

Incomplete Information Games with Ambiguity Averse Players*

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

This version: July 21, 2017

Abstract

We study incomplete information games of perfect recall involving players who perceive ambiguity and may be ambiguity averse. Our focus is on equilibrium concepts satisfying *sequential optimality* – each player’s strategy must be optimal at each stage given the strategies of the other players and the player’s conditional beliefs. We show that, for the purpose of identifying strategy profiles that are part of a sequential optimum, it is without loss of generality to restrict attention to a particular generalization of Bayesian updating. We examine comparative statics in ambiguity aversion and results and examples on belief robustness and new strategic behavior.

1 Introduction

Dynamic games of incomplete information are the subject of a large literature, both theory and application, with diverse fields including models of firm competition, agency theory, auctions, search, insurance and many others. In such games, how players perceive and react to

*We thank Itzhak Gilboa, Frederic Koessler, Fabio Maccheroni, Jean-Marc Tallon and participants at GAMES 2016, FUR 2016, RUD 2015, SING11-GTM2015, the Paris conferences “Ambiguity in Games and Mechanisms”, “Time, Uncertainties and Strategies” and D-TEA 2014, and the workshop “Knightian Uncertainty in Strategic Interactions and Markets” at the ZiF, Center for Interdisciplinary Research at Bielefeld University for helpful comments. Klibanoff and Mukerji additionally thank the Robust Finance research group at the ZiF, Bielefeld University for helpful comments and hospitality during their time as visiting fellows in Spring/Summer 2015. This research was partially supported by grant No. 2014350 from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel to Hanany and Klibanoff.

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

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

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

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 (2003), Maccheroni and Marinacci (2005) and Klibanoff, Marinacci and Mukerji (2009). We propose equilibrium notions for incomplete information games involving ambiguity about parameters or players' types. This allows us to examine effects of introducing ambiguity aversion in strategic settings, static and dynamic.

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

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

where Θ is the parameter space modeling the incomplete information, μ_i is a subjective probability over $\Delta(\Theta)$ (i.e., a second-order probability over Θ), $U_i(\sigma, \theta)$ is i 's expected payoff from σ given θ , and ϕ_i is an increasing function, the concavity of which reflects ambiguity aversion. All else equal, as ϕ_i becomes more concave, player i becomes more ambiguity averse (see e.g., Theorem 2 in Klibanoff, Marinacci, Mukerji 2005). The presence of ambiguity is captured by non-degeneracy of μ_i . In the smooth ambiguity model it is possible to hold the players' information fixed (by fixing μ_i) while varying their ambiguity attitude from aversion to neutrality (i.e., replacing a more concave ϕ_i with an affine one, which reduces preferences to expected utility). This facilitates a natural way to understand the effect of introducing ambiguity aversion into a strategic environment. Our focus is on extensive form games, specifically multistage games with perfect recall, and on equilibrium notions capturing perfection analogous to those in standard theories for ambiguity neutral players, such as subgame perfect equilibrium (Selten, 1965), sequential equilibrium (Kreps and Wilson, 1982) and perfect Bayesian equilibrium (PBE) (e.g., Fudenberg and Tirole, 1991a,b).

We first define an ex-ante (Nash) equilibrium concept allowing for aversion to ambiguity about parameters, a special case of which are players' types. 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*

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

requirement – each player i 's strategy must be optimal at each stage given the strategies of the other players and i 's conditional beliefs. When there is no parameter uncertainty, sequential optimality reduces to subgame perfection. Sequential optimality and our subsequent analysis and extensions of it are the main contributions of the paper.

We find that sequential optimality has a number of attractive properties including cutting through the vexing issue of what update rule to impose in dynamic games with ambiguity aversion. We show that for the purpose of identifying strategy profiles that are part of a sequential optimum, it is *without loss of generality* to restrict attention to belief systems updated using a generalization of Bayesian updating for smooth ambiguity preferences, called the smooth rule (Hanany and Klibanoff 2009), which coincides with Bayes' rule under ambiguity neutrality.² In this sense, the sequential optimality solution concept identifies a method for updating without arbitrarily assuming one a priori. Furthermore, an important method facilitating analysis of dynamic games with standard preferences is the sufficiency of checking only one-stage deviations (as opposed to general deviations) when verifying optimality at each information set. We show that this method retains its validity when applied to sequential optimality: a strategy profile is part of a sequential optimum if and only if, for each player and each information set, there are no profitable one-stage deviations with respect to beliefs updated according to the smooth rule.

Consider the game in Figure 1.1, illustrating some behavior that would be allowed under, for example, the no profitable one-stage deviation criterion with respect to beliefs updated according to Bayes' rule, but is ruled out by sequential optimality. To analyze this example, let us consider player 2 who is privately informed of $\theta \in \{I, II, III\}$ at the beginning of the game. Observe that for each θ , player 2 has a strictly dominant strategy if given the move: types I and II play U , and type III plays D . Given this strategy for player 2, observe that for player 1, the payoff to playing i followed by d if U is, θ -by- θ , strictly higher than the payoff to playing o followed by anything. Thus, no strategy involving playing o with positive probability can be a best reply to 2's optimal strategy no matter how player 1 perceives and treats the uncertainty about θ . This implies that o is not part of any ex-ante (Nash) equilibrium, let alone a sequentially optimal strategy profile.

In contrast, it is easy to specify preferences under which player 1 playing o with positive probability and playing u with probability $1/2$ if given the move together with player 2 playing her strictly dominant strategy if given the move is consistent with the no profitable one-stage deviation criterion under Bayesian updating. Notice that from the perspective

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).

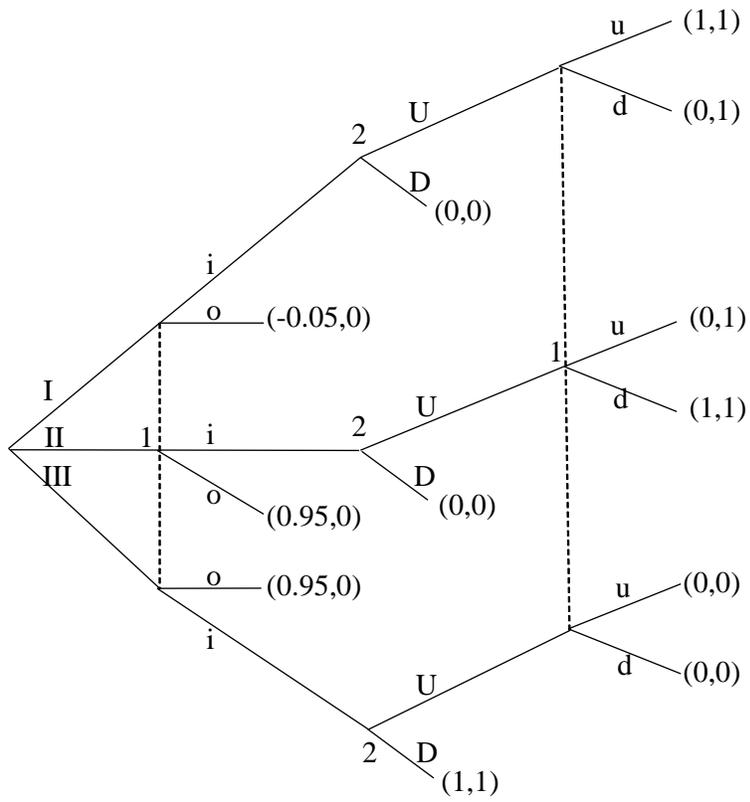


Figure 1.1: Game illustrating violation of sequential optimality.
 The vectors give utility payoffs for players 1 and 2, in that order, for each path.

of player 1, taking as given playing u with probability $1/2$, i does worse than o under II and better than o otherwise. If there is ambiguity concerning the probability of types II versus III and player 1 is sufficiently ambiguity averse, when considering playing i with probability 1 and evaluating moving toward o , that change will be evaluated using weights tilted sufficiently toward II to make that movement attractive. The details may be found in Appendix C.³ From a decision theoretic perspective, the vast majority of atemporal models of ambiguity averse preferences adopt state-by-state monotonicity as a fundamental building block. Thus, another reason for interest in sequential optimality as a solution concept for dynamic games is that it rules out the play of type-by-type (iteratively) strictly dominated strategies in dynamic settings.

Recent literature (e.g., Greenberg (2000), Bade (2011), Riedel and Sass (2013), Bose and Renou (2014), Kellner and Le Quement (2015), Ayouni and Koessler (2017) and di Tillio, Kos and Messner (2017)) identifies strategic use of ambiguity – intentionally choosing strategies that will be perceived as ambiguous by others – as a phenomenon of interest. This literature demonstrates that strategic use of ambiguity can be strictly valuable in both static and dynamic settings. However, its value in dynamic settings has only been shown for strategies *violating* sequential optimality. Not only do we show that our framework accommodates strategic use of ambiguity, but also that it can be valuable under sequential optimality. Sequential optimality may be viewed in the context of communication games and mechanism design 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 (and, not, as in the example in Figure 1.1, where without sequential optimality opting out can occur).

Sequential optimality does not restrict player i 's beliefs at information sets immediately following other players' deviations. 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. We show that SEA exists for any finite multistage game with perfect recall and incomplete information, and for any specification of players' ambiguity aversions and initial beliefs.

Section 2.5 provides results on comparative statics of the equilibrium set in ambiguity aversion. First, for fixed beliefs, ambiguity aversion may change the equilibrium set in a variety of ways – it can expand, shrink or simply change the set of equilibria. Second, we

³There 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, and, more generally, remains weaker than sequential optimality.

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. For a fixed game form, 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). With unrestricted heterogeneous beliefs (i.e., running over all possible μ_i), however, ambiguity aversion leaves unchanged the set of equilibria. However, if we were to limit attention to pure strategies (both in terms of the equilibrium profile and 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. We also discuss the relationship of these results to existing literature.

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 2.6 defines robustness of an equilibrium to increases in ambiguity aversion and shows that this is related to a type of belief robustness. The following example illustrates these notions of robustness and their relationship. Consider a two player one-stage game, where each player has a choice between two actions A and B . The parameter space about which there is ambiguity is $\Theta = \{\theta_1, \theta_2\}$. Both players have beliefs μ such that $\mu(\pi_1) = \mu(\pi_2) = \frac{1}{2}$, where $\pi_1(\theta_1) = \frac{2}{3}$ and $\pi_2(\theta_1) = \frac{1}{2}$, and do not learn anything about θ before choosing their action. Payoffs as a function of the actions and θ are as follows:

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

(1.1)

Observe that given θ_1 , A is strictly dominant for each player, while given θ_2 , B is strictly dominant. Under ambiguity neutrality, i.e. ϕ affine, both (A, A) and (B, B) are equilibrium strategy profiles. We claim that (A, A) is robust to increased ambiguity aversion (i.e., remains an equilibrium when ϕ becomes more concave), but (B, B) is not. To see that (A, A) is robust, note that, assuming her opponent plays A , a player evaluates the mixed strategy $\lambda A + (1 - \lambda)B$ according to $\frac{1}{2}\phi(2\lambda) + \frac{1}{2}\phi(4 - \lambda)$, which is maximized at $\lambda = 1$ for any concave ϕ . To see that (B, B) is not robust, note that for example, if $\phi(x) = -e^{-\alpha x}$ with $\alpha > \ln(\frac{1+\sqrt{5}}{2}) \approx 0.48$, it is profitable to deviate to A .

There is another sense in which (A, A) is robust. Consider the set of weights μ on π_1 and π_2 that support (A, A) as an equilibrium. Such weights are those satisfying $\mu(\pi_1) \geq \frac{\phi'(3)}{2\phi'(2)+\phi'(3)}$. Notice that as ϕ becomes more concave, $\frac{\phi'(3)}{2\phi'(2)+\phi'(3)}$ decreases, and thus increases in ambiguity aversion expand the set of weights μ supporting (A, A) . The fact that ambiguity aversion may expand the set of beliefs supporting (A, A) is not special to this example. We show, under some general conditions, that equilibria that are robust to increased ambiguity

aversion *must* have their sets of supporting beliefs expand under some increases in ambiguity aversion (see Theorem 2.5). We refer to this as ambiguity aversion making an equilibrium more belief robust. Ambiguity naturally arises for a potential entrant assessing an unfamiliar market. Section 3 contains an example of a Milgrom and Roberts (1982)-style limit pricing entry game with ambiguity about the incumbent’s cost. We show that limit pricing arises in an SEA. We provide conditions under which this limit pricing is robust to increased ambiguity aversion and thus, applying Theorem 2.5, under which ambiguity aversion makes limit pricing more belief robust compared to the usual setting of ambiguity neutrality.

As a running example (beginning in Section 2), to illustrate concepts as we introduce them, we consider a principal, multi-agent communication game. The principal is shown to strictly benefit from conditioning his cheap talk message to the agents on a payoff-irrelevant ambiguous event. The existing literature on the value of ambiguous communication in dynamic settings uses violations of sequential optimality in an essential way. Left open was the question of whether such value existed when respecting sequential optimality. Our analysis in the running example establishes that valuable strategic use of ambiguity can be sequentially optimal, and, in fact, occurs as part of an SEA.

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. A number of previous papers have analyzed incomplete information games with ambiguity sensitive preferences in settings without dynamics, including Salo and Weber (1995), Ozdenoren and Levin (2004), Kajii and Ui (2005), Bose, Ozdenoren and Pape (2006), Chen, Katuscak and Ozdenoren (2007), Lopomo, Rigotti and Shannon (2010), Azrieli and Teper (2011), Bade (2011), Bodoh-Creed (2012), Riedel and Sass (2013), Wolitzky (2013, 2016), Kellner (2015), Auster (2016) and di Tillio, Kos, Messner (2017). Ellis (2016) examines dynamic consistency between the ex-ante and interim stages of these static games and shows that a version of the folk result that dynamic consistency plus consequentialism implies probabilistic beliefs in an individual choice context can be extended to this game setting. There have been only a very few papers investigating aspects of dynamic games with ambiguity aversion (e.g., Lo 1999, Eichberger and Kelsey 1999, 2004, Bose and Daripa 2009, Kellner and Le Quement 2015, 2017, Bose and Renou 2014, Muraviev, Riedel and Sass 2017, Battigalli et al. 2015a,b, Dominiak and Lee 2017). Instead of sequential optimality, these other papers involving dynamic games take a variety of approaches. These include, e.g., optimality under consistent planning in the spirit of Strotz (1955-56), the notion of no profitable one-stage deviations, or taking a purely ex-ante perspective. Thus, equilibria described by these approaches may fail sequential optimality. Additionally, each of these papers limits attention to forms of updating that generally rule

out some (or all) sequentially optimal strategies.

Finally, Section 4 discusses some possible extensions of our approach, including to other models of ambiguity averse players' preferences.

2 Model, Equilibrium and Analysis

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 a finite set of players.
- H is a finite set of histories, each of which is a finite sequence of length $T + 2$ of the form $h = (h_{-1}, (h_{0,i})_{i \in N}, \dots, (h_{T,i})_{i \in N})$.
For $0 \leq t \leq T + 1$, let $H^t \equiv \{h^t \equiv (h_{-1}, (h_{0,i})_{i \in N}, \dots, (h_{t-1,i})_{i \in N}) \mid h \in H\}$ be the set of partial histories up to (but not including) stage t . The set of all partial histories is $\mathcal{H} \equiv \{\emptyset\} \cup \bigcup_{0 \leq t \leq T+1} H^t$. For each $i \in N$, $0 \leq t \leq T$ and $h^t \in H^t$, $A_i(h^t) \equiv \{\hat{h}_{t,i} \mid \hat{h} \in H, \hat{h}^t = h^t\}$ is the set of actions available to player i at h^t . The set of initial partial histories, $\Theta \equiv H^0$, is called the set of "parameters" or "types".
- $\mathcal{I}_i \equiv \bigcup_{0 \leq t \leq T} \mathcal{I}_i^t$ are the information sets for player i , where each \mathcal{I}_i^t is a partition of H^t such that, for all $h^t, \hat{h}^t \in H^t$, $\hat{h}^t \in I_i(h^t)$ implies $A_i(h^t) = A_i(\hat{h}^t)$ (where $I_i(h^t)$ is the unique element of \mathcal{I}_i^t containing h^t).
Perfect recall means: for each player i , stage $0 \leq t \leq T$ and partial histories $h^t, \hat{h}^t \in H^t$, $I_i(h^t) = I_i(\hat{h}^t)$ implies $R_i(h^t) = R_i(\hat{h}^t)$, where, for each partial history $\bar{h}^t \in H^t$, $R_i(\bar{h}^t)$

is the ordered list of information sets i encounters and the actions i takes under \bar{h}^t .⁴

- $u_i : H \rightarrow \mathbb{R}$ is the utility payoff of player i given the history.⁵
- 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 .⁶
- $\phi_i : \text{co}(u_i(H)) \rightarrow \mathbb{R}$ is a continuously differentiable, (weakly) concave and strictly increasing function.

The first three bullet points above describe the game form, while the rest describe preferences. The only notable non-standard 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 having multiple probability measures in its support.

We want to point out that the parameter space Θ may include both payoff-relevant and 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. This use of ambiguity may apply to strategies involving either payoff-relevant or payoff-irrelevant actions.

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. The example is a novel game in which deliberately introducing ambiguity about actions *without* payoff consequences (“ambiguous cheap talk”) proves valuable in equilibrium. Furthermore, in this game there is strict value in a player choosing to make his message depend on payoff-*irrelevant* ambiguous aspects of the parameters. This latter aspect is how our framework captures the idea that choosing to introduce uncertainty into one’s strategy through ambiguity can have value over and above the classical game theoretic notion of introducing strategic uncertainty via mixing or conditioning on randomizing devices.

Running Example: There are three players, a principal, P , and two agents, r (ow) and c (olumn). 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

⁴Formally, $R_i(\bar{h}^t) \equiv ((I_i(\bar{h}^s), \bar{h}_{s,i})_{0 \leq s < t}, I_i(\bar{h}^t))$. For future reference, note that we extend both A_i and R_i to information sets in the natural way.

⁵As 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.

⁶All of our results (except for Theorem B.3) also hold if the class of games is restricted to those with a common μ such that $\mu_i = \mu$ for all players i . None of our examples rely on differences in the μ_i .

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 $\Theta = \{IU, ID, IIU, IID\}$. At stage $t = 0$, only the principal 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 whether to start an independent b (usiness) or w (ork) for the principal. Thus a history h consists of a parameter value, a message, and an action for each agent. For example, one history is $(IU, \alpha, (w, b))$. The principal is privately informed of $\theta \in \Theta$ before sending his message, thus an information set at $t = 0$ for the principal consists of θ , and at $t = 1$ consists of θ together with the message sent. The message is publicly observed by both agents before they choose their actions. Thus, the only information set at $t = 0$ for an agent is the set Θ , and at $t = 1$, an agent's information set is the cross-product of Θ with the message the principal sent. Payoffs (the u_i) depend only on the payoff-relevant part of θ and the agents' actions and are given in the following matrices, where each cell lists the payoff to P , r , and c in that order:

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

Notice that the principal's message is cheap talk in that it does not have any direct effect on payoffs. 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 . The principal offers a fixed payment (equal to his revenue under partial success) to be split evenly among the agent(s) choosing to work for him. 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 , in addition to the benefit from her share of P 's payment, she incurs an effort cost that is higher when she works alone than if both agents work together. This reduction in cost when working together more than compensates for the lower payment share (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 agent c 's business idea.

Agent r is ambiguity averse with $\phi_r(x) = -e^{-11x}$. The exact specification of ϕ_P and ϕ_c will not be important for our analysis of the game.

The beliefs μ for all players are $\frac{1}{2}$ - $\frac{1}{2}$ over distributions π_1 and π_2 given by:

	IU	ID	IIU	IID
π_1	$\frac{1}{4}$	$\frac{1}{2}$	$\frac{1}{12}$	$\frac{1}{6}$
π_2	$\frac{1}{20}$	$\frac{3}{20}$	$\frac{1}{5}$	$\frac{3}{5}$

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 $\gamma \in \{\gamma_1, \gamma_2\}$ about which there is ambiguity (reflected in π_1 vs. π_2), and both the payoff-relevant, $\{I, II\}$, and payoff-irrelevant, $\{U, D\}$, parts of θ are determined as conditionally independent stochastic functions of γ with $\text{Prob}(I|\gamma_1) = 3/4$, $\text{Prob}(I|\gamma_2) = 1/5$, $\text{Prob}(U|\gamma_1) = 1/3$, $\text{Prob}(U|\gamma_2) = 1/4$. For instance, γ might be some scientific principle that is not well understood, and it influences both the functioning of the technology (I vs. II) and the findings of the laboratory experiment (U vs. D) not affecting any of the players' business ventures. ■

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

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

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

Given a strategy profile σ , history h and $0 \leq r \leq t \leq T + 1$, the probability of reaching h^t starting from h^r is $p_\sigma(h^t|h^r) \equiv \prod_{j \in N} \prod_{r \leq s < t} \sigma_j(I_j(h^s))(h_{s,j})$.⁷ 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 formally the assumption that players' ex-ante preferences over strategies are smooth ambiguity preferences (Klibanoff, Marinacci and Mukerji 2005) with the u_i, ϕ_i and μ_i as specified by the game.

Assumption 2.1 (Ex-ante Preferences) Fix a game Γ . Ex-ante (i.e., given the empty

⁷If $r = t$, so that the product is taken over an empty set, invoke the convention that a product over an empty set is 1.

partial history), each player i ranks strategy profiles σ according to

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

2.1 Ex-ante Equilibrium

Using the ex-ante preferences we define ex-ante (Nash) equilibrium:

Definition 2.2 (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 $\sigma'_i \in \Sigma_i$.

When there is no type uncertainty, the definition reduces to Nash equilibrium with expected utility preferences. Thus, in a game with complete information, we have nothing new to say compared to the standard theory. Also, in the case where the ϕ_i are linear (subjective expected utility) and $\mu_i = \mu$ for all players i , the definition reduces to the usual (ex-ante) Bayesian Nash Equilibrium definition.

Running Example continued: Returning to the principal-multi-agent game in the running example (p. 9), for the sake of comparison with what we will show occurs when, as specified, agent r is ambiguity averse, let us begin by considering the case when all players are ambiguity neutral (i.e., ϕ affine). Notice that 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 whether the payoff-relevant component of the parameter is I or II , 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 for sure, otherwise P sends message β for sure; after message α , agent r plays w for sure and agent c plays b for sure, while after message β , agent r plays b for sure and agent c plays w for sure. Observe that under this strategy, P gets his maximal payoff of 2 when I occurs, but gets 0 when II occurs. Why can't P do 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 agent 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 for sure, otherwise P sends message β for sure; after either message, agent r plays w for sure; after message α , agent c plays b for sure, and after β , c plays w for sure. Observe that P is making use of the payoff-irrelevant component of the parameter. We show that this strategy profile is an equilibrium (Proposition 2.1), and, that the principal does strictly better than if he were not allowed to use the payoff-irrelevant component (Proposition 2.2). The proof of these and all subsequent results in the paper may be found in the Appendices.⁸

Proposition 2.1 *The strategy profile σ^* is an ex-ante equilibrium. In this equilibrium, P attains his maximum possible payoff for each parameter.*

Remark 2.1 The strategy profile σ^* remains an equilibrium for any ϕ_r more concave than the one in the example.

Why is σ^* an equilibrium? Consider first agent c . By playing according to σ^* , c gets expected utility 2.75 under π_1 and 2.15 under π_2 . Deviating away from b after message α lowers utility under both π , no matter what c plays after β . Deviating by playing b after both messages, c would get 3.75 under π_1 and 1 under π_2 . Since σ_c^* has higher μ -average expected utility ($2.5 > 2.375$), is less exposed to ambiguity in that it has smaller variation across the π , and the two strategies order the π the same so that mixing between them does not help hedge against ambiguity, c has no incentive to deviate away from σ_c^* .

Agent r gets expected utility 1.75 under π_1 and 1.95 under π_2 . As with c , deviating after message α is never profitable for r . Deviating to b after message β yields 1.5 under π_1 and 4.05 under π_2 . Now there is a conflict: this deviation has a higher μ -average expected utility ($2.775 > 2.5$) but exposes r to more ambiguity. In the example, r is sufficiently ambiguity averse to prefer not to deviate.

Finally, given the strategies of the agents, P gets his highest possible payoffs and thus is playing optimally.

⁸Note that while in this example there is only one payoff irrelevant component available, if one wanted to explicitly model the principal choosing to condition on a payoff irrelevant component having the ‘‘optimal’’ ambiguity about it, one would do this 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. Furthermore, our point that the principal will strictly want to condition his message on such an ambiguous component is robust to any such enlargement.

Proposition 2.2 *If P were not allowed to make his strategy depend on the payoff-irrelevant component of the parameter (i.e., U or D), there would be no ex-ante equilibrium yielding P the maximum possible payoff for each parameter.*

One lesson from the proof of Proposition 2.2 is that, fixing the payoff-relevant component of the parameter, ambiguity about the payoff-irrelevant component is *necessary* for the principal to benefit from making use of it. The reason for this is that any strategy depending on payoff-irrelevant events that, conditional on the payoff relevant component, are unambiguous (i.e., are given the same conditional probability by all π in the support of the beliefs) may be replicated by using appropriate mixtures over messages conditional on the payoff-relevant component without changing the best responses of the agents.

How does playing the ambiguous communication strategy σ_P^* help the principal do better in the example? It allows P to expose agent r to more ambiguity in equilibrium than r would be exposed to under the optimal communication strategy that does not make use of U vs. D . To see this it may help to know what the best strategy for P that does not depend on U vs. D is: if I then send message α with probability ρ , and otherwise send message β for sure. Holding the play of agent c fixed, by playing w in response to both messages agent r would get expected utility $2 - \frac{3}{4}\rho$ under π_1 and $2 - \frac{1}{5}\rho$ under π_2 . For comparison, under the optimal ambiguous communication strategy σ_P^* the corresponding expressions are: $2 - \frac{3}{4}\pi_1(U)$ under π_1 and $2 - \frac{1}{5}\pi_2(U)$ under π_2 . Notice that more ambiguity can be generated with the ambiguous strategy, since the most that could be introduced without it is when $\rho = 1$, resulting in 1.25 under π_1 and 1.8 under π_2 while, for example, if $\pi_1(U) = 1$ and $\pi_2(U) = 0$, the ambiguous strategy would result in 1.25 under π_1 and 2 under π_2 . Even when the values of $\pi_1(U)$ and $\pi_2(U)$ and P 's optimal ρ are not as extreme, any $\pi_1(U) > \pi_2(U)$ with $\pi_1(U) \geq \rho$, as is the case in our example, can play a similar ambiguity-inducing role and allow the principal to leverage r 's ambiguity aversion to a greater extent than with randomization alone – resulting in greater ability to provide incentives for r to play w as P desires.

Remark 2.2 If agent r becomes sufficiently more ambiguity averse, Proposition 2.2 no longer holds: in addition to the equilibrium in Proposition 2.1, there will be an equilibrium where P conditions his play only on I vs. II and also obtains his maximum possible payoff for each parameter. Intuitively, with enough ambiguity aversion on the part of r , the additional ambiguity generated by conditioning on U vs. D is no longer needed.

Examples with strategic ambiguity involving actions that have payoff consequences may be found in the literature. For instance, Greenberg's (2000) peace negotiation example in

which he argues that a powerful country mediating peace negotiations between two smaller countries would wish to introduce ambiguity about which small country will suffer from worse relations with the powerful country if negotiations break down, has been discussed in Mukerji and Tallon (2004) and modeled as an equilibrium in Riedel and Sass (2013). It is straightforward to construct similar examples as equilibria in our framework. di Tillio, Kos and Messner (2017) consider a mechanism design problem where the designer may choose the mapping from participants' actions to the mechanism outcomes to be ambiguous. In our framework, this corresponds to allowing mechanism outcomes to be conditioned on the payoff-irrelevant ambiguous types in addition to participants' actions. Ayouni and Koessler (2017) examine settings with hard evidence and show that ambiguous auditing/certification strategies may be beneficial. Unlike our running example, all of these rely on ambiguity about payoff-relevant actions. The most closely related examples to ours in the literature, also having ambiguity about "cheap talk" actions, are analyzed in Bose and Renou (2014) and Kellner and Le Quement (2015). We discuss these papers further when we revisit the running example in Section 2.3. ■

2.2 Interim Preferences

Now turn to defining preferences beyond the ex-ante stage. We write down i 's preferences at an information set I_i , and devote the remainder of this subsection to formally defining and explaining any new notation involved in doing so. This is preparation for Section 2.3 that describes sequential optimality and an update rule for beliefs.

Assumption 2.2 (Interim Preferences) *Fix a game Γ and a strategy profile σ . Any player i at information set I_i ranks strategies $\hat{\sigma}_i$ according to*

$$V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i}) \equiv \sum_{\pi \in \Delta(H^{m_i(I_i)} | \pi(f_i^{\sigma_{-i}}(I_i)) > 0)} \phi_i \left(\sum_{h|h^t \in I_i} u_i(h) p_{(\hat{\sigma}_i, \sigma_{-i})}(h|h^t) \pi_{I_i, \sigma_{-i}}(h^t) \right) \nu_{i,I_i}(\pi), \quad (2.2)$$

where $t = s(I_i)$.

Compared to the ex-ante preferences given in (2.1), the conditional preferences (2.2) differ only in that (1) the beliefs may have changed in light of I_i and σ (i.e., μ_i is replaced by some interim belief ν_{i,I_i}), (2) each π is conditioned on I_i, σ_{-i} , which is the subset of I_i reachable under σ_{-i} from $f_i(I_i)$, the most recent information set for i not reachable from its immediate predecessor, and (3) 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.

In order to begin explaining the new notation used in interim preferences, it is helpful to define, having fixed a strategy profile σ , when one information set of player i is reachable from another without requiring a deviation from σ by players other than i .

Notation 2.1 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 .

Definition 2.3 (Reachability) Information set I_i is reachable from information set \hat{I}_i given strategies σ_{-i} via partial history $h^{s(I_i)} \in I_i$ if $h^{s(\hat{I}_i)} \in \hat{I}_i$ and $p_{-i, \sigma_{-i}}(h^{s(I_i)} | h^{s(\hat{I}_i)}) > 0$. Say that I_i is reachable from \hat{I}_i given σ_{-i} whenever such a partial history exists.

It is possible, for some information set, that no partial histories make it reachable from the beginning of the game given that other players follow σ_{-i} . Suppose \hat{I}_i is a first information set for which this happens (if it does). Again, it is possible that some subsequent information set is a first not reachable from \hat{I}_i given that other players follow σ_{-i} . It is useful to keep track of the most recent information set for player i that is not reachable from its immediate predecessor given σ_{-i} . The function f_i , defined recursively below, records this for each information set of player i , where we adopt the convention that if i has never yet encountered such an information set then f_i records player i 's initial information set.

Notation 2.2 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 2.4 Given information set I_i and strategies σ_{-i} , define $f_i(I_i)$ as follows: For $I_i \subseteq \Theta$, $f_i(I_i) = I_i$; for $I_i \not\subseteq \Theta$, if I_i is reachable from $f_i(I_i^{-1})$ given σ_{-i} , $f_i(I_i) = f_i(I_i^{-1})$, otherwise $f_i(I_i) = I_i$. Define $m_i(I_i) = s(f_i(I_i))$.

Next, define the set of partial histories in I_i that make it reachable from $f_i(I_i)$ given σ_{-i} , and the corresponding set of partial histories in $f_i(I_i)$:

Notation 2.3 $I_{i, \sigma_{-i}} \equiv \{h^{s(I_i)} \in I_i \mid I_i \text{ is reachable from } f_i(I_i) \text{ given } \sigma_{-i} \text{ via } h^{s(I_i)}\}$ and $f_i^{\sigma_{-i}}(I_i) \equiv \{h^{m_i(I_i)} \in f_i(I_i) \mid I_i \text{ is reachable from } f_i(I_i) \text{ given } \sigma_{-i} \text{ via } h^{s(I_i)}\}$.

The following is a defining property for interim (second-order) beliefs, ν_{i, I_i} , of player i given information set I_i and strategy profile σ : that they assign weight only to distributions that assign positive probability to the set of partial histories in $f_i(I_i)$ that have at least one continuation reaching I_i given σ_{-i} . This seems a minimal condition for interim beliefs to be consistent with the given σ_{-i} .

Definition 2.5 (Interim Belief) An interim belief for player i in a game Γ given information set I_i and strategy profile σ is a finite support probability measure ν_{i,I_i} over $\Delta(H^{m_i(I_i)})$ such that

$$\nu_{i,I_i}(\{\pi \mid \pi(f_i^{\sigma_{-i}}(I_i)) > 0\}) = 1. \quad (2.3)$$

Given a strategy profile σ , an interim belief system $\nu \equiv (\nu_{i,I_i})_{i \in N, I_i \in \mathcal{I}_i}$ is an interim belief for each player at each of that player's information sets.

Furthermore, the distributions $\pi \in \Delta(H^{m_i(I_i)})$ in the support of an interim belief should enter interim preferences only through their conditionals (given I_i and σ_{-i}). Such conditionals, $\pi_{I_i, \sigma_{-i}} \in \Delta(I_i)$, are related to π via the following Bayes' formula:

Definition 2.6 (Conditional π) For $I_i = f_i(I_i)$,

$$\pi_{I_i, \sigma_{-i}}(h^t) = \frac{\pi(h^t)}{\sum_{\hat{h}^t \in I_i, \sigma_{-i}} \pi(\hat{h}^t)} \text{ if } h^t \in I_i, \sigma_{-i} \text{ and 0 otherwise,} \quad (2.4)$$

and, for $I_i \neq f_i(I_i)$,

$$\pi_{I_i, \sigma_{-i}}(h^t) = \frac{p_{-i, \sigma_{-i}}(h^t | h^{t-1}) \pi_{I_i, \sigma_{-i}}^{-1}(h^{t-1})}{\sum_{\hat{h}^t \in I_i, \sigma_{-i}} p_{-i, \sigma_{-i}}(\hat{h}^t | \hat{h}^{t-1}) \pi_{I_i, \sigma_{-i}}^{-1}(\hat{h}^{t-1})} \text{ if } h^t \in I_i, \sigma_{-i} \text{ and 0 otherwise,} \quad (2.5)$$

where $t = s(I_i)$.

These Bayesian conditionals, $\pi_{I_i, \sigma_{-i}}$, differ across the π 's in the support of ν_{i, I_i} only in the relative weights assigned to elements of I_i, σ_{-i} . This fact, together with multiplication by $p_{-i, \sigma_{-i}}(h | h^t)$ in (2.5), reflects that there is no ambiguity about the part of σ_{-i} specifying behavior from $f_i(I_i)$ onward (i.e., the part of σ_{-i} that has not yet been contradicted). Thus, the ambiguity possibly impacting interim preferences concerns only the parameter θ and/or opponents' actions prior to $f_i(I_i)$ that are not uniquely identified given I_i, σ_{-i} .

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 $I_i \in R_i(\hat{I}_i)$. Interim preferences may be equivalently thought of as ranking continuation strategies.

2.3 Sequential Optimality

Fundamental to our theory will be sequential optimality. It requires that each player plays optimally at each information set given the strategies of the others. This optimality is

required even when the information set is before-the-fact viewed as unreachable according to the given strategy profile combined with the beliefs of the player.⁹

Using these preferences, we may now define sequential optimality – in words, each player i 's strategy must be optimal at each stage given the strategies of the other players and i 's interim beliefs:

Definition 2.7 Fix a game Γ . A pair (σ, ν) consisting of a strategy profile and interim belief system is sequentially optimal if, for all players i and all information sets I_i ,

$$V_i(\sigma) \geq V_i(\sigma'_i, \sigma_{-i}) \quad (2.6)$$

and

$$V_{i,I_i}(\sigma) \geq V_{i,I_i}(\sigma'_i, \sigma_{-i}) \quad (2.7)$$

for all $\sigma'_i \in \Sigma_i$, where the V_i and V_{i,I_i} are as specified in (2.1) and (2.2).¹⁰

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.

Assuming a common μ , sequential optimality implies subgame perfection adapted to allow for smooth ambiguity preferences. The only proper subgames start from partial histories where all players have singleton information sets (i.e., all uncertainty (if any) about the past has been resolved). For any such proper subgame, (2.7) ensures that the continuation strategy profile derived from σ forms an ex-ante equilibrium of the subgame. For the overall game, (2.6) ensures σ is an ex-ante equilibrium.

Sequential optimality identifies a set of strategy profiles. Each such profile is sequentially optimal with respect to *some* interim belief system. Recall that we have placed little restriction on how beliefs ν_{i,I_i} at different points in the game relate to one another and to the ex-ante beliefs μ_i . We now show (Theorem 2.1) that every such profile is sequentially optimal with respect to an interim belief system generated by one particular update rule. This generalization of Bayes' rule was proposed by Hanany and Klibanoff (2009) and is called the *smooth rule*, defined as follows:

Definition 2.8 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 $t = s(I_i)$:

⁹As is usual for refinements, when there are no such unreachable information sets, all ex-ante equilibria are also sequentially optimal (see Theorem A.1).

¹⁰Note that since $V_{i,I_i}(\tilde{\sigma}_i, \sigma_{-i}) = V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i})$ if $\tilde{\sigma}_i^{I_i} = \hat{\sigma}_i^{I_i}$, requiring the inequalities for the V_{i,I_i} to hold as i changes only her continuation strategy given I_i would result in an equivalent definition.

If $I_i \subseteq \Theta$, then for all π such that $\pi(I_i) > 0$,

$$\nu_{i,I_i}(\pi) \propto \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi_{I_i, \sigma_{-i}}(h^t) \right)} \pi(I_i) \mu_i(\pi); \quad (2.8)$$

and, if $f_i(I_i) \neq I_i$ and $\nu_{i,I_i^{-1}}(\pi) > 0$ for some π such that $\pi(f_i^{\sigma_{-i}}(I_i)) > 0$, then for all π such that $\pi(f_i^{\sigma_{-i}}(I_i)) > 0$,

$$\begin{aligned} \nu_{i,I_i}(\pi) \propto & \frac{\phi'_i \left(\sum_{h|h^{t-1} \in I_i^{-1}} u_i(h) p_\sigma(h|h^{t-1}) \pi_{I_i, \sigma_{-i}}(h^{t-1}) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi_{I_i, \sigma_{-i}}(h^t) \right)} \\ & \cdot \left(\sum_{h^t \in I_i, \sigma_{-i}} p_{-i, \sigma_{-i}}(h^t|h^{t-1}) \pi_{I_i, \sigma_{-i}}(h^{t-1}) \right) \nu_{i,I_i^{-1}}(\pi). \end{aligned} \quad (2.9)$$

In common with Bayesian updating, the smooth rule imposes no restrictions at information sets that are not reachable from their immediately preceding information set given σ_{-i} , i.e., information sets $I_i \not\subseteq \Theta$ such that $f_i(I_i) = I_i$. Under ambiguity neutrality (ϕ_i linear, which is expected utility), ϕ'_i is constant, and thus the ϕ'_i terms appearing in the formula cancel and the smooth rule becomes standard Bayesian updating of μ_i and the $\nu_{i,I_i^{-1}}$, applied whenever possible. More generally, the ϕ'_i ratio terms, which reflect changes in the motive to hedge against ambiguity (see Hanany and Klibanoff 2009 and Baliga, Hanany and Klibanoff 2013), are the only difference from Bayesian updating. These changes can be motivated via dynamic consistency. For ambiguity averse preferences, Bayesian updating does not ensure dynamic consistency. The smooth rule is dynamically consistent for all ambiguity averse smooth ambiguity preferences when σ is an ex-ante equilibrium (Hanany and Klibanoff 2009 established this for individual decisions).

We now show that, for the purposes of identifying sequentially optimal strategy profiles, restricting attention to beliefs updated according to the smooth rule is *without loss of generality*. Specifically, considering only interim belief systems satisfying the smooth rule yields the entire set of sequentially optimal strategy profiles.

Theorem 2.1 *Fix a game Γ . Suppose (σ, ν) is sequentially optimal. Then, there exists an interim belief system $\hat{\nu}$ satisfying the smooth rule using σ such that $(\sigma, \hat{\nu})$ is sequentially optimal.*

Why is it enough to consider smooth rule updating to identify sequentially optimal profiles? Consider any player i , updating upon moving from information set I_i^{-1} to information set I_i , given a sequentially optimal strategy profile σ . Three ingredients are key to the argument. First, 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 the ambiguity averse preferences used in this paper, this translates as follows: σ_i an interim best response to σ_{-i} for player i at information set I_i given ϕ_i and ν_{i,I_i} is equivalent to (see (A.5)) σ_i maximizing the following expected utility,

$$\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{s(I_i)}) q^{(\sigma,\nu),i,I_i}(h) \quad (2.10)$$

where $q^{(\sigma,\nu),i,I_i}$ is i 's (σ, ν) -local measure given I_i , defined for each $h \in H$ such that $h^{s(I_i)} \in I_i$ by,

$$q^{(\sigma,\nu),i,I_i}(h) \equiv \sum_{\pi | \pi(J_i^{\sigma_{-i}}(I_i)) > 0} \phi'_i \left(\sum_{\hat{h}|\hat{h}^{s(I_i)} \in I_i} u_i(\hat{h}) p_\sigma(\hat{h}|\hat{h}^{s(I_i)}) \pi_{I_i,\sigma_{-i}}(\hat{h}^{s(I_i)}) \right) \cdot p_{-i,\sigma_{-i}}(h|h^{s(I_i)}) \pi_{I_i,\sigma_{-i}}(h^{s(I_i)}) \nu_{i,I_i}(\pi). \quad (2.11)$$

Notice that i 's ambiguity aversion leads this local measure to tilt towards histories weighed more by conditional $\pi_{I_i,\sigma_{-i}}$'s under which i expects to fare less well under σ . As ambiguity aversion increases, this tilting becomes more severe.

Second, perfect recall allows us to conclude that σ_i maximizing the analogue of (2.10) for I_i^{-1} implies that σ_i (and thus also the continuation of σ_i from I_i onward) maximizes

$$\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{s(I_i)}) q^{(\sigma,\nu),i,I_i^{-1}}(h).$$

Third, substituting for $\hat{\nu}_{i,I_i}(\pi)$ using the smooth rule formula in (2.9) yields

$$q^{(\sigma,\nu),i,I_i}(h) \propto q^{(\sigma,\nu),i,I_i^{-1}}(h) \text{ for } \{h \mid h^{s(I_i)} \in I_i\},$$

where this proportionality says that $q^{(\sigma,\nu),i,I_i}(h)$ is the Bayesian update of $q^{(\sigma,\nu),i,I_i^{-1}}(h)$ given I_i and σ . Applying the first ingredient again shows that σ_i is optimal given I_i under smooth rule updating. Thus, wherever Bayes' rule has bite, using smooth rule updating preserves optimality because it implies updating of the local measure ($q^{(\sigma,\nu),i,I_i^{-1}}(h)$) by Bayes' rule. Where Bayes' rule cannot be applied, the smooth rule is unrestrictive and thus any belief

delivering optimality at that point is also consistent with the smooth rule.

Note that Theorem 2.1 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 in that case we would be applying Bayes' rule directly to the interim beliefs which would *not* correspond to Bayesian updating of the local measure, q . A Bayesian version of the theorem is true, however, if we restrict attention to expected utility preferences, for in that case the smooth rule specializes to Bayes' rule. The version of perfect Bayesian equilibrium (PBE) in, for example, Gibbons (1992) imposes sequential optimality (defined using only expected utility preferences) and also that beliefs are related via Bayesian updating wherever possible. From our theorem, it follows that in the expected utility case, sequential optimality alone (i.e., without additionally requiring Bayesian updating) identifies the same set of strategy profiles as sequential optimality plus Bayesian updating. This Bayesian version of the result was shown by Shimoji and Watson (1998) in the context of defining extensive form rationalizability.

Running Example continued: The strategy profile σ^* that we identified as an ex-ante equilibrium in the example is also sequentially optimal. We can make use of Theorem 2.1 to construct interim beliefs supporting this sequential optimality. 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). The theorem tells us that we can, without loss of generality, calculate these updated beliefs using the smooth rule. Doing so gives the following beliefs for agent r after observing β : $\nu_{r,\{ID,IIU,IID\}\times\{\beta\}}(\pi_1) \approx 0.877$ (the explicit calculation is $\frac{\nu_{r,\{ID,IIU,IID\}\times\{\beta\}}(\pi_1)}{\nu_{r,\{ID,IIU,IID\}\times\{\beta\}}(\pi_2)} = \frac{\frac{1}{2} \frac{3}{4} \frac{\phi'_r(1.75)}{\phi'_r(2)}}{\frac{1}{2} \frac{19}{20} \frac{\phi'_r(1.95)}{\phi'_r(2)}}$, recalling that 1.75 and 1.95 are r 's expected utilities under π_1 and π_2 respectively and 2 is r 's expected utility conditional on β under either π). Bayesian updating corresponds to $\nu_{\{ID,IIU,IID\}\times\{\beta\}}^B(\pi_1) = \frac{15}{34} \approx 0.441$. For agent c , who may have any degree of ambiguity aversion including neutrality, $\nu_{c,\{ID,IIU,IID\}\times\{\beta\}}(\pi_1) \leq \frac{15}{34} \approx 0.441$ (the explicit calculation is $\frac{\nu_{c,\{ID,IIU,IID\}\times\{\beta\}}(\pi_1)}{\nu_{c,\{ID,IIU,IID\}\times\{\beta\}}(\pi_2)} = \frac{\frac{1}{2} \frac{3}{4} \frac{\phi'_c(2.75)}{\phi'_c(2)}}{\frac{1}{2} \frac{19}{20} \frac{\phi'_c(2.15)}{\phi'_c(2)}} \leq \frac{15}{19}$; notice, the more ambiguity averse agent c is, the lower $\nu_{c,\{ID,IIU,IID\}\times\{\beta\}}(\pi_1)$ becomes). Recall that π_1 puts more weight on I than does π_2 . Since σ^* performs worse for r under I than under II , agent r does worse under π_1 than under π_2 in equilibrium. Therefore, ambiguity aversion leads r 's smooth rule update to place more weight on π_1 than the Bayesian update. 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 update places weight on I that is (weakly) less than under Bayesian updating. To ensure that both players coordinate on playing w after β , differing updated beliefs are crucial: if they shared a common updated belief, at least one agent would deviate. Thus ambiguity helps the principal implement σ^* by leading the agents' updated beliefs to move apart.¹¹

As the discussion of Theorem 2.1 indicated, one may think of smooth rule updating as ensuring Bayesian updating of the local measure. If we focus on the update after β , these updated local measures give the following marginal probabilities on I vs. II : For agent r , I is given probability ≈ 0.604 , leading w to indeed be better than b following β since $2 > 5(1 - 0.604)$. For agent c , the probability of I is at most $\frac{13}{34} \approx 0.382$, leading w to indeed be better than b following β since $2 > 5(0.382)$.¹² These show directly that r places more weight on I than c after observing β and verify the interesting part of sequential optimality for σ^* .

Remark 2.3 Aside from the specifics of the game, an important difference from previous analyses in our example is in the dynamics – demonstrating that strategic ambiguity occurs as part of a sequential optimum. In particular, the strategies involving strategic ambiguity in Bose and Renou (2014) and Kellner and Le Quement (2015), papers containing the most closely related examples, satisfy weaker concepts such as no profitable one-stage deviations but fail to be sequentially (and even ex-ante) optimal.

■

We next establish that smooth rule updating ensures (σ, ν) is sequentially optimal if, 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. Thus we establish that the important method from standard theory of checking only one-stage deviations (as opposed to general deviations) when verifying optimality retains its validity when applied to sequential optimality. Formally, the absence of these profitable one-stage deviations is the following property:

¹¹In the context of individual decisions, such belief polarization under ambiguity is explored in Baliga, Hanany and Klibanoff (2013).

¹²The explicit calculations are

$$\frac{q^{(\sigma^*, \nu), r, \beta}(ID, \beta, w, w)}{\sum_{\theta \in \{ID, IIU, IID\}} q^{(\sigma^*, \nu), r, \beta}(\theta, \beta, w, w)} = \frac{\frac{1}{2} \frac{1}{4} \phi'_r(1.75) + \frac{1}{2} \frac{3}{40} \phi'_r(1.95)}{\frac{1}{2} \frac{3}{8} \phi'_r(1.75) + \frac{1}{2} \frac{19}{40} \phi'_r(1.95)}$$

and

$$\frac{q^{(\sigma^*, \nu), c, \beta}(ID, \beta, w, w)}{\sum_{\theta \in \{ID, IIU, IID\}} q^{(\sigma^*, \nu), c, \beta}(\theta, \beta, w, w)} = \frac{\frac{1}{2} \frac{1}{4} \phi'_c(2.75) + \frac{1}{2} \frac{3}{40} \phi'_c(2.15)}{\frac{1}{2} \frac{3}{8} \phi'_c(2.75) + \frac{1}{2} \frac{19}{40} \phi'_c(2.15)} \leq \frac{13}{34}.$$

Definition 2.9 *The pair (σ, ν) satisfies the no profitable one-stage deviation property 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 .*

The potential complication in extending the sufficiency of no profitable one-stage deviations to our setting is the fact that players' preferences are non-separable across information sets. Smooth rule updating overcomes this, as the next theorem shows. As was true for Theorem 2.1, it is the combination of perfect recall, the existence of local linear approximations representing optima and Bayesian updating of the corresponding local measures that ensures the sufficiency of no profitable one-stage deviations.

Theorem 2.2 *If (σ, ν) satisfies the no profitable one-stage deviation property and ν satisfies smooth rule updating using σ then (σ, ν) is sequentially optimal.*

It follows that a strategy profile is part of a sequential optimum *if and only if* there are no profitable one-stage deviations with respect to beliefs updated according to the smooth rule. As a result, fixing a strategy profile σ , at any information set one may check for any player i whether σ_i is an optimal strategy for i for the remainder of the game using a “folding back” algorithm. Note that, unlike for consistent planning where folding back is built in by construction, this needed to be established for sequential optimality. It works as follows – for each information set for player i at the final stage, calculate an optimal mixture over the actions she has available at that stage given the information set, the strategies of the other players and i 's beliefs at that information set. Then, holding these optimal mixtures fixed, repeat this process for information sets one stage earlier in the game. Continue backwards iteratively. The only thing that (in common with the standard approach under ambiguity neutrality) cannot be calculated via folding back are the beliefs at each information set. These may be determined by updating according to the smooth rule (which reduces to Bayes' rule under ambiguity neutrality) using σ . Recall that smooth rule updating is without loss of generality for the purposes of identifying sequentially optimal strategies.

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

2.4 Sequential Equilibrium with Ambiguity

To define sequential equilibrium with ambiguity (SEA), we consider an auxiliary condition, *smooth rule consistency*, that imposes requirements on beliefs even at those points where sequential optimality has no implications for updating. These points are those for player i

that are not reachable from their immediately preceding information set given σ_{-i} .¹³ 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 (and adapting from ν using σ – see Definition 2.10). 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 Bayesian updating, then sequentially optimal strategies might fail to exist under ambiguity aversion.

In order to take limits of smooth rule updates derived from a sequence of strategy profiles, these updated beliefs need to be expressed on a common space. This is the purpose of the next definition, which transforms an interim belief ν_{i,I_i} (a distribution over $\Delta(H^{m_i(I_i)})$, a space that depends, via $m_i(I_i)$, on the strategy profile) into a distribution $\hat{\nu}_{i,I_i}$ over $\Delta(I_i)$ (a space independent of the strategy profile).

Definition 2.10 (Belief Adaptation) *For a strategy profile σ and an interim belief system ν , say that $\hat{\nu} \equiv (\hat{\nu}_{i,I_i})_{i \in N, I_i \in \mathcal{I}_i}$ is adapted from ν using σ if, for each player i , information set $I_i \in \mathcal{I}_i$ and $\hat{\pi} \in \Delta(I_i)$, $\hat{\nu}_{i,I_i}(\hat{\pi}) = \sum_{\pi \in \text{supp } \nu_{i,I_i} | \pi_{I_i, \sigma_{-i}} = \hat{\pi}} \nu_{i,I_i}(\pi)$.¹⁴*

Using limits of adapted beliefs, smooth rule consistency is defined as follows:

Definition 2.11 (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 ν is determined by smooth rule updating using σ , and there exists a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$, with $\lim_{k \rightarrow \infty} \sigma^k = \sigma$, such that $\hat{\nu} = \lim_{k \rightarrow \infty} \hat{\nu}^k$, where $\hat{\nu}$ is adapted from ν using σ , each $\hat{\nu}^k$ is adapted from ν^k using σ^k , and each ν^k is the interim belief system determined by smooth rule updating 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 2.12 (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.*

¹³If there are no such information sets, any sequentially optimal strategy profile is also part of an SEA (see Theorem A.3). Observe that in our running example this implies that the previously identified equilibrium profile σ^* is also part of an SEA.

¹⁴supp denotes the support of the measure in its argument.

We 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.

Theorem 2.3 *An SEA exists for any game Γ .*

We may use Theorem 2.2 to conclude that replacing sequential optimality in the definition of SEA by the no profitable one-stage deviation property would not change the set of equilibrium strategies.

Corollary 2.1 *(σ, ν) satisfies the no profitable one-stage deviation property and smooth rule consistency if and only if (σ, ν) is an SEA.*

Finally, we show that, for the purposes of identifying SEA strategy profiles, restricting attention to beliefs defined according to a version of the smooth rule using limits of likelihoods (see Theorem A.2 for details) is *without loss of generality*, just as using the smooth rule was in identifying sequentially optimal strategy profiles. This provides a constructive method for determining when a sequential optimum is part of an SEA.

Theorem 2.4 *Fix a game Γ and a strategy profile σ . If, for some interim belief system ν , (σ, ν) is an SEA (with corresponding sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$), then $(\sigma, \bar{\nu})$, for $\bar{\nu}$ as described in Theorem A.2, is also an SEA with the same corresponding sequence.*

The proof uses the fact that a strategy profile is part of an SEA to establish the existence of the limits of likelihoods needed to invoke Theorem A.2. The latter theorem is then used to prove smooth rule consistency for the constructed interim belief system $\bar{\nu}$. Furthermore, all interim belief systems satisfying smooth rule consistency with respect to the same sequence of completely mixed strategy profiles are shown to deliver the same interim preferences, implying that $(\sigma, \bar{\nu})$ is sequentially optimal as well.

2.5 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: Do changes in ambiguity aversion affect the set of equilibrium strategy profiles (and play paths) for fixed beliefs? The answer is yes they can, as was true in our running example. Such changes may lead to entirely non-overlapping sets of equilibrium strategy profiles (see Theorem B.1 for a formal result).

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? Formally, this corresponds to taking the union over beliefs μ of the set of equilibria generated if the common belief were μ . How do such predictions change when ambiguity aversion is introduced? We show that ambiguity aversion makes more equilibria (in the superset sense) possible, and that at least for some games, this expansion is strict (see Theorem B.2). 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 a simple three stage game in which this occurs, see the example in Appendix B.1.

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 (see Theorem B.3).

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

It is interesting to note that Battigalli et al. (2015a) emphasize a different result (their Theorem 1 together with an example of strict inclusion), 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). For the sake of comparison, if we were also to limit attention to pure strategies in both these respects, an analogous result would apply to our equilibria. In other words, we show that under the pure strategy limitation, increases in ambiguity aversion make more equilibria (in the superset sense) possible, and that this expansion may be strict (see Theorem B.4).

2.6 Robustness

Here we relate ambiguity aversion to a type of robustness of equilibria in the sense of the range of beliefs that support an equilibrium. An equilibrium supported for a wider range of beliefs is in a natural sense more robust.

We propose definitions of robustness to increased ambiguity aversion and of belief robustness and show that robustness to increased ambiguity aversion implies more belief robustness. 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 2.13 *For a game Γ , an ex-ante equilibrium σ is ex-ante robust to increased ambiguity aversion if it remains an ex-ante 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 more belief robust if at least some increases in players' ambiguity aversion, holding the π 's in the supports of players' beliefs $(\mu_i)_{i \in N}$ fixed, expand the set of beliefs supporting that profile as an equilibrium:

Definition 2.14 *For a game Γ , consider an ex-ante equilibrium σ . Ambiguity aversion makes σ ex-ante more belief robust if there exist, for each i , a continuum of $\hat{\phi}_i$ strictly more concave than ϕ_i such that, for each $(\hat{\mu}_i)_{i \in N}$ with the same supports as the $(\mu_i)_{i \in N}$, if σ is an ex-ante equilibrium of the game with $(\hat{\mu}_i)_{i \in N}$ and $(\phi_i)_{i \in N}$, then σ is also an ex-ante equilibrium of the game with $(\hat{\mu}_i)_{i \in N}$ and $(\hat{\phi}_i)_{i \in N}$. It makes σ ex-ante strictly more belief robust if, in addition, for each such $\hat{\phi}_i$, there exist $(\tilde{\mu}_i)_{i \in N}$ with the same supports as the $(\mu_i)_{i \in N}$ such that σ is not an ex-ante equilibrium of the game with $(\tilde{\mu}_i)_{i \in N}$ and $(\phi_i)_{i \in N}$, but is an ex-ante equilibrium of the game with $(\tilde{\mu}_i)_{i \in N}$ and $(\hat{\phi}_i)_{i \in N}$.*

Sequentially optimal robust to increased ambiguity aversion and *Sequentially optimal (strictly) more belief robust* are defined analogously, replacing ex-ante equilibrium with sequentially optimal in the above two definitions.

SEA robust to increased ambiguity aversion and *SEA (strictly) more belief robust* are defined by, in addition to replacing ex-ante equilibrium with SEA in each definition, requiring that there exists 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 for the assumed SEAs are simultaneously satisfied.

The equilibrium σ^* we analyzed in the running example is ex-ante, sequentially optimal and SEA robust to increased ambiguity aversion. It is also (all three varieties of) strictly more belief robust. In light of the robustness to increased ambiguity aversion, the fact that

σ^* is more belief robust is not an accident. The next result shows, under some general conditions, that equilibria that are robust to increased ambiguity aversion *must* have their sets of supporting beliefs expand under some increases in ambiguity aversion.

Theorem 2.5 *Fix a game Γ . If an equilibrium σ is ex-ante (resp. sequentially optimal or SEA) robust to increased ambiguity aversion and, for each player i , over the support of μ_i , $\sum_{h \in H} u_i(h)p_\sigma(h|h^0)\pi(h^0)$ either has a unique minimizing π or is constant in π , then ambiguity aversion makes σ ex-ante (resp. sequentially optimal or SEA) more belief robust. If, additionally, for some player i , $\sum_{h \in H} u_i(h)p_\sigma(h|h^0)\pi(h^0)$ can be strictly ordered across the π in the support of μ_i and there exists a $\tilde{\mu}_i$ with the same support as μ_i such that σ_i is not a best response to σ_{-i} given $\tilde{\mu}_i$ and ϕ_i , then ambiguity aversion makes σ ex-ante (resp. sequentially optimal or SEA) strictly more belief robust.*

An outline of the proof for the first (non-strict) part of the theorem in the ex-ante case is as follows: (1) establish via the local linear approximation of i 's objective function around σ_i (i.e., the ex-ante analogue of the expression in (2.10)) that σ_i an ex-ante best response to σ_{-i} for player i given any $\hat{\phi}_i$ and $\hat{\mu}_i$ is equivalent to

$$\sigma_i \in \arg \max_{\sigma'_i} \sum_{\pi} \left(\sum_{h \in H} u_i(h)p_{(\sigma'_i, \sigma_{-i})}(h|h^0)\pi(h^0) \right) \hat{\phi}'_i \left(\sum_{h \in H} u_i(h)p_\sigma(h|h^0)\pi(h^0) \right) \hat{\mu}_i(\pi); \quad (2.12)$$

(2) if $\sum_{h \in H} u_i(h)p_\sigma(h|h^0)\pi(h^0)$ is constant in π , then the $\hat{\phi}'_i$ term is constant in (2.12) and increases in ambiguity aversion leave the set of $\hat{\mu}_i$ supporting σ_i unchanged; (3) in all other cases, it is possible to increase ambiguity aversion in a way that is equivalent to placing more weight on the lowest value of $\sum_{h \in H} u_i(h)p_\sigma(h|h^0)\pi(h^0)$ in the maximand in (2.12) while maintaining the relative weights in that expression across the remaining expected utilities (doing so corresponds to applying any of the concave transformations ψ_i^1 – see Definition A.1 in the Appendix – to ϕ_i); (4) the rest of the argument assumes a unique minimizing π , denoted π_i^1 , and it then follows from step 3 that robustness to increased ambiguity aversion implies that σ_i maximizes i 's expected utility under π_i^1 given σ_{-i} ; and, finally, (5) for any $\hat{\mu}_i$ that supports σ_i as a best response to σ_{-i} , increasing ambiguity aversion by applying any of the continuum of concave transformations ψ_i^1 maintains σ_i as a best response because the maximand in (2.12) may be shown to be a linear combination of i 's expected utility under π_i^1 and the maximand in (2.12) using the original ϕ_i .

Note that without the assumption of a unique minimizing π or constancy in π the first part of the theorem would be false. To see this, observe that with two minimizing π 's, it might be that one supports σ_i as a best response and the other does not. Then, without constancy

in π , increasing ambiguity aversion using ψ_i^1 , could well, depending on the particular $\hat{\mu}_i$, push enough weight onto the non-supporting minimizing π relative to the supporting one that σ_i is no longer a best response. If there were only one other π (so a total of three) in the support of μ_i , all increases in ambiguity aversion shift weights in the same way as the transformations ψ_i^1 and no alternative increases in ambiguity aversion that might rescue the result are available.

Similarly, even with the unique minimizer assumption, without the strict ordering assumption the second part of the theorem would be false. Without it, the manner in which more ambiguity aversion can shift weights among the expected utilities in the maximand in (2.12) may be too limited to make new beliefs support σ . For example, suppose there are exactly three π 's in the support of μ_i , and two of them are tied for generating the highest expected utility under σ – in such a case, increasing ambiguity aversion can only push weight to the minimizing π while maintaining the relative weights among the top two. If σ_i were tied with another strategy at the minimizing π , this other strategy gave a higher expected utility than σ_i under one of the top two π and lower under the other, and convex combinations of these strategies form the feasible set for i , then increasing weight on the π where σ_i does worse is the only way to generate beliefs $\hat{\mu}_i$ that make σ_i non-optimal. Given any such $\hat{\mu}_i$, increasing ambiguity aversion would not be able to improve the relative standing of σ_i over the other strategy, because it cannot affect the relative weights on the top two π .

3 Example: Limit Pricing under Ambiguity

In this section, we use a parametric class of games based on the 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 refines ex-ante equilibrium and we find conditions under which this ambiguity aversion makes limit pricing behavior more robust compared to the standard case with ambiguity neutrality.

The game is as follows: An incumbent monopolist has private information concerning his per-unit production costs c_I (which is one of $c_L < c_M < c_H$). Thus the parameter space is $\Theta = \{L, M, H\}$.¹⁵ 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

¹⁵The use of at least three costs is necessary to have non-trivial updating on the equilibrium path under pure strategy limit pricing. With only two possible costs, pure limit pricing strategies involve full pooling.

same market demand and costs as in the first stage, and the entrant gets a payoff of zero. If entry is chosen, the entrant pays a fixed cost $K \geq 0$, the incumbent's cost is learned by the entrant, and in the final stage the two firms compete in a complete information Cournot duopoly with the same market demand. To make this a finite game, suppose a finite set of feasible quantities \mathcal{Q} (including at least the monopoly quantities for each possible production cost and the complete information Cournot quantities).¹⁶ 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 3.1 $a, b > 0$, $K \geq 0$, $c_H > c_M > c_L \geq 0$, $c_E \geq 0$, $a > c_H$, $a + c_E - 2c_H > 0$ and $a + c_L - 2c_E > 0$.

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

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

¹⁶The 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.

Proposition 3.1 *Under Assumption 3.1, the limit pricing strategy profile σ^{LP} is part of an SEA if*

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

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

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

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

Moreover, under the same conditions, σ^{LP} is SEA robust to increased ambiguity aversion, and ambiguity aversion makes it SEA more belief robust. If, additionally some $\pi \in \text{supp } \mu$ makes entry conditional on $\{L, M\}$ strictly profitable, then ambiguity aversion makes σ^{LP} SEA strictly more belief robust.

The proof uses the method provided by Theorem 2.4 to establish that σ^{LP} is part of an SEA and uses Theorem 2.5 to establish belief robustness. One observation following from the above result is 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, then there exists a strictly increasing and twice continuously differentiable concave function ϕ such that the entrant would be deterred by the limit price. In this way, increasing ambiguity aversion leads to expansion in the set of μ that can support such a limit pricing SEA. Here is a numerical example that illustrates this by having the equilibrium supported for the given μ under sufficient ambiguity aversion but not under ambiguity neutrality: μ puts equal weight on type distributions $\pi_1 = (1/6, 1/3, 1/2)$ and $\pi_2 = (1/2, 1/3, 1/6)$, where the vector notation gives the probabilities of L, M, H respectively, $a = 2, b = \frac{7}{128}, c_H = \frac{3}{2}, c_M = \frac{11}{8}, c_L = 1, c_E = \frac{5}{4}$ and $K = 1$. With these parameters, the entrant's Cournot profits net of entry cost as a function of the incumbent's type are: $w_L = -\frac{31}{63}, w_M = \frac{35}{63}$ and $w_H = \frac{65}{63}$. Under ambiguity neutrality, since $\frac{w_L + w_M}{2} > 0$ the entrant would want to enter after observing the limit price. If, however, $\phi(x) = -e^{-\alpha x}$, with $\alpha > \frac{189}{65} \log\left(\frac{39}{23}\right) \approx 1.54$, the limit price would deter entry and σ^{LP} would be an equilibrium.

4 Extensions

4.1 Other models of ambiguity averse players

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

By modifying our framework to specify V_i rather than μ_i and ϕ_i the definition of ex-ante preferences and equilibrium are easily adapted. However, since such preferences do not necessarily have separately specified beliefs and ambiguity aversion, the notion of interim belief system would need to be replaced by an interim preference system (i.e., an interim preference for each player and information set). Given that change, sequential optimality could be defined. Based on our proof of Theorem 2.1, we conjecture the following version of this result would be true: for the purpose of identifying sequentially optimal strategies, it is without loss of generality to limit attention to interim preference systems 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 . Observe that there are two key differences from our current result: 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). Similarly, under such updating, we conjecture that our results on the sufficiency of checking for no profitable one-stage deviations should also go through.

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.

4.2 Implementation of mixed actions

Recall that the objects of choice of a player are behavior strategies, which, for each type of the player, specify a mixture over the available actions at each point in the game where the player has an opportunity to move. Suppose at some point a player's strategy specifies a non-degenerate mixture, and, as can happen under ambiguity aversion, this strategy is strictly better than any specifying a pure action. If such a mixture is to be implemented by means of playing pure actions contingent on the outcome of a (possibly existing in the player's mind only) randomization device, then an additional sequential optimality concern beyond that formally reflected in Definition 2.7 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 2.7 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.

References

- [1] S. AUSTER, Robust contracting under common value uncertainty, *Theoretical Economics*, *forthcoming*.
- [2] M. AYOUNI, F. KOESSLER, Hard Evidence and Ambiguity Aversion, *Theory and Decision*, **82** (2017), 327-339.

- [3] Y. AZRIELI, R. TEPER, Uncertainty aversion and equilibrium existence in games with incomplete information, *Games and Economic Behavior* **73(2)** (2011), 310-317.
- [4] S. BADE, Ambiguous act equilibria, *Games and Economic Behavior* **71(2)** (2011), 246-260.
- [5] S. BADE, Dynamic Semi-Consistency, mimeo. (2016).
- [6] S. BALIGA, E. HANANY, P. KLIBANOFF, Polarization and Ambiguity, *American Economic Review* **103(7)** (2013), 3071-3083.
- [7] P. BATTIGALLI , S. CERREIA-VIOGLIO, F. MACCHERONI, M. MARINACCI, Self-confirming Equilibrium and Model Uncertainty, *American Economic Review* **105(2)** (2015a), 646-77.
- [8] P. BATTIGALLI , S. CERREIA-VIOGLIO, F. MACCHERONI, M. MARINACCI, Selfconfirming Equilibrium and Sequential Rationality with Ambiguity Averse Agents, mimeo. (2015b).
- [9] A.L. BODOH-CREED, Ambiguous beliefs and mechanism design, *Games and Economic Behavior* **75(2)** (2012), 518-537.
- [10] S. BOSE, A. DARIPA, A dynamic mechanism and surplus extraction under ambiguity, *Journal of Economic Theory* **144(5)** (2009), 2084–2114.
- [11] S. BOSE, E. OZDENOREN, A. PAPE, Optimal Auctions with Ambiguity, *Theoretical Economics* **1(4)** (2006), 411-438.
- [12] S. BOSE, L. RENO, Mechanism Design With Ambiguous Communication Devices, *Econometrica* **82** (2014), 1853–1872.
- [13] S. CERREIA-VIOGLIO, F. MACCHERONI, M. MARINACCI, L. MONTRUCCHIO, Uncertainty Averse Preferences, *Journal of Economic Theory* **146(4)** (2011), 1275-1330.
- [14] Y. CHEN, P. KATUSCAK, E. OZDENOREN, Sealed Bid Auctions with Ambiguity: Theory and Experiments, *Journal of Economic Theory*, **136(1)** (2007), 513–535.
- [15] A. DOMINIAK, M. S. LEE, Coherent Dempster-Shafer Equilibrium and Ambiguous Signals, *Journal of Mathematical Economics*, **68** (2017), 42–54.
- [16] J. EICHBERGER, D. KELSEY, Education-Signalling and Uncertainty, in Machina, M.J., and Munier, B.R. (eds.), *Beliefs, Interactions and Preferences in Decision Making* (1999), 135-157, Kluwer.

- [17] J. EICHBERGER, D. KELSEY, Sequential Two-Player Games with Ambiguity, *International Economic Review* **45** (2004), 1229-1261.
- [18] A. ELLIS, On Dynamic Consistency in Ambiguous Games, mimeo., (2016).
- [19] L. G. EPSTEIN, M. SCHNEIDER, IID: independently and indistinguishably distributed, *Journal of Economic Theory*, **113(1)** (2003), 32–50.
- [20] D. FUDENBERG, J. TIROLE, *Game Theory* (1991a), MIT Press.
- [21] D. FUDENBERG, J. TIROLE, Perfect Bayesian Equilibrium and Sequential Equilibrium, *Journal of Economic Theory*, **53** (1991b), 236–260.
- [22] P. GHIRARDATO, Defining Ambiguity and Ambiguity Aversion, *Uncertainty in Economic Theory: A Collection of Essays in Honor of David Schmeidler's 65th Birthday*, I. Gilboa. et al. (Eds.), (2004), London: Routledge.
- [23] P. GHIRARDATO, M. SINISCALCHI, Ambiguity in the small and in the large, *Econometrica* **80** (2012), 2827-2847.
- [24] R. GIBBONS, *Game Theory for Applied Economists* (1992), Princeton University Press.
- [25] I. GILBOA, M. MARINACCI, Ambiguity and the Bayesian Paradigm, *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society*, D. Acemoglu, M. Arellano, and E. Dekel (Eds.), (2013), New York: Cambridge University Press.
- [26] I. GILBOA, D. SCHMEIDLER, Maxmin Expected Utility with a Non-Unique Prior, *Journal of Mathematical Economics*, **18** (1989), 141-153.
- [27] I.L. GLICKSBERG, A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points, *Proceedings of the American Mathematical Society* **3(1)** (1952), 170-174.
- [28] J. GREENBERG, The Right to Remain Silent, *Theory and Decision*, **48** (2000), 193–204.
- [29] E. HANANY, P. KLIBANOFF, Updating Preferences with Multiple Priors, *Theoretical Economics* **2** (2007), 261-298.
- [30] E. HANANY, P. KLIBANOFF, Updating Ambiguity Averse Preferences, *The BE Journal of Theoretical Economics (advances)* **9(1)** (2009), article 37.

- [31] L. P. HANSEN, Beliefs, Doubts and Learning: Valuing Macroeconomic Risk, *American Economic Review* **97(2)** (2007), 1-30.
- [32] A. KAJII, T. UI, Incomplete Information Games with Multiple Priors, *Japanese Economic Review* **56(3)** (2005), 332-351.
- [33] C. KELLNER, Tournaments as a Response to Ambiguity Aversion in Incentive Contracts, *Journal of Economic Theory*, **159(A)** (2015), 627-655.
- [34] C. KELLNER, M. LE QUEMENT, Endogenous ambiguity in cheap talk, mimeo. University of Bonn (2015).
- [35] C. KELLNER, M. LE QUEMENT, Modes of ambiguous communication, *Games and Economic Behavior* **104** (2017), 271-292.
- [36] P. KLIBANOFF, M. MARINACCI, S. MUKERJI, A Smooth Model of Decision Making under Ambiguity, *Econometrica* **73** (2005), 1849-1892.
- [37] P. KLIBANOFF, M. MARINACCI, S. MUKERJI, Recursive Smooth Ambiguity Preferences, *Journal of Economic Theory* **144** (2009), 930-976.
- [38] D.M. KREPS, R. WILSON, Sequential Equilibria, *Econometrica* **50(4)** (1982), 863-894.
- [39] C. KUZMICS, Abraham Wald's complete class theorem and Knightian uncertainty, *Games and Economic Behavior*, *forthcoming*.
- [40] K.C. LO, Extensive Form Games with Uncertainty Averse Players, *Games and Economic Behavior* **28** (1999), 256-270.
- [41] G. LOPOMO, L. RIGOTTI, C. SHANNON, Uncertainty in Mechanism Design, mimeo. (2010).
- [42] F. MACCHERONI, M. MARINACCI, A Strong Law of Large Numbers for Capacities, *Annals of Probability* **33** (2005), 1171-1178.
- [43] F. MACCHERONI, M. MARINACCI, A. RUSTICHINI, Ambiguity aversion, robustness, and the variational representation of preferences, *Econometrica* **74(6)** (2006), 1447-1498.
- [44] P. MILGROM, J. ROBERTS, Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis, *Econometrica* **50(2)** (1982), 443-459.

- [45] I. MURAVIEV, F. RIEDEL, L. SASS, Kuhn's Theorem for Extensive Form Ellsberg Games, *Journal of Mathematical Economics*, **68** (2017), 26-41.
- [46] S. MUKERJI, J-M. TALLON, An overview of economic applications of David Schmeidler's models of decision making under uncertainty. In Gilboa, I. (ed.), *Uncertainty in Economic Theory* (2004), 283–302, Routledge, London.
- [47] R.B. MYERSON, *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA (1991).
- [48] E. OZDENOREN, D. LEVIN, Auctions with Uncertain Numbers of Bidders, *Journal of Economic Theory*, **118(2)**, (2004) 229-251.
- [49] F. RIEDEL, L. SASS, Ellsberg games, *Theory and Decision* **76** (2013), 469-509.
- [50] L. RIGOTTI, C. SHANNON, T. STRZALECKI, Subjective Beliefs and Ex Ante Trade, *Econometrica* **76** (2008), 1167-1190.
- [51] A. SALO, M. WEBER, Ambiguity Aversion in First Price Sealed Bid Auctions, *Journal of Risk and Uncertainty* **11** (1995), 123-137.
- [52] R. SELTEN, Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit, *Zeitschrift für die gesamte Staatswissenschaft* **121**, (1965), 301-324 and 667-689.
- [53] M. SHIMOJI, J. WATSON, Conditional Dominance, Rationalizability, and Game Forms, *Journal of Economic Theory* **83** (1998), 161-195.
- [54] R.H. STROTZ, Myopia and Inconsistency in Dynamic Utility Maximization, *The Review of Economic Studies* **23(3)** (1955-56), 165-180.
- [55] A. DI TILLIO, N. KOS, M. MESSNER, The Design of Ambiguous Mechanisms, *The Review of Economic Studies* **84(1)** (2017), 237-276.
- [56] A. WOLITZKY, Bilateral Trading with Maxmin Agents, mimeo. Stanford University (2013).
- [57] A. WOLITZKY, Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade, *Theoretical Economics* **11** (2016), 971-1004.

A Appendix: Proofs

Proof of Propositions 2.1. We first establish that the given strategies form an ex-ante equilibrium. P 's strategy is an ex-ante best response because it leads to payoff 2 for all parameters, which is the highest feasible payoff for this player. Let γ_m be the probability with which agent r plays w after message $m \in \{\alpha, \beta\}$, and similarly let δ_m be the corresponding probabilities for agent c . The proposed strategies correspond to $\gamma_\alpha = \gamma_\beta = \delta_\beta = 1$ and $\delta_\alpha = 0$. We now verify that these are ex-ante best responses. Denoting $\pi_k(IIU) + \pi_k(IID)$ by $\pi_k(II)$, given the strategies of the others, r maximizes

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

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

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

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

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

Since this function is strictly decreasing in δ_α , it is clearly maximized at $\delta_\alpha = 0$. The first derivative with respect to δ_β evaluated at $\delta_\alpha = 0$ and $\delta_\beta = 1$ is

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

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

Proof of Proposition 2.2.

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

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

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

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

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

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

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

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

Case 2: When $0 < \rho_{II} < 1$, since under II , P sends both messages with positive probability, it is necessary that w, w is played with probability 1 after both messages in order that the principal always receive the maximal payoff of 2. A necessary condition for this to be a best response for c is that the partial derivatives of (A.2) with respect to $\delta_\alpha, \delta_\beta$ are non-negative at $\gamma_\alpha = \gamma_\beta = \delta_\alpha = \delta_\beta = 1$. This is, respectively, equivalent to $14\rho_{II} \geq 19\rho_I$ and $14(1 - \rho_{II}) \geq 19(1 - \rho_I)$, which implies $14 \geq 19$, a contradiction.

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

for the π_k is,

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

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

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

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

Lemma A.1 *Fix a game Γ , a (σ, ν) such that σ is an ex-ante equilibrium, a player i and an information set I_i . If ν_{i,I_i} is derived from $\nu_{i,f_i(I_i)}$ with $\nu_{i,f_i(I_i)}(\pi) > 0$ for some π such that $\pi(f_i^{\sigma^{-i}}(I_i)) > 0$ (or, if $s(I_i) = 0$, from μ_i) via the smooth rule using σ and, for all $\sigma'_i \in \Sigma_i$,*

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

(or, if $s(I_i) = 0$, given ex-ante optimality), then, for all $\sigma'_i \in \Sigma_i$,

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

It is useful for the proof of this and other results to consider a player's "local ambiguity neutral measure" at an information set: Given (σ, ν) , for any player i , let $q^{\sigma,i}(h)$ denote i 's ex-ante σ -local measure, defined for each $h \in H$ by

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

additionally, as in (2.11), for any information set I_i , $q^{(\sigma, \nu), i, I_i}$ denotes i 's (σ, ν) -local measure given I_i , defined for each $h \in H$ such that $h^{s(I_i)} \in I_i$ by,

$$q^{(\sigma, \nu), i, I_i}(h) \equiv \sum_{\pi \in \Delta(H^{m_i(I_i)}) | \pi(f_i^{\sigma^{-i}}(I_i)) > 0} \phi'_i \left(\sum_{\hat{h} | \hat{h}^{s(I_i)} \in I_i} u_i(\hat{h}) p_\sigma(\hat{h} | \hat{h}^{s(I_i)}) \pi_{I_i, \sigma_{-i}}(\hat{h}^{s(I_i)}) \right) \cdot p_{-i, \sigma_{-i}}(h | h^{s(I_i)}) \pi_{I_i, \sigma_{-i}}(h^{s(I_i)}) \nu_{i, I_i}(\pi). \quad (\text{A.4})$$

Proof of Lemma A.1. The inequalities $V_{i,I_i}(\sigma) \geq V_{i,I_i}(\sigma'_i, \sigma_{-i})$ (respectively, $V_i(\sigma) \geq V_i(\sigma'_i, \sigma_{-i})$) for all σ'_i are equivalent (see Hanany and Klibanoff 2009, Lemma A.1) to the

condition that $\sigma'_i = \sigma_i$ maximizes

$$\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{s(I_i)}) q^{(\sigma,\nu),i,I_i}(h), \quad (\text{A.5})$$

where $q^{(\sigma,\nu),i,I_i}$ is i 's (σ,ν) - local measure given I_i (defined in (A.4)), (respectively, $\sigma'_i = \sigma_i$ maximizes

$$\sum_h u_i(h) p_{i,\sigma'_i}(h|h^0) q^{\sigma,i}(h), \quad (\text{A.6})$$

where $q^{\sigma,i}$ is i 's ex-ante σ - local measure (defined in (A.3)).

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 (A.5).

Consider the case where $s(I_i) > 0$ (the case where $s(I_i) = 0$ is similar, using (A.6) instead of (A.7), and is omitted). By assumption in the statement of the lemma, $V_{i,f_i(I_i)}(\sigma) \geq V_{i,f_i(I_i)}(\sigma'_i, \sigma_{-i})$ for all $\sigma'_i \in \Sigma_i$. As in (A.5), this is equivalent to the condition that $\sigma'_i = \sigma_i$ maximizes

$$\sum_{h|h^{m_i(I_i)} \in f_i(I_i)} u_i(h) p_{i,\sigma'_i}(h|h^{m_i(I_i)}) q^{(\sigma,\nu),i,f_i(I_i)}(h). \quad (\text{A.7})$$

Notice that, since i 's strategy is a function only of i 's information sets and, by perfect recall, $R_i(h^{s(I_i)}) = R_i(I_i)$ for any h such that $h^{s(I_i)} \in I_i$, $p_{i,\sigma'_i}(h^{s(I_i)}|h^{m_i(I_i)})$ is the same for any such h . Thus, the objective function in (A.7) can be equivalently written as

$$\begin{aligned} & \sum_{\substack{h|h^{m_i(I_i)} \in f_i(I_i) \\ \text{and } h^{s(I_i)} \notin I_i}} u_i(h) p_{i,\sigma'_i}(h|h^{m_i(I_i)}) q^{(\sigma,\nu),i,f_i(I_i)}(h) \\ & + p_{i,\sigma'_i}(\bar{h}^{s(I_i)}|\bar{h}^{m_i(I_i)}) \sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{s(I_i)}) q^{(\sigma,\nu),i,f_i(I_i)}(h) \end{aligned}$$

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

$$\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{s(I_i)}) q^{(\sigma,\nu),i,f_i(I_i)}(h) \quad (\text{A.8})$$

is affected by the specification of σ'_i from I_i onward and no other part of σ'_i affects (A.8). Therefore (A.7) implies that σ_i maximizes (A.8). For that to imply σ_i maximizes (A.5), it is sufficient to show that $q^{(\sigma,\nu),i,I_i}(h) \propto q^{(\sigma,\nu),i,f_i(I_i)}(h)$ holds for $\{h \mid h^{s(I_i)} \in I_i\}$. This proportionality may be shown by using the local measure definition (A.4), applying the smooth rule iteratively to substitute for $\nu_{i,I_i}(\pi)$ for all $\pi \in \Delta(H^{m_i(I_i)})$ such that $\pi(f_i^{\sigma-i}(I_i)) > 0$ (as $\nu_{i,I_i}(\pi) = 0$ for other π) and then using the definitions of $\pi_{I_i,\sigma_{-i}}$ and $\pi_{f_i(I_i),\sigma_{-i}}$ and cancelling

terms. ■

Proof of Theorem 2.1. We show that $(\sigma, \hat{\nu})$, where, for all i, I_i , if $I_i \subseteq \Theta$ or $[f_i(I_i) \neq I_i$ and $\nu_{i, f_i(I_i)}(\pi) > 0$ for some π such that $\pi(f_i^{\sigma^{-i}}(I_i)) > 0]$, $\hat{\nu}_{i, I_i}$ is derived via the smooth rule, and $\hat{\nu}_{i, I_i} = \nu_{i, I_i}$ everywhere else, is sequentially optimal. First, observe that $\hat{\nu}$ does not enter into the function V_i , so the fact that (σ, ν) is sequentially optimal directly implies that $V_i(\sigma) \geq V_i(\sigma'_i, \sigma_{-i})$ for all $\sigma'_i \in \Sigma_i$. Second, by construction, $\hat{\nu}$ satisfies the smooth rule using σ except, possibly, for i, I_i not satisfying the criteria $I_i \subseteq \Theta$ or $[f_i(I_i) \neq I_i$ and $\nu_{i, f_i(I_i)}(\pi) > 0$ for some π such that $\pi(f_i^{\sigma^{-i}}(I_i)) > 0]$. However, from the definition of the smooth rule (Definition 2.8), observe that it is exactly in such cases where the smooth rule allows any interim beliefs. Thus $\hat{\nu}$ satisfies the smooth rule using σ . Finally, to see that $(\sigma, \hat{\nu})$ satisfies $V_{i, I_i}(\sigma) \geq V_{i, I_i}(\sigma'_i, \sigma_{-i})$ for all $\sigma'_i \in \Sigma_i$, observe that (a) where $\hat{\nu}_{i, I_i} = \nu_{i, I_i}$, it directly inherits this from (σ, ν) and (b) everywhere else, Lemma A.1 shows that smooth rule updating ensures the required optimality. ■

Theorem A.1 Fix a game Γ . Suppose σ is an ex-ante equilibrium and, for each player i , each of i 's information sets is reachable from some information set in \mathcal{I}_i^0 given σ_{-i} (i.e., $m_i(I_i) = 0$ for all I_i). Then, there exists an interim belief system ν such that (σ, ν) is sequentially optimal.

Proof of Theorem A.1. By ex-ante optimality of σ , (2.6) in the definition of sequential optimality is satisfied. Choose a ν satisfying the smooth rule using σ . Since $m_i(I_i) = 0$ for all i and I_i and $\sum_{\pi \in \Delta(\Theta)} \mu_i(\pi) \pi(\theta) > 0$, either $I_i \subseteq \Theta$ or $f_i(I_i) \neq I_i$ and $\nu_{i, f_i(I_i)}(\pi) > 0$ for some π such that $\pi(f_i^{\sigma^{-i}}(I_i)) > 0$, implying that ν is pinned down completely by the smooth rule. By Lemma A.1, (2.7) in the definition of sequential optimality is satisfied for all i and I_i . ■

Proof of Theorem 2.2. Suppose that (σ, ν) satisfies the no profitable one-stage deviation property and ν satisfies smooth rule updating using σ . First, for each player i , the no profitable one-stage deviation property implies optimality of σ_i according to V_{i, I_i} for all $I_i \in \mathcal{I}_i^T$. Next we 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 (A.5))

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

Since ν satisfies smooth rule updating using σ , for all such J_i that are reachable from I_i given σ_{-i} for which the smooth rule has bite, $q^{(\sigma,\nu),i,J_i}(h) \propto q^{(\sigma,\nu),i,I_i}(h)$ holds for $\{h \mid h^t \in J_i\}$ (see the argument for this near the end of the proof of Lemma A.1). After substituting in (A.9) for $q^{(\sigma,\nu),i,J_i}$, cancelling the constant of proportionality and multiplying by $p_{i,\sigma'_i}(h^t|h^{t-1})$, which is constant for any h such that $h^t \in J_i$ because i 's strategy is a function only of i 's information sets and, due to perfect recall, $R_i(h^t) = R_i(J_i)$ for any h such that $h^t \in J_i$, (A.9) becomes

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

Notice that if either $J_i \in \mathcal{I}_i^t$ is not reachable from I_i given σ_{-i} , or $J_i \in \mathcal{I}_i^t$ is reachable from I_i given σ_{-i} but the smooth rule lacks bite, it holds that, for all h with $h^t \in J_i$, $q^{(\sigma,\nu),i,I_i}(h) = 0$. In the former case this follows since $p_{-i,\sigma_{-i}}(h|h^{t-1}) = 0$, while in the latter case it follows from $\pi_{I_i,\sigma_{-i}}(h^{s(I_i)}) = 0$. Thus, summing (A.10) for all $J_i \in \mathcal{I}_i^t$ reachable from I_i given σ_{-i} is the same as summing for all $J_i \in \mathcal{I}_i^t$ such that $J_i^{-1} = I_i$, yielding:

$$\begin{aligned} & \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h) \\ & \geq \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h). \end{aligned} \tag{A.11}$$

The no profitable one-stage deviation property implies σ_i is optimal according to V_{i,I_i} among all strategies deviating only at I_i . By the optimality representation invoked in (A.9) applied to I_i and restricted to such deviations,

$$\begin{aligned} & \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma_i}(h|h^{t-1}) q^{(\sigma,\nu),i,I_i}(h) \\ & \geq \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i,\sigma'_i}(h^t|h^{t-1}) p_{i,\sigma_i}(h|h^t) q^{(\sigma,\nu),i,I_i}(h). \end{aligned} \tag{A.12}$$

Combining (A.12) and (A.11) implies

$$\begin{aligned} & \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i, \sigma_i}(h|h^{t-1}) q^{(\sigma, \nu), i, I_i}(h) \\ \geq & \sum_{h|h^{t-1} \in I_i} u_i(h) p_{i, \sigma'_i}(h|h^{t-1}) q^{(\sigma, \nu), i, I_i}(h). \end{aligned} \quad (\text{A.13})$$

Since (A.13) holds for any σ'_i , it is the same as (A.9) 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 interim optimality conditions (2.7) in the definition of sequentially optimal.

It remains to show that σ also satisfies the ex-ante optimality conditions (2.6). Since ν satisfies smooth rule updating using σ , for all $I_i \subseteq \Theta_i$, $q^{(\sigma, \nu), i, I_i}(h) \propto q^{\sigma, i}(h)$ holds for $\{h \mid h^0 \in I_i\}$. Using this to substitute for $q^{(\sigma, \nu), i, I_i}$ in (A.13) with $t = 1$, cancelling the constant of proportionality and summing for all I_i , yields:

$$\begin{aligned} & \sum_h u_i(h) p_{i, \sigma_i}(h|h^0) q^{\sigma, i}(h) \\ \geq & \sum_h u_i(h) p_{i, \sigma'_i}(h|h^0) q^{\sigma, i}(h). \end{aligned} \quad (\text{A.14})$$

Since (A.14) holds for any σ'_i , σ satisfies (A.6) which is equivalent to the ex-ante optimality condition (2.6). ■

Proof of Theorem 2.3. 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 $\varepsilon_I^k \leq \frac{1}{|A_i(I_i)|}$ for all players i and information sets I_i . For any k , let Γ^k be the restriction of the game Γ defined such that the set of feasible strategy profiles is the set of all completely mixed σ^k satisfying $\sigma_i^k(I_i)(a_i) \geq \varepsilon_{I_i}^k$ for all i , I_i and actions $a_i \in A_i(I_i)$. Consider the agent normal form G^k of the game Γ^k (see e.g., Myerson, 1991, p.61). Since the payoff functions are concave and the set of strategies of each player in G^k is non-empty, compact and convex, G^k has an ex-ante equilibrium by Glicksberg (1952). Let $\hat{\sigma}^k$ be the strategy profile in the game Γ^k corresponding to this equilibrium. Then $\hat{\sigma}^k$ is an ex-ante equilibrium of Γ^k . By Theorem A.1, since all information sets are on the equilibrium path, there exists an interim belief system ν^k such that $(\hat{\sigma}^k, \nu^k)$ is sequentially optimal. By Theorem 2.1, 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

adaptations from $\hat{\nu}^k$ using $\hat{\sigma}^k$ converges to a limit, which itself is adapted using $\hat{\sigma}$ from some interim belief system. Denote this interim belief system by ν . 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 the adaptation from $\hat{\nu}_{i,I_i}^k$ using $\hat{\sigma}_i^{k,I_i}$ converges to the adaptation from ν_{i,I_i} using $\hat{\sigma}_i^{I_i}$, 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 ν_{i,I_i} . Therefore $(\hat{\sigma}, \nu)$ satisfies sequential optimality. Finally, observe that $(\hat{\sigma}, \nu)$ satisfies smooth rule consistency (since it is explicitly constructed as the limit of an appropriate sequence). Therefore $(\hat{\sigma}, \nu)$ is an SEA of Γ . ■

Proof of Corollary 2.1. It is enough to show that the no profitable one-stage deviation property and smooth rule consistency imply (σ, ν) is sequentially optimal. This follows directly from Definition 2.11 and Theorem 2.2. ■

Theorem 2.4 makes use of a result we state next, providing an explicit formula for interim beliefs satisfying smooth rule consistency. This smooth rule 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 .¹⁷

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

$\text{supp } \mu_i$, define $\bar{\pi}^\pi \in \Delta(H^{m_i(I_i)})$ as follows: if $f_i(I_i) \subseteq \Theta$, $\bar{\pi}^\pi(h^0) \equiv \pi(h^0)$, otherwise, $\bar{\pi}^\pi(h^{m_i(I_i)}) \equiv \bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0) \pi(h^0)$ for all $h^{m_i(I_i)} \in f_i(I_i)$ and is constant for all other $h^{m_i(I_i)}$.

Theorem A.2 Fix a game Γ , a strategy profile σ and a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$ such that $\lim_{k \rightarrow \infty} \sigma^k = \sigma$. Assume that $\bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0)$ exists for each player i , each information set $I_i \not\subseteq \Theta$ and each $h^{m_i(I_i)} \in f_i(I_i)$. Then $(\sigma, \bar{\nu})$ satisfies smooth

¹⁷More precisely, the only parts of this likelihood that need to be derived via a limit concern the partial histories up to $f_i(I_i)$, as the continuations from $f_i(I_i)$ to I_i can be (and are) calculated directly from σ_{-i} .

rule consistency, where $\bar{\nu}$ is defined, for each i, I_i , via the formula

$$\bar{\nu}_{i,I_i}(\hat{\pi}) \propto \sum_{\pi \in \text{supp } \mu_i | \bar{\pi}^\pi = \hat{\pi}} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0) \right)}{\phi'_i \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_\sigma(h|h^{s(I_i)}) \hat{\pi}_{I_i, \sigma_{-i}}(h^{s(I_i)}) \right)} \cdot \left(\sum_{h^{s(I_i)} \in I_i, \sigma_{-i}} p_{-i, \sigma_{-i}}(h^{s(I_i)} | h^{m_i(I_i)}) \hat{\pi}(h^{m_i(I_i)}) \right) \mu_i(\pi), \quad (\text{A.15})$$

when $\hat{\pi}_{I_i, \sigma_{-i}}$ exists and is zero otherwise.

As is useful in the proof of this result, rather than the “one-step ahead” formulation of smooth rule updating in (2.9) for an information set from an immediately prior one, one could alternatively (and equivalently) write the smooth rule as updating $\nu_{i, f_i(I_i)}$ to ν_{i, I_i} all at once:

For each information set $I_i \not\subseteq \Theta$, if $\nu_{i, f_i(I_i)}(\pi) > 0$ for some π such that $\pi(f_i^{\sigma_{-i}}(I_i)) > 0$, then, for all π such that $\pi(f_i^{\sigma_{-i}}(I_i)) > 0$,

$$\nu_{i, I_i}(\pi) \propto \frac{\phi'_i \left(\sum_{h|h^{m_i(I_i)} \in f_i(I_i)} u_i(h) p_\sigma(h|h^{m_i(I_i)}) \pi_{f_i(I_i), i, \sigma_{-i}}(h^{m_i(I_i)}) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi_{I_i, \sigma_{-i}}(h^t) \right)} \cdot \left(\sum_{h^t \in I_i, \sigma_{-i}} p_{-i, \sigma_{-i}}(h^t | h^{m_i(I_i)}) \pi_{f_i(I_i), i, \sigma_{-i}}(h^{m_i(I_i)}) \right) \nu_{i, f_i(I_i)}(\pi), \quad (\text{A.16})$$

where

$$\pi_{I_i, \sigma_{-i}}(h^t) = \frac{p_{-i, \sigma_{-i}}(h^t | h^{m_i(I_i)}) \pi_{f_i(I_i), i, \sigma_{-i}}(h^{m_i(I_i)})}{\sum_{\hat{h}^t \in I_i, \sigma_{-i}} p_{-i, \sigma_{-i}}(\hat{h}^t | \hat{h}^{m_i(I_i)}) \pi_{f_i(I_i), i, \sigma_{-i}}(\hat{h}^{m_i(I_i)})} \text{ if } h^t \in I_i, \sigma_{-i} \text{ and 0 otherwise} \quad (\text{A.17})$$

and $t = s(I_i)$.

Proof of Theorem A.2. We establish that $(\sigma, \bar{\nu})$ satisfies smooth rule consistency. Fix i, I_i . In what follows, when we include π in an expression, the expression is understood to apply for all $\pi \in \Delta(\Theta)$ such that $\pi(\{h^0 | h^{s(I_i)} \in I_i\}) > 0$. Since σ^k is completely mixed, $\{h^0 | h^{s(I_i)} \in I_i\} = f_i^{\sigma^k}(I_i)$ and the smooth rule has bite at each information set and so, applying the “all-at-once” formulation of the smooth rule in (A.16) and the initial smooth

rule step (2.8) from μ_i , for each k , the interim belief $\bar{\nu}_{i,I_i}^k$ determined by smooth rule updating using σ^k satisfies

$$\bar{\nu}_{i,I_i}^k(\pi) \propto \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma^k}(h|h^0) \pi(h^0) \right)}{\phi'_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_{\sigma^k}(h|h^s(I_i)) \pi_{I_i, \sigma^k_{-i}}(h^s(I_i)) \right)} \left(\sum_{h^s(I_i) \in I_i, \sigma^k_{-i}} p_{-i, \sigma^k_{-i}}(h^s(I_i)|h^0) \pi(h^0) \right) \mu_i(\pi). \quad (\text{A.18})$$

Using (A.17) for I_i, σ^k_{-i} and $f_i(I_i)_{i, \sigma^k_{-i}}$, where $f_i(I_i)$ and $m_i(I_i)$ are determined from σ , and observing that $I_i, \sigma^k_{-i} = I_i$ since σ^k is completely mixed,

$$\pi_{I_i, \sigma^k_{-i}}(h^s(I_i)) = \frac{p_{-i, \sigma^k_{-i}}(h^s(I_i)|h^{m_i(I_i)}) \pi_{f_i(I_i)_{i, \sigma^k_{-i}}}(h^{m_i(I_i)})}{\sum_{\hat{h}^s(I_i) \in I_i} p_{-i, \sigma^k_{-i}}(\hat{h}^s(I_i)|\hat{h}^{m_i(I_i)}) \pi_{f_i(I_i)_{i, \sigma^k_{-i}}}(\hat{h}^{m_i(I_i)})} \text{ for all } h^s(I_i) \in I_i$$

and

$$\pi_{f_i(I_i)_{i, \sigma^k_{-i}}}(h^{m_i(I_i)}) = \frac{p_{-i, \sigma^k_{-i}}(h^{m_i(I_i)}|h^0) \pi(h^0)}{\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)_{i, \sigma^k_{-i}}} p_{-i, \sigma^k_{-i}}(\hat{h}^{m_i(I_i)}|\hat{h}^0) \pi(\hat{h}^0)} \text{ for all } h^{m_i(I_i)} \in f_i(I_i)_{i, \sigma^k_{-i}}.$$

Substituting the latter into the former yields

$$\pi_{I_i, \sigma^k_{-i}}(h^s(I_i)) = \frac{p_{-i, \sigma^k_{-i}}(h^s(I_i)|h^0) \pi(h^0)}{\sum_{\hat{h}^s(I_i) \in I_i} p_{-i, \sigma^k_{-i}}(\hat{h}^s(I_i)|\hat{h}^0) \pi(\hat{h}^0)} \text{ for all } h^s(I_i) \in I_i.$$

The adaptation (Definition 2.10) of $\bar{\nu}_{i,I_i}^k$ using σ^k is: for $\tilde{\pi} \in \Delta(I_i)$,

$$\begin{aligned} \bar{\bar{\nu}}_{i,I_i}^k(\tilde{\pi}) \propto & \sum_{\pi \in \text{supp } \bar{\nu}_{i,I_i}^k | \pi_{I_i, \sigma^k_{-i}} = \tilde{\pi}} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma^k}(h|h^0) \pi(h^0) \right)}{\phi'_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_{\sigma^k}(h|h^s(I_i)) \pi_{I_i, \sigma^k_{-i}}(h^s(I_i)) \right)} \quad (\text{A.19}) \\ & \cdot \left(\sum_{h^s(I_i) \in I_i, \sigma^k_{-i}} p_{-i, \sigma^k_{-i}}(h^s(I_i)|h^0) \pi(h^0) \right) \mu_i(\pi). \end{aligned}$$

We next show that the limit of $\bar{\bar{\nu}}_{i,I_i}^k(\tilde{\pi})$ equals the adaptation from $\bar{\nu}_{i,I_i}$ (as defined in (A.15))

using σ . Divide (A.19) by $\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} | \hat{h}^0) > 0$, which is constant with respect to π . Then the limit of $\bar{\nu}_{i, I_i}^{\sigma_{-i}^k}(\tilde{\pi})$ is proportional to:

$$\begin{aligned} & \lim_{k \rightarrow \infty} \sum_{\pi \in \text{supp } \bar{\nu}_{i, I_i}^{\sigma_{-i}^k} | \pi_{I_i, \sigma_{-i}^k} = \tilde{\pi}} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma^k}(h | h^0) \pi(h^0) \right)}{\phi'_i \left(\sum_{h | h^s(I_i) \in I_i} u_i(h) p_{\sigma^k}(h | h^s(I_i)) \pi_{I_i, \sigma_{-i}^k}(h^s(I_i)) \right)} \\ & \cdot \left(\sum_{h^s(I_i) \in I_i, \sigma_{-i}^k} p_{-i, \sigma_{-i}^k}(h^s(I_i) | h^{m_i(I_i)}) \frac{p_{-i, \sigma_{-i}^k}(h^{m_i(I_i)} | h^0)}{\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} | \hat{h}^0)} \pi(h^0) \right) \mu_i(\pi). \end{aligned} \quad (\text{A.20})$$

Recall that $\bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0) \equiv \lim_{k \rightarrow \infty} \frac{p_{-i, \sigma_{-i}^k}(h^{m_i(I_i)} | h^0)}{\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} | \hat{h}^0)}$ exists for all $h^{m_i(I_i)} \in f_i(I_i)$ (either by assumption, or, if $I_i \subseteq \Theta$, because it is constant in k). Then (A.20) is proportional (in $\tilde{\pi} \in \Delta(I_i)$) to

$$\begin{aligned} & \sum_{\pi \in \text{supp } \mu_i | \bar{\pi}_{I_i, \sigma_{-i}}^{\pi} = \tilde{\pi}} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h | h^0) \pi(h^0) \right)}{\phi'_i \left(\sum_{h | h^s(I_i) \in I_i} u_i(h) p_{\sigma}(h | h^s(I_i)) \bar{\pi}_{I_i, \sigma_{-i}}^{\pi}(h^s(I_i)) \right)} \\ & \cdot \left(\sum_{h^s(I_i) \in I_i, \sigma_{-i}} p_{-i, \sigma_{-i}}(h^s(I_i) | h^{m_i(I_i)}) \bar{\pi}^{\pi}(h^{m_i(I_i)}) \right) \mu_i(\pi), \end{aligned} \quad (\text{A.21})$$

$$\begin{aligned} & \text{since, whenever it exists, } \lim_{k \rightarrow \infty} \pi_{I_i, \sigma_{-i}^k}(h^s(I_i)) = \lim_{k \rightarrow \infty} \frac{\frac{p_{-i, \sigma_{-i}^k}(h^s(I_i) | h^{m_i(I_i)}) p_{-i, \sigma_{-i}^k}(h^{m_i(I_i)} | h^0) \pi(h^0)}{\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} | \hat{h}^0)}}{\sum_{\hat{h}^s(I_i) \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^s(I_i) | \hat{h}^{m_i(I_i)}) p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} | \hat{h}^0) \pi(\hat{h}^0)}} = \\ & \frac{\sum_{\hat{h}^s(I_i) \in I_i} p_{-i, \sigma_{-i}}(h^s(I_i) | h^{m_i(I_i)}) \bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0) \pi(h^0)}{\sum_{\hat{h}^s(I_i) \in I_i} p_{-i, \sigma_{-i}}(\hat{h}^s(I_i) | \hat{h}^{m_i(I_i)}) \bar{p}_{-i, \sigma_{-i}}(\hat{h}^{m_i(I_i)} | \hat{h}^0) \pi(\hat{h}^0)}} = \frac{p_{-i, \sigma_{-i}}(h^s(I_i) | h^{m_i(I_i)}) \bar{\pi}^{\pi}(h^{m_i(I_i)})}{\sum_{\hat{h}^s(I_i) \in I_i} p_{-i, \sigma_{-i}}(\hat{h}^s(I_i) | \hat{h}^{m_i(I_i)}) \bar{\pi}^{\pi}(\hat{h}^{m_i(I_i)})} = \bar{\pi}_{I_i, \sigma_{-i}}^{\pi}(h^s(I_i)). \end{aligned}$$

By definition, the adaptation of an interim belief ν_{i,I_i} using σ is, for each $\tilde{\pi} \in \Delta(I_i)$,

$$\sum_{\hat{\pi} \in \text{supp } \nu_{i,I_i} | \hat{\pi}_{I_i, \sigma_{-i}} = \tilde{\pi}} \nu_{i,I_i}(\hat{\pi}),$$

for $\hat{\pi} \in \Delta(H^{m_i(I_i)})$. Substituting for $\nu_{i,I_i}(\hat{\pi})$ using $\bar{\nu}_{i,I_i}(\hat{\pi})$ from (A.15) yields

$$\begin{aligned} \sum_{\hat{\pi} \in \text{supp } \bar{\nu}_{i,I_i} | \hat{\pi}_{I_i, \sigma_{-i}} = \tilde{\pi}} \sum_{\pi \in \text{supp } \mu_i | \bar{\pi}^\pi = \hat{\pi}} \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0) \right)}{\phi'_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_\sigma(h|h^s(I_i)) \bar{\pi}_{I_i, \sigma_{-i}}^\pi(h^s(I_i)) \right)} \\ \cdot \left(\sum_{h^s(I_i) \in I_i, \sigma_{-i}} p_{-i, \sigma_{-i}}(h^s(I_i) | h^{m_i(I_i)}) \bar{\pi}^\pi(h^{m_i(I_i)}) \right) \mu_i(\pi). \end{aligned}$$

Let $B(\pi)$ denote the summand. Observe that

$$\begin{aligned} \sum_{\hat{\pi} \in \text{supp } \bar{\nu}_{i,I_i} | \hat{\pi}_{I_i, \sigma_{-i}} = \tilde{\pi}} \sum_{\pi \in \text{supp } \mu_i | \bar{\pi}^\pi = \hat{\pi}} B(\pi) &= \sum_{\hat{\pi} \in \Delta(H^{m_i(I_i)}) | \hat{\pi}_{I_i, \sigma_{-i}} = \tilde{\pi}} \sum_{\pi \in \text{supp } \mu_i | \bar{\pi}^\pi = \hat{\pi}} B(\pi) \\ &= \sum_{\pi \in \text{supp } \mu_i | \bar{\pi}_{I_i, \sigma_{-i}}^\pi = \tilde{\pi}} B(\pi). \end{aligned}$$

The first equality follows because whenever $\bar{\pi}_{I_i, \sigma_{-i}}^\pi$ exists, $B(\pi) = 0$ for all π such that $\bar{\pi}^\pi = \hat{\pi}$ when $\hat{\pi} \notin \text{supp } \bar{\nu}_{i,I_i}$. The second equality holds because $B(\pi)$ does not depend on $\hat{\pi}$, each π appears in at most one term in the double summation on the left-hand side since π uniquely determines $\bar{\pi}^\pi$, and the π that appear are exactly those in $\text{supp } \mu_i$ for which $\bar{\pi}_{I_i, \sigma_{-i}}^\pi = \tilde{\pi}$. Thus the limit of $\bar{\nu}_{i,I_i}^k(\tilde{\pi})$ equals the adaptation from $\bar{\nu}_{i,I_i}$ using σ .

It remains to show that the interim belief system $\bar{\nu}$ satisfies smooth rule updating using σ . This follows by inspection of (A.15) and the formulation of the smooth rule in (A.16) and (2.8). In particular, applying the ‘‘all-at-once’’ formulation of the smooth rule in (A.16) to $\bar{\nu}_{i,f_i(I_i)}$ and simplifying delivers $\bar{\nu}_{i,I_i}$. Key in verifying this is noticing that in (A.15), the only terms that depend directly on π rather than $\hat{\pi}$ are the ϕ'_i term in the numerator and $\mu_i(\pi)$. Therefore $(\sigma, \bar{\nu})$ satisfies smooth rule consistency in Γ . ■

The following lemma is used in proving Theorem 2.4: If σ satisfies smooth rule consistency for *some* interim belief system then the assumption in Theorem A.2 necessarily holds. Formally:

Lemma A.2 *Fix a game Γ and a strategy profile σ . If, for some interim belief system*

ν , (σ, ν) satisfies smooth rule consistency (with corresponding sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$), then $\bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0)$ exists for each player i , each information set $I_i \notin \Theta$ and each $h^{m_i(I_i)} \in f_i(I_i)$.

Proof of Lemma A.2. Since (σ, ν) satisfies smooth rule consistency with sequence $\{\sigma^k\}_{k=1}^\infty$, $\lim_{k \rightarrow \infty} \sigma^k = \sigma$ and $\hat{\nu} = \lim_{k \rightarrow \infty} \hat{\nu}^k$, where $\hat{\nu}$ is adapted from ν using σ , each $\hat{\nu}^k$ is adapted from ν^k using σ^k , and each ν^k is the interim belief system determined by smooth rule updating using σ^k (see (A.18) for the formula). Fix such sequences and any i, I_i . The existence of $\lim_{k \rightarrow \infty} \hat{\nu}^k$ implies $\lim_{k \rightarrow \infty} \frac{p_{-i, \sigma_{-i}^k}(h^{m_i(I_i)} | h^0)}{\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} | \hat{h}^0)}$ exists for each

$h^{m_i(I_i)} \in f_i(I_i)$. ■

Proof of Theorem 2.4. By Lemma A.2 and Theorem A.2, $(\sigma, \bar{\nu})$ satisfies smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$. Since (σ, ν) and $(\sigma, \bar{\nu})$ both satisfy smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$, the adaptations from ν and $\bar{\nu}$ using σ must be equal (denote this adaptation by $\hat{\nu}$). Since, by Definition 2.10, for any I_i and $\hat{\pi} \in \Delta(I_i)$, $\hat{\nu}_{i, I_i}(\hat{\pi}) = \sum_{\pi \in \text{supp } \nu_{i, I_i} | \pi_{I_i, \sigma_{-i}} = \hat{\pi}} \nu_{i, I_i}(\pi)$,

ν and $\bar{\nu}$ must have the same interim valuations:

$$\begin{aligned} & \sum_{\pi \in \Delta(H^{m_i(I_i)}) | \pi(f_i^{\sigma_{-i}}(I_i)) > 0} \phi_i \left(\sum_{h | h^t \in I_i} u_i(h) p_{(\hat{\sigma}_i, \sigma_{-i})}(h | h^t) \pi_{I_i, \sigma_{-i}}(h^t) \right) \nu_{i, I_i}(\pi) \\ = & \sum_{\hat{\pi} \in \Delta(I_i)} \phi_i \left(\sum_{h | h^t \in I_i} u_i(h) p_{(\hat{\sigma}_i, \sigma_{-i})}(h | h^t) \hat{\pi}(h^t) \right) \hat{\nu}_{i, I_i}(\hat{\pi}) \\ = & \sum_{\pi \in \Delta(H^{m_i(I_i)}) | \pi(f_i^{\sigma_{-i}}(I_i)) > 0} \phi_i \left(\sum_{h | h^t \in I_i} u_i(h) p_{(\hat{\sigma}_i, \sigma_{-i})}(h | h^t) \pi_{I_i, \sigma_{-i}}(h^t) \right) \bar{\nu}_{i, I_i}(\pi). \end{aligned}$$

Therefore sequential optimality holds for ν if and only if it holds for $\bar{\nu}$, implying that $(\sigma, \bar{\nu})$ is an SEA with the corresponding sequence $\{\sigma^k\}_{k=1}^\infty$. ■

Theorem A.3 Fix a game Γ . Suppose (σ, ν) is sequentially optimal and for each player i , each of i 's information sets is reachable from some information set in \mathcal{I}_i^0 given σ_{-i} (i.e., $m_i(I_i) = 0$ for all I_i). Then, there exists an interim belief system $\hat{\nu}$ such that $(\sigma, \hat{\nu})$ is an SEA.

Proof of Theorem A.3. Consider a sequence of a completely mixed strategy profiles converging to σ . Since the limits $\bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0)$ in Theorem A.2 exist at on-path information sets, Theorem A.2 implies that $(\sigma, \bar{\nu})$ satisfies smooth rule consistency, and is thus an SEA. ■

Parts of the next proof (of Theorem 2.5) make use of the following particularly convenient set, parametrized by $l \geq 1$ and $b > 1$, of $\hat{\phi}_i$ strictly more concave than ϕ_i :

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

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

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

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

Proof of Theorem 2.5. This proof makes use of the E_Γ , Q_Γ and S_Γ notations for sets of equilibria given in Notation B.1 (see p. 70). Fixing an equilibrium $\sigma \in E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ (resp. Q_Γ with associated interim belief system ν or S_Γ with associated ν and sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$) and a player i , say that ambiguity aversion makes σ_i *ex-ante* (resp. *sequentially optimal or SEA*) *more belief robust* if there exists a continuum of $\hat{\phi}_i$ strictly more concave than ϕ_i such that, for each $\hat{\mu}_i$ with the same support, Π_i , as μ_i , σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i$ and $\hat{\mu}_i$ (resp. and an interim best response to σ_{-i} at any information set I_i given $\hat{\phi}_i$ and ν_{i,I_i} or (for SEA) that plus satisfying the part for player i of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$) whenever it is an ex-ante best response to σ_{-i} given ϕ_i and $\hat{\mu}_i$ (resp. and an interim best response to σ_{-i} at any information set I_i given ϕ_i and $\hat{\nu}_{i,I_i}$ or (for SEA) that plus satisfying the part for player i of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$). Similarly, say that ambiguity aversion makes σ_i *ex-ante* (resp. *sequentially optimal or SEA*) *strictly more belief robust* if, in addition, for each such $\hat{\phi}_i$, there exists $\tilde{\mu}_i$ with support Π_i such that σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i$ and $\tilde{\mu}_i$ (resp. and an interim best response to σ_{-i} at any information set I_i given $\hat{\phi}_i$ and ν_{i,I_i} or (for SEA) that plus satisfying the part for player i of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$) but is not an ex-ante best response to σ_{-i} given ϕ_i and $\tilde{\mu}_i$ (resp., for all interim beliefs, fails either the ex-ante or interim best response to σ_{-i} property given ϕ_i and $\tilde{\mu}_i$, or (for SEA) fails the same property for all interim beliefs satisfying smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$). To prove that ambiguity aversion makes σ ex-ante (resp. sequentially optimal or SEA) (strictly) more belief robust, it is sufficient to show, for each player i , that ambiguity aversion makes σ_i ex-ante (resp. sequentially optimal or SEA) more belief robust (and that there is some

player i for whom ambiguity aversion makes σ_i ex-ante (resp. sequentially optimal or SEA) strictly more belief robust). The argument is the same for each player, so for the remainder of the proof fix a player i . Also assume for the remainder of the argument that $|\Pi_i| > 1$, as otherwise the result follows immediately because there is only one possible belief with that support.

We begin by proving that ambiguity aversion makes σ_i ex-ante more belief robust under the assumption that the antecedent of the second part of the theorem holds for this player i . Using the local linear approximation, (A.6), of i 's objective function around σ_i it follows that σ_i is an ex-ante best response to σ_{-i} for player i given any $\hat{\phi}_i$ and $\hat{\mu}_i$ if and only if (2.12). Observe that any strategies σ'_i that are weakly worse than σ_i (in terms of ex-ante expected utility, $\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi(h^0)$) for all $\pi \in \Pi_i$ can never interfere with optimality of σ_i and will thus, without loss of generality, be ignored whenever making statements about strategies other than σ_i in what follows. For each l , denote by π_i^l the unique $\pi \in \Pi_i$ under which σ_i gives e_i^l , the l^{th} lowest distinct ex-ante expected utility generated by Π_i . By the existence of $\tilde{\mu}_i$ as posited in the antecedent of the second part of the theorem, there exists a strategy $\hat{\sigma}_i$ with strictly higher ex-ante expected utility than σ_i under π_i^l for some $1 \leq l \leq |\Pi_i|$.

We next show that all strategies must give weakly lower ex-ante expected utility than σ_i under π_i^1 . To see this, suppose, to the contrary, there exists a strategy $\hat{\sigma}_i$ with strictly higher ex-ante expected utility than σ_i under π_i^1 . Since σ is ex-ante robust to increased ambiguity aversion, σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^1$ and μ_i (and in particular is at least as good as $\hat{\sigma}_i$). Using (2.12) and applying the definition of $\hat{\phi}_i^1$ (Definition A.1), this implies

$$\begin{aligned} & \sum_{\pi \in \Pi_i} \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi(h^0) - \sum_{h \in H} u_i(h) p_{(\hat{\sigma}_i, \sigma_{-i})}(h|h^0) \pi(h^0) \right) \phi_i' \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi(h^0) \right) \mu_i(\pi) \\ & + (b-1) \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^1(h^0) - \sum_{h \in H} u_i(h) p_{(\hat{\sigma}_i, \sigma_{-i})}(h|h^0) \pi_i^1(h^0) \right) \\ & \geq 0. \end{aligned}$$

Since the value of the first line is bounded and the term multiplying $(b-1)$ on the second line is strictly negative, taking b large enough generates a contradiction.

Let \hat{M}_i denote the set of probabilities $\hat{\mu}_i$ with support contained in Π_i for which (2.12) holds with ϕ_i replacing $\hat{\phi}_i$. Let $m < |\Pi_i|$ be the smallest number l for which there exists a strategy that gives a strictly higher ex-ante expected utility than σ_i under π_i^{l+1} . By the previous paragraph, $m \geq 1$. Using (2.12) with ϕ_i replacing $\hat{\phi}_i$ and $\hat{\mu}_i \in \hat{M}_i$ together with

the fact that all strategies give weakly lower ex-ante expected utility than σ_i under π_i^l for all $1 \leq l \leq m$, conclude that, for any $b > 1$,

$$\begin{aligned} \sigma_i \in \arg \max_{\sigma'_i} & \sum_{\pi \in \Pi_i} \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi(h^0) \right) \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi(h^0) \right) \hat{\mu}_i(\pi) \\ & + (b-1) \sum_{l=1}^m \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi_i^l(h^0) \right) \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^l(h^0) \right) \hat{\mu}_i(\pi_i^l). \end{aligned} \quad (\text{A.22})$$

Applying the definition of $\hat{\phi}_i^m$, we see that (A.22) is equivalent to (2.12) with $\hat{\phi}_i^m$ replacing $\hat{\phi}_i$. This shows that ambiguity aversion makes σ_i ex-ante more belief robust.

To complete the second part of the theorem in the ex-ante case, it remains to show the strict part of the result. By Lemma A.3, there exists a continuum of $b > 1$ and corresponding $\tilde{\mu}_i \notin \hat{M}_i$ with support Π_i such that σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. This optimality together with (2.12), where $\hat{\phi}_i^m$ with this b replaces $\hat{\phi}_i$ and for all $\hat{\mu}_i \in \hat{M}_i$, as was established above, implies that ambiguity aversion makes σ_i ex-ante strictly more belief robust.

Consider now sequential optimality in the second part of the theorem. Consider m as defined above. Since σ is sequentially optimal robust to increased ambiguity aversion, σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i^m$ (for any $b > 1$) and μ_i , and for any information set I_i there exists an interim belief ν_{i,I_i} such that σ_i is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ and ν_{i,I_i} . Consider any $\hat{\mu}_i$ with the same supports as the μ_i for which σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i$ and $\hat{\mu}_i$. By the ex-ante equilibrium argument above, σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i^m$ and $\hat{\mu}_i$. Given $\hat{\phi}_i^m$ and $\hat{\mu}_i$, where the smooth rule has bite, derive $\hat{\nu}_{i,I_i}$ for information sets I_i reachable from some information set in \mathcal{I}_i^0 given σ_{-i} via the smooth rule using σ . By Theorem 2.1, limiting attention to such interim beliefs is without loss of generality. By Lemma A.1, σ_i an ex-ante best response to σ_{-i} given $\hat{\phi}_i^m$ and $\hat{\mu}_i$ implies σ_i is an interim best response to σ_{-i} at these information sets given $\hat{\phi}_i^m$ and $\hat{\nu}_{i,I_i}$. Extend $\hat{\nu}$ by setting $\hat{\nu}_{i,I_i} = \nu_{i,I_i}$ elsewhere. Thus σ_i is an interim best response to σ_{-i} also at these remaining information sets given $\hat{\phi}_i^m$ and $\hat{\nu}_{i,I_i}$, as this fact is not affected by the shift from μ_i to $\hat{\mu}_i$. This shows that ambiguity aversion makes σ_i sequentially optimal more belief robust. It remains to show the strict part of the result. By the argument above for the ex-ante equilibrium case using Lemma A.3 to produce a continuum of $b > 1$ and corresponding $\tilde{\mu}_i$, σ_i is an ex-ante best response to σ_{-i} given, for any such b , $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. By an analogous construction as for $\hat{\nu}_{i,I_i}$, construct interim beliefs $\tilde{\nu}_{i,I_i}$ for each I_i given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$ such that σ_i is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ and $\tilde{\nu}_{i,I_i}$. Thus ambiguity aversion makes σ_i sequentially optimal strictly more belief robust.

Finally turn to SEA under the assumptions of the second part of the theorem. We establish the existence of interim beliefs so that both interim best response and i 's part of smooth rule consistency are satisfied. Given any $\hat{\phi}_i$, for each I_i , construct an interim belief, $\hat{\nu}_{i,I_i}$, to equal $\bar{\nu}_{i,I_i}$ as defined in (A.15) except with $\hat{\mu}_i$ replacing μ_i and $\hat{\phi}_i$ replacing ϕ_i . Then, since (σ, ν) satisfies smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$, Lemma A.2 yields existence of $\bar{p}_{-i,\sigma_{-i}}(h^{m_i(I_i)}|h^0)$. Theorem A.2 applied using $\hat{\mu}_i$ and $\hat{\phi}_i$ (and noting that $\bar{p}_{-i,\sigma_{-i}}(h^{m_i(I_i)}|h^0)$ is independent of the choice of $\hat{\mu}_i$ and $\hat{\phi}_i$) implies that σ together with $(\hat{\nu}_{i,I_i})_{I_i \in \mathcal{I}_i}$ satisfies player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ given $\hat{\mu}_i$ and $\hat{\phi}_i$.

It remains to show that σ_i is an interim best response to σ_{-i} for player i at I_i given $\hat{\phi}_i$ and $\hat{\nu}_{i,I_i}$, which is equivalent to showing that

$$\sigma_i \in \arg \max_{\sigma'_i} \sum_{\pi | \pi(\{h^0 | h^{s(I_i)} \in I_i\}) > 0} \left(\sum_{h | h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h | h^{s(I_i)}) \bar{\pi}_{I_i, \sigma_{-i}}^\pi(h^{s(I_i)}) \right) \cdot \hat{\phi}'_i \left(\sum_{h | h^{s(I_i)} \in I_i} u_i(h) p_\sigma(h | h^{s(I_i)}) \bar{\pi}_{I_i, \sigma_{-i}}^\pi(h^{s(I_i)}) \right) \hat{\nu}_{i,I_i}(\bar{\pi}^\pi)$$

as can be seen by rearranging terms in (A.5) with $\hat{\phi}_i$ replacing ϕ_i and $\hat{\nu}_{i,I_i}$ replacing ν_{i,I_i} and observing that each element in the support of $\hat{\nu}_{i,I_i}$ is $\bar{\pi}^\pi$ for some π in the support of $\hat{\mu}_i$ such that $\pi(\{h^0 | h^{s(I_i)} \in I_i\}) > 0$. Substituting for $\bar{\pi}_{I_i, \sigma_{-i}}^\pi$ using (A.17, p. 46) and for $\hat{\nu}_{i,I_i}$ and $\bar{\pi}^\pi$ using the (A.15) construction, using the notation $\tilde{p}_{-i,\sigma_{-i}}(h^{s(I_i)}|h^0) \equiv p_{-i,\sigma_{-i}}(h^{s(I_i)}|h^{m_i(I_i)}) \bar{p}_{-i,\sigma_{-i}}(h^{m_i(I_i)}|h^0)$ for any h such that $h^{s(I_i)} \in I_i$, noticing that $\pi(\{h^0 | h^{s(I_i)} \in I_i\}) = 0$ implies $\pi(h^0) = 0$ for all h such that $h^{s(I_i)} \in I_i$, so that we may add terms multiplied by $\pi(h^0)$ to the sum in the objective function without changing its value, and, finally, simplifying, yields

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

Since σ is SEA robust to increased ambiguity aversion, σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i^1$ and μ_i , and for any I_i there exists an interim belief ν_{i,I_i} , constructed as was $\hat{\nu}_{i,I_i}$ at the beginning of the SEA part of the proof except now using $\hat{\phi}_i^1$ and μ_i , such that σ_i is an interim best response to σ_{-i} given $\hat{\phi}_i^1$ and ν_{i,I_i} . From the definition of $\hat{\phi}_i^1$ and the assumption that π_i^1 is well-defined, (A.23) with $\hat{\mu}_i = \mu_i$ and $\hat{\phi}_i = \hat{\phi}_i^1$ with b large enough

(i.e., $\hat{\phi}_i^1$ sufficiently concave) implies that, for each I_i , $\sigma'_i = \sigma_i$ must maximize the interim expected utility under π_i^1 ,

$$\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^{s(I_i)}) \tilde{p}_{-i, \sigma_{-i}}(h^{s(I_i)}|h^0) \pi_i^1(h^0), \quad (\text{A.24})$$

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

For each I_i , let \hat{M}_{i, I_i} denote the set of probabilities $\hat{\mu}_i$ with support contained in Π_i for which (A.23) holds with ϕ_i replacing $\hat{\phi}_i$. Let $1 \leq m < |\Pi_i|$ be the smallest number l for which there exists a strategy that, under π_i^{l+1} , gives either a strictly higher ex-ante expected utility than σ_i or a strictly higher interim expected utility than σ_i at some I_i . Consider any $\hat{\mu}_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i, I_i} \right)$ with support Π_i . By the argument above for ex-ante equilibrium, σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i^m$ and $\hat{\mu}_i$. Analogous to the construction of $\hat{\nu}_{i, I_i}$ at the beginning of the SEA part of the proof, for each I_i construct $\hat{\nu}_{i, I_i}^m$ such that σ_i is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ and $\hat{\nu}_{i, I_i}^m$. Substituting $\hat{\phi}_i^m$ for $\hat{\phi}_i$, (A.23) is equivalent to

$$\begin{aligned} \sigma_i \in \arg \max_{\sigma'_i} & \sum_{\pi \in \Pi_i} \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^{s(I_i)}) \tilde{p}_{-i, \sigma_{-i}}(h^{s(I_i)}|h^0) \pi(h^0) \right) \\ & \cdot \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi(h^0) \right) \hat{\mu}_i(\pi) \\ & + (b-1) \sum_{l=1}^m \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^{s(I_i)}) \tilde{p}_{-i, \sigma_{-i}}(h^{s(I_i)}|h^0) \pi_i^l(h^0) \right) \\ & \cdot \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^l(h^0) \right) \hat{\mu}_i(\pi_i^l). \end{aligned} \quad (\text{A.25})$$

That (A.25) holds follows from the fact that $\sigma'_i = \sigma_i$ maximizes each of the two terms – the first by the fact that $\hat{\mu}_i \in \bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i, I_i}$, and the second since, by definition of m , under π_i^l for all $1 \leq l \leq m$, all strategies give weakly lower ex-ante expected utility than σ_i and weakly lower interim expected utility at each I_i than σ_i . Thus σ_i is an interim best response to σ_{-i} for player i at information set I_i given $\hat{\phi}_i^m$ and $\hat{\nu}_{i, I_i}^m$. Furthermore, by construction, σ together with interim beliefs $\hat{\nu}_{i, I_i}^m$ for player i satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^{\infty}$ given $\hat{\phi}_i^m$ and $\hat{\mu}_i$. Therefore ambiguity aversion makes σ_i SEA more belief robust.

It remains to show the strict part of the SEA result. The conclusion follows immediately when no l exists for defining m , as then σ_i is an ex-ante best response to σ_{-i} for player i given

any $\bar{\phi}_i$ and any $\bar{\mu}_i$ and is an interim best response to σ_{-i} given any $\bar{\phi}_i$ and any $\bar{\nu}_{i,I_i}$ for each I_i , and σ together with any interim beliefs $\bar{\nu}_{i,I_i}$ for player i , constructed analogously to the construction of $\hat{\nu}_{i,I_i}$ at the beginning of the SEA part of the proof, satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$. If there exists a l for defining m , then by Lemma A.4 there exists a continuum of $b > 1$ and corresponding $\tilde{\mu}_i \notin \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ with support Π_i and interim beliefs $\tilde{\nu}_{i,I_i}$ for each I_i such that, for each such b , σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ (with this b) and corresponding $\tilde{\mu}_i$ and is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ (with that b) and $\tilde{\nu}_{i,I_i}$ for each I_i , and σ together with the interim beliefs $\tilde{\nu}_{i,I_i}$ for player i satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ given $\hat{\phi}_i^m$ (with this b) and $\tilde{\mu}_i$. These facts together with the SEA more belief robustness results already proven above given $\hat{\phi}_i^m$ (for any such b) and for all $\hat{\mu}_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ with support Π_i , imply that ambiguity aversion makes σ SEA strictly more belief robust.

We conclude the proof by showing that ambiguity aversion makes σ_i ex-ante (resp. sequentially optimal or SEA) more belief robust under the assumption that the antecedent in only the first part of the theorem holds for this player i . If $\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0)$ is constant in π , then the $\hat{\phi}_i^l$ term is constant in (2.12) and increases in ambiguity aversion leave the set of $\hat{\mu}_i$ supporting σ_i unchanged, establishing the desired conclusion in this case. Therefore, consider that $\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0)$ has a unique minimizing π over the support of μ_i . Under this assumption, π_i^1 is well defined but π_i^l for $l > 1$ may not be. That ambiguity aversion makes σ_i ex-ante (resp. sequentially optimal or SEA) more belief robust now follows from the argument above for the ex-ante (resp. sequentially optimal or SEA) case in the second part of the theorem, with the exception that $\hat{\phi}_i^1$ should be used in place of $\hat{\phi}_i^m$ everywhere and excluding entirely the arguments used only for the strictly more belief robust part of the conclusions.

Lemma A.3 *Consider an equilibrium σ that is ex-ante robust to increased ambiguity aversion. Fix a player i such that $\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0)$ can be strictly ordered across the π in the support of μ_i . If there exists $1 \leq m < |\Pi_i|$ such that m is the smallest number l for which there exists a strategy that gives a strictly higher ex-ante expected utility than σ_i under π_i^{l+1} , then there exists a continuum of $b > 1$ and, for each b , a $\tilde{\mu}_i \notin \hat{M}_i$ with support Π_i such that σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ (with that b) and $\tilde{\mu}_i$.*

■

Proof of Lemma. Under the assumptions of the lemma, the definition of m implies that each σ_i^l gives a weakly lower ex-ante expected utility than σ_i under each π_i^l for $l = 1, \dots, m$. For each $\hat{\sigma}_i$ that gives a strictly higher ex-ante expected utility than σ_i under π_i^{m+1} , $\hat{\sigma}_i$ gives a strictly lower ex-ante expected utility than σ_i for some π_i^l for $l = 1, \dots, m$. This is necessary

to avoid $\hat{\sigma}_i$ being used together with $\hat{\phi}_i^{m+1}$ and μ_i to generate a contradiction to σ_i being ex-ante robust to increased ambiguity aversion. Consider the following function of σ'_i for $\lambda \in [0, 1]$,

$$\begin{aligned} \xi(\sigma'_i; \lambda) \equiv & \sum_{l=1}^m \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi_i^l(h^0) \right) \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^l(h^0) \right) \frac{\mu_i(\pi_i^l)}{\sum_{j=1}^m \mu_i(\pi_i^j)} \lambda \\ & + \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi_i^{m+1}(h^0) \right) \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^{m+1}(h^0) \right) (1 - \lambda). \end{aligned}$$

Observe that $\sigma_i \notin \arg \max_{\sigma'_i} \xi(\sigma'_i; 0)$ and $\sigma_i \in \arg \max_{\sigma'_i} \xi(\sigma'_i; 1)$. Moreover, for each $\hat{\sigma}_i$ that gives a strictly higher ex-ante expected utility than σ_i under π_i^{m+1} , $\xi(\sigma_i; 0) < \xi(\hat{\sigma}_i; 0)$ and $\xi(\sigma_i; 1) > \xi(\hat{\sigma}_i; 1)$. Thus, for each such $\hat{\sigma}_i$, there exists some $\underline{\lambda}(\hat{\sigma}_i) \in (0, 1)$ such that, for $\lambda \in [0, 1]$, $\lambda \geq \underline{\lambda}(\hat{\sigma}_i)$ if and only if

$$\sigma_i \in \arg \max_{\sigma'_i \in \{\hat{\sigma}_i, \sigma_i\}} \xi(\sigma'_i; \lambda).$$

Since $\xi(\sigma'_i; \lambda)$ is linear in the mixing weights in σ'_i , to determine optimality of σ_i it is sufficient to compare it with all of player i 's pure strategies, of which there are a finite number. Define $\underline{\lambda} \in (0, 1)$ as the maximal $\underline{\lambda}(\hat{\sigma}_i)$ over all pure strategies $\hat{\sigma}_i$ giving a strictly higher ex-ante expected utility than σ_i under π_i^{m+1} . Thus, for $\lambda \in [0, 1]$, $\lambda \geq \underline{\lambda}$ if and only if

$$\sigma_i \in \arg \max_{\sigma'_i} \xi(\sigma'_i; \lambda).$$

Since $\mu_i \in \hat{M}_i$ and \hat{M}_i is closed and convex, there exists $\alpha, \varepsilon \in (0, 1)$ such that, for each $\lambda \in (0, \varepsilon \underline{\lambda})$

$$\tilde{\mu}_i(\pi_i^\ell) \equiv \begin{cases} \alpha \frac{\mu_i(\pi_i^\ell)}{\sum_{j=1}^m \mu_i(\pi_i^j)} \lambda + (1 - \alpha) \mu_i(\pi_i^\ell) & , \quad \ell = \{1, \dots, m\} \\ \alpha(1 - \lambda) + (1 - \alpha) \mu_i(\pi_i^\ell) & , \quad \ell = m + 1 \\ (1 - \alpha) \mu_i(\pi_i^\ell) & , \quad \ell \in \{m + 2, \dots, |\Pi_i|\} \end{cases}$$

satisfies $\tilde{\mu}_i \notin \hat{M}_i$. By construction, $\tilde{\mu}_i$ has the same support as μ_i . By the properties of $\underline{\lambda}$ and convexity of \hat{M}_i , for any $\gamma \in [0, 1]$,

$$\check{\mu}_i(\pi_i^\ell) \equiv \begin{cases} \gamma \frac{\mu_i(\pi_i^\ell)}{\sum_{j=1}^m \mu_i(\pi_i^j)} \underline{\lambda} + (1 - \gamma) \mu_i(\pi_i^\ell) & , \quad \ell = \{1, \dots, m\} \\ \gamma(1 - \underline{\lambda}) + (1 - \gamma) \mu_i(\pi_i^\ell) & , \quad \ell = m + 1 \\ (1 - \gamma) \mu_i(\pi_i^\ell) & , \quad \ell \in \{m + 2, \dots, |\Pi_i|\} \end{cases}$$

satisfies $\check{\mu}_i \in \hat{M}_i$. Observe that, for each $\lambda \in (0, \varepsilon\lambda)$, if $\gamma = \frac{(1-\lambda)\alpha}{(1-\lambda)(1-\alpha)+(1-\lambda)\alpha}$ then, for $\beta = \frac{\gamma-\alpha}{1-\alpha} \in (0, 1)$,

$$\check{\mu}_i(\pi_i^\ell) = \begin{cases} \beta \frac{\mu_i(\pi_i^\ell)}{\sum_{j=1}^m \mu_i(\pi_i^j)} + (1-\beta)\tilde{\mu}_i(\pi_i^\ell) & , \quad \ell = \{1, \dots, m\} \\ (1-\beta)\tilde{\mu}_i(\pi_i^\ell) & , \quad \ell \in \{m+1, \dots, |\Pi_i|\}. \end{cases}$$

One can verify that for

$$b = 1 + \frac{\beta}{(1-\beta)\sum_{j=1}^m \tilde{\mu}_i(\pi_i^j)} = 1 + \frac{\beta}{(1-\beta)(\alpha\lambda + (1-\alpha)\sum_{j=1}^m \mu_i(\pi_i^j))},$$

the version of the objective function in (2.12) given ϕ_i and $\check{\mu}_i$ is, up to normalization, equal to the version of the objective function in (2.12) given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. Thus σ_i being an ex-ante best response to σ_{-i} for player i given ϕ_i and $\check{\mu}_i$ is equivalent to σ_i being an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ (for all such b) and $\tilde{\mu}_i$. The former holds since $\check{\mu}_i \in \hat{M}_i$, and thus we have established the latter. This proves that all such b and corresponding $\check{\mu}_i$ satisfy the conclusion of the lemma. ■

Lemma A.4 *Consider an equilibrium σ that is SEA robust to increased ambiguity aversion. Fix a player i such that $\sum_{h \in H} u_i(h)p_\sigma(h|h^0)\pi(h^0)$ can be strictly ordered across the π in the support of μ_i . If there exists $1 \leq m < |\Pi_i|$ such that m is the smallest number l for which there exists a strategy that, under π_i^{l+1} , gives either a strictly higher ex-ante expected utility than σ_i or a strictly higher interim expected utility than σ_i at some I_i , then there exists a continuum of $b > 1$ and, for each b , $\check{\mu}_i \notin \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i}\right)$ with support Π_i and interim beliefs $\tilde{\nu}_{i,I_i}$ for each I_i such that σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ (with that b) and $\check{\mu}_i$ and is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ (with that b) and $\tilde{\nu}_{i,I_i}$ for each I_i , and σ together with the interim beliefs $\tilde{\nu}_{i,I_i}$ for player i satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ given $\hat{\phi}_i^m$ (with that b) and $\check{\mu}_i$.*

Proof. Under the assumptions of the lemma, the definition of m implies that, under each π_i^l for $l = 1, \dots, m$, each σ_i^l gives a weakly lower ex-ante expected utility than σ_i and a weakly lower interim expected utility than σ_i at each I_i . Define $\xi(\sigma_i^l; \lambda)$ as in the proof of A.3. For any $\hat{\sigma}_i$ that gives a strictly higher ex-ante expected utility than σ_i under π_i^{m+1} , there exists a $\underline{\lambda}(\hat{\sigma}_i)$ as defined in the proof of Lemma A.3. Consider any I_i for which there exists some $\hat{\sigma}_i$ that gives a strictly higher interim expected utility than σ_i at I_i under π_i^{m+1} . Any such $\hat{\sigma}_i$ gives a strictly lower interim expected utility than σ_i at I_i for some π_i^l for $l = 1, \dots, m$. This is necessary to avoid I_i and $\hat{\sigma}_i$ being used together with $\hat{\phi}_i^{m+1}$ and μ_i to generate a contradiction to σ_i being SEA robust to increased ambiguity aversion. Consider

the function $\xi_{i,I_i}(\sigma'_i; \lambda)$ of σ'_i for $\lambda \in [0, 1]$, defined similarly to $\xi(\sigma'_i; \lambda)$ with interim expected utilities replacing ex-ante expected utilities. Observe that $\sigma_i \notin \arg \max_{\sigma'_i} \xi_{i,I_i}(\sigma'_i; 0)$ and $\sigma_i \in \arg \max_{\sigma'_i} \xi_{i,I_i}(\sigma'_i; 1)$. Moreover, for each $\hat{\sigma}_i$ that gives a strictly higher interim expected utility than σ_i at I_i under π_i^{m+1} , $\xi_{i,I_i}(\sigma_i; 0) < \xi_{i,I_i}(\hat{\sigma}_i; 0)$ and $\xi_{i,I_i}(\sigma_i; 1) > \xi_{i,I_i}(\hat{\sigma}_i; 1)$. Thus, for each such $\hat{\sigma}_i$, there exists some $\underline{\lambda}_{i,I_i}(\hat{\sigma}_i) \in (0, 1)$ such that, for $\lambda \in [0, 1]$, $\lambda \geq \underline{\lambda}_{i,I_i}(\hat{\sigma}_i)$ if and only if

$$\sigma_i \in \arg \max_{\sigma'_i \in \{\hat{\sigma}_i, \sigma_i\}} \xi_{i,I_i}(\sigma'_i; \lambda).$$

Since $\xi(\sigma'_i; \lambda)$ and $\xi_{i,I_i}(\sigma'_i; \lambda)$ are linear in the mixing weights in σ'_i , to determine optimality of σ_i according to $\xi(\sigma'_i; \lambda)$ or $\xi_{i,I_i}(\sigma'_i; \lambda)$ it is sufficient to compare it with all of player i 's pure strategies, of which there are a finite number. Define $\underline{\lambda} \in (0, 1)$ as the maximal $\underline{\lambda}(\hat{\sigma}_i)$ or $\underline{\lambda}_{i,I_i}(\hat{\sigma}_i)$ over all pure strategies $\hat{\sigma}_i$ giving, under π_i^{m+1} , a strictly higher ex-ante expected utility than σ_i or a strictly higher interim expected utility at I_i , respectively. Thus, for $\lambda \in [0, 1]$, $\lambda \geq \underline{\lambda}$ if and only if

$$\sigma_i \in \arg \max_{\sigma'_i} \xi(\sigma'_i; \lambda) \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \arg \max_{\sigma'_i} \xi_{i,I_i}(\sigma'_i; \lambda) \right).$$

Since $\mu_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ and $\hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ is closed and convex, there exists $\alpha, \varepsilon \in (0, 1)$ such that, for each $\lambda \in (0, \varepsilon \underline{\lambda})$ the $\tilde{\mu}_i$ defined in the proof of Lemma A.3 satisfies $\tilde{\mu}_i \notin \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$. By construction, $\tilde{\mu}_i$ has the same support as μ_i . By the properties of $\underline{\lambda}$ and convexity of \hat{M}_i , for any $\gamma \in [0, 1]$, the $\check{\mu}_i$ defined in the proof of Lemma A.3 satisfies $\check{\mu}_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$. One can verify that for all b defined in the proof of Lemma A.3, the version of the objective functions in (2.12) and (A.23) given ϕ_i and $\check{\mu}_i$ are, up to normalization, correspondingly equal to the versions of the objective functions in (2.12) and (A.23) given $\hat{\phi}_i^m$ (with that b) and $\tilde{\mu}_i$. For each b defined in the proof of Lemma A.3, using Lemma A.2 and the formula (A.15), for each I_i construct interim beliefs $\check{\nu}_{i,I_i}$ given $\hat{\phi}_i^m$ (with this b) and corresponding $\tilde{\mu}_i$ and also construct using the same formula interim beliefs $\check{\nu}_{i,I_i}$ given ϕ_i and $\check{\mu}_i$. Thus σ_i being an ex-ante best response to σ_{-i} for player i given ϕ_i and $\check{\mu}_i$ and an interim best response to σ_{-i} given ϕ_i and $\check{\nu}_{i,I_i}$ for each I_i is equivalent to σ_i being an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ (with this b) and corresponding $\tilde{\mu}_i$ and an interim best response to σ_{-i} given $\hat{\phi}_i^m$ (with this b) and $\hat{\nu}_{i,I_i}^m$ for each I_i . The former holds since $\check{\mu}_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$, and thus we have established the latter. By Theorem A.2, σ together with interim beliefs $\check{\nu}_{i,I_i}^m$ for player i satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. This proves that all such b and corresponding $\tilde{\mu}_i$ satisfy the conclusion of the lemma. ■

The next result relates to analysis of the limit pricing example. Denote the entrant's

Cournot profit net of entry costs when facing an incumbent of type τ by $w_\tau \equiv b\left(\frac{a+c_\tau-2c_E}{3b}\right)^2 - K$.

Lemma A.5 *Under Assumption 3.1, σ^{LP} is an ex-ante equilibrium if and only if (ICH for I), (ICM for I), $w_H \geq 0$ and*

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

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

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

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

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

This is equivalent to (ICH for I).

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

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

This is equivalent to (ICM for I).

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

It remains to examine the entry decision of the entrant. As a best-response to the

incumbent's strategy, ex-ante the entrant wants to maximize

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

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

Before turning to the proof of Proposition 3.1, we remark that we actually prove a slightly stronger result, allowing for the possibility that $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M = 0\}) = 1$ (i.e., that the entrant unambiguously believes that it will exactly break even if it enters conditional on the incumbent's type being in $\{L, M\}$). This appears in the proof only in the proof of Lemma A.6.

Proof of Proposition 3.1. Consider the limit pricing strategy profile σ^{LP} .

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

Next, we construct an interim belief system that, together with σ^{LP} , satisfies smooth rule consistency. Consider a sequence of completely mixed strategy profiles, σ^k , where $\gamma_{\tau,q}^k > 0$ is the probability that type τ of the incumbent chooses first period quantity q , $\lambda_q^k > 0$ is the probability that the entrant enters after observing quantity q , $\delta_{\tau,(q,enter,r)}^k > 0$ and $\delta_{(q,enter,r)}^k > 0$ are the probabilities of second period quantity r being chosen by, respectively, type τ of the incumbent and the entrant, after observing first period quantity q followed by entry and revelation of τ , and $\delta_{\tau,(q,no\ entry,r)}^k > 0$ is the probability of second period quantity r being chosen by type τ of the incumbent after observing first period quantity q followed

by no entry. Specifically, let $\gamma_{\tau,q}^k \equiv \frac{\beta_{\tau,q}^k}{\sum_{\hat{q} \in \mathcal{Q}} \beta_{\tau,\hat{q}}^k}$ for $k = 1, 2, \dots$, where $\beta_{\tau,q}^k$ is defined by

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

, λ_q^k converge to 1 as $k \rightarrow \infty$ when $q < q_L$ and converge to 0 otherwise, $\delta_{\tau,(q,enter,r)}^k$ converge to 1 as $k \rightarrow \infty$ when r is the Cournot quantity for type τ and converge to 0 otherwise, $\delta_{(q,enter,r)}^k$ converge to 1 as $k \rightarrow \infty$ when r is the Cournot quantity for the entrant and converge to 0 otherwise, and $\delta_{\tau,(q,no\ entry,r)}^k$ converge to 1 as $k \rightarrow \infty$ when r is the monopoly quantity for type τ and converge to 0 otherwise. Note that σ^k converges to σ^{LP} .

Theorem A.2 provides a formula (A.15) for an interim belief $\bar{\nu}$ satisfying smooth rule consistency under the assumption that a particular limit exists for information sets where $I_i \not\subseteq \Theta$. We now show this assumption is met. First, for those information sets of the entrant that are on-path according to σ_{-i}^{LP} , $f_i(I_i) = \{L, M, H\}$ and $m_i(I_i) = 0$, thus

$$\lim_{k \rightarrow \infty} \frac{p_{-i,\sigma_{-i}^k}(h^{m_i(I_i)}|h^0)}{\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i,\sigma_{-i}^k}(\hat{h}^{m_i(I_i)}|\hat{h}^0)} = \lim_{k \rightarrow \infty} \frac{1}{3} = \frac{1}{3}.$$

Next consider off-path information sets of the entrant following observation of first period quantity q . If $q \neq q_H, q_L$ (with or without also having entered and learned the incumbent's type is τ), $f_i(I_i) = \{L, M, H\} \times \{q\}$ and $m_i(I_i) = 1$, thus

$$\lim_{k \rightarrow \infty} \frac{p_{-i,\sigma_{-i}^k}(h^{m_i(I_i)}|h^0)}{\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i,\sigma_{-i}^k}(\hat{h}^{m_i(I_i)}|\hat{h}^0)} = \lim_{k \rightarrow \infty} \frac{\gamma_{\tau,q}^k}{\sum_{\hat{\tau} \in \{L,M,H\}} \gamma_{\hat{\tau},q}^k}.$$

For $q_H \neq q < q_L$, calculation shows this limit is 0 if $\tau \in \{L, M\}$ and 1 if $\tau = H$. For $q > q_L$, calculation shows this limit is 0 if $\tau \in \{M, H\}$ and 1 if $\tau = L$.

If $q = q_H$ and the entrant learns the incumbent's type is $\tau \in \{L, M\}$ or $q = q_L$ and the entrant learns the incumbent's type is $\tau = H$, $f_i(I_i) = I_i$, which is a singleton, and $m_i(I_i) = 3$, thus

$$\lim_{k \rightarrow \infty} \frac{p_{-i,\sigma_{-i}^k}(h^{m_i(I_i)}|h^0)}{\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i,\sigma_{-i}^k}(\hat{h}^{m_i(I_i)}|\hat{h}^0)} = \lim_{k \rightarrow \infty} \frac{p_{-i,\sigma_{-i}^k}(h^{m_i(I_i)}|h^0)}{p_{-i,\sigma_{-i}^k}(h^{m_i(I_i)}|h^0)} = 1.$$

There are no other information sets for the entrant before the end of the game. Since, after the ex-ante stage, the incumbent knows its type and this is the only source of incomplete information in the game, the interim beliefs for the incumbent are irrelevant and thus we will not consider them here. Having established the necessary limits exist, Theorem A.2 delivers interim beliefs, $\bar{\nu}$, such that $(\sigma^{\text{LP}}, \bar{\nu})$ satisfies smooth rule consistency.

The final step in the proof is to verify that $(\sigma^{\text{LP}}, \bar{\nu})$ satisfies the interim optimality conditions (2.7) of sequential optimality. By Corollary 2.1, it is enough to check one-stage deviations. The Cournot strategies in the last stage given entry are optimal because all distributions over type become degenerate when conditioned on the entrant learning the incumbent's type. The fact that $w_L < 0$ plus $w_H \geq 0$ implies that it is optimal for the entrant to stay out if its interim objective function places all weight on type L and to enter if its interim objective function places all weight on type H . We now verify that when $q \neq q_L$ the interim objective function does exactly this when entry/no entry are supposed to occur according to σ^{LP} . Entry is supposed to occur if and only if $q < q_L$. When $q = q_H$, since $\pi_{I_i, \sigma^{\text{LP}}_i}$ is the degenerate distribution on type H for all π that may be so conditioned, it is optimal to enter. When $q_H \neq q < q_L$, for $\tau \in \{L, M, H\}$, $\bar{\pi}_{I_i, \sigma^{\text{LP}}_i}(\tau, q) = \frac{\bar{p}_{-i, \sigma^{\text{LP}}_i}(\tau, q | \tau) \pi(\tau)}{\sum_{\hat{\tau} \in \{L, M, H\}} \bar{p}_{-i, \sigma^{\text{LP}}_i}(\hat{\tau}, q | \hat{\tau}) \pi(\hat{\tau})}$ which is 1 if $\tau = H$ and 0 otherwise. Since (A.15) implies that all $\hat{\pi}$ in the support of $\bar{\nu}_{E, q}$ are such that $\hat{\pi} = \bar{\pi}^\pi$, each $\hat{\pi}_{I_i, \sigma^{\text{LP}}_i}$ puts weight only on type H , it is again optimal to enter. Similarly, when $q > q_L$, $\bar{\pi}_{I_i, \sigma^{\text{LP}}_i}(\tau, q)$ is 1 if $\tau = L$ and 0 otherwise, implying that all $\hat{\pi}$ in the support of $\bar{\nu}_{E, q}$ are such that $\hat{\pi}_{I_i, \sigma^{\text{LP}}_i}$ puts weight only on type L , so it is optimal not to enter.

When $q = q_L$, not entering being optimal from an interim perspective is equivalent to the following:

$$\sum_{\pi \in \Delta(\Theta) | \pi(H) < 1} \bar{\nu}_{E, q_L}(\pi) \frac{1}{1 - \pi(H)} (\pi(L)w_L + \pi(M)w_M) \phi'(0) \leq 0. \quad (\text{A.27})$$

Using the formula (A.15) to substitute for $\bar{\nu}_{E, q_L}(\pi)$ in (A.27) yields that not entering remaining optimal is equivalent to (ICL for E). Therefore $(\sigma^{\text{LP}}, \bar{\nu})$ satisfies the interim optimality conditions (2.7) of sequential optimality as long as the entrant's ϕ is at least as concave as the $\hat{\phi}$ identified from Lemma A.6. For such sufficiently concave ϕ , having shown $(\sigma^{\text{LP}}, \bar{\nu})$ is sequentially optimal and satisfies smooth rule consistency, it is therefore an SEA.

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

We next verify that the other conditions in the antecedents of Theorem 2.5 are satisfied.

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

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

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

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

By Assumption 3.1 and (ICM for I), the expression multiplied by $\pi(L)$ is strictly larger than the one multiplied by $\pi(M)$, which is, in turn, strictly larger than the one multiplied by $\pi(H)$. Thus, likelihood-ratio ordering of the support of μ implies strict ordering of $\sum_{h \in H} u_i(h) p_{\sigma^{\text{LP}}}(h|h^0) \pi(h^0)$.

By the first part of Theorem 2.5, ambiguity aversion makes σ^{LP} SEA more belief robust. Furthermore, if $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M > 0\}) > 0$, as implied by the additional assumption in the proposition that some $\pi \in \text{supp } \mu$ makes entry conditional on $\{L, M\}$ strictly profitable, then any $\tilde{\mu}$, with the same support as μ , that puts sufficient weight on the π for which $\pi(L)w_L + \pi(M)w_M$ is most positive, will fail to support σ^{LP} as an ex-ante equilibrium. Using the second part of Theorem 2.5 yields that ambiguity aversion makes σ^{LP} SEA strictly more belief robust. ■

Lemma A.6 *Under the assumptions of Proposition 3.1 there exists an $\alpha > 0$ such that if ϕ is at least as concave as $-e^{-\alpha x}$ then (ICL for E) is satisfied.*

Proof. Assume the conditions of the proposition. We show that (ICL for E) is satisfied for concave enough ϕ . The assumption in the proposition that some $\pi \in \text{supp } \mu$ makes entry conditional on $\{L, M\}$ strictly unprofitable means $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M < 0\}) > 0$. If $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M \leq 0\}) = 1$ then (ICL for E) is trivially satisfied for any ϕ . For the remainder of the proof, therefore, suppose that $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M > 0\}) >$

0. Let $\Pi^- \equiv \{\pi \mid \pi(L)w_L + \pi(M)w_M < 0\}$, $\Pi^+ \equiv \{\pi \mid \pi(L)w_L + \pi(M)w_M > 0\}$, $N \equiv \sum_{\pi \in \Pi^-} \mu(\pi)(\pi(L)w_L + \pi(M)w_M)$, and $P \equiv \sum_{\pi \in \Pi^+} \mu(\pi)(\pi(L)w_L + \pi(M)w_M)$. Let $\pi^- \in \arg \max_{\pi \in \Pi^-} \pi(H)$ and $\pi^+ \in \arg \min_{\pi \in \Pi^+} \pi(H)$. The left-hand side of (ICL for E) can be bounded from above as follows:

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

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

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

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

To complete the proof, fix α satisfying this inequality and consider ϕ such that $\phi(x) = h(-e^{-\alpha x})$ for all x with h concave and strictly increasing on $(-\infty, 0)$. We show that (ICL for E) holds. Observe that $\phi'(x) = h'(-e^{-\alpha x})\alpha e^{-\alpha x}$. Since $\pi^-(H) - \pi^+(H) < 0$ and $w_H > 0$, we have

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

and, by concavity of h ,

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

Therefore the upper bound derived above satisfies

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

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

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

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

B Appendix: Comparative Statics

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

B.1 Example: New Strategic Behavior in Equilibrium

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

There are three players: 1, 2 and 3. First, it is determined whether player 1 is of type I or type II and 1 observes her own type. Players 2 and 3 have only one type, so there is complete information about them. The payoff triples in Figure B.1 describe vNM utility payoffs given players' actions and players' types (i.e., (u_1, u_2, u_3) means that player i receives u_i). Players 2 and 3 have ambiguity about player 1's type and have smooth ambiguity preferences with an associated $\phi_2 = \phi_3 = \phi$ and $\mu_2 = \mu_3 = \mu$. Player 1 also has smooth ambiguity preferences, but nothing in what follows depends on either ϕ_1 or μ_1 . Player 1's first and only move in the game is to choose between action P (lay) which leads to players 2 and 3 playing a simultaneous move game in which their payoffs depend on 1's type, and action Q (uit), which ends the game (equivalently think of it leading to a stage where all players have only one action).¹⁸

¹⁸Note that to eliminate any possible effects of varying players' risk aversion, think of the payoffs being generated using lotteries over two "physical" outcomes, the better of which has utility u normalized to 5/2

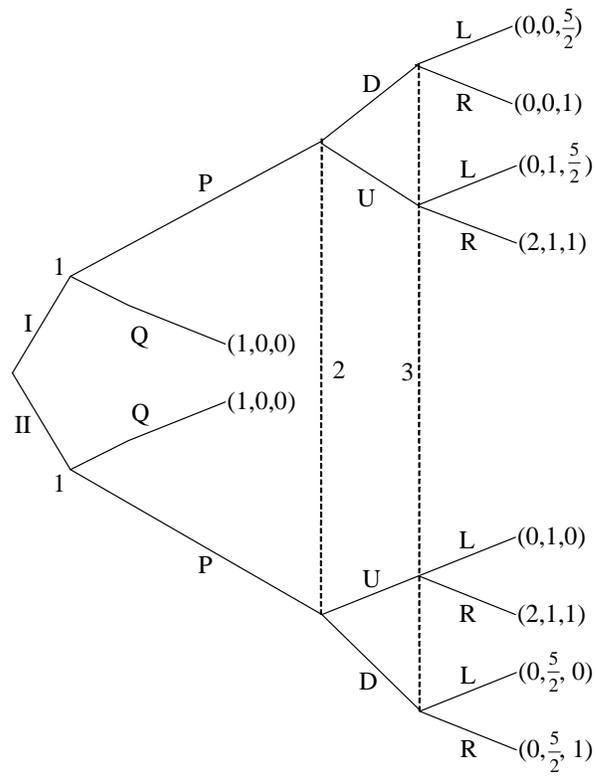


Figure B.1: New equilibrium behavior with ambiguity aversion

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

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

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

which is equivalent to

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

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

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

which is equivalent to

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

Since (B.1) and (B.2) cannot both be satisfied when $p_I + p_{II} > 0$ (i.e., P is played with positive probability), σ must specify that the history (P, U, R) is never realized. This implies that player 1 has an ex-ante profitable deviation to the strategy of always playing Q , contradicting the assumption that σ is an ex-ante equilibrium. ■

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

and the worse of which has u normalized to 0. So, for example, the payoff 1 can be thought of as generated by the lottery giving the better outcome with probability $2/5$ and the worse outcome with probability $3/5$.

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

Proof of Proposition B.2. Let μ put probability $\frac{1}{2}$ on π_0 and $\frac{1}{2}$ on π_1 , where $\pi_0(I) = 1$ and $\pi_1(I) = 0$.¹⁹ Let $\phi(x) \equiv -e^{-x}$.²⁰ Let σ be a strategy profile specifying that both types of player 1 play P with probability 1, player 2 plays U with probability λ^* if given the move and player 3 plays R with probability λ^* if given the move, where $\lambda^* = 1 - \frac{2}{5} \ln(3/2)$. Notice that according to σ , the history (P, U, R) occurs with probability $(1 - \frac{2}{5} \ln(3/2))^2 > \frac{7}{10}$. Observe that player 1 strictly prefers ex-ante to play P with probability 1 for both types if and only if, after the play of P , (U, R) will be played with probability greater than $\frac{1}{2}$. The same is true for each type of player 1 after her type is realized as well. Player 2 ex-ante chooses the probability, $\lambda \in [0, 1]$, with which to play U if given the move to maximize

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

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

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

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

Now consider the following sequence of completely mixed strategies with limit σ : σ^k has each type of player 1 play P with probability $1 - \frac{1}{2k}$, and leaves the strategies otherwise the same as in σ . Theorem A.2 provides a formula (A.15) for an interim belief \bar{v} satisfying smooth rule consistency under the assumption that a particular limit exists for information sets where $I_i \not\subseteq \Theta$. We now show this assumption is met. Recall that only player 1 has more than one possible type, and player 1 learns his type at the beginning of the game. Thus we need only be concerned with the beliefs of players 2 and 3. For the information sets of player 2 and 3 after observing P , $f_i(I_i) = \{I, II\}$ and $m_i(I_i) = 0$, thus

$$\lim_{k \rightarrow \infty} \frac{p_{-i, \sigma_{-i}^k}(h^{m_i(I_i)} | h^0)}{\sum_{\hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} | \hat{h}^0)} = \lim_{k \rightarrow \infty} \frac{1}{2} = \frac{1}{2}.$$

¹⁹The degeneracy of the π in the support of μ is not necessary for the argument to go through – it merely shortens some calculations and reduces the ambiguity aversion required.

²⁰Any more concave ϕ will also work, as will any ϕ more concave than $-e^{-\alpha x}$ for $\alpha = \frac{-4(\ln(2/3))}{5(2-\sqrt{2})} \approx 0.554$.

Therefore $(\sigma, \bar{\nu})$ satisfies smooth rule consistency. It remains to show $(\sigma, \bar{\nu})$ is sequentially optimal. Since

$$\bar{\nu}_{2,\{(I,P),(II,P)\}}(\pi_0) = \bar{\nu}_{3,\{(I,P),(II,P)\}}(\pi_1) \propto \frac{\phi'_i(\lambda^*)}{\phi'_i(\lambda^*)} \frac{1}{2} = \frac{1}{2}$$

and

$$\bar{\nu}_{2,\{(I,P),(II,P)\}}(\pi_1) = \bar{\nu}_{3,\{(I,P),(II,P)\}}(\pi_0) \propto \frac{\phi'_i(\lambda^* + \frac{5}{2}(1 - \lambda^*))}{\phi'_i(\lambda^* + \frac{5}{2}(1 - \lambda^*))} \frac{1}{2} = \frac{1}{2},$$

σ remains optimal for players 2 and 3 following the play of P given $\bar{\nu}$. Thus, $(\sigma, \bar{\nu})$ is sequentially optimal. It is therefore an SEA. ■

As the proof of Proposition B.2 mentions, the example μ is chosen for simplicity, and degeneracy of the measures in its support is not necessary for the result.

B.2 Formal Comparative Statics in Ambiguity Aversion

Notation B.1 For a game $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$, let $E_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ denote the set of all ex-ante equilibria of the game $\hat{\Gamma} = (N, H, (\mathcal{I}_i)_{i \in N}, (\hat{\mu}_i)_{i \in N}, (u_i, \hat{\phi}_i)_{i \in N})$ differing from Γ only in ambiguity aversions and beliefs. Let $Q_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ denote the analogous set of sequentially optimal strategy profiles and $S_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ denote the analogous set of SEA strategy profiles.

Notation B.2 Denote the identity function by ι .

Theorem B.1 There exists a game Γ and $(\hat{\phi}_i)_{i \in N}$ such that

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

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

Examination of the proof shows that, fixing beliefs, not only are the equilibrium strategies distinct under ambiguity aversion compared to ambiguity neutrality, but it can also be that the strategies under ambiguity aversion generate paths of play that do not occur in equilibrium under ambiguity neutrality. An analogue of Theorem B.1 is true for sequential

optima, SEA and any other refinement of ex-ante equilibria as well, as they are all ex-ante equilibria. Thus, with fixed beliefs, change in ambiguity aversion can impact the set of equilibrium strategies and realized play.

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

Theorem B.2 *For all games Γ and $(\hat{\phi}_i)_{i \in N}$, $\bigcup_{\hat{\mu}} E_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} E_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$, and the same holds when Q or S replaces E . There exists a game Γ and $(\hat{\phi}_i)_{i \in N}$ such that all these inclusions are strict and some of the new equilibrium strategies induce new paths of play.*

Proof of Theorem B.2. That $\bigcup_{\hat{\mu}} E_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} E_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ follows by considering only degenerate beliefs on the left-hand side and choosing them to have the same reduced measure as the right-hand side beliefs. $\bigcup_{\hat{\mu}} Q_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} Q_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ follows using the same construction and additionally taking the left-hand side interim beliefs at each information set to be degenerate with the same reduced measure as the right-hand side interim beliefs at the corresponding information set and noting that this preserves optimality at each information set. $\bigcup_{\hat{\mu}} S_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} S_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ follows using the same construction as for sequential optima, observing that the left-hand side degenerate beliefs satisfy smooth rule consistency since the right-hand side beliefs do so. As shown by Propositions B.1 and B.2, Section B.1 provides an example where the inclusion is strict and the new strategies generate new paths of play. ■

For each equilibrium strategy profile of a game, when ambiguity aversion(s) change, there always exist modified beliefs $\hat{\mu}_i$ under which the strategy profile remains an equilibrium. Formally the result is:

Theorem B.3 *Fix a game Γ . For all $(\hat{\phi}_i)_{i \in N}$,*

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

and the same holds when Q or S replaces E .

The constructive proof of this result shows that the beliefs $\hat{\mu}_i$ and interim beliefs $\hat{\nu}_{i, I_i}$ that work for a given equilibrium profile σ are related to the beliefs μ_i and ν_{i, I_i} in the game

with the original ambiguity aversion(s) by the formulae in (B.3) and (B.4) where the ϕ_i are the original and $\hat{\phi}_i$ the new specifications of ambiguity aversions.

Proof of Theorem B.3. Fix Γ and let $\sigma \in E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$. Ex-ante equilibrium is equivalent to ex-ante optimality for all players i of σ_i according to i 's preferences given σ_{-i} . This ex-ante optimality is equivalent to (see (A.6)) $\sigma'_i = \sigma_i$ maximizing

$$\sum_h u_i(h) p_{i, \sigma'_i}(h|h^0) q^{\sigma, i}(h)$$

with respect to σ'_i , where $q^{\sigma, i}(h)$ is i 's ex-ante σ -local measure as in (A.3). Let $\hat{\mu}_i$ be the probability measure such that

$$\hat{\mu}_i(\pi) \propto \frac{\phi'_i(\sum_h u_i(h) p_\sigma(h|h^0) \pi(h^0))}{\hat{\phi}'_i(\sum_h u_i(h) p_\sigma(h|h^0) \pi(h^0))} \mu_i(\pi). \quad (\text{B.3})$$

Using $\hat{\phi}_i$ and $\hat{\mu}_i$, ex-ante optimality for player i as a function of σ'_i is equivalent to $\sigma'_i = \sigma_i$ maximizing

$$\begin{aligned} & \sum_h u_i(h) p_{i, \sigma'_i}(h|h^0) \sum_\pi \hat{\phi}'_i \left(\sum_{\hat{h}} u_i(\hat{h}) p_\sigma(\hat{h}|\hat{h}^0) \pi(\hat{h}^0) \right) p_{-i, \sigma_{-i}}(h|h^0) \pi(h^0) \hat{\mu}_i(\pi) \\ & \propto \sum_h u_i(h) p_{i, \sigma'_i}(h|h^0) \sum_\pi \phi'_i \left(\sum_{\hat{h}} u_i(\hat{h}) p_\sigma(\hat{h}|\hat{h}^0) \pi(\hat{h}^0) \right) p_{-i, \sigma_{-i}}(h|h^0) \pi(h^0) \mu_i(\pi) \\ & = \sum_h u_i(h) p_{i, \sigma'_i}(h|h^0) q^{\sigma, i}(h). \end{aligned}$$

Thus σ_i is ex-ante optimal for player i given $\hat{\phi}_i$, $\hat{\mu}_i$ and σ_{-i} . As this is true for each player i , $\sigma \in E_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

Turn now to sequentially optimal strategy profiles. Suppose $\sigma \in Q_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (σ, ν) is sequentially optimal for Γ . Let $\hat{\mu}_i$ be defined as in (B.3) and for each I_i , $\hat{\nu}_{i, I_i}$ be the probability measure such that

$$\hat{\nu}_{i, I_i}(\pi) \propto \frac{\phi'_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_\sigma(h|h^s(I_i)) \pi_{I_i, \sigma_{-i}}(h^s(I_i)) \right)}{\hat{\phi}'_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_\sigma(h|h^s(I_i)) \pi_{I_i, \sigma_{-i}}(h^s(I_i)) \right)} \nu_{i, I_i}(\pi). \quad (\text{B.4})$$

By the argument in the ex-ante equilibrium part of this proof, $\sigma \in E_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. Using $\hat{\phi}_i$ and $\hat{\nu}_{i, I_i}$, interim optimality for player i at I_i as a function of σ'_i is equivalent (see

(A.5) and (A.4)) to $\sigma'_i = \sigma_i$ maximizing

$$\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{s(I_i)}) \sum_{\pi|\pi(f_i^{\sigma-i}(I_i))>0} \hat{\phi}'_i \left(\sum_{\hat{h}|\hat{h}^{s(I_i)} \in I_i} u_i(\hat{h}) p_\sigma(\hat{h}|\hat{h}^{s(I_i)}) \pi_{I_i,\sigma-i}(\hat{h}^{s(I_i)}) \right) \cdot p_{-i,\sigma-i}(h|h^{s(I_i)}) \pi_{I_i,\sigma-i}(h^{s(I_i)}) \hat{\nu}_{i,I_i}(\pi). \quad (\text{B.5})$$

Since (B.5) is proportional to

$$\begin{aligned} & \sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{s(I_i)}) \sum_{\pi|\pi(f_i^{\sigma-i}(I_i))>0} \phi'_i \left(\sum_{\hat{h}|\hat{h}^{s(I_i)} \in I_i} u_i(\hat{h}) p_\sigma(\hat{h}|\hat{h}^{s(I_i)}) \pi_{I_i,\sigma-i}(\hat{h}^{s(I_i)}) \right) \\ & \quad \cdot p_{-i,\sigma-i}(h|h^{s(I_i)}) \pi_{I_i,\sigma-i}(h^{s(I_i)}) \nu_{i,I_i}(\pi) \\ = & \sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{s(I_i)}) q^{(\sigma,\nu),i,I_i}(h), \end{aligned}$$

optimality of σ_i at such an I_i in game Γ implies (see (A.5)) the same in game $\hat{\Gamma}$. This is true for each player i . Thus, $(\sigma, \hat{\nu})$ is sequentially optimal in $\hat{\Gamma}$.

We now extend the argument to SEA. Suppose $\sigma \in S_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (σ, ν) is an SEA for Γ . As above, let $\hat{\mu}_i$ be as in (B.3) and for each I_i , $\hat{\nu}_{i,I_i}$ be defined as in (B.4). By our previous arguments, $(\sigma, \hat{\nu})$ is sequentially optimal in $\hat{\Gamma}$. Since ν is determined by smooth rule updating using σ , it follows from inspection of (B.4) and the formulation of the smooth rule in (A.16) and (2.8) that $\hat{\nu}$ satisfies smooth rule updating using σ . It remains to show that $(\sigma, \hat{\nu})$ satisfies the remaining parts of smooth rule consistency.

Since (σ, ν) satisfies 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 $\hat{\nu} = \lim_{k \rightarrow \infty} \hat{\nu}^k$, where $\hat{\nu}$ is adapted from ν using σ , each $\hat{\nu}^k$ is adapted from ν^k using σ^k , and each ν^k is the interim belief system determined by smooth rule updating using σ^k . Thus we can invoke Lemma A.2 and Theorem A.2 (respectively, Theorem A.2) to conclude that $(\sigma, \bar{\nu})$ (respectively, $(\sigma, \bar{\hat{\nu}})$) with $\bar{\nu}$ (respectively, $\bar{\hat{\nu}}$) defined according to the formula in (A.15) with ϕ_i and μ_i (respectively, with $\hat{\phi}_i$ and $\hat{\mu}_i$) satisfies smooth rule consistency in Γ (respectively, $\hat{\Gamma}$). Since (σ, ν) and $(\sigma, \bar{\nu})$ satisfy smooth rule consistency using the same sequence $\{\sigma^k\}_{k=1}^\infty$, the adaptations from ν and $\bar{\nu}$ using σ must be equal (to $\hat{\nu}$). By using (B.3) and (B.4) and observing that the expectation terms in (B.4) are constant with respect to the summation in Definition 2.10, one can establish that the adaptations from $\hat{\nu}$ and $\bar{\hat{\nu}}$ using σ must also be equal to each other. It follows that $(\sigma, \hat{\nu})$ satisfies smooth rule consistency in $\hat{\Gamma}$. Therefore $(\sigma, \hat{\nu})$ is an SEA of $\hat{\Gamma}$.

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

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

Definition B.1 For a game $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$, let $\tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ be the set of all ex-ante equilibria with respect to pure strategies of a game $\hat{\Gamma} = (N, H, (\mathcal{I}_i)_{i \in N}, (\hat{\mu}_i)_{i \in N}, (u_i, \hat{\phi}_i)_{i \in N})$ differing from Γ only in ambiguity aversions and beliefs. Let $\tilde{Q}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ and $\tilde{S}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ be the analogous respective sets of sequentially optimal and SEA strategy profiles with respect to pure strategies.

The formal result delivering the pure strategy comparative statics results referenced in Section 2.5 is:

Theorem B.4 Fix a game Γ . For all $(\hat{\phi}_i)_{i \in N}$ such that, for each i , $\hat{\phi}_i$ is at least as concave as ϕ_i , $\tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$, and the same holds when \tilde{Q} or \tilde{S} replaces \tilde{E} . There exists a game Γ and $(\hat{\phi}_i)_{i \in N}$ such that for each i , $\hat{\phi}_i$ is at least as concave as ϕ_i , all these inclusions are strict and some of the new equilibrium strategies induce new paths of play.

Proof of Theorem B.4. Fix a game Γ . Suppose $\varsigma \in \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$, for each i , $\hat{\phi}_i = \chi_i(\phi_i)$ for some increasing, differentiable and concave χ_i (note that differentiability of χ_i is implied by the continuous differentiability of $\hat{\phi}_i$ in the class of games considered in this paper) and $\hat{\mu}_i$ is the probability measure such that

$$\hat{\mu}_i(\pi) \propto \frac{\mu_i(\pi)}{\chi_i'(\phi_i(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0)))}. \quad (\text{B.6})$$

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

$$\sum_\pi \phi_i \left(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0) \right) \mu_i(\pi) \geq \sum_\pi \phi_i \left(\sum_h u_i(h)p_{(\varsigma'_i, \varsigma_{-i})}(h|h^0)\pi(h^0) \right) \mu_i(\pi). \quad (\text{B.7})$$

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

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

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

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

where the last inequality follows from B.7. Since this is true for each i and each $\varsigma'_i, \varsigma \in \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. This shows $\tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

Turn now to the part of the theorem about sequentially optimal strategy profiles. Suppose $\varsigma \in \tilde{Q}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (ς, ν) is sequentially optimal for Γ with respect to pure strategies. Further suppose that for each i , $\hat{\phi}_i = \chi_i(\phi_i)$ for some increasing, differentiable and concave χ_i , $\hat{\mu}_i$ is defined as in (B.6), and for each I_i , $\hat{\nu}_{i, I_i}$ is the probability measure such that

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

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

$$\begin{aligned} & \sum_\pi \phi_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_\varsigma(h|h^s(I_i)) \pi_{I_i, \varsigma_{-i}}(h^s(I_i)) \right) \nu_{i, I_i}(\pi) \tag{B.8} \\ & \geq \sum_\pi \phi_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_{(\varsigma'_i, \varsigma_{-i})}(h|h^s(I_i)) \pi_{I_i, \varsigma_{-i}}(h^s(I_i)) \right) \nu_{i, I_i}(\pi). \end{aligned}$$

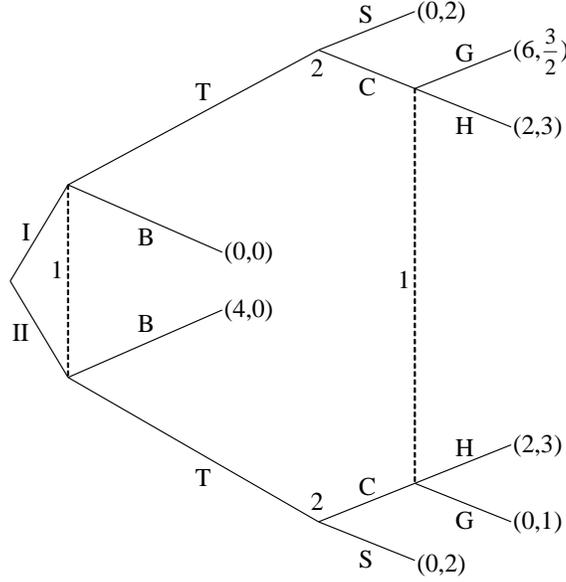


Figure B.2: Ambiguity aversion generates new equilibria with respect to pure strategies

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

Finally, turn to the part of the theorem about SEA strategy profiles. Suppose $\varsigma \in \tilde{S}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (ς, ν) is an SEA for Γ with respect to pure strategies, where the sequence used in satisfying smooth rule consistency is $\{\sigma^k\}_{k=1}^\infty$. Further suppose that for each i , $\hat{\phi}_i = \chi_i(\phi_i)$ for some increasing, differentiable and concave χ_i and $\hat{\mu}_i$ is defined as in (B.6). Invoking Lemma A.2 and Theorem A.2, $(\varsigma, \bar{\nu})$ (respectively, $(\sigma, \bar{\nu})$) with $\bar{\nu}$ (respectively, $\bar{\nu}$) defined according to the formula in (A.15) with ϕ_i and μ_i (respectively, with $\hat{\phi}_i$ and $\hat{\mu}_i$) satisfies smooth rule consistency in Γ (respectively, $\hat{\Gamma}$) with corresponding sequence $\{\sigma^k\}_{k=1}^\infty$. Inspection of the expressions for $\bar{\nu}_{i, I_i}$ and $\bar{\nu}_{i, I_i}$ yields that for each I_i ,

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

By the argument in the ex-ante equilibrium part of this proof, $\varsigma \in \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. By the argument in the sequentially optimal part of the proof, $\varsigma \in \tilde{Q}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. Thus $\varsigma \in \tilde{S}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. This shows $\tilde{S}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{S}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

To prove that strict inclusions may happen, consider the game depicted in Figure B.2. There are two players, 1 and 2. First, it is determined whether player 2 is of type I or

type II and 2 observes his own type. Player 1 does not observe the type. The payoff pairs in Figure B.2 describe vNM utility payoffs given players' actions and players' types (i.e., (u_1, u_2) means that player i receives u_i). Player 1 has ambiguity about player 2's type and has smooth ambiguity preferences with an associated ϕ_1 and μ_1 . Player 2 also has smooth ambiguity preferences, with ϕ_2 and μ_2 . Player 1's first move in the game is to choose between action T (wo) which gives the move to player 2 and action B (et) (i.e., betting that player 2 is of type II) which ends the game. If T , then player 2's move is a choice between C (ontinue) which leads to player 1 again being given the move, and S (top) which ends the game. If C , then player 1 has a choice between G (amble) and H (edge) after which the game ends.

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

By the weak inclusions already shown, and since SEA implies sequentially optimal, which in turn implies ex-ante equilibrium, it is enough to show that for some strictly concave $\hat{\phi}_1$ there is an SEA strategy profile with respect to pure strategies not contained in $\bigcup_{(\mu_i)_{i \in N}} \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\iota)_{i \in N})$. To this end, suppose $\hat{\phi}_1(x) \equiv -e^{-2x}$, $\hat{\phi}_2 \equiv \iota$ and $\mu_1(\pi_1) = \mu_1(\pi_2) = \frac{1}{2}$, where $\pi_1(I) = \frac{1}{4}$ and $\pi_2(I) = \frac{3}{4}$. Consider the pure strategy profile $(T, (C, C), H)$ and a sequence of completely mixed strategy profiles approaching it where the k^{th} element of the sequence has player 1 and each type of player 2 playing the action not assigned by $(T, (C, C), H)$ with probability $\frac{1}{k}$ at any point they are given the move. Since all information sets are on-path given $(T, (C, C), H)$, $\bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0)$ always exists and is constant for each player. Let interim beliefs \bar{v}_1 for player 1 be defined according to the formula in (A.15) with $\hat{\phi}_1$ and μ_1 . By calculation, $\bar{v}_{1, \{(I, T, C), (II, T, C)\}}(\pi_1) = \bar{v}_{1, \{(I, T, C), (II, T, C)\}}(\pi_2) = \frac{1}{2}$. By Theorem A.2, \bar{v} satisfies smooth rule consistency. Since $\hat{\phi}_1(2) > \frac{1}{2}(\hat{\phi}_1(3) + \hat{\phi}_1(1))$ and $\hat{\phi}_1(2) > \frac{1}{2}(\hat{\phi}_1(\frac{3}{2}) + \hat{\phi}_1(\frac{9}{2}))$, and since C is always a best response for player 2 given 1 plays T and then H , $(T, (C, C), H)$ is sequentially optimal with respect to pure strategies given \bar{v} . Therefore $(T, (C, C), H) \in \tilde{S}_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ and the proof is complete. ■

C Appendix: Details on Figure 1.1 and the comparison with no profitable one-stage deviations and consistent planning

A strengthening of the no profitable one-stage deviations criterion appearing in some of the existing literature investigating games with ambiguity is the following condition, describing a consistent planning requirement in the spirit of Strotz (1955-56):

Definition C.1 Fix a game Γ and a pair (σ, ν) consisting of a strategy profile and interim belief system. Specify V_i and V_{i,I_i} as in (2.1) and (2.2). For each player i and information set $I_i \in \mathcal{I}_i^T$, let

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

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

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

Finally, let

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

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

$$\sigma_i \in CP_i.$$

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

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

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

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

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

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

If (σ, ν) is sequentially optimal then it is also optimal under consistent planning. However, if (σ, ν) is optimal under consistent planning it may fail to be sequentially optimal (even when limiting attention to ambiguity neutrality). For such a failure to occur, the optimal strategy from player i 's point of view at some earlier stage must have a continuation that fails to be optimal from the viewpoint of some later reachable stage. This is what makes the extra constraints imposed in the optimization inequalities under consistent planning bind. Just as with no profitable one-stage deviations, when updating is according to the smooth rule, (σ, ν) optimal under consistent planning implies (σ, ν) is sequentially optimal, making the two equivalent under smooth rule updating. This follows from Theorem 2.2 and the fact that consistent planning implies no profitable one-stage deviations. Thus, under smooth rule updating, sequential optimality, optimality under consistent planning, and no profitable one-stage deviations are *equivalent*. These observations are generalizations of the fact that updating according to Bayes' rule makes all three concepts equivalent for expected utility preferences.

Recall that the example in Figure 1.1 in the Introduction showed how the no profitable one-stage deviation criterion under Bayesian updating allowed strategy profiles that are not even ex-ante (Nash) equilibria of a game (and thus clearly not sequentially optimal). Replacing no profitable one-stage deviations by consistent planning does not change this fact. It is easy to specify ϕ_1 , μ and an interim belief system for player 1 such that o can be played with positive probability while satisfying consistent planning and no profitable one-stage deviations. For example, this is the case if $\phi_1(x) = -e^{-10x}$, μ is $1/2$ on $(1/3, 1/9, 5/9)$ and $1/2$ on $(1/3, 5/9, 1/9)$, and 1's beliefs after seeing U are given by Bayes' rule applied to μ : $1/3$ on $(3/4, 1/4, 0)$ and $2/3$ on $(3/8, 5/8, 0)$. With these parameters and beliefs, the following strategy profile satisfies no profitable one-stage deviations and consistent planning: player 1 plays o with probability $1 - \frac{9}{20} \ln(\frac{29}{11}) \approx 0.564$ and mixes evenly between u and d if U , while player 2 plays her strictly dominant strategy if given the move. Notice, if we consider any more concave ϕ_1 , playing o with even higher probability will be consistent with consistent planning or no profitable one-stage deviations given these beliefs. In the limit where the decision maker is Maxmin EU with set of priors equal to the convex combinations of $(1/3, 1/9, 5/9)$ and $(1/3, 5/9, 1/9)$ and applies Bayes' rule to each measure in the set, playing o with probability 1 is consistent with consistent planning and no profitable one-stage deviations.