. Profits if it instead pools with M,L at monopoly quantity for L and deters entry are $\frac{a-c_L}{2b}(a - \frac{a-c_L}{2} - c_H) + b(\frac{a-c_H}{2b})^2$. H is at least as well off not pooling if and only if

$$b(\frac{a+c_E-2c_H}{3b})^2 \geq \frac{a-c_L}{2b}(a - \frac{a-c_L}{2} - c_H).$$

This is equivalent to (ICH for I).

When does type M not prefer to produce the monopoly quantity for M and fail to deter entry? Profits for M in the conjectured equilibrium are $\frac{a-c_L}{2b}(a - \frac{a-c_L}{2} - c_M) + b(\frac{a-c_M}{2b})^2$. If it instead produced at the monopoly quantity for M and fails to deter entry, profits are $b(\frac{a-c_M}{2b})^2 + b(\frac{a+c_E-2c_M}{3b})^2$. M is at least as well off pooling with L if and only if

$$\frac{a-c_L}{2b}(a - \frac{a-c_L}{2} - c_M) \geq b(\frac{a+c_E-2c_M}{3b})^2.$$

This is equivalent to (ICM for I).

Type L is playing optimally since its monopoly quantity also deters entry.

It remains to examine the entry decision of the entrant. As a best-response to the incumbent's strategy, ex-ante the entrant wants to maximize

$$\sum_{\pi} \mu(\pi) \phi [\lambda_L(\pi(L)w_L + \pi(M)w_M) + \lambda_H\pi(H)w_H] \quad (\text{A.26})$$

with respect to $\lambda_H, \lambda_L \in [0, 1]$, where λ_H and λ_L are the mixed-strategy probabilities of entering contingent on seeing the monopoly quantity for H and the monopoly quantity for L , respectively. When is this maximized at $\lambda_H = 1$ and $\lambda_L = 0$? Notice, by monotonicity, some maximum involves $\lambda_H = 1$ if and only if $w_H \geq 0$, and $w_H > 0$ is equivalent to $\lambda_H = 1$ being part of every maximum. This says that entering against a known high cost incumbent is profitable. Assuming this is satisfied, so that $\lambda_H = 1$ is optimal, then $\lambda_L = 0$ is optimal if and only if the derivative of (A.26) with respect to λ_L evaluated at $\lambda_L = 0$ and $\lambda_H = 1$ is non-positive, which yields (ICL for E). ■

Before turning to the proof of Proposition 5.3, we remark that we actually prove a slightly stronger result, allowing for the possibility that $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M = 0\}) = 1$ (i.e., that the entrant unambiguously believes that it will exactly break even if it enters conditional on the incumbent's type being in $\{L, M\}$). This appears in the proof only in the proof of Lemma A.6.

Proof of Proposition 5.3. Consider the limit pricing strategy profile σ^{LP} .

By Lemma A.6, under the assumptions of the proposition there exists a $\hat{\phi}$ such that if the entrant's ϕ is at least as concave as $\hat{\phi}$, then (ICL for E) is satisfied. By Lemma A.5, the assumptions of the proposition together with (ICL for E) are sufficient for σ^{LP} to be an ex-ante equilibrium.

Next, we construct an interim belief system that, together with σ^{LP} , satisfies smooth rule consistency. Consider a sequence of completely mixed strategy profiles, σ^k , where $\gamma_{\theta,q}^k > 0$ is the probability that type θ of the incumbent chooses first period quantity q , $\lambda_q^k > 0$ is the probability that the entrant enters after observing quantity q , $\delta_{\theta,(q,\text{enter},r)}^k > 0$ and $\delta_{(q,\text{enter},r)}^k > 0$ are the probabilities of second period quantity r being chosen by, respectively, type θ of the incumbent and the entrant, after observing first period quantity q followed by entry and revelation of θ , and $\delta_{\theta,(q,\text{no entry},r)}^k > 0$ is the probability of second period quantity r being chosen by type θ of the incumbent after observing first period quantity q followed by no entry. Specifically, let $\gamma_{\theta,q}^k \equiv \frac{\beta_{\theta,q}^k}{\sum_{\hat{q} \in \mathcal{Q}} \beta_{\theta,\hat{q}}^k}$ for $k = 1, 2, \dots$, where $\beta_{\theta,q}^k$ is defined by

θ	$q \in \mathcal{Q}$			
	$q = q_H$	$q = q_L$	$q_H \neq q < q_L$	$q > q_L$
L	1	k^2	1	k
M	1	k^2	1	1
H	k^2	1	k	1

, λ_q^k converge to 1 as $k \rightarrow \infty$ when $q < q_L$ and converge to 0 otherwise, $\delta_{\theta,(q,\text{enter},r)}^k$ converge to 1 as $k \rightarrow \infty$ when r is the Cournot quantity for type θ and converge to 0 otherwise, $\delta_{(q,\text{enter},r)}^k$ converge to 1 as $k \rightarrow \infty$ when r is the Cournot quantity for the entrant and converge to 0 otherwise, and $\delta_{\theta,(q,\text{no entry},r)}^k$ converge to 1 as $k \rightarrow \infty$ when r is the monopoly quantity for type θ and converge to 0 otherwise. Note that σ^k converges to σ^{LP} . By Lemma A.4, Theorem A.2 delivers an interim belief system ν such that $(\sigma^{\text{LP}}, \nu)$ satisfies smooth rule consistency.

The final step in the proof is to verify that $(\sigma^{\text{LP}}, \nu)$ satisfies the optimality conditions (3.2) at all information sets. By Theorem 3.3, for optimality, it is sufficient to check against one-stage deviations, and therefore only at information sets where the player has a non-trivial move. The Cournot strategies in the last stage given entry are optimal because all distributions over type become degenerate when conditioned on the entrant learning the incumbent's type. The fact that $w_L < 0$ plus $w_H \geq 0$ implies that it is optimal for the entrant to stay out if its objective function after observing q places all weight on type L and to enter if that objective function places all weight on type H . We now verify that when $q \neq q_L$ this objective function does exactly that when entry/no entry are supposed to occur according to σ^{LP} . Entry is supposed to occur if and only if $q < q_L$. When $q = q_H$, since π_{I_i}

is the degenerate distribution on type H for all π that may be so conditioned, it is optimal to enter. When $q_H \neq q < q_L$, since $\bar{p}_{-i, \sigma_{-i}^{\text{LP}}}(\theta, q|\theta)$ places all weight on (H, q) , (A.12) implies that all π in the support of $\nu_{E, \Theta \times \{q\}}$ puts weight only on (H, q) , and so it is again optimal to enter. Similarly, when $q > q_L$, since $\bar{p}_{-i, \sigma_{-i}^{\text{LP}}}(\theta, q|\theta)$ places all weight on (L, q) , all π in the support of $\nu_{E, \Theta \times \{q\}}$ puts weight only on (L, q) , and so it is optimal not to enter.

Not entering being optimal after observing $q = q_L$ is equivalent (see 4.1) to the following:

$$\sum_{\pi \in \Delta(\Theta \times \{q_L\})} (\pi(L, q_L)w_L + \pi(M, q_L)w_M)\phi'(0)\nu_{E, \{L, M\} \times \{q_L\}}(\pi) \leq 0. \quad (\text{A.27})$$

Using the formula (A.12) to substitute for $\nu_{E, q_L}(\pi)$ in (A.27) yields that not entering remaining optimal is equivalent to (ICL for E). Therefore $(\sigma^{\text{LP}}, \nu)$ satisfies the optimality conditions (3.2) at all information sets as long as the entrant's ϕ is at least as concave as the $\hat{\phi}$ identified from Lemma A.6. For such sufficiently concave ϕ , having shown $(\sigma^{\text{LP}}, \nu)$ is sequentially optimal and satisfies smooth rule consistency, it is therefore an SEA.

Since the only assumption on ϕ made in the above argument that σ^{LP} is part of an SEA was that it was sufficiently concave for the entrant, the argument goes through in its entirety for all $\tilde{\phi}$ at least as concave as ϕ . Furthermore, the same sequence $\{\sigma^k\}_{k=1}^\infty$ may be used for all $\tilde{\phi}$. Thus, σ^{LP} is SEA robust to increased ambiguity aversion.

We next verify that the other conditions in the antecedents of Theorem 4.5 are satisfied. We begin by showing that, for each player, $\sum_{h \in H} u_i(h)p_{\sigma^{\text{LP}}}(h|h^0)\pi(h^0)$ can be strictly ordered across the π in the support of μ . For the entrant,

$$\sum_{h \in H} u_i(h)p_{\sigma^{\text{LP}}}(h|h^0)\pi(h^0) = \pi(H)w_H.$$

Thus, strict ordering corresponds to strict ordering by $\pi(H)$. The assumption that the support of μ can be ordered in the likelihood-ratio ordering ensures the latter, as it implies that for any two distinct $\pi, \pi' \in \text{supp } \mu$, $\pi(H) \neq \pi'(H)$. To see this, suppose to the contrary that $\pi(H) = \pi'(H)$. By distinctness and that weights must sum to one, $\pi(M) \neq \pi'(M)$, $\pi(L) \neq \pi'(L)$ and $\pi(M) > \pi'(M)$ if and only if $\pi(L) < \pi'(L)$, a violation of likelihood-ratio ordering. For the incumbent,

$$\begin{aligned} \sum_{h \in H} u_i(h)p_{\sigma^{\text{LP}}}(h|h^0)\pi(h^0) &= \pi(L)2b\left(\frac{a - c_L}{2b}\right)^2 \\ &\quad + \pi(M)\left[\frac{a - c_L}{2b}\left(a - \frac{a - c_L}{2} - c_M\right) + b\left(\frac{a - c_M}{2b}\right)^2\right] \\ &\quad + \pi(H)\left[b\left(\frac{a - c_H}{2b}\right)^2 + b\left(\frac{a + c_E - 2c_H}{3b}\right)^2\right]. \end{aligned}$$

By Assumption 5.1 and (ICM for I), the expression multiplied by $\pi(L)$ is strictly larger than the one multiplied by $\pi(M)$, which is, in turn, strictly larger than the one multiplied by $\pi(H)$. Thus, likelihood-ratio ordering of the support of μ implies strict ordering of $\sum_{h \in H} u_i(h) p_{\sigma^{\text{LP}}}(h|h^0) \pi(h^0)$. By Theorem 4.5, ambiguity aversion makes σ^{LP} SEA belief robust. ■

Lemma A.6 *Under the assumptions of Proposition 5.3 there exists an $\alpha > 0$ such that if ϕ is at least as concave as $-e^{-\alpha x}$ then (ICL for E) is satisfied.*

Proof. Assume the conditions of the proposition. We show that (ICL for E) is satisfied for concave enough ϕ . The assumption in the proposition that some $\pi \in \text{supp } \mu$ makes entry conditional on $\{L, M\}$ strictly unprofitable means $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M < 0\}) > 0$. If $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M \leq 0\}) = 1$ then (ICL for E) is trivially satisfied for any ϕ . For the remainder of the proof, therefore, suppose that $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M > 0\}) > 0$. Let $\Pi^- \equiv \{\pi \mid \pi(L)w_L + \pi(M)w_M < 0\}$, $\Pi^+ \equiv \{\pi \mid \pi(L)w_L + \pi(M)w_M > 0\}$, $N \equiv \sum_{\pi \in \Pi^-} \mu(\pi)(\pi(L)w_L + \pi(M)w_M)$, and $P \equiv \sum_{\pi \in \Pi^+} \mu(\pi)(\pi(L)w_L + \pi(M)w_M)$. Let $\pi^- \in \arg \max_{\pi \in \Pi^-} \pi(H)$ and $\pi^+ \in \arg \min_{\pi \in \Pi^+} \pi(H)$. The left-hand side of (ICL for E) can be bounded from above as follows:

$$\begin{aligned} & \sum_{\pi \in \Pi^-} \mu(\pi)(\pi(L)w_L + \pi(M)w_M) \phi'(\pi(H)w_H) + \sum_{\pi \in \Pi^+} \mu(\pi)(\pi(L)w_L + \pi(M)w_M) \phi'(\pi(H)w_H) \\ & \leq \sum_{\pi \in \Pi^-} \mu(\pi)(\pi(L)w_L + \pi(M)w_M) \phi'(\pi^-(H)w_H) + \sum_{\pi \in \Pi^+} \mu(\pi)(\pi(L)w_L + \pi(M)w_M) \phi'(\pi^+(H)w_H) \\ & = N \phi'(\pi^-(H)w_H) + P \phi'(\pi^+(H)w_H). \end{aligned}$$

Consider $\phi(x) = -e^{-\alpha x}$, $\alpha > 0$. The upper bound above becomes

$$\alpha N e^{-\alpha \pi^-(H)w_H} + \alpha P e^{-\alpha \pi^+(H)w_H}.$$

We show that this upper bound is non-positive for sufficiently large α , implying (ICL for E). The upper bound is non-positive if and only if $P e^{-\alpha \pi^+(H)w_H} \leq -N e^{-\alpha \pi^-(H)w_H}$ if and only if $e^{\alpha(\pi^-(H) - \pi^+(H))w_H} \leq -\frac{N}{P}$ if and only if $\alpha(\pi^-(H) - \pi^+(H))w_H \leq \ln(-\frac{N}{P})$. Since $\pi^-(L)w_L + \pi^-(M)w_M < 0 < \pi^+(L)w_L + \pi^+(M)w_M$ and $c_L < c_M$, we have $w_L < 0 < w_M$. Thus, $\frac{\pi^-(L)}{\pi^-(M)} > -\frac{w_M}{w_L} > \frac{\pi^+(L)}{\pi^+(M)}$. By our assumption on the support of μ and Lemma A.7, $\frac{\pi^-(L)}{\pi^-(M)} > \frac{\pi^+(L)}{\pi^+(M)}$ implies $\pi^-(H) < \pi^+(H)$. Therefore, $\alpha(\pi^-(H) - \pi^+(H))w_H \leq \ln(-\frac{N}{P})$ if and only if $\alpha \geq \frac{\ln(-\frac{N}{P})}{(\pi^-(H) - \pi^+(H))w_H}$.

To complete the proof, fix α satisfying this inequality and consider ϕ such that $\phi(x) = h(-e^{-\alpha x})$ for all x with h concave and strictly increasing on $(-\infty, 0)$. We show that (ICL

for E) holds. Observe that $\phi'(x) = h'(-e^{-\alpha x})\alpha e^{-\alpha x}$. Since $\pi^-(H) - \pi^+(H) < 0$ and $w_H > 0$, we have

$$-e^{-\alpha\pi^-(H)w_H} \leq -e^{-\alpha\pi^+(H)w_H}$$

and, by concavity of h ,

$$h'(-e^{-\alpha\pi^-(H)w_H}) \geq h'(-e^{-\alpha\pi^+(H)w_H}).$$

Therefore the upper bound derived above satisfies

$$\begin{aligned} & N\phi'(\pi^-(H)w_H) + P\phi'(\pi^+(H)w_H) \\ = & \alpha N e^{-\alpha\pi^-(H)w_H} h'(-e^{-\alpha\pi^-(H)w_H}) + \alpha P e^{-\alpha\pi^+(H)w_H} h'(-e^{-\alpha\pi^+(H)w_H}) \\ \leq & (\alpha N e^{-\alpha\pi^-(H)w_H} + \alpha P e^{-\alpha\pi^+(H)w_H}) h'(-e^{-\alpha\pi^-(H)w_H}) \leq 0 \end{aligned}$$

by the first part of the proof and the assumption on α . This implies (ICL for E). ■

Lemma A.7 *If the support of μ can be ordered in the likelihood-ratio ordering, then, for any $\pi, \pi' \in \text{supp } \mu$, $\frac{\pi(L)}{\pi(M)} > \frac{\pi'(L)}{\pi'(M)}$ implies $\pi(H) < \pi'(H)$.*

Proof. Suppose the support of μ can be so ordered. Fix any $\pi, \pi' \in \text{supp } \mu$. Suppose $\frac{\pi(L)}{\pi(M)} > \frac{\pi'(L)}{\pi'(M)}$. Then $\frac{\pi'(L)}{\pi(L)} < \frac{\pi'(M)}{\pi(M)}$, and thus, by likelihood-ratio ordering, $\frac{\pi'(L)}{\pi(L)} < \frac{\pi'(M)}{\pi(M)} \leq \frac{\pi'(H)}{\pi(H)}$. This implies $\pi'(H) > \pi(H)$ since the last two ratios cannot be less than or equal to 1 without violating the total probability summing to 1. ■

B Appendix: Details on the analysis of the game in Figure 3.1 and the comparison with no profitable one-stage deviations and consistent planning

A strengthening of no profitable one-stage deviations used in some of the existing literature investigating games with ambiguity is the following condition, describing a consistent planning requirement in the spirit of Strotz (1955-56) (for a formal decision theoretic treatment see Siniscalchi 2011):

Definition B.1 *Fix a game Γ and a pair (σ, ν) consisting of a strategy profile and interim belief system. Specify V_i and V_{i,I_i} as in (2.1) and (2.4). For each player i and information set $I_i \in \mathcal{I}_i^T$, let*

$$CP_{i,I_i} \equiv \operatorname{argmax}_{\hat{\sigma}_i \in \Sigma_i} V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i}).$$

Then, inductively, for $0 \leq t \leq T - 1$, and $I_i \in \mathcal{I}_i^t$ let

$$CP_{i,I_i} \equiv \operatorname{argmax}_{\hat{\sigma}_i \in \bigcap_{\hat{I}_i \in \mathcal{I}_i^{t+1} | \hat{I}_i^{-1} = I_i} CP_{i,\hat{I}_i}} V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i}).$$

Finally, let

$$CP_i \equiv \operatorname{argmax}_{\hat{\sigma}_i \in \bigcap_{\hat{I}_i \in \mathcal{I}_i^0} CP_{i,\hat{I}_i}} V_i(\hat{\sigma}_i, \sigma_{-i}).$$

(σ, ν) is optimal under consistent planning if, for all players i ,

$$\sigma_i \in CP_i.$$

Equivalently, (σ, ν) is such that for all players i ,

$$V_i(\sigma) \geq V_i(\hat{\sigma}_i, \sigma_{-i}) \text{ for all } \hat{\sigma}_i \in \bigcap_{\hat{I}_i \in \mathcal{I}_i^0} CP_{i,\hat{I}_i}$$

and, for all information sets $I_i \in \mathcal{I}_i^t$, $0 \leq t \leq T - 1$,

$$V_{i,I_i}(\sigma) \geq V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i}) \text{ for all } \hat{\sigma}_i \in \bigcap_{\hat{I}_i \in \mathcal{I}_i^{t+1} | \hat{I}_i^{-1} = I_i} CP_{i,\hat{I}_i}$$

and, for all information sets $I_i \in \mathcal{I}_i^T$,

$$V_{i,I_i}(\sigma) \geq V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i}) \text{ for all } \hat{\sigma}_i \in \Sigma_i.$$

If (σ, ν) is sequentially optimal then it is also optimal under consistent planning. However, if (σ, ν) is optimal under consistent planning it may fail to be sequentially optimal (even when limiting attention to ambiguity neutrality). For such a failure to occur, the optimal strategy from player i 's point of view at some earlier stage must have a continuation that fails to be optimal from the viewpoint of some later reachable stage. This is what makes the extra constraints imposed in the optimization inequalities under consistent planning bind. Just as with no profitable one-stage deviations, when updating is according to the smooth rule, (σ, ν) optimal under consistent planning implies (σ, ν) is sequentially optimal, making the three equivalent under smooth rule updating.

Recall that the example in Figure 3.1 in the Introduction showed how the no profitable one-stage deviation criterion under Bayesian updating allowed strategy profiles that are

not even ex-ante (Nash) equilibria of a game (and thus clearly not sequentially optimal). Replacing no profitable one-stage deviations by consistent planning does not change this fact. It is easy to specify ϕ_1 , μ and an interim belief system for player 1 such that o can be played with positive probability while satisfying consistent planning and no profitable one-stage deviations. For example, this is the case if $\phi_1(x) = -e^{-10x}$, μ is $1/2$ on $(1/3, 1/9, 5/9)$ and $1/2$ on $(1/3, 5/9, 1/9)$, and 1's beliefs after seeing U are given by Bayes' rule applied to μ : $1/3$ on $(3/4, 1/4, 0)$ and $2/3$ on $(3/8, 5/8, 0)$. With these parameters and beliefs, the following strategy profile satisfies no profitable one-stage deviations and consistent planning: player 1 plays o with probability $1 - \frac{9}{20} \ln(\frac{29}{11}) \approx 0.564$ and mixes evenly between u and d if U , while player 2 plays her strictly dominant strategy if given the move. Notice, if we consider any more concave ϕ_1 , playing o with even higher probability will be consistent with consistent planning or no profitable one-stage deviations given these beliefs. In the limit where the decision maker is Maxmin EU with set of priors equal to the convex combinations of $(1/3, 1/9, 5/9)$ and $(1/3, 5/9, 1/9)$ and applies Bayes' rule to each measure in the set, playing o with probability 1 is consistent with consistent planning and no profitable one-stage deviations.

School of Economics and Finance



This working paper has been produced by
the School of Economics and Finance at
Queen Mary University of London

Copyright © 2018 Eran Hanany, Peter
Klibanoff and Sujoy Mukerji all rights
reserved

School of Economics and Finance
Queen Mary University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 7356
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/research/workingpapers/