

Incomplete Information Games with Ambiguity Averse Players[†]

By ERAN HANANY, PETER KLIBANOFF, AND SUJOY MUKERJI*

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. (JEL C73, D81, D83)

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 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 modeling 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.

*Hanany: Faculty of Engineering, Tel Aviv University, Tel Aviv 69978, Israel (email: hananye@tau.ac.il); Klibanoff: Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston, IL 60208 (email: peterk@kellogg.northwestern.edu); Mukerji: School of Economics and Finance, Queen Mary University of London, London E1 4NS, United Kingdom (email: s.mukerji@qmul.ac.uk). John Asker was coeditor for this article. 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.

[†]Go to <https://doi.org/10.1257/mic.20180302> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

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

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, and Mukerji 2005). The presence of ambiguity is captured by nondegeneracy 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 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 toward 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 their strategies an unpredictably changing nature that Bewley (1988, 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's (2012) 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 IVB) 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's (1982) 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 noncredible 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's (1982) 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 exists for any finite multistage game with perfect recall and incomplete information, and for any specification of players' ambiguity aversions and ex ante beliefs.

Section IIIA provides results on comparative statics of the equilibrium set in ambiguity aversion that apply to any of the three notions of equilibrium mentioned above. First, for fixed beliefs, ambiguity aversion may change the equilibrium set

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 Klibanoff (2009); Ghirardato and Siniscalchi (2012).

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, 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 IIIB defines robustness of an equilibrium to increases in ambiguity aversion and shows that this is related to a type of belief robustness (Theorem 9).

Section IV contains an example of a Milgrom and Roberts-style (1982) 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 9 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 IV 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. Our analysis of this example establishes that increasing the ambiguity of communication 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 V discusses some possible extensions of our approach, including to other models of ambiguity averse players' preferences. Finally, in addition to the discussion of alternative approaches following Theorem 2 in Section IIB, Section VI discusses closely related literature (especially Battigalli et al. 2019 and Pahlke 2018) and compares it to our theory. The print Appendix contains all proofs for results from Section II and some further analysis. The proofs for results from the remaining sections and some further analysis are contained in an online Appendix.

I. 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 $h^t \in H^t$, $A_i(h^t) \equiv \{\hat{h}_{t,i} \mid \hat{h}^t \in H, \hat{h}^t = h^t\}$ 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 $h^t, \hat{h}^t \in H^t$, $\hat{h}^t \in \mathcal{I}_i^t(h^t)$ implies $A_i(h^t) = A_i(\hat{h}^t)$ (where $\mathcal{I}_i^t(h^t)$ 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 $h^t, \hat{h}^t \in H^t$, $\mathcal{I}_i(h^t) = \mathcal{I}_i(\hat{h}^t)$ 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, $R_i(\bar{h}^t) \equiv ((\mathcal{I}_i(\bar{h}^s), \bar{h}_{s,i})_{0 \leq s < t}, \mathcal{I}_i(\bar{h}^t))$. 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 7) 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 .

The first three bullet points describe the *game form*, while the rest describe *preferences*. Observe that at each partial history, each player (not just those with nontrivial moves at this point) has an information set. The nonstandard 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 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 nontrivial move, which is either to c (cooperate) or to d (defect) in the negotiations. If player 1 plays c , then player 2 observes this and may choose either to c (cooperate) or to d (defect) 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

⁶ $\text{supp}(\cdot)$ denotes the support of a measure.




FIGURE 1. A PEACE NEGOTIATION GAME

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

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 half-half 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, we have the following definition.

DEFINITION 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 $p_{i,\sigma_i}(h^t|h^r) \equiv \prod_{r \leq s < t} \sigma_i(I_i(h^s))(h_{s,i})$ 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

⁷To formally write some histories, one may need to include some of the “dummy” actions chosen from singleton action sets needed to create the multistage structure of the game. These actions are not depicted in the game tree. For example, if player 1 plays d , to keep the multistage 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 nontrivial 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.

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 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

$$(1) \quad V_i(\sigma'_i, \sigma_{-i}) \equiv \sum_{\pi \in \Delta(\Theta)} \phi_i \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi(h^0) \right) \mu_i(\pi).$$

If there were only a single distribution π over parameters, and so no ambiguity, μ_i would be degenerate and ϕ_i irrelevant and (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 nonlinear 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 (1), σ_i , being 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:

$$(2) \quad \sum_{h \in H} u_i(h) p_{i, \sigma'_i}(h|h^0) q^{(\sigma, \mu_i), i}(h),$$

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,

$$(3) \quad p_{-i, \sigma_{-i}}(h|h^0) \sum_{\pi \in \Delta(\Theta)} \pi(h^0) \phi'_i \left(\sum_{\hat{h} \in H} u_i(\hat{h}) p_{\sigma}(\hat{h}|\hat{h}^0) \pi(\hat{h}^0) \right) \mu_i(\pi).$$

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, toward 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. This is defined formally as follows.

DEFINITION 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 (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*

$$(4) \quad V_{i,I_i}(\sigma'_i, \sigma_{-i}) \equiv \sum_{\pi \in \Delta(I_i)} \phi_i \left(\sum_{h|h^t \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^t) \pi(h^t) \right) \nu_{i,I_i}(\pi),$$

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

Compared to the ex ante preferences given in (1), the preferences (4) at I_i differ only in that 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 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:

$$(5) \quad \sum_{h|h^t \in I_i} u_i(h) p_{i,\sigma_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h),$$

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

$$(6) \quad p_{-i,\sigma_{-i}}(h|h^t) \sum_{\pi \in \Delta(I_i)} \phi_i \left(\sum_{\hat{h}|\hat{h}^t \in I_i} u_i(\hat{h}) p_{\sigma}(\hat{h}|\hat{h}^t) \pi(\hat{h}^t) \right) \pi(h^t) \nu_{i,I_i}(\pi).$$

Now, i ’s ambiguity aversion leads the marginal of this local measure on I_i to tilt toward 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 II 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 assumes no particular connection between beliefs at different information sets, the ex ante beliefs, and the strategy profile. Our Theorem 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).

II. Equilibrium

A. Ex Ante Equilibrium

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

DEFINITION 3 (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 (2014); Kellner (2015); and Grant, Meneghel, and Tourky (2016).

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's (1974) 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. (2019).

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 (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 $a \geq \ln(5/4)$), peace with probability 1 is an ex ante equilibrium outcome. Formally, we have the following result.

PROPOSITION 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 player 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 (3) that ambiguity aversion would then, at this strategy profile, tilt the marginals of their local measures on parameters in different directions—1's toward I and 2's toward 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 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.

B. 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's (1982) *sequential rationality*. Sequential optimality is defined formally as follows.

DEFINITION 4: 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 ,

$$(7) \quad V_i(\sigma) \geq V_i(\sigma'_i, \sigma_{-i})$$

and

$$(8) \quad V_{i,I_i}(\sigma) \geq V_{i,I_i}(\sigma'_i, \sigma_{-i})$$

for all $\sigma'_i \in \Sigma_i$, where the V_i and V_{i,I_i} are as specified in (1) and (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 noncredible 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 10 in Appendix AB). How does sequential optimality relate to familiar concepts? Sequential optimality with a common μ implies (and is stronger

⁹Note that since $V_{i,I_i}(\tilde{\sigma}_i, \sigma_{-i}) = V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i})$, if $\tilde{\sigma}_i^I = \hat{\sigma}_i^I$, 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.

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 (8) 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, (7) ensures σ is an ex ante equilibrium. When preferences are expected utility and there is a common μ , we show (see Corollary 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 nonsingleton 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 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 noncredible 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.

PROPOSITION 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.*

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 nontrivial move at that point), ex ante equilibrium plus sequential rationality need not imply subgame perfection (see, e.g., Mas-Colell et al. 1995, 290, 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.

REMARK 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 $\phi_1(x) \equiv -e^{-x \ln(5/3)} = -(5/3)^{-x}$ (so that $\phi_1'(x)/\phi_1'(y) = (5/3)^{y-x}$) and player 2 is ambiguity neutral, then, as is shown in the proof of Proposition 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 p_1 and p_2 . 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.

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

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 (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 nonconstant 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.

¹¹ More 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).

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

We defer further discussion of the rule and suggest that the reader may also wish to defer parsing its formal definition.

NOTATION 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 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 5: 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., $t = s(I_i)$):

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

$$(9) \quad \nu_{i,I_i}(\pi) \propto \sum_{\hat{\pi} \in \Delta(\Theta) | \hat{\pi}_{I_i} = \pi} \frac{\phi'_i(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \hat{\pi}(h^0))}{\phi'_i(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi(h^t))} \hat{\pi}(I_i) \mu_i(\hat{\pi}),$$

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

$$(10) \quad \nu_{i,I_i}(\pi) \propto \sum_{\hat{\pi} \in \Delta(I_i^{-1}) | \hat{\pi}_{I_i} = \pi} \frac{\phi'_i(\sum_{h|h^{t-1} \in I_i^{-1}} u_i(h) p_\sigma(h|h^{t-1}) \hat{\pi}(h^{t-1}))}{\phi'_i(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi(h^t))} \times \left(\sum_{h^t \in I_i} p_{-i,\sigma_{-i}}(h^t|h^{t-1}) \hat{\pi}(h^{t-1}) \right) \nu_{i,I_i^{-1}}(\hat{\pi}),$$

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 (10) 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 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 (5). Second, perfect recall and the linearity of this objective allows us to conclude that σ_i maximizes (5) implies that, for each J_i , σ_i maximizes the part of the summation in (5) taken only over J_i . Finally, the smooth rule formula in (10) 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 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 it is “non-consequentialist” in the sense that its formula includes expected payoffs under σ at all (terminal) histories reachable from the immediately preceding information set, as well as ambiguity aversion via ϕ_i . Importantly however, it is consequentialist in the weaker sense 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 1 says that this belief *may be determined by updating according to the smooth rule using* $\sigma = (\sigma_i, \sigma_{-i})$.

Non-consequentialist updating more generally has been criticized in the context of dynamic decision-making under ambiguity by, e.g., Siniscalchi (2009) as violating the spirit of ambiguity being purely an “informational” phenomenon, in the sense that it may generate updated perceptions of ambiguity that depend on contextual factors such as payoffs and attitudes. While there are types of analysis that become problematic given non-consequentialism, such as studies of learning or inference carried out in isolation, separate from any decision or game

context, an important takeaway from our analysis (including Theorem 2 below on the sufficiency of checking one-stage deviations) is that the non-consequentialism of smooth rule updating doesn't introduce any new conceptual difficulties into the analysis of dynamic games. For example, just as with Bayesian updating, once one fixes some beliefs at an information set, to proceed with analysis of play at that and successor information sets does not require any knowledge of parts of the game outside of this continuation sub-tree.

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

COROLLARY 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's (1982) 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 6: *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 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.*

¹³Shimoji 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.

Comparison with an Alternative Approach.—Sequential optimality refines ex ante equilibrium to deal with noncredible threats by requiring players to be optimizing over continuation strategies at each information set given the strategies of the other players. It follows that for any strategy profile that is not sequentially optimal, there must be at least some information set at which some player could strictly improve by shifting to a different continuation strategy. However, this has not been the approach to noncredible threats generally taken in the small body of existing literature on dynamic games with ambiguity aversion that addresses the issue. These alternative approaches instead require no profitable one-stage deviations (or the slightly stronger Strotzian consistent planning, see online Appendix B4) with respect to *exogenously imposed* particular update rules.^{14,15} They start from the principle that, at any information set, players choose only current actions (or mixtures over actions) rather than continuation strategies. Thus, when best responding, players take as given not only other players' strategies, as in our approach, but also their *own* future continuation strategy. While such approaches emphasize the conflicting views of optimality at different information sets that may arise for a player under ambiguity aversion when using particular update rules, sequential optimality emphasizes the unity of a player in requiring that strategies are optimal from the point of view of all of the player's information sets simultaneously, while taking a more agnostic or endogenous view of updating.

How do such approaches compare with sequential optimality? Without exogenous assumptions on updating, sequential optimality is a stronger requirement than consistent planning, which is, in turn, stronger than requiring no profitable one-stage deviations. It then follows from Theorem 2 that all three concepts are equivalent under strong smooth rule updating. This finding generalizes the fact that under expected utility, the three are equivalent under strong Bayes' rule updating. However, the no profitable one-stage deviations and similar approaches under ambiguity aversion have been most commonly applied together with different updating. In the context of smooth ambiguity preferences, Bayesian updating is often assumed. As noted earlier, because of the conflicts this generates between preferences at a player's different information sets, this may rule out some (or all) sequential optima. This may occur even when all information sets are on-path. Thus, differently from sequential optimality, this approach may eliminate an ex ante equilibrium by restricting on-path as well as off-path behavior.

Such approaches also differ from sequential optimality in the other direction—under Bayesian updating, some strategy profiles that are not even ex ante equilibria may satisfy no profitable one-stage deviations. One type of such behavior that has been a source of criticism often discussed in the literature (see, e.g., Machina 1989, Siniscalchi 2011) is a player in a dynamic decision problem strictly preferring not to obtain freely available information. As Siniscalchi (2011) explains, information is typically bundled together with the future ability to condition one's

¹⁴Examples of such literature include Bose and Renou (2014); Kellner and Le Quement (2018); Battigalli et al. (2019); and Beauchêne, Li, and Li (2019).

¹⁵A second approach appearing in the literature is based on recursive preferences. We relate sequential optimality to this approach in the latter part of Section VI.

action on that information, and therefore increases the flexibility given to the player at future information sets. When there are conflicts between the current and future objectives of the player, such flexibility can be strictly costly for the player at the current information set. Thus what might appear to be negative value of information is reinterpreted as valuable information outweighed by a costly increase in future flexibility (or, as more commonly referred to, loss of commitment power). Such behavior cannot occur under sequential optimality, as future flexibility is never costly given the unified agreement on optimality it entails. The following example illustrates another consequence of this difference between the approaches.

Example: Consider the game in Figure 2. 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*. Let $\phi_1(x) = -e^{-10x}$ and the common μ be half-half on $(1/3, 1/9, 5/9)$ and $(1/3, 5/9, 1/9)$. Then the unique strategy profile satisfying no profitable one-stage deviations combined with updating according to strong Bayes' rule has 1 playing *o* with positive probability and then the mixture $(1/2)u + (1/2)d$ if given the move, and 2 playing her dominant strategy. However, the unique sequential optimum is player 1 playing *i* and then *d* if given the move, together with 2 playing her dominant strategy.

What drives this difference in 1's behavior? Playing *i* rather than *o* gives player 1 flexibility at the final information set. With Bayesian updating, 1 uses this flexibility to play $(1/2)u + (1/2)d$ rather than *d*, and this is costly from the perspective of player 1 at the initial information set. Since, for some θ , 1's payoffs from *o* are higher than those from *i* followed by $(1/2)u + (1/2)d$, and the reverse is true for other θ , 1 is motivated to hedge against ambiguity by mixing between *o* and *i* at the initial information set. The details may be found in online Appendix B4.¹⁶ In contrast, observe that for player 1, given 2's dominant strategy, the payoff to playing *i* followed by *d* is, θ -by- θ , strictly higher than the payoff to playing *o*. Thus, no strategy involving playing *o* with positive probability can be an ex ante best reply to 2's optimal strategy no matter how player 1 perceives and treats the uncertainty about θ . This is why *o* is not part of any ex ante equilibrium, let alone a sequential optimum.

More generally, sequential optimality always rules out the play of type-by-type (iteratively) strictly dominated strategies in dynamic settings while no profitable one-stage deviations combined with Bayesian updating does not.

As observed in this section, an ambiguity averse player who exogenously adopts, for example, Bayesian updating will generally have disagreement (across information sets) on the optimal strategy, and therefore may suffer from future flexibility. In contrast, a player who, either ex ante or at an early information set, both recognizes updating as the generator of these costly disagreements and is able to influence *their own* future information processing would generally be better off departing

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.




FIGURE 2. GAME CONTRASTING THE APPROACHES

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

from Bayesian updating in order to remove these difficulties. One possible view of sequential optimality is as modeling the outcomes of strategic interaction of such players.

C. 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’s (1982) *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 1, since, by that result, sequential optima are as if players are responding

¹⁷If there are no such information sets, any sequentially optimal strategy profile is also part of an SEA (see Theorem 12).

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.

DEFINITION 7 (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 8 (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 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: 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 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: 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 4: *An SEA exists for any game Γ .*

In Theorem 11 in the Appendix, 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.

III. Effects of Ambiguity Aversion on Equilibria

A. 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 ex ante 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 5: *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 6: *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 7: *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) as follows:

$$(11) \quad \sum_{\pi \in \Delta(\Theta)} \left(\sum_{h \in H} u_i(h) p(\sigma'_i, \sigma_{-i})(h|h^0) \pi(h^0) \right) \phi'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0) \right) \mu_i(\pi).$$

Equation (11) 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 (11) 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 6, this is what may generate new equilibria. In Theorem 7, 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 7 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 (2017). 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, 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. (2019) 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 8: *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.*

B. 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 $\mu(\pi_1) = \mu(\pi_2) = 1/2$, where $\pi_1(\theta_1) = 2/3$ and $\pi_2(\theta_1) = 1/2$, 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 in Figure 3.

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 $(1/2)\phi(2\lambda) + (1/2)\phi(4 - \lambda)$, 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 $\alpha > \ln((1 + \sqrt{5})/2) \approx 0.48$, 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 $\mu(\pi_1) \geq \phi'(3)/(2\phi'(2) + \phi'(3))$. Notice that as ϕ becomes more concave, $\phi'(3)/(2\phi'(2) + \phi'(3))$ decreases, approaching 0 in the limit as $\phi'(3)/\phi'(2)$ 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

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

FIGURE 3. ROBUSTNESS EXAMPLE

supporting set of beliefs may be made as large as desired (see Theorem 9 and Remark 2). 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 9: *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 10: *For a game Γ , consider an equilibrium σ . Ambiguity aversion makes σ belief robust if, for each i and $\varepsilon_i \in (0, 1/|\text{supp}(\mu_i)|)$, 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 9: *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 9 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 2: Suppose ϕ_i is twice continuously differentiable with $\phi_i' > 0$ everywhere. Then Theorem 9 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 2 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, 1865–66).

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 IVB, we apply Theorem 9 in analyzing the effect of the entrant's ambiguity aversion on the robustness of limit-pricing equilibria.

IV. Examples

A. 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) for a principal communicating to two agents. The equilibrium we identify would not be a solution under no profitable one-stage deviations (or Strotzian consistent planning) together with Bayesian updating. It is only under the latter approach that existing literature (see, e.g., Bose and Renou 2014; Beauchêne, Li, and Li 2019; Kellner and Le Quement 2018) has been able to establish, through arguments relying in an essential way on violations of sequential

optimality, a value for ambiguous cheap talk. As was discussed following Theorem 2 in Section IIB, such approaches may lead ambiguity averse players to reject freely available, relevant information, and thus it is not clear whether the value of ambiguous communication could survive in the absence of such an effect. Thus, our analysis of this example establishes a new reason why ambiguous cheap talk can be valuable: like the ambiguity about payoff-relevant actions in our running example, increasing the ambiguity of the communication can enhance the ability to provide incentives to both agents simultaneously by moving their effective beliefs further apart in the desired directions.

There are three players, a principal, P , and two agents, r (row) and c (column). Principal P wishes to induce the agents to w (work) for him rather than start their own b (businesses). 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 . Principal 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 Figure 4, where each cell lists the payoff to P , r , and c in that order.

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 half-half over distributions π_1 and π_2 given in Figure 5 (where the numbers on the top and left edges are the corresponding marginals). 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

<i>I</i>	<i>b</i>	<i>w</i>
<i>b</i>	0, 0, 5	0, 0, 1
<i>w</i>	2, 1, 5	2, 2, 2

<i>II</i>	<i>b</i>	<i>w</i>
<i>b</i>	0, 5, 0	0, 5, 1
<i>w</i>	0, 1, 0	2, 2, 2

FIGURE 4. PAYOFFS FOR AMBIGUOUS CHEAP TALK EXAMPLE

		3/4	1/4
	π_1	<i>I</i>	<i>II</i>
1/3	<i>U</i>	3/12	1/12
2/3	<i>D</i>	6/12	2/12

		1/5	4/5
	π_2	<i>I</i>	<i>II</i>
1/4	<i>U</i>	1/20	4/20
3/4	<i>D</i>	3/20	12/20

FIGURE 5. BELIEFS FOR AMBIGUOUS CHEAP TALK EXAMPLE

scientific principle that is not well understood, which influences both the functioning of the technology (*I* versus *II*) as well as the findings of the laboratory experiment (*U* versus *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 4), and, that the principal does strictly better than if he were not allowed to use the payoff-irrelevant component (Proposition 5). Formally, we have the following results.

PROPOSITION 4: *The strategy profile σ^* is an SEA. In this equilibrium, *P* attains his maximum possible payoff for each parameter.*

REMARK 3: *The strategy profile σ^* remains an equilibrium for any ϕ_r more concave than the one in the example.*

PROPOSITION 5: *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 versus D . To understand this, first note that the best strategy, $\hat{\sigma}_P^*$, for P that does not depend on U versus D is: if I then send message α with probability $\rho \approx 0.267$, and otherwise send message β . One can then show that the only ambiguity that is relevant to r 's payoffs under each of these strategies is that concerning the event that the message α is sent. This event is assigned probabilities $\pi_k(I)\rho$ under $\hat{\sigma}_P^*$, and $\pi_k(I)\pi_k(U)$ under σ_P^* . As $\pi_1(U) > \rho > \pi_2(U)$, there is more ambiguity under σ_P^* . This additional ambiguity about the event where α is sent 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 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 4: *If agent r becomes sufficiently more ambiguity averse, Proposition 5 no longer holds: in addition to the equilibrium in Proposition 4, there will be an equilibrium where P conditions his play only on I versus 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 versus D is no longer needed.*

B. Example: Limit Pricing under Ambiguity

In this section, we use a parametric class of games based on Milgrom and Roberts' (1982) 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

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

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 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.

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 3: $a, b > 0$, $K \geq 0$, $c_H > c_M > c_L \geq 0$, $c_E \geq 0$, $a > c_H$, $a + c_E - 2c_H > 0$ 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 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

¹⁹The use of at least three costs is necessary to have nontrivial 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.

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 6: *Under Assumption 3, the limit pricing strategy profile σ^{LP} is part of an SEA if*

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

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

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

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 11 to establish that (σ^{LP}, ν) is an SEA and uses Theorem 9 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.

V. Extensions

A. 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 is essentially what Cerreia-Vioglio et al. (2011) call Uncertainty Averse preferences, and includes smooth ambiguity preferences along with many other models from the literature, some of which are recursive. Note that Maxmin Expected Utility (Gilboa and Schmeidler 1989) is a subclass of Uncertainty Averse preferences, 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 an 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 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 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.

B. Implementation of Mixed Actions

Players choose behavior strategies, which, for each type of the player, specify a mixture over the available actions at each information set. 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 4 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 4 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.

C. 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.

VI. 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

There have been very few papers investigating general dynamic games with incomplete information and ambiguity aversion. The two most closely related to ours are Battigalli et al. (2019) and Pahlke (2018). Battigalli et al. (2019) 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, by building on self-confirming equilibrium they are able to tightly link their solution concept to steady-state learning. In contrast, the Nash-like and stronger equilibrium notions we build on are not as linked to learning foundations. While the Introduction offered some thoughts about how restrictions on learning opportunities might relate to our equilibria, whether and how such ideas can be formalized and their lessons for the equilibrium concepts we explore is left for future research. Second, 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. Third, 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 Fourth, 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. (2019) 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 8 showing

²¹In addition to the papers we mention when discussing *ex ante* equilibrium in Section IIA, 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, Katuščák, and Ozdenoren (2007); Lopomo, Rigotti, and Shannon (2014); Bodoh-Creed (2012); Wolitzky (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); Muraviev, Riedel, and Sass (2017); and Eichberger, Grant, and Kelsey (2018).

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.

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 IIIA). 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 7).

Subsequent to our paper, Pahlke (2018) explored 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 nonstandard 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 $\mathcal{P} \subseteq \Delta(\Theta)$ that 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 nonsingleton, i.e., when there is nontrivial ambiguity. This heterogeneity explains how Pahlke's approach overcomes the specialization of the results of Ellis (2018) to MEU preferences, which imply 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.²³ We make four observations comparing Pahlke (2018) to our framework and approach when adapted to MEU preferences (see Section VA). 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 (2019); Auster and Kellner (2018), and all focus on behavior that cannot occur under ambiguity neutrality. Their approach is to assume

²³Grant, Meneghel, and Tourky (2016, section 5) also provide an example illustrating that in Bayesian games with recursive strictly ambiguity averse preferences, when the commonality condition in Ellis (2018) is relaxed, an ex ante equilibrium may exist. We observe that the combination of no off-path information sets in Bayesian games and such recursion ensure that any ex ante equilibrium will also be a sequential optimum with respect to the interim preferences aggregated in the recursion.

MEU preferences and prior-by-prior Bayesian updating and use as an equilibrium concept optimality under consistent planning in the spirit of Strotz (1955) (see online Appendix B4 for a definition). Thus, as was discussed and illustrated in Section IIB, 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.

APPENDIX A: PROOFS FOR SECTION II AND FURTHER ANALYSIS

A. Proofs Related to the Running Example

PROOF OF PROPOSITION 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

$$\begin{aligned} V_2(\sigma) &= \sigma_1(c)\sigma_2(c)(5\rho(f2) + 4\rho(f1)) \\ &\quad + (1 - \sigma_1(c)\sigma_2(c))(6\rho(n) + 10(1 - \rho(n) - \rho(p2))). \end{aligned}$$

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 to $\sigma_2(c) = 0$ 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

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

This is equivalent to

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

which is

$$(A1) \quad \frac{\phi'_i(4)}{\phi'_i(5)} \geq \frac{5}{4}.$$

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

Suppose one player is sufficiently ambiguity averse that (A1) 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

$$\begin{aligned} 1 \in \operatorname{argmax}_{\lambda_1 \in [0,1]} & \frac{1}{2} \phi_1 \left((4\sigma_3(I)(f2) + 5(1 - \sigma_3(I)(f2))) \lambda_1 \right. \\ & \left. + (10\sigma_3(I)(p2) + 6\sigma_3(I)(n))(1 - \lambda_1) \right) \\ & + \frac{1}{2} \phi_1 \left((4\sigma_3(II)(f2) + 5(1 - \sigma_3(II)(f2))) \lambda_1 \right. \\ & \left. + (10\sigma_3(II)(p2) + 6\sigma_3(II)(n))(1 - \lambda_1) \right). \end{aligned}$$

This is equivalent to

$$\begin{aligned} (A2) \quad & (4\sigma_3(I)(f2) + 5(1 - \sigma_3(I)(f2)) - 10\sigma_3(I)(p2) - 6\sigma_3(I)(n)) \\ & \times \phi'_1(4\sigma_3(I)(f2) + 5(1 - \sigma_3(I)(f2))) \\ & + (4\sigma_3(II)(f2) + 5(1 - \sigma_3(II)(f2)) - 10\sigma_3(II)(p2) - 6\sigma_3(II)(n)) \\ & \times \phi'_1(4\sigma_3(II)(f2) + 5(1 - \sigma_3(II)(f2))) \\ & \geq 0. \end{aligned}$$

Player 2 is ex ante best responding if and only if

$$\begin{aligned} 1 \in \operatorname{argmax}_{\lambda_2 \in [0,1]} & \frac{1}{2} \phi_2 \left((5\sigma_3(I)(f2) + 4(1 - \sigma_3(I)(f2))) \lambda_2 \right. \\ & \left. + (10(1 - \sigma_3(I)(n) - \sigma_3(I)(p2)) + 6\sigma_3(I)(n))(1 - \lambda_2) \right) \\ & + \frac{1}{2} \phi_2 \left((5\sigma_3(II)(f2) + 4(1 - \sigma_3(II)(f2))) \lambda_2 \right. \\ & \left. + (10(1 - \sigma_3(II)(n) - \sigma_3(II)(p2)) + 6\sigma_3(II)(n))(1 - \lambda_2) \right). \end{aligned}$$

This is equivalent to

$$\begin{aligned}
 \text{(A3)} \quad & (5\sigma_3(I)(f2) + 4(1 - \sigma_3(I)(f2))) \\
 & - 10(1 - \sigma_3(I)(n) - \sigma_3(I)(p2)) - 6\sigma_3(I)(n) \\
 & \quad \times \phi'_2(5\sigma_3(I)(f2) + 4(1 - \sigma_3(I)(f2))) \\
 & + (5\sigma_3(II)(f2) + 4(1 - \sigma_3(II)(f2))) \\
 & \quad - 10(1 - \sigma_3(II)(n) - \sigma_3(II)(p2)) - 6\sigma_3(II)(n) \\
 & \quad \times \phi'_2(5\sigma_3(II)(f2) + 4(1 - \sigma_3(II)(f2))) \\
 & \geq 0.
 \end{aligned}$$

Suppose, for example, that player 2 is not ambiguity averse enough to satisfy (A1). 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 (A2) and (A3) 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 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 $\hat{\sigma}_3(I, d)(p1) = \sigma_3(I, d)(p1) + \sigma_3(I, d)(n)/2$, $\hat{\sigma}_3(I, d)(p2) = \sigma_3(I, d)(p2) + \sigma_3(I, d)(n)/2$, 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 $10/2 < 6$, 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 2, it is sufficient to check optimality against one-stage deviations, and therefore only at information sets where the player has a nontrivial move. Begin with player 3. Following the play of 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 nonsingleton 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 $1/2$ 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 (A1) is satisfied, the proof of part (ii) of Proposition 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 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 (A2) and (A3) are necessary for players 1 and 2 to be best responding. We now show that, under (A3), the maximal value of the left-hand side of (A2) is negative, contradicting (A2). 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 (A2) and positive in (A3), the maximal value of the left-hand side of (A2) is

obtained when (A3) is binding. This equality under ambiguity neutrality of player 2 is equivalent to

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

Substituting using this and $\phi_1(x) = -e^{-x \ln(5/3)}$, the left-hand side of (A2) becomes (up to multiplication by a positive constant)

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

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 $\sigma_3(I)(n) = (1/2) - (\sigma_3(I)(f2) + \sigma_3(II)(f2) + 10\sigma_3(I)(p2))/4$ and $\sigma_3(I)(p1) = (1/2) + (\sigma_3(I)(f2) + \sigma_3(II)(f2) + 6\sigma_3(I)(p2))/4$. 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 $-81/1,250$. Therefore (A2) and (A3) cannot be simultaneously satisfied.

The $x \leq 0.5$ statement follows from the proof of Proposition 2. ■

PROOF OF PROPOSITION 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 11 implies that formula (A12) 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 (A12) implies that beliefs are the same at both information sets where 3 has a nontrivial 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 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 4 and Theorem 11, formula (A12) identifies an interim belief system $\hat{\nu}$ satisfying smooth rule consistency. By Theorem 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 nontrivial 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 nontrivial move, their beliefs $\hat{\nu}_{i, I_i}$ place probability $1/2$ on each of the two degenerate π on the corresponding information set, where the $1/2$ 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. ■

B. Proofs of Results in Section IIB

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

DEFINITION 11 (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 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 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 (5) (respectively, $\sigma'_i = \sigma_i$ maximizes (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 (5).

Let $t = s(I_i)$. Consider the case where $t > 0$ (the case where $t = 0$ is similar, using (2) instead of (A4), 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 (5), this is equivalent to the condition that $\sigma'_i = \sigma_i$ maximizes

$$(A4) \quad \sum_{h|h^{t-1} \in I_i^{-1}} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i^{-1}}(h).$$

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 $h^t \in I_i$, $p_{i,\sigma'_i}(h|h^{t-1})$ is the same for any such h . Thus, the objective function in (A4) can be equivalently written as

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

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

$$(A5) \quad \sum_{h|h^t \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^t) q^{(\sigma,\nu),i,I_i^{-1}}(h)$$

is affected by the specification of σ'_i from I_i onward and no other part of σ'_i affects (A5). Therefore, since reachability implies $p_{i,\sigma'_i}(\bar{h}^t|\bar{h}^{t-1}) > 0$, σ_i maximizes (A4) implies that σ_i maximizes (A5). For that to imply σ_i maximizes (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 $\{h|h^t \in I_i\}$. This proportionality may be shown by using the local measure definition (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 10: 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 10:

By ex ante optimality of σ , (7) 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 1, (8) in the definition of sequential optimality is satisfied for all i and I_i . ■

PROOF OF THEOREM 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 5), 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 (i) where $\hat{\nu}_{i, I_i} = \nu_{i, I_i}$, it directly inherits this from (σ, ν) and (ii) everywhere else, Lemma 1 shows that smooth rule updating ensures the required optimality. ■

PROOF OF THEOREM 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 (5))

$$(A6) \quad \sum_{h|h^t \in J_i} u_i(h) p_{i, \sigma_i}(h|h^t) q^{(\sigma, \nu), i, J_i}(h) \geq \sum_{h|h^t \in J_i} u_i(h) p_{i, \sigma'_i}(h|h^t) q^{(\sigma, \nu), i, J_i}(h).$$

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 $\{h | h^t \in J_i\}$. This proportionality follows from using the local measure definition (6), applying the strong smooth rule iteratively to substitute for ν_{i, I_i} and simplifying. After substituting in (A6) 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 $h^t \in J_i$, (A6) becomes

$$\begin{aligned}
 \text{(A7)} \quad & \sum_{h|h^t \in J_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma'_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h) \\
 & \geq \sum_{h|h^t \in J_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma'_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h) \\
 & = \sum_{h|h^t \in J_i} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h).
 \end{aligned}$$

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 (A7) 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:

$$\begin{aligned}
 \text{(A8)} \quad & \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma'_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h) \\
 & \geq \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h).
 \end{aligned}$$

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 (A6) applied to I_i and restricted to such deviations,

$$\begin{aligned}
 \text{(A9)} \quad & \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h) \\
 & \geq \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma'_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h).
 \end{aligned}$$

Combining (A9) and (A8) implies

$$\text{(A10)} \quad \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h) \geq \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h).$$

Since (A10) holds for any σ'_i , it is the same as (A6) 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 (8) in the definition of sequentially optimal.

It remains to show that σ also satisfies the ex ante optimality conditions (7). 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 $\{h|h^0 \in I_i\}$. Using this to substitute for $q^{(\sigma,\nu),i,I_i}$ in (A10) with $t = 1$, cancelling the constant of proportionality and summing for all I_i , yields:

$$\text{(A11)} \quad \sum_h u_i(h) p_{i,\sigma'_i}(h|h^0) q^{(\sigma,\mu_i),i}(h) \geq \sum_h u_i(h) p_{i,\sigma'_i}(h|h^0) q^{(\sigma,\mu_i),i}(h).$$

Since (A11) holds for any σ'_i , σ maximizes (2), which is equivalent to the ex ante optimality condition (7). ■

PROOF OF COROLLARY 1:

Assume ambiguity neutrality. From Theorem 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 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 . ■

C. Proofs of Results in Section IIC

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

PROOF OF LEMMA 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 (9) and (10) hold for all I_i . For any player i and I_i , if either $I_i \subseteq \Theta$ or $\sum_{h^t \in I_i} \sum_{\hat{\pi} \in \Delta(I_i^{-1})} p_{-i, \sigma^k}(h^t | h^{t-1}) \hat{\pi}(h^{t-1}) \nu_{i, I_i^{-1}}(\hat{\pi}) > 0$, then, by continuity in the strategy profile of the formulas (9) and (10), 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:

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 2 and Theorem 2. ■

PROOF OF THEOREM 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 $\varepsilon_I^k \leq 1/|A_i(I_i)|$ 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 $\sigma_i^k(I_i)(a_i) \geq \varepsilon_I^k$ 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, 61). Since the payoff functions are concave and the set of strategies of each player in G^k is nonempty, 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 10, 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 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 3: 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 nonsingleton 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 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 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 $h^t \in I_i$,

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

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 11. 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

$$\begin{aligned} \text{(A12)} \quad \nu_{i,I_i}(\pi) &\propto \sum_{\hat{\pi} \in \Delta(\Theta)|\hat{\pi}_{I_i}=\pi} \frac{\phi'_i(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \hat{\pi}(h^0))}{\phi'_i(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi(h^t))} \\ &\quad \times \left(\sum_{h^t \in I_i} \bar{p}_{-i,\sigma_{-i}}(h^t|h^0) \hat{\pi}(h^0) \right) \mu_i(\hat{\pi}) \end{aligned}$$

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 11:

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 $\pi \in \Delta(I_i)$,

$$(A13) \quad \nu_{i,I_i}^k(\pi) \propto \sum_{\hat{\pi} \in \Delta(\Theta) | \hat{\pi}_{I_i} = \pi} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma^k}(h | h^0) \hat{\pi}(h^0) \right)}{\phi'_i \left(\sum_{h|h' \in I_i} u_i(h) p_{\sigma^k}(h | h^t) \pi(h^t) \right)} \\ \times \left(\sum_{h^t \in I_i} p_{-i, \sigma_{-i}^k}(h^t | h^0) \hat{\pi}(h^0) \right) \mu_i(\hat{\pi}),$$

where

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

It remains to show that the limit of ν_{i,I_i}^k equals ν_{i,I_i} (as defined in (A12)). Divide (A13) 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

$$(A14) \quad \lim_{k \rightarrow \infty} \sum_{\hat{\pi} \in \Delta(\Theta) | \hat{\pi}_{I_i} = \pi} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma^k}(h | h^0) \hat{\pi}(h^0) \right)}{\phi'_i \left(\sum_{h|h' \in I_i} u_i(h) p_{\sigma^k}(h | h^t) \pi(h^t) \right)} \\ \times \left(\sum_{h^t \in I_i} \frac{p_{-i, \sigma_{-i}^k}(h^t | h^0)}{\sum_{\hat{h}^t \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^t | \hat{h}^0)} \hat{\pi}(h^0) \right) \mu_i(\hat{\pi}).$$

By Lemma 3, $\bar{p}_{-i, \sigma_{-i}}(h^t | h^0)$ exists. Then (A14) is proportional (in $\pi \in \Delta(I_i)$) to the right-hand side of (A12) since, whenever it exists, $\lim_{k \rightarrow \infty} \hat{\pi}_{I_i}(h^t) = \lim_{k \rightarrow \infty} \frac{p_{-i, \sigma_{-i}^k}(h^t | h^0) \hat{\pi}(h^0)}{\sum_{\hat{h}^t \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^t | \hat{h}^0) \hat{\pi}(\hat{h}^0)} = \frac{\bar{p}_{-i, \sigma_{-i}}(h^t | h^0) \hat{\pi}(h^0)}{\sum_{\hat{h}^t \in I_i} \bar{p}_{-i, \sigma_{-i}}(\hat{h}^t | \hat{h}^0) \hat{\pi}(\hat{h}^0)} = \bar{\pi}_{I_i}(h^t)$.

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 (A13) for the formula). Fix any i, I_i . The existence of $\lim_{k \rightarrow \infty} \nu^k$, as described in (A14), 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 (A12) 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 (A12) using $\{\sigma^k\}_{k=1}^\infty$. ■

THEOREM 12: Fix a game Γ . Suppose (σ, ν) is sequentially optimal and for each player i , and each information set $I_i \not\subseteq \Theta$, $\sum_{h^{s(i)} \in I_i} p_\sigma(h^{s(i)} | h^0) > 0$. Then there exists $\hat{\nu}$ satisfying (A12) such that $(\sigma, \hat{\nu})$ is an SEA.

PROOF OF THEOREM 12:

Consider a sequence of a completely mixed strategy profiles converging to σ . Since the limits $\bar{p}_{-i, \sigma_{-i}}(h^{s(i)} | h^0)$ in Theorem 11 exist if $\sum_{h^{s(i)} \in I_i} p_\sigma(h^{s(i)} | h^0) > 0$ for all $I_i \not\subseteq \Theta$, there exists $\hat{\nu}$ satisfying (A12). Theorem 11 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 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 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

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

REFERENCES

Aumann, Robert J. 1974. "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics* 1 (1): 67–96.

Auster, Sarah. 2018. "Robust Contracting under Common Value Uncertainty." *Theoretical Economics* 13 (1): 175–204.

Auster, Sarah, and Christian Kellner. 2018. "Robust Bidding and Revenue in Descending Price Auctions." <https://pdfs.semanticscholar.org/9287/59ff9222735616cbf48801425393ae3aa8d4.pdf>.

Ayouni, Mehdi, and Frédéric Koessler. 2017. "Hard Evidence and Ambiguity Aversion." *Theory and Decision* 82 (3): 327–39.

Azrieli, Yaron, and Roei Teper. 2011. "Uncertainty Aversion and Equilibrium Existence in Games with Incomplete Information." *Games and Economic Behavior* 73 (2): 310–17.

Bade, Sophie. 2011. "Ambiguous Act Equilibria." *Games and Economic Behavior* 71 (2): 246–60.

Bade, Sophie. 2016. "Dynamic Semi-Consistency." <https://econ.biu.ac.il/files/economics/seminars/dc8.pdf>.

Baliga, Sandeep, Eran Hanany, and Peter Klibanoff. 2013. "Polarization and Ambiguity." *American Economic Review* 103 (7): 3071–83.

- Battigalli, P., E. Catonini, G. Lanzani, and M. Marinacci.** 2019. "Ambiguity Attitudes and Self-Confirming Equilibrium in Sequential Games." *Games and Economic Behavior* 115: 1–29.
- Battigalli, Pierpaolo, Simone Cerreia-Vioglio, Fabio Maccheroni, and Massimo Marinacci.** 2015. "Self-Confirming Equilibrium and Model Uncertainty." *American Economic Review* 105 (2): 646–77.
- Beauchêne, Dorian, Jian Li, and Ming Li.** 2019. "Ambiguous Persuasion." *Journal of Economic Theory* 179: 312–65.
- Bewley, Truman F.** 1988. "Knightian Decision Theory and Econometric Inference." Cowles Foundation Discussion Paper 868.
- Bodoh-Creed, Aaron L.** 2012. "Ambiguous Beliefs and Mechanism Design." *Games and Economic Behavior* 75 (2): 518–37.
- Bose, Subir, and Arup Daripa.** 2009. "A Dynamic Mechanism and Surplus Extraction under Ambiguity." *Journal of Economic Theory* 144 (5): 2084–114.
- Bose, Subir, Emre Ozdenoren, and Andreas Pape.** 2006. "Optimal Auctions with Ambiguity." *Theoretical Economics* 1 (4): 411–38.
- Bose, Subir, and Ludovic Renou.** 2014. "Mechanism Design with Ambiguous Communication Devices." *Econometrica* 82 (5): 1853–72.
- Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio.** 2011. "Uncertainty Averse Preferences." *Journal of Economic Theory* 146 (4): 1275–1330.
- Chen, Yan, Peter Katsušák, and Emre Ozdenoren.** 2007. "Sealed Bid Auctions with Ambiguity: Theory and Experiments." *Journal of Economic Theory* 136 (1): 513–35.
- di Tillio, Alfredo, Nenad Kos, and Matthias Messner.** 2017. "The Design of Ambiguous Mechanisms." *Review of Economic Studies* 84 (1): 237–76.
- Dominiak, Adam, and Min Suk Lee.** 2017. "Coherent Dempster–Shafer Equilibrium and Ambiguous Signals." *Journal of Mathematical Economics* 68: 42–54.
- Eichberger, Jürgen, Simon Grant, and David Kelsey.** 2018. "Ambiguity and the Centipede Game: Strategic Uncertainty in Multi-stage Games with Almost Perfect Information." <https://people.exeter.ac.uk/dk210/centipede-dk.pdf>.
- Eichberger, Jürgen, and David Kelsey.** 1999. "Education Signalling and Uncertainty." In *Beliefs, Interactions and Preferences in Decision Making*, edited by Mark Machina and Bertrand Munier, 135–57. New York: Springer.
- Eichberger, Jürgen, and David Kelsey.** 2000. "Non-additive Beliefs and Strategic Equilibria." *Games and Economic Behavior* 30 (2): 183–215.
- Eichberger, Jürgen, and David Kelsey.** 2004. "Sequential Two-Player Games with Ambiguity." *International Economic Review* 45 (4): 1229–61.
- Ellis, Andrew.** 2018. "On Dynamic Consistency in Ambiguous Games." *Games and Economic Behavior* 111: 241–49.
- Epstein, Larry G., and Martin Schneider.** 2003a. "Recursive Multiple-Priors." *Journal of Economic Theory* 113 (1): 1–31.
- Epstein, Larry G., and Martin Schneider.** 2003b. "IID: Independently and Indistinguishably Distributed." *Journal of Economic Theory* 113 (1): 32–50.
- Fudenberg, Drew, and Jean Tirole.** 1991a. *Game Theory*. Cambridge, MA: MIT Press.
- Fudenberg, Drew, and Jean Tirole.** 1991b. "Perfect Bayesian Equilibrium and Sequential Equilibrium." *Journal of Economic Theory* 53 (2): 236–60.
- Ghirardato, Paolo.** 2004. "Defining Ambiguity and Ambiguity Attitude." In *Uncertainty in Economic Theory: A Collection of Essays in Honor of David Schmeidler's 65th Birthday*, edited by Itzhak Gilboa, 36–45. London: Routledge.
- Ghirardato, Paolo, and Marciano Siniscalchi.** 2012. "Ambiguity in the Small and in the Large." *Econometrica* 80 (6): 2827–47.
- Gibbons, Robert.** 1992. *Game Theory for Applied Economists*. Princeton: Princeton University Press.
- Gilboa, Itzhak, and Massimo Marinacci.** 2013. "Ambiguity and the Bayesian Paradigm." In *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress, Volume 1, Economic Theory*, edited by Daron Acemoglu, Manuel Arellano, and Eddie Dekel, 179–242. New York: Cambridge University Press.
- Gilboa, Itzhak, and David Schmeidler.** 1989. "Maxmin Expected Utility with a Non-unique Prior." *Journal of Mathematical Economics* 18 (2): 141–53.
- Glicksberg, I.L.** 1952. "A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points." *Proceedings of the American Mathematical Society* 3 (1): 170–74.

- Grant, Simon, Idione Meneghel, and Rabee Tourky. 2016. "Savage Games." *Theoretical Economics* 11: 641–82.
- Greenberg, Joseph. 2000. "The Right to Remain Silent." *Theory and Decision* 48 (2): 193–204.
- Hanany, Eran, and Peter Klibanoff. 2007. "Updating Preferences with Multiple Priors." *Theoretical Economics* 2: 261–98.
- Hanany, Eran, and Peter Klibanoff. 2009. "Updating Ambiguity Averse Preferences." *BE Journal of Theoretical Economics* 9 (1).
- Hansen, Lars Peter. 2007. "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk." *American Economic Review* 97 (2): 1–30.
- Hendon, Ebbe, Hans Jørgen Jacobsen, and Birgitte Sloth. 1996. "The One-Shot-Deviation Principle for Sequential Rationality." *Games and Economic Behavior* 12 (2): 274–82.
- Kajii, Atsushi, and Takashi Ui. 2005. "Incomplete Information Games with Multiple Priors." *Japanese Economic Review* 56 (3): 332–51.
- Kellner, Christian. 2015. "Tournaments as a Response to Ambiguity Aversion in Incentive Contracts." *Journal of Economic Theory* 159 (A): 627–55.
- Kellner, Christian, and Mark T. Le Quement. 2017. "Modes of Ambiguous Communication." *Games and Economic Behavior* 104: 271–92.
- Kellner, Christian, and Mark T. Le Quement. 2018. "Endogenous Ambiguity in Cheap Talk." *Journal of Economic Theory* 173: 1–17.
- Klibanoff, Peter. 1996. "Uncertainty, Decision, and Normal Form Games." <https://pdfs.semanticscholar.org/d517/476dd78103e43e670baed408d84f5e43ddd9.pdf>.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. "A Smooth Model of Decision Making under Ambiguity." *Econometrica* 73 (6): 1849–92.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2009. "Recursive Smooth Ambiguity Preferences." *Journal of Economic Theory* 144 (3): 930–76.
- Kreps, David M., and Robert Wilson. 1982. "Sequential Equilibria." *Econometrica* 50 (4): 863–94.
- Kuzmics, Christoph. 2017. "Abraham Wald's Complete Class Theorem and Knightian Uncertainty." *Games and Economic Behavior* 104: 666–73.
- Lehrer, Ehud. 2012. "Partially Specified Probabilities: Decisions and Games." *American Economic Journal: Microeconomics* 4 (1): 70–100.
- Levin, Dan, and Emre Ozdenoren. 2004. "Auctions with Uncertain Numbers of Bidders." *Journal of Economic Theory* 118 (2): 229–51.
- Lo, Kin Chung. 1999. "Extensive Form Games with Uncertainty Averse Players." *Games and Economic Behavior* 28 (2): 256–70.
- Lopomo, Giuseppe, Luca Rigotti, and Chris Shannon. 2014. "Uncertainty in Mechanism Design." <https://pdfs.semanticscholar.org/fc5a/7e60328f7fc1519d06699e7e0b980633ca47.pdf>.
- Maccheroni, Fabio, and Massimo Marinacci. 2005. "A Strong Law of Large Numbers for Capacities." *Annals of Probability* 33 (3): 1171–78.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini. 2006. "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences." *Econometrica* 74 (6): 1447–98.
- Machina, Mark J. 1989. "Dynamic Consistency and Non-expected Utility Models of Choice under Uncertainty." *Journal of Economic Literature* 27 (4): 1622–68.
- Marinacci, Massimo. 2000. "Ambiguous Games." *Games and Economic Behavior* 31 (2): 191–219.
- Milgrom, Paul, and John Roberts. 1982. "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis." *Econometrica* 50 (2): 443–59.
- Muraviev, Igor, Frank Riedel, and Linda Sass. 2017. "Kuhn's Theorem for Extensive Form Ellsberg Games." *Journal of Mathematical Economics* 68: 26–41.
- Mukerji, Sujoy, and Jean-Marc Tallon. 2004. "An Overview of Economic Applications of David Schmeidler's Models of Decision Making under Uncertainty." In *Uncertainty in Economic Theory: A Collection of Essays in Honor of David Schmeidler's 65th Birthday*, edited by Itzhak Gilboa, 283–302. London: Routledge.
- Myerson, Roger B. 1991. *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press.
- Pahlke, Marieke. 2018. "Dynamic Consistency in Incomplete Information Games with Multiple Priors." Center for Mathematical Economics Working Paper 599, Bielefeld University.
- Perea, Andrés. 2002. "A Note on the One-Deviation Property in Extensive Form Games." *Games and Economic Behavior* 40 (2): 322–38.
- Riedel, Frank, and Linda Sass. 2014. "Ellsberg Games." *Theory and Decision* 76 (4): 469–509.
- Rigotti, Luca, Chris Shannon, and Tomasz Strzalecki. 2008. "Subjective Beliefs and Ex Ante Trade." *Econometrica* 76 (5): 1167–90.

- Salo, Ahtia, and Martin Weber.** 1995. "Ambiguity Aversion in First-Price Sealed-Bid Auctions." *Journal of Risk and Uncertainty* 11 (2): 123–37.
- Selten, Reinhard.** 1965. "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit." *Journal of Institutional and Theoretical Economics* 121 (2): 301–24 and 667–89.
- Shimoji, Makoto, and Joel Watson.** 1998. "Conditional Dominance, Rationalizability, and Game Forms." *Journal of Economic Theory* 83 (2): 161–95.
- Siniscalchi, Marciano.** 2009. "Two out of Three Ain't Bad: A Comment on 'The Ambiguity Aversion Literature: A Critical Assessment.'" *Economics and Philosophy* 25 (3): 335–56.
- Siniscalchi, Marciano.** 2011. "Dynamic Choice under Ambiguity." *Theoretical Economics* 6 (3): 379–421.
- Strotz, R.H.** 1955. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23 (3): 165–80.
- Wolitzky, Alexander.** 2016. "Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade." *Theoretical Economics* 11: 971–1004.