

Incomplete Information Games with Ambiguity Averse Players*

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

This version: Feb 6, 2016

Abstract

We study incomplete information games involving players who perceive ambiguity about the types of others and may be ambiguity averse as modeled through smooth ambiguity preferences (Klibanoff, Marinacci and Mukerji, 2005). Our focus is on multi-stage games with observed actions and 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 about types. 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 beliefs generated using a particular generalization of Bayesian updating. We propose and analyze two strengthenings of sequential optimality. Examples illustrate new strategic behavior that can arise under ambiguity aversion. Our concepts and framework are also suitable for examining the strategic use of ambiguity.

*We thank Pierpaolo Battigalli, Itzhak Gilboa, Frederic Koessler, Fabio Maccheroni, Jean-Marc Tallon and participants at 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@post.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

1 Introduction

Dynamic games of incomplete information are the subject of a large literature, both theory and application, with diverse fields including models of firm competition, agency theory, auctions, search, insurance and many others. In such games, how players perceive and react to uncertainty, and the way it evolves over the course of the game, is of central importance. In 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). We propose equilibrium notions for incomplete information games involving ambiguity about 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. In the smooth ambiguity model it is possible to hold the players' information fixed while varying their ambiguity attitude from aversion to neutrality (i.e., 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 multi-stage games with observed actions, and on equilibrium notions capturing perfection analogous to those in standard theories for ambiguity neutral players, such as Sequential equilibrium (Kreps and Wilson, 1982) and Perfect Bayesian equilibrium (PBE) (e.g., Fudenberg and Tirole, 1991a,b).

We first define an ex-ante equilibrium concept allowing for aversion to ambiguity about players' types. When there is no type uncertainty, this collapses to Nash equilibrium. When there are common beliefs and ambiguity neutrality, it becomes Bayesian Nash equilibrium. Next, we refine ex-ante equilibrium by imposing perfection in the form of a *sequential optimality* requirement – each player i 's strategy must be optimal at each stage given the strategies of the other players and i 's conditional beliefs about types. When there is no type 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 along with the potential to cut 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 dynamically consistent generalization of Bayesian updating for smooth ambiguity preferences, called the smooth rule (Hanany and Klibanoff 2009). An important method facilitating analysis of dynamic games with standard preferences is

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

the sufficiency of checking only one-stage deviations (as opposed to general deviations) when verifying optimality. 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 there are no profitable one-stage deviations with respect to beliefs updated according to the smooth rule.

Sequential optimality places little restriction on player i 's beliefs about types at stage games that could be reached only immediately following a deviation of at least one other player $j \neq i$. We propose two refinements of sequential optimality restricting such beliefs: Sequential Equilibrium with Ambiguity (SEA) and Perfect Equilibrium with Ambiguity (PEA). In addition to sequential optimality, SEA imposes a generalization of Kreps and Wilson's (1982) consistency condition from their definition of sequential equilibrium. PEA replaces this additional requirement of SEA with the property that, under certain conditions, any player i 's conditional beliefs about player $j \neq i$'s type remain the same as they were at the previous stage if player j had no choice (or only one action) available at that stage. These strengthenings are nested: any SEA is shown to be a PEA. We show that even the strongest of these concepts, SEA, exists for any multi-stage game with observed actions and incomplete information, and for any specification of players' ambiguity aversion and initial beliefs.

In Section 3, we provide several examples that apply our equilibrium notions. First, we present a game with a path that is played in an SEA given sufficient ambiguity aversion, but is never an ex-ante equilibrium (and thus also not sequentially optimal, a PEA, an SEA, a PBE or a sequential equilibrium) given ambiguity neutrality. This example is truly strategic in that it relies on one player recognizing the ambiguity aversion of others and changing play because of it. In fact, we show (Theorem 2.7) that strategic interaction is necessary for ambiguity aversion to generate new equilibrium behavior in any example. Second, we consider a game with a path that is played in an SEA given sufficient ambiguity aversion, but is never played in a PEA (or PBE) given ambiguity neutrality. This example illustrates how differential screening based on beliefs and ambiguity attitudes can give rise to strategic behavior ruled out under ambiguity neutrality. Third, we present an example of a Milgrom and Roberts (1982)-style limit pricing entry game with an SEA involving limit pricing and non-trivial smooth rule updating on the equilibrium path that departs from Bayes' rule. We provide conditions under which ambiguity aversion makes limit pricing more robust.

In Section 4, building on ideas of Aumann (1974) and Bade (2011), we demonstrate that our equilibrium notions and framework can also be used to model strategic ambiguity through strategies that are optimally chosen to be contingent on payoff irrelevant types about which there is ambiguity. We present an example in which a principal strictly benefits from conditioning her cheap talk message to her agents on such payoff-irrelevant ambiguous types.

Our analysis establishes that this strategic use of ambiguity occurs as part of a sequential optimum. This feature is missing in the analyses in recent literature on the role of ambiguous communication (e.g., Bose and Renou, 2014 and Kellner and Le Quement, 2015).

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), di Tillio, Kos, Messner (2012), Auster (2013), Riedel and Sass (2013), Wolitzky (2013, 2014) and Kellner (2015). In contrast, 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 2013, 2015, Bose and Renou 2014, Mouraviev, Riedel and Sass 2015, Battigalli et al. 2015a,b, Dominiak and Lee 2015). 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. In Section 5, we define optimality under consistent planning and say more about how these approaches relate to ours, and also discuss some possible extensions, including to Maxmin expected utility (Gilboa and Schmeidler, 1989) preferences.

2 Model

We begin by defining the central domain of the paper, multi-stage games with observed actions and incomplete information (cf. Fudenberg and Tirole, 1991a, Chapter 8.2.3) where players have (weakly) ambiguity averse smooth ambiguity preferences. It is on this domain that we will develop and apply our equilibrium concepts. While this class of games is broad enough to cover many applications to economics and elsewhere, it does embody some limitations. In such games, the only observation that a player may see while others do not is what is revealed to her at the start of the game by nature. There are no private observations as the game proceeds. 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 (i.e., $\mathcal{T} = \{0\}$).

Definition 2.1 *A (finite) extensive-form multi-stage game with observed actions and incomplete information and (weakly) ambiguity averse smooth ambiguity preferences, Γ , is*

a tuple $(N, \mathcal{T}, (A_i^t)_{i \in N, t \in \mathcal{T}}, (\Theta_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$ where N is a finite set of players, $\mathcal{T} = \{0, 1, \dots, T\}$ is the set of stages, $A_i^t(\eta^t)$ gives the finite set of actions (possibly singleton) available to player i in stage t as a function of the partial history $\eta^t \in H^t$ of action profiles up to (but not including) time t , where the sets of partial histories are defined by $H^0 = \{\emptyset\}$ and, for $1 \leq t \leq T + 1$, $H^t \equiv \{(\eta^{t-1}, a) \mid \eta^{t-1} \in H^{t-1}, a \in \prod_{j \in N} A_j^{t-1}(\eta^{t-1})\}$, Θ_i is the finite set of possible “types” for player i , 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 $\Theta \subseteq \prod_{j \in N} \Theta_j$ and $\Delta(\Theta)$ is the set of all probability measures over Θ , $u_i : H \times \Theta \rightarrow \mathbb{R}$ is the utility payoff of player i given the history of actions ($H \equiv H^{T+1}$) and the type of each player, and $\phi_i : u_i(H \times \Theta) \rightarrow \mathbb{R}$ is a continuously differentiable, concave and strictly increasing function.

All of the definitions and formal results of this paper continue to hold if restricted to the class of games with a common μ such that $\mu_i = \mu$ for all players i . Furthermore, none of our examples or the conclusions we draw from them will rely on differences in the μ_i .

To interpret u_i in this definition, one can think of this utility function as coming from the composition of two more fundamental functions. The first function $c_i : H \times \Theta \rightarrow Z$ is a consequence function determined by the structure of the game – for each history and type profile, it specifies the consequence or prize or outcome $z \in Z$ received by player i . The second function is a vNM utility over consequences, $w_i : Z \rightarrow \mathbb{R}$.² Assume that Z is big enough so that $u_i(H \times \Theta)$ is interior in $w_i(Z)$.³

Given a history $h \in H$ and stage $t \leq T + 1$, $h^t \in \prod_{j \in N} \prod_{s < t} A_j^s(h^s)$ is the partial history up to but not including t specified by h . It is useful to define a strategy for player i as specifying the distribution over i 's actions conditional on each possible partial history and each possible type of player i . Formally:

Definition 2.2 A (behavior) strategy for player i in a game Γ is a function σ_i such that $\sigma_i(h^t, \theta_i) \in \Delta(A_i^t(h^t))$ for each type θ_i , history h and stage t .

Given a strategy σ_i for player i , the continuation strategy at stage t given partial history h^t , $\sigma_i^{h^t}$, is the restriction of σ_i to the set of all partial histories starting with h^t . 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. Similarly, the associated continuation strategy profile at stage t given partial history h^t is $\sigma^{h^t} \equiv (\sigma_i^{h^t})_{i \in N}$.

²Notice that w_i is independent of θ , so that even though it might appear that $u_i = w_i \circ c_i$ is state-dependent, this does not mean that decision-theoretically we are in a state-dependent setting, since the dependence is only via the usual dependence of the consequence of an act on the state.

³This will be convenient for some later optimality characterizations, the proofs of Theorem 2.3 and Lemma A.1 in particular.

For a history h , the action taken at stage s by player j is denoted by $h_{s,j}$. Given a strategy profile σ , type profile θ , history h and $0 \leq r < t \leq T + 1$, the probability of reaching h^t starting from h^r is $p_{\sigma,\theta}(h^t|h^r) \equiv \prod_{j \in N} \prod_{r \leq s < t} \sigma_j(h^s, \theta_j)(h_{s,j})$. It will be useful in what follows to separate this probability into a part affected only by σ_i and θ_i and a part affected only by σ_{-i} and θ_{-i} . These are $p_{i,\sigma,\theta}(h^t|h^r) \equiv \prod_{r \leq s < t} \sigma_i(h^s, \theta_i)(h_{s,i})$ and $p_{-i,\sigma,\theta}(h^t|h^r) \equiv \prod_{j \neq i} \prod_{r \leq s < t} \sigma_j(h^s, \theta_j)(h_{s,j})$ respectively, with $p_{i,\sigma,\theta}(h^t|h^r)p_{-i,\sigma,\theta}(h^t|h^r) = p_{\sigma,\theta}(h^t|h^r)$. 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 *Fix a game Γ . Ex-ante (before own-types are known), each player i ranks strategy profiles σ according to*

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

Using these preferences we define ex-ante equilibrium:

Definition 2.3 *Fix a game Γ . A strategy profile σ^* is an ex-ante equilibrium if, for all players i ,*

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

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

When there is no type uncertainty, the definition collapses to Nash equilibrium. 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 (expected utility) and $\mu_i = \mu$ for all players i , the definition reduces to the usual (ex-ante) Bayesian Nash Equilibrium definition.

We next turn to defining preferences beyond the ex-ante stage. Given a player i of type τ_i a partial history h^t , and a strategy profile σ , consider the set of type profiles consistent with τ_i that make h^t reachable without requiring a deviation from σ by players other than i . It is possible that this set might be empty (i.e., h^t can be reached only by some player(s) deviating from σ_{-i}). For this reason, consider the furthest point back from h^t for which there is some θ_{-i} such that getting to h^t from that point requires no deviation from σ_{-i} . Note that such a point always exists, as h^t is always reachable from h^t itself. We will be interested in the set of type profiles consistent with τ_i that make h^t reachable from such a point without requiring a deviation from σ by players other than i . Formally this set is the following:

Notation 2.1 $\Theta_{i,\tau_i,h^t} \equiv \{\theta \in \Theta \mid \theta_i = \tau_i \text{ and } p_{-i,\sigma,\theta}(h^t|h^{m_i(h^t)}) > 0\}$, where

$$m_i(h^t) \equiv \min_{\theta} \bigcup \{r \in \{0, \dots, t\} \mid p_{-i,\sigma,\theta}(h^t|h^r) > 0\}.$$

Using $m_i(h^t)$ we can make precise what it means for one partial history to be reachable from another:

Definition 2.4 Given a strategy profile σ , player i views partial history h^t with $t \geq 1$ as reachable from h^s (where $0 \leq s \leq t$) if $m_i(h^t) \leq s$.

The following expresses a defining property for interim (second-order) beliefs of player i of type τ_i given partial history h^t and strategy profile σ : that they assign weight only to type distributions that assign positive probability to type profiles in Θ_{i,τ_i,h^t} .

Definition 2.5 An interim belief for player i of type τ_i in a game Γ given partial history h^t and strategy profile σ is a finite support probability measure ν_{i,τ_i,h^t} over $\Delta(\Theta)$ such that

$$\nu_{i,\tau_i,h^t}(\{\pi \in \Delta(\Theta) \mid \pi(\Theta_{i,\tau_i,h^t}) > 0\}) = 1. \quad (2.2)$$

Given a strategy profile σ , an interim belief system $\nu \equiv (\nu_{i,\tau_i,h^t})_{i \in N, \tau_i \in \Theta_i, h^t \in H^t}$ is an interim belief for each type of each player at each partial history. The associated interim belief profile given partial history h^t is $\nu^{h^t} \equiv (\nu_{i,\tau_i,h^t})_{i \in N, \tau_i \in \Theta_i}$.

Fundamental to our equilibrium notion will be sequential optimality. It requires that each player plays optimally for each partial history and each own-type realization given the strategies of the others. This optimality is required even when the partial history is before-the-fact viewed as a null event according to the given strategy profile combined with the beliefs of the player. In order to describe optimality for i given τ_i and h^t , we need to write i 's conditional preferences. These make use of interim beliefs.

Assumption 2.2 Fix a game Γ and an interim belief system ν . Any player i of type τ_i at partial history h^t ranks strategy profiles σ according to

$$V_{i,\tau_i,h^t}(\sigma; \nu) \equiv \sum_{\pi \in \Delta(\Theta) \mid \pi(\Theta_{i,\tau_i,h^t}) > 0} \phi_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H \mid \hat{h}^t = h^t} u_i(\hat{h}, \hat{\theta}) p_{i,\sigma,\hat{\theta}}(\hat{h} \mid h^t) p_{-i,\sigma,\hat{\theta}}(\hat{h} \mid h^t) \pi_{\Theta_{i,\tau_i,h^t}}(\hat{\theta}) \right) \nu_{i,\tau_i,h^t}(\pi), \quad (2.3)$$

where

$$\pi_{\Theta_{i,\tau_i,h^t}}(\theta) = \frac{p_{-i,\sigma,\theta}(h^t|h^{m_i(h^t)})\pi(\theta)}{\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} p_{-i,\sigma,\hat{\theta}}(h^t|h^{m_i(h^t)})\pi(\hat{\theta})} \text{ if } \theta \in \Theta_{i,\tau_i,h^t} \text{ and 0 otherwise.} \quad (2.4)$$

Compared to the ex-ante preferences given in (2.1), the conditional preferences (2.3) differ only in that (1) the beliefs may have changed in light of τ_i and h^t (μ_i is replaced by ν_{i,τ_i,h^t} and π by $\pi_{\Theta_{i,\tau_i,h^t}}$), and (2) the probabilities of reaching various histories according to the strategy profile are now calculated starting from h^t rather than from the beginning of the game. Observe that, while $\pi_{\Theta_{i,\tau_i,h^t}}$ is calculated using Bayes' formula (see the Remark below), there is no restriction placed at this point on the ν other than (2.2).

Remark 2.1 To see that the conditioning formula for π in (2.4) is the usual Bayes' formula, note that, because $\theta_i = \tau_i$ for all $\theta \in \Theta_{i,\tau_i,h^t}$,

$$\frac{p_{-i,\sigma,\theta}(h^t|h^{m_i(h^t)})\pi(\theta)}{\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} p_{-i,\sigma,\hat{\theta}}(h^t|h^{m_i(h^t)})\pi(\hat{\theta})} = \frac{p_{\sigma,\theta}(h^t|h^{m_i(h^t)})\pi(\theta)}{\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} p_{\sigma,\hat{\theta}}(h^t|h^{m_i(h^t)})\pi(\hat{\theta})}$$

if $\theta \in \Theta_{i,\tau_i,h^t}$ and $p_{i,\sigma,\theta}(h^t|h^{m_i(h^t)}) > 0$.

Furthermore, as long as $m_i(h^t) < t$ (so that one may go back at least one stage from h^t without a deviation by players other than i and therefore $m_i(h^t) = m_i(h^{t-1})$), such conditional probabilities are also related by the one-step-ahead Bayes' formula

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

Using these preferences, we may now define sequential optimality:

Definition 2.6 Fix a game Γ . A pair (σ^P, ν^P) consisting of a strategy profile and interim belief system is sequentially optimal if, for all players i , all types τ_i and all partial histories h^t ,

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

and

$$V_{i,\tau_i,h^t}(\sigma^P; \nu^P) \geq V_{i,\tau_i,h^t}((\sigma'_i, \sigma^P_{-i}); \nu^P) \quad (2.7)$$

for all $\sigma'_i \in \Sigma_i$, where the V_i and V_{i,τ_i,h^t} are as specified in (2.1) and (2.3).

Note that since $V_{i,\tau_i,h^t}(\sigma;\nu) = V_{i,\tau_i,h^t}(\hat{\sigma};\nu)$ if $\sigma^{h^t} = \hat{\sigma}^{h^t}$ for type τ_i , requiring the inequalities for the V_{i,τ_i,h^t} to hold as i changes only her continuation strategy given h^t and τ_i would result in an equivalent definition. 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.

While sequential optimality does not place restrictions on the interim belief system beyond (2.5), observe that V_{i,τ_i,h^t} does entail the assumption that, even at stages immediately following a deviation by some player other than i , player i continues to assume that the players $-i$ will play according to σ_{-i} from the current stage onward.

Assuming a common μ , sequential optimality implies subgame perfection adapted to allow for smooth ambiguity preferences. In multistage games with observed actions, the only proper subgames occur at stages where all type uncertainty (if any) has been resolved. For any such proper subgame, (2.7) ensures that the continuation strategy profile derived from σ^P forms a Nash equilibrium of the subgame. For the overall game, (2.6) ensures σ^P is an ex-ante equilibrium, which, with common μ , is the natural extension of Nash equilibrium to allow for smooth ambiguity preferences.

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 a player i 's beliefs 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 update rule was proposed by Hanany and Klibanoff (2009) and is called the *smooth rule*. The smooth rule is defined as follows:

Definition 2.7 *An interim belief system ν satisfies the smooth rule using σ as the ex-ante equilibrium if the following holds for each player i and type τ_i : First, for all $\pi \in \Delta(\Theta)$ such that $\pi(\Theta_{i,\tau_i,\emptyset}) > 0$,*

$$\nu_{i,\tau_i,\emptyset}(\pi) \propto \frac{\phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | \hat{h}^0) \pi(\hat{\theta}) \right)}{\phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | \hat{h}^0) \pi_{\Theta_{i,\tau_i,\emptyset}}(\hat{\theta}) \right)} \pi(\Theta_{i,\tau_i,\emptyset}) \mu_i(\pi);$$

Second, for each partial history h^t such that i views h^t as reachable from h^{t-1} , for all $\pi \in$

$\Delta(\Theta)$ such that $\pi(\Theta_{i,\tau_i,h^t}) > 0$,

$$\nu_{i,\tau_i,h^t}(\pi) \propto \frac{\phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H | \hat{h}^{t-1} = h^{t-1}} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | h^{t-1}) \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \right)}{\phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H | \hat{h}^t = h^t} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | h^t) \pi_{\Theta_{i,\tau_i,h^t}}(\hat{\theta}) \right)} \cdot \left(\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} p_{-i,\sigma,\hat{\theta}}(h^t | h^{t-1}) \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \right) \nu_{i,\tau_i,h^{t-1}}(\pi). \quad (2.8)$$

Note that 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 coincides with standard Bayesian updating. 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 (Hanany and Klibanoff 2009).

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. The proof of this and all subsequent results in the paper may be found in the Appendix.

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

Why is it enough to consider smooth rule updating to identify sequentially optimal profiles? For each player individually, sequential optimality reduces to, essentially, dynamic consistency, and it is this property of smooth rule updating that ensures its sufficiency. 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. A Bayesian version of the theorem is true, however, if we restrict attention to expected utility preferences, for in that case the smooth rule and Bayes' rule agree. Perfect Bayesian Equilibrium (PBE) imposes sequential optimality (defined using only expected utility preferences) and also that beliefs are related via

Bayesian updating wherever possible (plus some auxiliary conditions, which we ignore for now). 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 first shown by Shimoji and Watson (1998) in the context of defining extensive form rationalizability.

The next result shows that, in common with most refinements of Nash equilibrium under ambiguity neutrality, the power of sequential optimality to refine ex-ante equilibrium comes from the presence of off-path actions. In particular, all ex-ante equilibria for which all partial histories before the last stage are on-path can be supported as sequentially optimal.

Theorem 2.2 *Fix a game Γ . Suppose σ is an ex-ante equilibrium, all players view all partial histories h^t with $1 \leq t \leq T$ as reachable from h^0 . Then, there exists an interim belief system ν such that (σ, ν) is sequentially optimal.*

To verify that a pair (σ, ν) where ν obeys smooth rule updating is sequentially optimal it is sufficient to check one-stage deviations. To show this, we formally define what it means to have no profitable one-stage deviations and what we mean by obeying the smooth rule in this context.

Definition 2.8 *The pair (σ, ν) satisfies the no profitable one-stage deviation property if for each player i and each partial history h^t , σ_i is optimal for i given interim belief ν_{i, τ_i, h^t} among all σ'_i satisfying $\sigma'_i(\eta, \tau_i) = \sigma_i(\eta, \tau_i)$ for all $\eta \in \bigcup_{t \in T} H^t$ such that $\eta \neq h^t$.*

For our next result, it is helpful to apply the smooth rule given a strategy profile σ , even when σ is not necessarily an ex-ante equilibrium. We say that an interim belief system ν satisfies *extended smooth rule updating* using σ as the strategy profile if ν satisfies the conditions in Definition 2.7 with respect to σ . With extended smooth rule updating, the absence of profitable one-stage deviations implies sequential optimality:

Theorem 2.3 *If (σ, ν) satisfies the no profitable one-stage deviation property and ν satisfies extended smooth rule updating using σ as the strategy profile 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. Also notice that, fixing the strategies of i 's opponents, at any stage in the game where i has a non-trivial strategy choice, i may identify an optimal strategy for the remainder of the

game using a “folding back” algorithm. It works as follows – for each partial history leading to a final stage where i may move, i calculates an optimal mixture over the actions she has available at that stage given the partial history, the strategies of the other players and i ’s beliefs about types. Then, holding these fixed, i repeats this process for partial histories one stage earlier in the game, and so on. The only thing that (in common with the standard approach under ambiguity neutrality) cannot be calculated via folding back are the beliefs about types at each partial history. These may be determined by updating according to the smooth rule (which reduces to Bayes’ rule under ambiguity neutrality). 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 two sections we explore refinements of sequential optimality, SEA and PEA. We show existence for these refinements, thus implying existence of sequential optima.

2.1 Sequential Equilibrium with Ambiguity

To describe our proposed Sequential Equilibrium with Ambiguity (SEA), we consider an auxiliary condition that imposes requirements on beliefs even at those points where sequential optimality has no implications for updating. These points are those a player i thinks are only reached immediately following deviation(s) of other player(s). Formally, these are partial histories h^t such that $t > 0$ and $m_i(h^t) = t$. The condition is an extension of Kreps and Wilson’s (1982) *consistency* condition that they use in defining Sequential Equilibrium. We modify consistency in order to accommodate ambiguity aversion by replacing Bayes’ rule in their definition with the (extended) smooth rule. Recall that if we simply limited attention to Bayesian updating then sequentially optimal strategies might fail to exist. Also, observe that this is a true extension of Kreps and Wilson’s consistency because Bayes’ rule and the smooth rule coincide under ambiguity neutrality.

Definition 2.9 *Fix a game Γ . A pair (σ^S, ν^S) consisting of a strategy profile and interim belief system satisfies smooth rule consistency if there exists a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$, with $\lim_{k \rightarrow \infty} \sigma^k = \sigma^S$, such that $\nu^S = \lim_{k \rightarrow \infty} \nu^k$, where ν^k is determined by extended smooth rule updating using σ^k as the strategy profile.*

Definition 2.10 *A sequential equilibrium with ambiguity (SEA) of a game Γ is a pair (σ^S, ν^S) consisting of a strategy profile and interim belief system such that (σ^S, ν^S) is sequentially optimal and satisfies smooth rule consistency.*

By definition, any SEA is sequentially optimal. In general, a sequential optimum might not be an SEA. However, if all actions before the last stage of the game are on the equilibrium path, any strategy profile that is part of a sequential optimum is also part of an SEA. Thus the SEA refinement has bite only through restricting off-path beliefs. Formally:

Theorem 2.4 *Fix a game Γ . Suppose (σ, ν) is sequentially optimal and all players view all partial histories h^t with $1 \leq t \leq T$ as reachable from h^0 . Then, there exists an interim belief system $\hat{\nu}$ such that $(\sigma, \hat{\nu})$ is an SEA.*

We show that every game Γ has at least one SEA (and thus also at least one sequential optimum). 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 2.5 *An SEA exists for any game Γ .*

Finally, we show that smooth rule consistency implies extended smooth rule updating in the limit, and use this plus Theorem 2.3 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 equilibria.

Lemma 2.1 *If (σ, ν) satisfies smooth rule consistency, then ν satisfies extended smooth rule updating using σ as the strategy profile.*

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.*

2.2 Perfect Equilibrium with Ambiguity

In this section, we consider a weaker auxiliary condition than smooth rule consistency. Though weaker, it has the advantage of not invoking limits of sequences of strategies and beliefs. This auxiliary condition relates to beliefs exactly at those points where sequential optimality has no implications for updating. The condition requires that, absent considerations related to hedging against ambiguity, if players' types are viewed as independent, there should be no updating of player i 's belief about player j 's type immediately following a partial history at which player j has no choice (i.e., only one action) available. This reflects an idea present in versions of PBE (see e.g., Fudenberg and Tirole 1991b, p. 241) that when players' types are independent, only player j has information to reveal about her own type and so i 's beliefs about player j 's type should not be affected by another player's deviation.

When j has only one action, she has no means to reveal anything, and so, absent reasons related to hedging against ambiguity, player i should not change her marginal on j 's type.

To formalize this in our setting we need to define i 's marginal on j 's type, as well as a condition ensuring that no change in ambiguity hedging concerns occurs in moving from h^{t-1} to h^t .

Definition 2.11 *Given an interim belief system ν , player i 's marginal on player j 's type at partial history h^t is*

$$\sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^t}}(\{\{\theta_j\} \times \Theta_{-j}\}) \nu_{i,\tau_i,h^t}(\pi).$$

Definition 2.12 *Given a strategy profile σ and an interim belief system ν , if player i does not view a partial history h^t with $t \geq 1$ as reachable from h^{t-1} , i has no costly ambiguity exposure under σ at h^{t-1} and h^t if*

$$\phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H | \hat{h}^{t-1} = h^{t-1}} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | h^{t-1}) \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \right)$$

is constant for all π in the support of $\nu_{i,\tau_i,h^{t-1}}$, and

$$\phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H | \hat{h}^t = h^t} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | h^t) \pi_{\Theta_{i,\tau_i,h^t}}(\hat{\theta}) \right)$$

is constant for all π in the support of ν_{i,τ_i,h^t} .

Observe that there are essentially two ways that a player i could have no costly ambiguity exposure under σ at h^{t-1} and h^t – strategies might be such that i is fully hedged against ambiguity from the points of view of the two partial histories (i.e., the conditional expected utility arguments of ϕ'_i in the definition do not vary with π) or, where i is exposed to fluctuations in these conditional expected utilities, ϕ'_i is constant (i.e., i is ambiguity neutral over some range) so the ambiguity exposure is not costly. We can now state our auxiliary condition:

Definition 2.13 *Fix a game Γ . A pair (σ^P, ν^P) consisting of a strategy profile and interim belief system naturally extends updating if, for all players $i, j \neq i$, all types τ_i and all partial histories h^t with $t \geq 1$, if*

- (a) *player i does not view h^t as reachable from h^{t-1} ,*
- (b) *$\sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}} \nu_{i,\tau_i,h^{t-1}}^P(\pi) \in \Delta(\Theta)$ is a product measure,*

(c) player j has no choice (i.e., only one action) available at h^{t-1} and
(d) i has no costly ambiguity exposure under σ at h^{t-1} and h^t ,
then

i 's marginal on player j 's type at partial history h^t must remain the same as it would be at h^{t-1} if the smooth rule using σ^P as the ex-ante equilibrium were used to derive $\nu_{i,\tau_i,h^{t-1}}$ from $\nu_{i,\tau_i,h^{m_i(h^{t-1})}}^P$.

In the case where players are ambiguity neutral, Definition 2.13 is implied by Fudenberg and Tirole's (1991b) PBE requirement that Bayes' rule is used to update beliefs whenever possible (see Fudenberg and Tirole (1991b) condition (1) of Definition 3.1, p.242).

Adding this condition to sequential optimality leads to the following equilibrium definition:

Definition 2.14 *A perfect equilibrium with ambiguity (PEA) of a game Γ is a pair (σ^P, ν^P) consisting of a strategy profile and interim belief system such that (σ^P, ν^P) is sequentially optimal and naturally extends updating.*

Theorem 2.1 showed that interim belief systems using smooth rule updating generate all sequentially optimal strategy profiles. Similarly, they generate all PEA strategy profiles.

Corollary 2.2 *Fix a game Γ . Suppose (σ^P, ν^P) is a PEA of Γ . Then, there exists an interim belief system $\hat{\nu}^P$ satisfying the smooth rule using σ^P as the ex-ante equilibrium such that $(\sigma^P, \hat{\nu}^P)$ is a PEA of Γ .*

The following result tells us that smooth rule consistency implies naturally extending updating, and thus any SEA is also a PEA.

Theorem 2.6 *Fix an SEA, (σ^S, ν^S) , of a game Γ . (σ^S, ν^S) is also a PEA of Γ .*

From this result and the existence of SEA (Theorem 2.5), it immediately follows that a PEA exists for any game Γ . That some PEA may not be SEA, and so SEA can be a strictly stronger concept for some games, follows from the observation that PEA does not restrict i 's beliefs at partial histories immediately following a deviation by another player unless conditions (b), (c), and (d) of naturally extends updating are satisfied. In contrast, smooth rule consistency does place restrictions on beliefs at such partial histories.

2.3 No new behavior without strategic interaction

Before we turn to examples and use of these new concepts, we show that strategic interaction is necessary for any new equilibrium behavior to occur under ambiguity aversion that would never occur under ambiguity neutrality. Specifically, we show that in any game where only one player, j , has non-trivial choice of strategies, any path played with positive probability in an ex-ante equilibrium (the weakest of the concepts we have defined) is also played with positive probability under ambiguity neutrality in an SEA (the strongest concept we have defined) for *some* initial beliefs $\hat{\mu}_j$. This shows that our preference generalization alone (i.e., allowing ambiguity aversion) is not what is generating any new behavior that we find, but rather it is the *combination* of ambiguity aversion with strategic interaction. 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).

Theorem 2.7 *Fix a game Γ such that all players other than a player, j , have singleton action sets following every partial history. If σ is an ex-ante equilibrium of Γ and h is a history such that $\sum_{\pi} \sum_{\theta} p_{\sigma, \theta}(h|h^0) \pi(\theta) \mu_j(\pi) > 0$, then there exists a game $\hat{\Gamma}$, identical to Γ except that (1) all players are ambiguity neutral and (2) the belief of player j ($\hat{\mu}_j$) may differ from that in Γ , such that σ is part of an SEA of $\hat{\Gamma}$ and $\sum_{\pi} \sum_{\theta} p_{\sigma, \theta}(h|h^0) \pi(\theta) \hat{\mu}_j(\pi) > 0$.*

Theorem 2.7 says that with only one player having real choices, no paths of play occur under ambiguity aversion that could not also occur for some beliefs under ambiguity neutrality. The converse – that no paths of play occur under ambiguity neutrality that could not also occur for some beliefs under ambiguity aversion – is also true: let the beliefs under ambiguity aversion, $\hat{\mu}_j$, place probability 1 on the reduced measure ($\sum_{\pi} \pi(\theta) \mu_j(\pi)$). In sum then, when there is only one player with choices, neither the presence nor the absence of ambiguity aversion affects what paths of play may occur under some beliefs. This is true despite the fact that, for fixed beliefs μ , changing ambiguity aversion generally does change the optimal strategy.

3 Examples

In this section, we present a number of examples designed to illustrate different aspects of our equilibrium concepts and compare with standard concepts that limit attention to ambiguity neutral players. All of our examples look at games where there is a common belief μ such




Figure 3.1: SEA Example

that $\mu_i = \mu$ for all players i . Thus the behavior in our examples is never driven by differences in ex-ante beliefs.

3.1 Ambiguity Aversion Generates New Strategic Behavior in Equilibrium

3.1.1 Example 1: Opting in and different hedges

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 3.1.

There are three players: 1,2 and 3. First, nature determines 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 3.1 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).⁴

Proposition 3.1 *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.*

Since σ being part of a sequentially optimal (σ, ν) implies σ is an ex-ante equilibrium, Proposition 3.1 immediately implies that none of the stronger concepts such as PEA, SEA, PBE or Sequential Equilibrium can admit the play of P with positive probability under ambiguity neutrality. The next result shows that the situation changes dramatically under sufficient ambiguity aversion.

Proposition 3.2 *There exist ϕ and μ such that in an SEA both types of player 1 play P with probability 1, and (U, R) is played with probability greater than $\frac{1}{2}$.*

If one modifies the example in Figure 3.1 by reversing the 0's and 2's in player 1's payoffs then, by similar reasoning as in the proof of Proposition 3.2, it should be possible to construct an SEA where Q is played with probability 1 by both types of player 1. There should not be such an SEA under ambiguity neutrality (and thus no such sequential equilibrium), though there would be such a PEA.

3.1.2 Example 2: Hedging plus Screening on beliefs

We now present an example of a 3-player game, with incomplete information about player 3, in which a path of play can occur as part of an SEA when player 2 is sufficiently ambiguity averse, but never occurs as part of a PEA (nor a PBE) when player 2 has expected utility preferences (and is thus ambiguity neutral). Furthermore, under the SEA we construct,

⁴Note that to eliminate any possible effects of varying players' risk aversion, think of the payoffs being generated using lotteries over two "physical" outcomes, the better of which has utility u normalized to 5/2 and the worse of which has u normalized to 0. So, for example, the payoff 1 can be thought of as generated by the lottery giving the better outcome with probability 2/5 and the worse outcome with probability 3/5.




Figure 3.2: 3-Player Game

player 3 achieves a higher expected payoff than under any PEA with player 2 having expected utility preferences. The example is constructed so that if player 2 is sufficiently ambiguity averse then 3 changes his strategy to allow an action by 2 that is favorable to 3. The role of player 1 is to effectively “screen” player 2 and prevent the part of the game that has the play path in question from being reached when 2 puts sufficiently high weight on player 3 being of a particular type (type II). This screening, by design, catches player 2 for a smaller range of parameters when 2 is more ambiguity averse. When 2 is ambiguity neutral, the screening works for a large enough range of parameters that the part of the game in question is reached only when player 2 does not have incentive to carry out the action favorable to player 3, thus 3 opts out of this portion of the game. The game is depicted in Figure 3.2.

There are three players: 1, 2 and 3. First, nature determines whether player 3 is of type I or type II and 3 observes his own type. Players 1 and 2 have only one type, so there is complete information about them. The payoff triples in Figure 3.2 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). Player 2 has ambiguity about player 3’s type and has smooth ambiguity preferences with an associated ϕ_2 and μ_2 . Players 1 and 3 also have smooth ambiguity preferences, but nothing in what follows depends on either ϕ_j or μ_j for $j = 1, 3$. Player 1’s first and only move

in the game is to choose between action T (wo) which gives the move to player 2 and action (th) R (ee), which gives the move to player 3. If T , then 2 makes a single move that ends the game, by choosing between F (ixed) and B (et) (i.e., player 2 effectively chooses between a fixed payoff and betting that player 3 is of type II). If R , then player 3's move is a choice between C (ontinue) which leads to player 2 being given the move, and S (top) which ends the game. If C , then player 2 has a choice between G (amble) and H (edge) after which the game ends.⁵

In any sequential optimum where only type I of player 3 plays C with positive probability, (2.3) requires player 2, following C , to put weight only on type I. Thus, 2 would then always play G if given the move in any such sequential optimum. Notice that if 2 plays G , player 3 is always better off playing S than C . Therefore no sequential optimum can have only type I of player 3 play C with positive probability. Similarly, no sequential optimum can have only type II of player 3 play C with positive probability as 2 would play H and type I would gain from deviating to C . Observe that for any pure strategy sequential optimum, player 3 plays (C, C) if and only if player 2 plays H . Thus $(C, C), H$ is part of a sequential optimum if and only if Player 2 is behaving optimally by playing H (with sufficiently high probability) from the point of view of both stage 1 and stage 2. From (2.3), the only difference in the point of view of these stages can be the beliefs and the event the π are conditioned on. For extreme beliefs such as putting all weight on player 3 being type II with probability 1, H is indeed optimal. For other beliefs, such as, putting all weight on player 3 being type I with probability 1, H is not optimal. Are there beliefs supporting it as part of an SEA?

Our first result shows that C may be played on the equilibrium path as part of an SEA.

Proposition 3.3 *There exist ϕ_2 and μ_2 such that C is played on the equilibrium path with probability 1 as part of an SEA.*

The proof of Proposition 3.3 relies on ambiguity aversion on the part of player 2. Our next result shows that this is essential:

Proposition 3.4 *Regardless of the beliefs of any player, if player 2 is ambiguity neutral (ϕ_2 affine), then no PEA results in C being played on the equilibrium path with positive probability.*

Thus, under ambiguity neutrality no PEA can ever result in play of C , while, when there is enough ambiguity aversion there are SEA involving the play of C with probability 1. Note

⁵Note that to eliminate any possible effects of varying 2's risk attitude, think of the payoffs of player 2 being generated using lotteries over two "physical" outcomes, the better of which has utility u_2 normalized to 6 and the worse of which has u_2 normalized to 0. So, for example, the payoff 2 can be thought of as generated by the lottery giving the better outcome with probability 1/3 and the worse with probability 2/3.

that if we were looking only for profiles satisfying sequential optimality then $(R, F, (C, C), H)$ could be the equilibrium strategies under ambiguity neutrality if 2's initial reduced probability that 3 is of type II were sufficiently high. This is compatible with 2's play of F if given the move by 1 by specifying (off-path) beliefs for 2 following T that place sufficient weight on type I. Such off-path beliefs are unrestricted by sequential optimality, but are not compatible with PEA or SEA.

3.2 Example 3: Limit Pricing under Ambiguity

In this section we present an example of a game with an SEA involving non-trivial updating on the equilibrium path that departs from Bayes rule. The example is based on the Milgrom and Roberts (1982) limit pricing entry model. In this example, an incumbent has private information concerning his production costs. The incumbent chooses a quantity, an entrant observes the quantity (or, equivalently, price) and decides whether or not to enter, in which case he pays a fixed cost $K > 0$. Then the private information is revealed and the last stage of game played, either by both firms competing in a Cournot duopoly or by the incumbent again being a monopolist.⁶ To make this a finite game, suppose there are three possible costs for the incumbent (H, M, L) and a finite set of feasible quantities (including at least the monopoly quantities for each possible production cost and the complete information Cournot quantities).^{7,8} We construct an SEA where in the first period, types M and L pool at the monopoly quantity for L , and type H plays the monopoly quantity for H . Then the entrant, with known cost, enters after observing any quantity strictly below the monopoly quantity for L and does not enter otherwise. If entry occurs, the firms play the complete information Cournot quantities in the second period. If no entry occurs, the incumbent plays its monopoly quantity in the second period. We will see that after observing the monopoly quantity for L (on-path), sequential optimality will require that the entrant's

⁶As stated, it might seem that, due to the revelation of private information following entry, this is not a multi-stage game with observed actions. Observe, however, that the following technique may be used to embed such revelation in any multi-stage game with observed actions: model the revelation of a player's private information as a stage in which that player has actions equal to the set of types (and no other player has any choice of actions) and payoffs are such that playing her realized type leads to the game continuing as described, while playing any other type ends the game with a payoff to the player worse than under any other outcome of the game (and any payoffs to the other players).

In any sequentially optimal strategy profile, the revealing player's strategy will involve playing its realized type with probability one at the revelation stage. Thus, it is safe to ignore the formal revelation stage and simply assume complete information about that player from that point forward in any analysis of sequential optima.

⁷The 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.

⁸The 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.

updating under ambiguity aversion departs from Bayes rule. Sequential optimality also ensures 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).

We assume that the inverse market demand is given by $P(Q) = a - bQ$, $a, b > 0$. Given the incumbent's cost c_I and quantity q_I and entrant cost c_E and quantity q_E , the complete information Cournot reaction functions are given by

$$q_E(q_I) = \arg \max(P(q_E + q_I) - c_E)q_E = \frac{a - c_E}{2b} - \frac{q_I}{2}$$

and

$$q_I(q_E) = \arg \max(P(q_E + q_I) - c_I)q_I = \frac{a - c_I}{2b} - \frac{q_E}{2}.$$

This yields equilibrium values:

$$q_I = \frac{a + c_E - 2c_I}{3b}, q_E = \frac{a + c_I - 2c_E}{3b}$$

and corresponding profits:

$$b\left(\frac{a + c_E - 2c_I}{3b}\right)^2, b\left(\frac{a + c_I - 2c_E}{3b}\right)^2.$$

Similarly, if there is only one firm in the market, with cost c_I and quantity q_I , the monopoly quantity is defined by

$$\arg \max_{q_I} (P(q_I) - c_I)q_I = \frac{a - c_I}{2b}.$$

Thus monopoly profits are

$$b\left(\frac{a - c_I}{2b}\right)^2.$$

For later reference, we collect here conditions on the parameters assumed explicitly or implicitly already plus restrictions equivalent to the monopoly and duopoly quantities above being non-negative:

Assumption 3.1 $a, b > 0$, $K \geq 0$, $c_H > c_M > c_L \geq 0$, $c_E \geq 0$, $a \geq c_H$, $a + c_E - 2c_H \geq 0$ and $a + c_L - 2c_E \geq 0$.

We proceed to check whether the strategies described above satisfy condition (2.6) of sequential optimality. First, we take the incumbent's point of view.

1. Check that type H does not prefer to pool with M,L at the monopoly quantity for L thus

detering entry. Profits for H in the conjectured equilibrium are $b(\frac{a-c_H}{2b})^2 + b(\frac{a+c_E-2c_H}{3b})^2$. Profits if it instead pools with M,L at monopoly quantity for L and deters entry are $\frac{a-c_L}{2b}(a - \frac{a-c_L}{2} - c_H) + b(\frac{a-c_H}{2b})^2$. For H to be better off not pooling, must have

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

This is equivalent to

$$(\frac{a+c_E-2c_H}{3})^2 \geq \frac{a-c_L}{2}(a - \frac{a-c_L}{2} - c_H) \quad (3.1)$$

2. Check that type M does not prefer to produce the monopoly quantity for M and fail to deter entry. Profits for M in the conjectured equilibrium are $\frac{a-c_L}{2b}(a - \frac{a-c_L}{2} - c_M) + b(\frac{a-c_M}{2b})^2$. If it instead produced at the monopoly quantity for M and fails to deter entry, profits are $b(\frac{a-c_M}{2b})^2 + b(\frac{a+c_E-2c_M}{3b})^2$. For M to be better off pooling with L, must have

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

This is equivalent to

$$\frac{a-c_L}{2}(a - \frac{a-c_L}{2} - c_M) \geq (\frac{a+c_E-2c_M}{3})^2. \quad (3.2)$$

3. Type L is playing optimally since its profit maximizing strategy in the absence of a potential entrant also deters entry.

Now for the entry decision of the entrant. We assume that the entrant views each of the three types (L, M, H) as non-null events ex-ante. This means $\sum_{\pi} \mu(\pi)\pi(t) > 0$ for $t \in \{L, M, H\}$. Denote the entrant's Cournot profit net of entry costs when facing an incumbent of type t by $w_t \equiv b(\frac{a+c_t-2c_E}{3b})^2 - K$. As a best-response to the incumbent's strategy, ex-ante the entrant wants to maximize

$$\sum_{\pi} \mu(\pi)\phi[\lambda_L(\pi(L)w_L + \pi(M)w_M) + \lambda_H\pi(H)w_H] \quad (3.3)$$

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, and $K \geq 0$ is the fixed cost of entry. We need it to be the case that this is maximized at $\lambda_H = 1$ and $\lambda_L = 0$. Notice, by monotonicity, some maximum involves $\lambda_H = 1$

if and only if

$$w_H \geq 0 \tag{3.4}$$

and the strict version of this 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 (3.3) with respect to λ_L evaluated at $\lambda_L = 0$ and $\lambda_H = 1$ is non-positive:

$$\sum_{\pi} \mu(\pi)(\pi(L)w_L + \pi(M)w_M)\phi'(\pi(H)w_H) \leq 0. \tag{3.5}$$

Since $\phi' > 0$, a necessary condition for (3.5) is $w_L < 0$ (i.e., entering against a known low cost incumbent is not profitable). To sum up, the equilibrium strategies we described will satisfy condition (2.6) of sequential optimality if and only if the four inequalities (3.1), (3.2), (3.4) and (3.5) are satisfied.

The following theorem provides sufficient conditions for the existence of a SEA of the form described above. One of the conditions is that the entrant is ambiguity averse enough. All else equal, as a player's ϕ becomes more concave, the player becomes more ambiguity averse (see e.g., Klibanoff, Marinacci, Mukerji (2005), Theorem 2). Thus, formally, when we say a player is *ambiguity averse enough* we mean that there exists a $\hat{\phi}$ such that the conclusion of the theorem holds if the player's ϕ is at least as concave as $\hat{\phi}$.

Proposition 3.5 *Suppose Assumption 3.1, (3.1), (3.2), and the strict version of (3.4) hold, and that μ is such that $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M < 0\}) > 0$ or $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M = 0\}) = 1$, and the support of μ can be ordered in the likelihood-ratio ordering. Then, if the entrant is ambiguity averse enough, the limit pricing strategy profile described above is part of an SEA.*

One observation following from the above result is that for any $\mu \in \Delta(\Delta(\{H, M, L\}))$ such that $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M < 0\}) > 0$ and (3.5) is violated when ϕ is linear, there exists a strictly increasing and twice continuously differentiable concave function ϕ such that (3.5) is satisfied. In this way, ambiguity aversion leads to an expansion in the set of μ that can support such a semi-pooling equilibrium. The numerical example below shows that this expansion can be strict. Similarly, increasing ambiguity aversion increases the set of μ that can support such a semi-pooling equilibrium. We also conjecture that in the limit as ambiguity aversion becomes infinite, the set of such μ approaches the set of all μ satisfying $\mu(\{\pi \mid w_L(\pi) < 0\}) > 0$ or $\mu(\{\pi \mid w_L(\pi) = 0\}) = 1$.

An example of parameters that satisfy the four inequalities so that the limit pricing strategies are part of a PEA are the following: μ puts equal weight on type distributions $\pi_0 =$

$(1/6, 1/3, 1/2)$ and $\pi_1 = (1/2, 1/3, 1/6)$, where the vector notation gives the probabilities of L, M, H respectively, $\phi(x) = -e^{-\alpha x}$, with $\alpha > \frac{189}{65} \log(\frac{39}{23}) \approx 1.53546$, $a = 2, b = \frac{7}{128}, c_H = \frac{3}{2}, c_M = \frac{11}{8}, c_L = 1, c_E = \frac{5}{4}$ and $K = 1$. With these parameters, $w_L = -\frac{31}{63}, w_M = \frac{35}{63}$ and $w_H = \frac{65}{63}$. Applying Bayesian updating after observing the monopoly price for L gives $\nu_{E,qL}^B(\pi_0) = \frac{3}{8}$. Applying the smooth rule, the updated beliefs after observing the monopoly price for L are $\nu_{E,qL}(\pi_0) = \frac{3e^{-\alpha(\frac{\nu_H}{2})}}{3e^{-\alpha(\frac{\nu_H}{2})} + 5e^{-\alpha(\frac{\nu_H}{6})}} = \frac{1}{1 + \frac{5}{3}e^{\alpha(\frac{65}{189})}} < \frac{3}{8}$. For example, when $\alpha = 2$, $\nu_{E,qL}(\pi_0) = \frac{1}{1 + \frac{5}{3}e^{\frac{130}{189}}} \approx 0.232$. If the entrant had used Bayesian updating in this example then these limit pricing strategies would not have been sequentially optimal. Specifically, after observing the monopoly quantity for L , the entrant would have wished to deviate by entering.

4 Modeling Strategic Ambiguity

At first glance, since the only source of ambiguity in our framework is ambiguity about players' types, one might think that this is too restrictive to address strategic ambiguity (i.e., ambiguity about the strategies of other players). We show that, in fact, ambiguous strategies may be modeled within our framework. The approach builds on that introduced by Bade (2011) in normal form games who in turn built on Aumann (1974). The basic idea is as follows: in the framework of this paper, players may perceive ambiguity about types. Thus, a type-contingent strategy of a single player may be viewed as ambiguous since the ambiguity about that player's type will translate through to ambiguity about that player's actions. However, since players' payoff functions may also depend on the types, this method of generating strategic ambiguity might in general be confounded with the desire to make actions type contingent due to this payoff dependence. To allow for "pure" strategic ambiguity, we can impose some structure on the type space so that some aspects of a players type are assumed not to affect payoffs, i.e., players' payoff functions are constant with respect to those aspects of the types. In such a game, if a player prefers, in equilibrium, to make his strategy responsive to the realization of such "action" types, the only reason for this can be a desire to affect the ambiguity that other players' perceive about his strategy. One may think of a mixed strategy as choosing a strategy contingent on the outcome of a roulette wheel or other randomizing device. Here, instead of a roulette wheel, there is an "Ellsberg urn" (or, more generally, a payoff irrelevant "natural event") and the player may make his strategy contingent on the draw from the urn(s). Observe that sequential optimality ensures that whenever a player chooses to condition her strategy on these payoff-irrelevant but ambiguous aspects of types, she also necessarily wants to carry out that conditioning

even after the type is realized. Thus our approach to strategic ambiguity is able to satisfy this important implementation issue that was raised in the context of Ellsberg strategies and normal form games by Riedel and Sass (2013).

Formally, consider placing the following structure on the finite set of types for each player i : $\Theta_i \equiv \Theta_i^{\mathcal{U}} \times \Theta_i^{\mathcal{A}}$ (with generic element $\theta_i \equiv (\theta_i^{\mathcal{U}}, \theta_i^{\mathcal{A}})$) where $\Theta_i^{\mathcal{A}} \equiv \prod_{t \in T} \prod_{\eta^t \in H^t} \Theta_{i, \eta^t}^{\mathcal{A}}$ and the utility payoff function for each player i depends only on the $\Theta^{\mathcal{U}} \equiv \prod_{j \in N} \Theta_j^{\mathcal{U}}$ component of the space of type profiles, i.e., for each history $h \in H$, $u_i(h, \theta) = u_i(h, \hat{\theta})$ for all $\theta, \hat{\theta} \in \Theta$ such that $\theta_j^{\mathcal{U}} = \hat{\theta}_j^{\mathcal{U}}$ for all players j .

We call a strategy for a player an Ellsberg strategy if it makes play depend on the (payoff-irrelevant) $\Theta_i^{\mathcal{A}}$ component of i 's type.

Definition 4.1 *A strategy for player i , σ_i , is an Ellsberg strategy if $\sigma_i(\eta^t, \theta_i) \neq \sigma_i(\eta^t, \hat{\theta}_i)$ for some $\theta_i, \hat{\theta}_i \in \Theta_i$ such that $\theta_i^{\mathcal{U}} = \hat{\theta}_i^{\mathcal{U}}$, some partial history $\eta^t \in H^t$, and some $t \in \{1, \dots, T + 1\}$.*

Notice that if μ makes $\Theta_i^{\mathcal{A}}$ ambiguous given some $\theta_i^{\mathcal{U}} \in \Theta_i^{\mathcal{U}}$ then an Ellsberg strategy allows player i to create ambiguity about his strategy even fixing the payoff relevant component of the type. Because ignoring $\theta_i^{\mathcal{A}}$ is always an option, if, in a sequential optimum/PEA/SEA, a player uses such an Ellsberg strategy it must be the case that she views choosing to create this strategic ambiguity (and follow through on it) as a best response. This is a key difference with the older literature on complete information games with ambiguity about others' strategies (e.g., Dow and Werlang 1994, Lo 1996, 1999, Klibanoff 1996, Eichberger and Kelsey 2000, Marinacci 2000, Mukerji and Shin 2002). In that literature, while each player is assumed to best respond to the ambiguity she has about the others' strategies, the others' strategies in the support of that ambiguity are not all required to be part of others' best responses. A notable exception is Lo (1996, 1999), which does require this best response property. Even in Lo (1996, 1999) however, there is no choice on the part of a player to create (or not) ambiguity about her play, as the ambiguity there is based on the set of a player's best responses and not just on the best response chosen by the player.

Mouraviev, Riedel and Sass (2015) define Ellsberg behavior strategies and Ellsberg mixed strategies and show that in games of perfect recall the two are not generally equivalent, violating Kuhn's theorem. Their Ellsberg behavior strategies allow only, for each node, conditioning of the mixture over actions at that node on an ambiguous urn for *that node*. This restricts the ability to connect the mixture used across nodes and generates their non-equivalence. Since our notion of Ellsberg (behavior) strategies allows conditioning on the

overall type space, the same issue does not arise for us.⁹

4.1 Example 4: Strategic Ambiguity

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 involving Ellsberg strategies in our framework. di Tillio, Kos and Messner (2012) 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 (2015) examine a principal-multi-agent auditing game and show that the principal may benefit from an ambiguous auditing strategy. All these examples have ambiguity about payoff-relevant actions. Most are also “static” in the sense that they may be analyzed using only an ex-ante level of optimality.

We present here a novel example where strategic ambiguity about actions *without* payoff consequences (“cheap talk”) proves valuable in equilibrium. The most closely related examples in the literature are analyzed in Bose and Renou (2014) and Kellner and Le Quement (2015). Aside from the specifics of the game, the most important difference from previous analyses is in the dynamics – specifically, demonstrating that strategic ambiguity occurs as part of a sequential optimum. In particular, Kellner and Le Quement (2015) do not show that ex-ante equilibrium obtains, while the strategic ambiguity in Bose and Renou (2014) may fail to be sequentially optimal.

The example is a game with three players, a principal, P , and two agents, a_1 and a_2 . The principal’s type has two components, a payoff-relevant component, which takes the value I or II , and a payoff-irrelevant component, which takes the value B or R . Thus the possible types for the principal are given by the set $\{IB, IR, IIB, IIR\}$. The principal is privately informed of his type. The agents have no private information. After learning his type, the principal publicly sends a message $m \in \{m_1, m_2\}$ that is seen by all players. This message is cheap talk in that it does not have any direct effect on payoffs. After seeing the message, the

⁹In Aryal and Stauber (2014), it is not Kuhn’s theorem itself, but optimality results that follow from it in the standard framework that are shown may fail under ambiguity aversion due to dynamic inconsistencies. The failures they identify cannot occur for any sequentially optimal strategies.

agents play a simultaneous-move game in which each agent chooses between the actions g and h . Payoffs for the three players contingent on the payoff-relevant part of the principal's type and the agents' actions are given in the following matrices:

I	g_2	h_2	II	g_2	h_2
g_1	0, 0, 5	0, 0, 1	g_1	0, 5, 0	0, 5, 1
h_1	2, 1, 5	2, 2, 2	h_1	0, 1, 0	2, 2, 2

Observe that the principal would always like $a1$ to play h and, in event I is indifferent to the play of $a2$, while in event II also wants $a2$ to play h . Each agent is indifferent to the other agent's play when herself playing g , and strictly prefers the other agent to play h when herself playing h . If informed of the state, $a1$ prefers to play h if I and g if II , while $a2$ prefers to play g if I and h if II . Suppose that $\phi_{a1}(x) = -e^{-11x}$. The specification of ϕ_P and ϕ_{a2} is not important for what follows.

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

	IB	IR	IIB	IIR
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 the principal's type and, fixing that component, ambiguity about the payoff-irrelevant component of the principal's type. This belief structure is, for example, consistent with there being an underlying parameter $\gamma \in \{\gamma_1, \gamma_2\}$ (generating π_1 vs. π_2) about which there is ambiguity, and both the payoff-relevant, $\{I, II\}$, and payoff-irrelevant, $\{R, B\}$, parts of the principal's type are determined as conditionally independent stochastic functions of γ with $\text{Prob}(I|\gamma_1) = 3/4$, $\text{Prob}(I|\gamma_2) = 1/5$, $\text{Prob}(B|\gamma_1) = 1/3$, $\text{Prob}(B|\gamma_2) = 1/4$. For instance, γ might be some scientific principle that is not well understood, and it influences both the functioning of a technology relevant for the task-at-hand (I vs. II) and the findings of a laboratory experiment (B vs. R) not affecting the task-at-hand.

Consider the following strategy for the principal: if his type is IB send message m_1 , otherwise send message m_2 . Observe that this strategy makes use of the payoff-irrelevant component of the principal's type, and is thus an Ellsberg strategy. We will show that this strategy is an equilibrium strategy for the principal (Proposition 4.1), and, that the principal does strictly better than if he were restricted to play a non-Ellsberg strategy (Proposition 4.2). This establishes that the strategic ambiguity is strictly valuable for the principal.

Proposition 4.1 *The following strategies are part of an SEA: P sends message m_1 if his type is IB and sends m_2 otherwise. a_1 plays h after both messages. a_2 plays g after m_1 and h after m_2 . In such an SEA, the principal attains his maximum possible payoff for each type.*

Remark 4.1 *The above strategies remain an SEA for any ϕ_P, ϕ_{a_2} , and, by Lemma A.5, for any ϕ_{a_1} more concave than the one stated above.*

Proposition 4.2 *If the principal were restricted to play a non-Ellsberg strategy, there would be no ex-ante equilibrium yielding the principal the maximum possible payoff for each type.*

One lesson from the proof of Proposition 4.2 is that, fixing I or II , ambiguity about B vs. R is *necessary* for the principal to do better by playing an Ellsberg strategy. Specifically, if π_1 and π_2 were modified so that the uncertainty about the payoff-relevant component of type remained the same as above, but, given the payoff-relevant component of the type, there were no ambiguity about the payoff-irrelevant component (i.e., $\pi_1(IB) + \pi_1(IR) = \frac{3}{4}$, $\pi_2(IB) + \pi_2(IR) = \frac{1}{5}$, $\pi_1(IB)/\pi_1(IR) = \pi_2(IB)/\pi_2(IR)$ and $\pi_1(IIB)/\pi_1(IIR) = \pi_2(IIB)/\pi_2(IIR)$) then any Ellsberg strategy could be replaced by a non-Ellsberg strategy using appropriate mixtures conditional on the payoff-relevant component of the type without changing the best responses of the agents.

Remark 4.2 *If a_1 becomes sufficiently more ambiguity averse, Proposition 4.2 no longer holds: in addition to being able to attain the maximum by using an Ellsberg strategy, there will be an equilibrium in non-Ellsberg strategies that allows the maximum possible payoff for each type of the principal. Intuitively, with enough ambiguity aversion, the additional ambiguity generated by the Ellsberg strategy is no longer needed to induce the principal's desired behavior.*

5 Extensions and Other Approaches

5.1 Other Approaches in the literature

In order to compare with some of the existing literature investigating dynamic games with ambiguity, the following condition, describing a consistent planning requirement in the spirit of Strotz (1955-56), is useful.

Definition 5.1 Fix a game Γ and a pair (σ^P, ν^P) consisting of a strategy profile and interim belief system. Specify V_i and V_{i,τ_i,h^t} as in (2.1) and (2.3). For each player i , type τ_i and history h , let

$$CP_{i,\tau_i,h^{T+1}} \equiv \Sigma_i.$$

Then, inductively, for $0 \leq t \leq T$, let

$$CP_{i,\tau_i,h^t} \equiv \underset{\hat{\sigma}_i \in \bigcap_{\hat{h} \in H | \hat{h}^t = h^t} CP_{i,\tau_i,\hat{h}^{t+1}}}{\operatorname{argmax}} V_{i,\tau_i,h^t}((\hat{\sigma}_i, \sigma_{-i}^P); \nu^P).$$

Finally, let

$$CP_i \equiv \underset{\hat{\sigma}_i \in \bigcap_{\hat{\tau}_i \in \Theta_i} CP_{i,\hat{\tau}_i,\emptyset}}{\operatorname{argmax}} V_i(\hat{\sigma}_i, \sigma_{-i}^P).$$

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

$$\sigma_i^P \in CP_i.$$

Equivalently, (σ^P, ν^P) is such that for all players i , all types τ_i and all partial histories h^t ,

$$V_i(\sigma^P) \geq V_i(\hat{\sigma}_i, \sigma_{-i}^P) \text{ for all } \hat{\sigma}_i \in \bigcap_{\hat{\tau}_i \in \Theta_i} CP_{i,\hat{\tau}_i,\emptyset}$$

and

$$V_{i,\tau_i,h^t}(\sigma^P; \nu^P) \geq V_{i,\tau_i,h^t}((\hat{\sigma}_i, \sigma_{-i}^P); \nu^P) \text{ for all } \hat{\sigma}_i \in \bigcap_{\hat{h} \in H | \hat{h}^t = h^t} CP_{i,\tau_i,\hat{h}^{t+1}}.$$

If (σ^P, ν^P) is sequentially optimal then it is also optimal under consistent planning. However, if (σ^P, ν^P) 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.

Sequential optimality is a feature of PEA and SEA. The above shows that sequential optimality is not generally satisfied by the consistent planning approach taken in much existing literature using extensive-form games of incomplete information with ambiguity. Sequential optimality is a relatively uncontroversial part of the main equilibrium concepts for extensive-form games with incomplete information under ambiguity neutrality, such as Perfect Bayesian Equilibrium and Sequential Equilibrium. Thus, it is both important and




Figure 5.1: Violation of sequential optimality

interesting to explore sequential optimality in the context of games with ambiguity.

Furthermore, when updating is according to the smooth rule, (σ^P, ν^P) optimal under consistent planning implies (σ^P, ν^P) is sequentially optimal, making the two equivalent under smooth rule updating. This follows from Theorem 2.3 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.

Under ambiguity aversion without smooth rule updating, when sequential optimality is a stronger requirement than just consistent planning or no profitable one-stage deviations, what kinds of behavior does sequential optimality rule out? Consider the following example (Figure 5.1), which shows how consistent planning or no profitable one-stage deviations allows strategy profiles that are not even ex-ante (Nash) equilibria of a game (and thus clearly not sequentially optimal).

To analyze the game, let us consider player 2. Observe that each type of 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, type-by-type, strictly higher than the payoff to playing o followed by anything. Thus no strategy involving o can be a best reply to 2's optimal strategy no matter what player 1's ambiguity attitude or beliefs about 2's type. This immediately implies that o is not part of any ex-ante equilibrium, let alone a sequentially optimal strategy profile.

In contrast, 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. 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.

Battigalli et al. (2015b) is another paper exploring dynamic games with smooth ambiguity preferences (building on Battigalli et al. (2015a), which analyzed games in strategic form and so took a purely ex-ante perspective). A key difference from our approach is that instead of sequential optimality, they require no profitable one-stage deviations plus Bayesian updating. Thus, while both approaches satisfy no profitable one-stage deviations, the equilibria described by their approach may fail both sequential optimality and optimality under consistent planning. Additionally, because restricting attention to interim belief systems satisfying Bayesian updating generally rules out some (or all) sequentially optimal strategies, equilibria we identify might fail to be equilibria according to their approach. A further difference is that they focus on a form of self-confirming equilibria while we concentrate on a form of sequential equilibria.

5.2 Extensions

5.2.1 Maxmin Expected Utility

We have assumed players have smooth ambiguity preferences (Klibanoff, Marinacci and Mukerji, 2005). This facilitated our analysis by allowing ambiguity aversion (via ϕ) and beliefs (via μ) to be separately and conveniently specified. Can our approach be applied to players with Maxmin expected utility (Gilboa and Schmeidler, 1989) preferences? We suggest one way to do so. If the set of probability measures in the Maxmin EU representation is taken to be the (convex hull of) the support of μ , then these preferences can be interpreted as a model of an infinitely ambiguity averse player with beliefs given by the support of μ . By modifying our framework to specify this set rather than μ , eliminate the specification of ϕ ,

and use the Maxmin EU functional rather than the smooth ambiguity functional to evaluate strategies, the key definitions of ex-ante equilibrium, interim belief system, and Sequential Optimality can all be naturally adapted. We conjecture the following version of Theorem 2.1 would be true: for the purposes of identifying such sequentially optimal strategies, it is without loss of generality to limit attention to interim belief systems derived according to any one of the dynamically consistent update rules described in Hanany and Klibanoff (2007) for Maxmin EU preferences. With this in hand, one could then explore analogues of the rest of our analysis and see to what extent they remain true with infinite ambiguity aversion and if any new phenomena arise.

5.2.2 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 2.6 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 2.6 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] G. ARYAL, R. STAUBER, A note on Kuhn's Theorem with ambiguity averse players, *Economics Letters* **125** (2014), 110–114.
- [2] R. AUMANN, Subjectivity and correlation in randomized strategies, *Journal of Mathematical Economics* **1** (1974), 67–96.
- [3] S. AUSTER, Adverse Selection Under Ambiguity, mimeo. (2013).

- [4] M. AYOUNI, F. KOESSLER, Ambiguous Audit, mimeo. (2015)
- [5] Y. AZRIELI, R. TEPER, Uncertainty aversion and equilibrium existence in games with incomplete information, *Games and Economic Behavior* **73(2)** (2011), 310-317.
- [6] S. BADE, Ambiguous act equilibria, *Games and Economic Behavior* **71(2)** (2011), 246-260.
- [7] S. BALIGA, E. HANANY, P. KLIBANOFF, Polarization and Ambiguity, *American Economic Review* **103(7)** (2013), 3071-3083.
- [8] P. BATTIGALLI, S. CERREIA VIOLIO, F. MACCHERONI, M. MARINACCI, Self-confirming Equilibrium and Model Uncertainty, *American Economic Review* **105(2)** (2015a), 646-77.
- [9] P. BATTIGALLI, S. CERREIA VIOLIO, F. MACCHERONI, M. MARINACCI, Selfconfirming Equilibrium and Sequential Rationality with Ambiguity Averse Agents, mimeo. (2015b).
- [10] A.L. BODOH-CREED, Ambiguous beliefs and mechanism design, *Games and Economic Behavior* **75(2)** (2012), 518-537.
- [11] S. BOSE, A. DARIPA, A dynamic mechanism and surplus extraction under ambiguity, *Journal of Economic Theory* **144(5)** (2009), 2084-2114.
- [12] S. BOSE, E. OZDENOREN, A. PAPE, Optimal Auctions with Ambiguity, *Theoretical Economics* **1(4)** (2006), 411-438.
- [13] S. BOSE, L. RENO, Mechanism Design With Ambiguous Communication Devices, *Econometrica* **82** (2014), 1853-1872.
- [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, mimeo. Virginia Tech (2015).
- [16] J. DOW, S.R.C WERLANG, Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction, *Journal of Economic Theory*, **64(2)** (1994), 305-324.
- [17] 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.

- [18] J. EICHBERGER, D. KELSEY, Non-Additive Beliefs and Strategic Equilibria, *Games and Economic Behavior* **30(2)** (2000), 183–215.
- [19] J. EICHBERGER, D. KELSEY, Sequential Two-Player Games with Ambiguity, *International Economic Review* 45 (2004), 1229-1261.
- [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] 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.
- [24] I. GILBOA, D. SCHMEIDLER, Maxmin Expected Utility with a Non-Unique Prior, *Journal of Mathematical Economics*, **18** (1989), 141-153.
- [25] 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.
- [26] J. GREENBERG, The Right to Remain Silent, *Theory and Decision*, **48** (2000), 193–204.
- [27] E. HANANY, P. KLIBANOFF, Updating Preferences with Multiple Priors, *Theoretical Economics* **2** (2007), 261-298.
- [28] E. HANANY, P. KLIBANOFF, Updating Ambiguity Averse Preferences, *The BE Journal of Theoretical Economics (advances)* **9(1)** (2009), article 37.
- [29] A. KAJII, T. UI, Incomplete Information Games with Multiple Priors, *Japanese Economic Review* **56(3)** (2005), 332-351.
- [30] C. KELLNER, Tournaments as a Response to Ambiguity Aversion in Incentive Contracts, *Journal of Economic Theory*, **159(A)** (2015), 627-655.
- [31] C. KELLNER, M. LE QUEMENT, Modes of ambiguous communication, mimeo. University of Bonn (2013).

- [32] C. KELLNER, M. LE QUEMENT, Endogenous ambiguity in cheap talk, mimeo. University of Bonn (2015).
- [33] P. KLIBANOFF, Uncertainty, decision, and normal form games, mimeo. Northwestern University (1996).
- [34] P. KLIBANOFF, M. MARINACCI, S. MUKERJI, A Smooth Model of Decision Making under Ambiguity, *Econometrica* **73** (2005), 1849-1892.
- [35] D.M. KREPS, R. WILSON, Sequential Equilibria, *Econometrica* **50(4)** (1982), 863-894.
- [36] C. KUZMICS, Abraham Wald's complete class theorem and Knightian uncertainty, mimeo. Bielefeld University (2015).
- [37] K.C. LO, Equilibrium in beliefs under uncertainty, *Journal of Economic Theory* **71(2)** (1996), 443-484.
- [38] K.C. LO, Extensive Form Games with Uncertainty Averse Players, *Games and Economic Behavior* **28** (1999), 256-270.
- [39] G. LOPOMO, L. RIGOTTI, C. SHANNON, Uncertainty in Mechanism Design, mimeo. (2010).
- [40] M. MARINACCI, Ambiguous Games, *Games and Economic Behavior* **31(2)** (2000), 191-219.
- [41] P. MILGROM, J. ROBERTS, Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis, *Econometrica* **50(2)** (1982), 443-459.
- [42] I. MOURAVIEV, F. RIEDEL, L. SASS, Kuhn's Theorem for Extensive Form Ellsberg Games, mimeo. (2015).
- [43] S. MUKERJI, H.S. SHIN, Equilibrium departures from common knowledge in games with non-additive expected utility, *Advances in Theoretical Economics* **2(1)** (2002).
- [44] 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.
- [45] R.B. MYERSON, *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA (1991).

- [46] E. OZDENOREN, D. LEVIN, Auctions with Uncertain Numbers of Bidders, *Journal of Economic Theory*, **118(2)**, (2004) 229-251.
- [47] F. RIEDEL, L. SASS, Ellsberg games, *Theory and Decision* **76** (2013), 469-509.
- [48] A. SALO, M. WEBER, Ambiguity Aversion in First Price Sealed Bid Auctions, *Journal of Risk and Uncertainty* **11** (1995), 123-137.
- [49] M. SHIMOJI, J. WATSON, Conditional Dominance, Rationalizability, and Game Forms, *Journal of Economic Theory* **83** (1998), 161-195.
- [50] R.H. STROTZ, Myopia and Inconsistency in Dynamic Utility Maximization, *The Review of Economic Studies* **23(3)** (1955-56), 165-180.
- [51] A. DI TILLO, N. KOS, M. MESSNER, The Design of Ambiguous Mechanisms, mimeo. Bocconi University (2012).
- [52] A. WOLITZKY, Bilateral Trading with Maxmin Agents, mimeo. Stanford University (2013).
- [53] A. WOLITZKY, Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade, mimeo. Stanford University (2014).

A Appendix: Proofs

We begin with a key lemma on the preservation of optimality under smooth rule updating:

Lemma A.1 *Fix a game Γ and (σ, ν) such that σ is an ex-ante equilibrium. For $t \neq 0$ and i, τ_i, h^t such that $m_i(h^t) \neq t$, or for $t = 0$, if ν_{i, τ_i, h^t} is derived from $\nu_{i, \tau_i, h^{m_i(h^t)}}$ (or, if $t = 0$, from μ_i) via the smooth rule using σ as the ex-ante equilibrium and σ_i is optimal for player i given σ_{-i} and given $\tau_i, h^{m_i(h^t)}$ (or, if $t = 0$, given ex-ante optimality), then σ_i is optimal for player i given σ_{-i}, τ_i and h^t : for all $\sigma'_i \in \Sigma_i$,*

$$V_{i, \tau_i, h^t}(\sigma; \nu) \geq V_{i, \tau_i, h^t}((\sigma'_i, \sigma_{-i}); \nu).$$

It is useful for the proof of this result (as well as that of Theorem 2.3) to refer to a player's "local ambiguity neutral measure" after some partial history in the game. Given (σ, ν) , for any player i and type τ_i , let $q^{\sigma, i}(h, \theta)$ denote i 's ex-ante σ -local measure, defined for each

$\theta \in \Theta$ and $h \in H$ by

$$q^{\sigma,i}(h, \theta) \equiv \sum_{\pi \in \Delta(\Theta)} \phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | \hat{h}^0) \pi(\hat{\theta}) \right) \cdot p_{-i, \sigma, \theta}(h | h^0) \pi(\theta) \mu_i(\pi); \quad (\text{A.1})$$

additionally, for any partial history $\eta \in \bigcup_{t \in \mathcal{T}} H^t$, let $q^{(\sigma, \nu), i, \tau_i, \eta}$ denote i 's (σ, ν) - local measure given τ_i and η , defined for each $\theta \in \Theta$ and $h \in H$ with $\theta_i = \tau_i$ and $h^t = \eta$ and by

$$q^{(\sigma, \nu), i, \tau_i, \eta}(h, \theta) \equiv \sum_{\pi \in \Delta(\Theta) | \pi(\Theta_{i, \tau_i, \eta}) > 0} \phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H | \hat{h}^t = \eta} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | \eta) \pi_{\Theta_{i, \tau_i, \eta}}(\hat{\theta}) \right) \cdot p_{-i, \sigma, \theta}(h | \eta) \pi_{\Theta_{i, \tau_i, \eta}}(\theta) \nu_{i, \tau_i, \eta}(\pi). \quad (\text{A.2})$$

Proof of Lemma A.1. Consider first the case where $t \neq 0$ and $m_i(h^t) \neq t$. By assumption, σ_i is optimal given $\tau_i, h^{m_i(h^t)}$ and σ_{-i} . This is equivalent (see Hanany and Klibanoff 2009, Lemma A.1) to the condition that σ_i solves

$$\max_{\sigma'_i \in \Sigma_i} \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^{m_i(h^t)} = h^{m_i(h^t)}} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | h^{m_i(h^t)}) q^{(\sigma, \nu), i, \tau_i, h^{m_i(h^t)}}(\bar{h}, \bar{\theta}), \quad (\text{A.3})$$

where $q^{(\sigma, \nu), i, \tau_i, h^{m_i(h^t)}}$ is i 's (σ, ν) - local measure given τ_i and $h^{m_i(h^t)}$ (defined in (A.2)). Notice that the objective function in (A.3) can be equivalently written as

$$\begin{aligned} & \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\substack{\bar{h} \in H | \bar{h}^{m_i(h^t)} = h^{m_i(h^t)} \\ \text{and } \bar{h}^t \neq h^t}} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | h^{m_i(h^t)}) q^{(\sigma, \nu), i, \tau_i, h^{m_i(h^t)}}(\bar{h}, \bar{\theta}) \\ & + p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(h^t | h^{m_i(h^t)}) \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | h^t) q^{(\sigma, \nu), i, \tau_i, h^{m_i(h^t)}}(\bar{h}, \bar{\theta}). \end{aligned}$$

The advantage of doing so is making clear that only the term

$$\sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | h^t) q^{(\sigma, \nu), i, \tau_i, h^{m_i(h^t)}}(\bar{h}, \bar{\theta}) \quad (\text{A.4})$$

is affected by $\sigma'_i(\eta, \theta_i)$ for partial histories η that include h^t and only those same components of σ'_i affect (A.4). Therefore (A.3) implies that σ_i maximizes (A.4).

We want to show that σ_i is optimal given τ_i, h^t and σ_{-i} . This is equivalent to the condition that σ_i solves

$$\max_{\sigma'_i \in \Sigma_i} \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | h^t) q^{(\sigma, \nu), i, \tau_i, h^t}(\bar{h}, \bar{\theta}).$$

where $q^{(\sigma, \nu), i, \tau_i, h^t}$ is the gradient of the interim indifference curve given τ_i and h^t at σ , which, for each $\bar{\theta} \in \Theta$ and $\bar{h} \in H$ with $\bar{\theta}_i = \tau_i$ and $\bar{h}^t = h^t$, is given by

$$q^{(\sigma, \nu), i, \tau_i, h^t}(\bar{h}, \bar{\theta}) \equiv \sum_{\pi \in \Delta(\Theta) | \pi_{\Theta_{i, \tau_i, h^t}} > 0} \phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H | \hat{h}^t = h^t} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | h^t) \pi_{\Theta_{i, \tau_i, h^t}}(\hat{\theta}) \right) \cdot p_{-i, \sigma, \bar{\theta}}(\bar{h} | h^t) \pi_{\Theta_{i, \tau_i, h^t}}(\bar{\theta}) \nu_{i, \tau_i, h^t}(\pi). \quad \text{A.5}$$

Thus it is sufficient to show that $q^{(\sigma, \nu), i, \tau_i, h^t}(\bar{h}, \bar{\theta}) \propto q^{(\sigma, \nu), i, \tau_i, h^{m_i(h^t)}}(\bar{h}, \bar{\theta})$ for each $\bar{\theta} \in \Theta$ and $\bar{h} \in H$ with $\bar{\theta}_i = \tau_i$ and $\bar{h}^t = h^t$. This follows by using the local measure definition (A.2), applying the smooth rule to substitute for $\nu_{i, \tau_i, h^t}(\pi)$ for all $\pi \in \Delta(\Theta)$ such that $\pi_{\Theta_{i, \tau_i, h^t}} > 0$ (as $\nu_{i, \tau_i, h^t}(\pi) = 0$ for other π) and then using the definitions of $\pi_{\Theta_{i, \tau_i, h^t}}(\theta)$ and $\pi_{\Theta_{i, \tau_i, h^{m_i(h^t)}}$ and cancelling terms.

The case where $t = 0$ is similar. ■

Proof of Theorem 2.1. We show that $(\sigma^P, \hat{\nu}^P)$, where, for all $i, \tau_i, \hat{\nu}_{i, \tau_i, h^t}^P = \nu_{i, \tau_i, h^t}^P$ whenever $(t > 0$ and $m_i(h^t) = t)$, and where, everywhere else, $\hat{\nu}_{i, \tau_i, h^t}^P$ is derived via the smooth rule, is sequentially optimal. First, observe that $\hat{\nu}^P$ does not enter into the function V_i , so the fact that (σ^P, ν^P) is sequentially optimal directly implies that $V_i(\sigma^P) \geq V_i(\sigma'_i, \sigma_{-i}^P)$ for all $\sigma'_i \in \Sigma_i$. Second, by construction, $\hat{\nu}^P$ satisfies the smooth rule using σ^P as the ex-ante equilibrium except, possibly, for i, τ_i, h^t where $(t > 0$ and $m_i(h^t) = t)$. However, from the definition of the smooth rule (Definition 2.7), observe that it is exactly for i, τ_i, h^t where $(t > 0$ and $m_i(h^t) = t)$ for which the smooth rule allows any interim beliefs. Thus $\hat{\nu}^P$ satisfies the smooth rule using σ^P as the ex-ante equilibrium. Finally, to see that $(\sigma^P, \hat{\nu}^P)$ satisfies $V_{i, \tau_i, h^t}(\sigma^P; \hat{\nu}^P) \geq V_{i, \tau_i, h^t}(\sigma'_i, \sigma_{-i}^P; \hat{\nu}^P)$ for all $\sigma'_i \in \Sigma_i$, observe that (a) for i, τ_i, h^t such that $(h^t \neq \emptyset$ and $m_i(h^t) = t)$, it directly inherits this from (σ^P, ν^P) and (b) everywhere else, Lemma A.1 shows that smooth rule updating ensures the required optimality. ■

Proof of Theorem 2.2. 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 as the ex-ante equilibrium. Since all players view all partial histories h^t with $1 \leq t \leq T$ as reachable from h^0 , ν is pinned down completely except possibly after partial histories h^{T+1} for which $m_i(h^{T+1}) < T + 1$ for at least one player i (i.e., except after histories involving off-path ac-

tion(s) in the final stage). By Lemma A.1, (2.7) in the definition of sequential optimality is satisfied for all i, τ_i, h^t with $t \leq T$. Note however, using the definition of $V_{i, \tau_i, h^{T+1}}$, optimality after the last stage is trivial because there are no actions remaining, and so how ν is specified after off-path play at $T + 1$ is irrelevant for sequential optimality. ■

Proof of Theorem 2.3. Suppose that (σ, ν) satisfies the no profitable one-stage deviation property and ν satisfies extended smooth rule updating using σ as the strategy profile. First, for each player i , the no profitable one-stage deviation property implies conditional optimality of σ_i according to ν_{i, τ_i, η^T} for all $\eta^T \in H^T$. 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 conditionally optimal according to ν_{i, τ_i, η^t} for all $\eta^t \in H^t$. We claim that, for each player i , σ_i is conditionally optimal according to $\nu_{i, \tau_i, \eta^{t-1}}$ for all $\eta^{t-1} \in H^{t-1}$. The argument for this is as follows: Fix $\eta^{t-1} \in H^{t-1}$ and fix a player i . Consider any strategy σ'_i for player i . For any $h \in H$ such that $h^{t-1} = \eta^{t-1}$ and i views h^t as reachable from η^{t-1} , the conditional optimality of σ_i according to ν_{i, τ_i, h^t} implies (see Hanany and Klibanoff 2009, Lemma A.1)

$$\begin{aligned} & \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, \sigma, \bar{\theta}}(\bar{h} | \bar{h}^t) q^{(\sigma, \nu), i, \tau_i, \bar{h}^t}(\bar{h}, \bar{\theta}) \\ & \geq \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | \bar{h}^t) q^{(\sigma, \nu), i, \tau_i, \bar{h}^t}(\bar{h}, \bar{\theta}). \end{aligned} \quad (\text{A.6})$$

Since i 's preferences satisfy extended smooth rule updating using σ , for all such h , $q^{(\sigma, \nu), i, \tau_i, \bar{h}^t}(\bar{h}, \bar{\theta}) \propto q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta})$ for all $\bar{\theta} \in \Theta$ and $\bar{h} \in H$ with $\bar{\theta}_i = \tau_i$ and $\bar{h}^t = h^t$. Thus, after substituting for $q^{(\sigma, \nu), i, \tau_i, \bar{h}^t}(\bar{h}, \bar{\theta})$, cancelling the constant of proportionality and multiplying by $p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h}^t | \eta^{t-1})$ (which is identical for all $\bar{\theta}$ with $\bar{\theta}_i = \tau_i$, as it depends only on τ_i , and for all \bar{h} with $\bar{h}^t = h^t$), (A.6) becomes

$$\begin{aligned} & \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h}^t | \eta^{t-1}) p_{i, \sigma, \bar{\theta}}(\bar{h} | \bar{h}^t) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}) \\ & \geq \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h}^t | \eta^{t-1}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | \bar{h}^t) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}) \\ & = \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | \eta^{t-1}) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}). \end{aligned}$$

Summing for all such h , yields:

$$\begin{aligned}
& \sum_{\substack{h \in H | h^{t-1} = \eta^{t-1} \\ \text{and } m_i(h^t) \leq t-1}} \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h}^t | \eta^{t-1}) p_{i, \sigma, \bar{\theta}}(\bar{h} | \bar{h}^t) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}) \\
&= \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\substack{\bar{h} \in H | \bar{h}^{t-1} = \eta^{t-1} \\ \text{and } m_i(\bar{h}^t) \leq t-1}} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h}^t | \eta^{t-1}) p_{i, \sigma, \bar{\theta}}(\bar{h} | \bar{h}^t) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}) \\
&= \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^{t-1} = \eta^{t-1}} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h}^t | \eta^{t-1}) p_{i, \sigma, \bar{\theta}}(\bar{h} | \bar{h}^t) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}) \quad (\text{A.7}) \\
&\geq \sum_{\substack{h \in H | h^{t-1} = \eta^{t-1} \\ \text{and } m_i(h^t) \leq t-1}} \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^t = h^t} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | \eta^{t-1}) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}) \\
&= \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^{t-1} = \eta^{t-1}} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | \eta^{t-1}) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}).
\end{aligned}$$

where the second and third equalities follow since, for all $\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i$ and all $\bar{h} \in H | \bar{h}^{t-1} = \eta^{t-1}$ with $m_i(\bar{h}^t) > t - 1$, $q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}) = 0$ because $p_{-i, \sigma, \bar{\theta}}(\bar{h} | \eta^{t-1}) = 0$.

Since (σ, ν) satisfies the no profitable one-stage deviation property, the conditional optimality of σ_i according to $\nu_{i, \tau_i, \eta^{t-1}}$ among all strategies deviating only at $t - 1$ implies

$$\begin{aligned}
& \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^{t-1} = \eta^{t-1}} u_i(\bar{h}, \bar{\theta}) p_{i, \sigma, \bar{\theta}}(\bar{h} | \eta^{t-1}) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}) \quad (\text{A.8}) \\
&\geq \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^{t-1} = \eta^{t-1}} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h}^t | \eta^{t-1}) p_{i, \sigma, \bar{\theta}}(\bar{h} | \bar{h}^t) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}).
\end{aligned}$$

Combining (A.7) and (A.8) implies

$$\begin{aligned}
& \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^{t-1} = \eta^{t-1}} u_i(\bar{h}, \bar{\theta}) p_{i, \sigma, \bar{\theta}}(\bar{h} | \eta^{t-1}) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}) \quad (\text{A.9}) \\
&\geq \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H | \bar{h}^{t-1} = \eta^{t-1}} u_i(\bar{h}, \bar{\theta}) p_{i, (\sigma'_i, \sigma_{-i}), \bar{\theta}}(\bar{h} | \eta^{t-1}) q^{(\sigma, \nu), i, \tau_i, \eta^{t-1}}(\bar{h}, \bar{\theta}).
\end{aligned}$$

Since (A.9) holds for any σ'_i , σ_i is conditionally optimal according to $\nu_{i, \tau_i, \eta^{t-1}}$. Since this conclusion holds for any $\eta^{t-1} \in H^{t-1}$, the induction step is completed. It follows that (σ, ν) satisfies the second set of inequalities in the definition of sequentially optimal.

To show it also satisfies the first set of inequalities, note that for any $\tau_i \in \Theta_i$, the conditional

optimality of σ_i according to $\nu_{i,\tau_i,\emptyset}$ implies

$$\begin{aligned} & \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H} u_i(\bar{h}, \bar{\theta}) p_{i,\sigma,\bar{\theta}}(\bar{h} | \bar{h}^0) q^{(\sigma,\nu),i,\tau_i,\emptyset}(\bar{h}, \bar{\theta}) \\ \geq & \sum_{\bar{\theta} \in \Theta | \bar{\theta}_i = \tau_i} \sum_{\bar{h} \in H} u_i(\bar{h}, \bar{\theta}) p_{i,(\sigma'_i, \sigma_{-i}),\bar{\theta}}(\bar{h} | \bar{h}^0) q^{(\sigma,\nu),i,\tau_i,\emptyset}(\bar{h}, \bar{\theta}). \end{aligned} \quad (\text{A.10})$$

Since i 's preferences satisfy extended smooth rule updating using σ , for all τ_i , $q^{(\sigma,\nu),i,\tau_i,\emptyset}(\bar{h}, \bar{\theta}) \propto q^{\sigma,i}(\bar{h}, \bar{\theta})$ for all $\bar{\theta} \in \Theta$ and $\bar{h} \in H$. Substituting back into (A.10), cancelling the constant of proportionality and summing for all τ_i , yields:

$$\begin{aligned} & \sum_{\bar{\theta} \in \Theta} \sum_{\bar{h} \in H} u_i(\bar{h}, \bar{\theta}) p_{i,\sigma,\bar{\theta}}(\bar{h} | \bar{h}^0) q^{\sigma,i}(\bar{h}, \bar{\theta}) \\ \geq & \sum_{\bar{\theta} \in \Theta} \sum_{\bar{h} \in H} u_i(\bar{h}, \bar{\theta}) p_{i,(\sigma'_i, \sigma_{-i}),\bar{\theta}}(\bar{h} | \bar{h}^0) q^{\sigma,i}(\bar{h}, \bar{\theta}). \end{aligned} \quad (\text{A.11})$$

Since (A.11) holds for any σ'_i , (σ, ν) also satisfies the first set of inequalities in the definition of sequentially optimal, thus it is sequentially optimal. ■

Proof of Theorem 2.4. By Theorem 2.1, there exists an interim belief system $\tilde{\nu}$ satisfying the smooth rule using σ as the ex-ante equilibrium such that $(\sigma, \tilde{\nu})$ is sequentially optimal. Consider a sequence of a completely mixed strategy profiles converging to σ and the corresponding sequence of interim belief systems determined by extended smooth rule updating using the strategies in the sequence as the strategy profile. Let $\hat{\nu}$ be the limit of this sequence of interim belief systems. By construction, $(\sigma, \hat{\nu})$ satisfies smooth rule consistency. It remains to show that $(\sigma, \hat{\nu})$ is sequentially optimal. By continuity of extended smooth rule updating in the strategy profile and the fact that all players view all partial histories h^t with $1 \leq t \leq T$ as reachable from h^0 , $\hat{\nu} = \tilde{\nu}$ except possibly after partial histories h^{T+1} for which $m_i(h^{T+1}) < T+1$ for at least one player i (i.e., except after histories involving off-path action(s) in the final stage). Note however, using the definition of $V_{i,\tau_i,h^{T+1}}$, optimality after the last stage is trivial because there are no actions remaining, and so how interim beliefs are specified after off-path play at $T+1$ is irrelevant for sequential optimality. Therefore, $(\sigma, \hat{\nu})$ inherits sequential optimality from $(\sigma, \tilde{\nu})$. ■

Proof of Theorem 2.5. Fix a sequence $\varepsilon^k = (\varepsilon_{\theta_j, \eta}^k)_{\theta_j \in \cup_{i \in N} \Theta_i, \eta \in \bigcup_{t \in T} H^t}$ of strictly positive vectors of dimension $\left| \bigcup_{t \in T} H^t \right| \sum_{j \in N} |\Theta_j|$, converging in the sup-norm to 0 and such that $\varepsilon_{\theta_j, h^t}^k \leq \frac{1}{|A_j^t(h^t)|}$ for all types θ_j and all histories h and all stages t . 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 $\sigma_j^k(h^t, \theta_j)(a_j^t) \geq \varepsilon_{\theta_j, h^t}^k$ for all j, θ_j, h^t and actions $a_j^t \in A_j^t(h^t)$. 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 . Let $\hat{\nu}^k$ be a belief system in Γ^k that satisfies the smooth rule using $\hat{\sigma}^k$ as the ex-ante equilibrium. Since the strategy of each player (h^t, θ_j) in G^k according to $\hat{\sigma}^k$ is an ex-ante best response to $\hat{\sigma}^k$, and since all partial histories h^t are on the equilibrium path, by Lemma A.1 smooth rule updating ensures 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 of the payoff functions, $\hat{\sigma}$ is an ex-ante equilibrium of Γ . Inspection of Definition 2.7 reveals that the beliefs generated by the smooth rule vary continuously in the ex-ante equilibrium σ , as σ enters continuously in $p_{i, \sigma, \theta}(h^t|h^r)$ and $p_{-i, \sigma, \theta}(h^t|h^r)$ and only affects $m_i(h^t)$ when the weight on some action hits zero, in which case the smooth rule becomes less restrictive and so the same beliefs can be maintained in that case. By this continuity of the smooth rule in the ex-ante equilibrium $\hat{\sigma}^k$, the associated sub-sequence of $\hat{\nu}^k$ converges to a limit denoted by $\hat{\nu}$. Given any partial history h^t and continuation strategy $\tilde{\sigma}_j^{h^t}$ of player j of type θ_j in Γ , let $\tilde{\sigma}_j^{k, h^t}$ be a feasible strategy in Γ^k for this player that is closest (in the sup-norm) to $\tilde{\sigma}_j^{h^t}$. Since, by sequential optimality of $(\hat{\sigma}^k, \hat{\nu}^k)$ for each k , $\hat{\sigma}_j^{k, h^t}$ is weakly better than $\tilde{\sigma}_j^{k, h^t}$ for player j of type θ_j given belief $\hat{\nu}_{j, \theta_j, h^t}^k$, and since, along the sub-sequence, $\tilde{\sigma}_j^{k, h^t}$ converges to $\tilde{\sigma}_j^{h^t}$ and $\hat{\nu}^k$ converges to $\hat{\nu}$, continuity of the payoff functions implies that $\hat{\sigma}_j^{h^t}$ is weakly better than $\tilde{\sigma}_j^{h^t}$ for this player given belief $\hat{\nu}_{j, \theta_j, h^t}$. Therefore $(\hat{\sigma}, \hat{\nu})$ satisfies sequential optimality. Finally, observe that $(\hat{\sigma}, \hat{\nu})$ satisfies smooth rule consistency (since it is explicitly constructed as the limit of an appropriate sequence). Therefore $(\hat{\sigma}, \hat{\nu})$ is an SEA of Γ . ■

Proof of Lemma 2.1. By smooth rule consistency and upper semi-continuity of the extended smooth rule in the strategy profile at i, τ_i, h^t such that i does not view h^t as reachable from h^{t-1} in the limit and continuity of the extended smooth rule in the strategy profile everywhere else, ν satisfies extended smooth rule updating using σ as the strategy profile. (The upper semi-continuity comes from the fact that at unreachable partial histories the extended smooth rule is less restrictive than at reachable partial histories.) ■

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 Lemma 2.1 and Theorem 2.3. ■

Proof of Corollary 2.2. We show that $(\sigma^P, \hat{\nu}^P)$, where, for all $i, \tau_i, \hat{\nu}_{i, \tau_i, h^t}^P = \nu_{i, \tau_i, h^t}^P$ whenever $(t > 0$ and $m_i(h^t) = t)$, and where, everywhere else, $\hat{\nu}_{i, \tau_i, h^t}$ is derived via the

smooth rule, is a PEA. By the proof of Theorem 2.1, $(\sigma^P, \hat{\nu}^P)$ is sequentially optimal. It remains to show that it naturally extends updating. This imposes restrictions on i 's beliefs only at partial histories h^t where i does *not* view h^t as reachable from h^{t-1} . By construction of $\hat{\nu}^P$, at all such h^t , $\hat{\nu}_{i,\tau_i,h^t}^P = \nu_{i,\tau_i,h^t}^P$. Thus $(\sigma^P, \hat{\nu}^P)$ naturally extends updating because (σ^P, ν^P) does. ■

Lemma A.2 *Any $(\hat{\sigma}, \hat{\nu})$ satisfying smooth rule consistency also naturally extends updating.*

Proof of Lemma A.2. Fix any (σ, ν) such that ν satisfies extended smooth rule updating using σ as the strategy profile. Consider players $i, j \neq i$, type τ_i and partial histories h^t that player i views as reachable from h^{t-1} and for which conditions (b)-(d) in the definition of naturally extends updating are satisfied. Since player i has no costly ambiguity exposure under σ at h^{t-1} and h^t , the smooth rule updating formula (2.8) simplifies to

$$\nu_{i,\tau_i,h^t}(\pi) = A \cdot \left(\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} p_{-i,\sigma,\hat{\theta}}(h^t|h^{t-1}) \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \right) \nu_{i,\tau_i,h^{t-1}}(\pi)$$

for π such that $\pi(\Theta_{i,\tau_i,h^t}) > 0$, where A is the normalization factor ensuring the left-hand side sums (over π) to 1.

Since player j has only one action at h^{t-1} , $p_{-i,\sigma,\theta}(h^t|h^{t-1}) = \prod_{k \neq i,j} \sigma_k(h^{t-1}, \theta_k)(h_{t-1,k})$. Thus, for each θ_j ,

$$\begin{aligned} & \pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) \nu_{i,\tau_i,h^t}(\pi) \\ = & A \pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) \left(\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} \left(\prod_{k \neq i,j} \sigma_k(h^{t-1}, \hat{\theta}_k)(h_{t-1,k}) \right) \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \right) \nu_{i,\tau_i,h^{t-1}}(\pi). \end{aligned}$$

Summing over the π yields,

$$\begin{aligned} & \sum_{\pi \text{ s.t. } \nu_{i,\tau_i,h^{t-1}}(\pi) > 0 \text{ and } \pi(\Theta_{i,\tau_i,h^t}) > 0} \pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) \nu_{i,\tau_i,h^t}(\pi) \quad (\text{A.12}) \\ = & A \sum_{\pi \text{ s.t. } \nu_{i,\tau_i,h^{t-1}}(\pi) > 0 \text{ and } \pi(\Theta_{i,\tau_i,h^t}) > 0} \pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) \nu_{i,\tau_i,h^{t-1}}(\pi) \\ & \cdot \left(\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} \left(\prod_{k \neq i,j} \sigma_k(h^{t-1}, \hat{\theta}_k)(h_{t-1,k}) \right) \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \right). \quad (\text{A.13}) \end{aligned}$$

By (2.2) applied to $\nu_{i,\tau_i,h^{t-1}}$, we can replace the left hand side of (A.12) with

$$\sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) \nu_{i,\tau_i,h^t}(\pi).$$

By definition of $\pi_{\Theta_{i,\tau_i,h^t}}$,

$$\pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) = \frac{\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t} \cap (\{\theta_j\} \times \Theta_{-j})} \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \left(\prod_{k \neq i,j} \sigma_k(h^{t-1}, \hat{\theta}_k)(h_{t-1,k}) \right)}{\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \left(\prod_{k \neq i,j} \sigma_k(h^{t-1}, \hat{\theta}_k)(h_{t-1,k}) \right)}.$$

Thus, (A.12) becomes,

$$\begin{aligned} & \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) \nu_{i,\tau_i,h^t}(\pi) \tag{A.14} \\ &= A \sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t} \cap (\{\theta_j\} \times \Theta_{-j})} \left(\prod_{k \neq i,j} \sigma_k(h^{t-1}, \hat{\theta}_k)(h_{t-1,k}) \right) \\ & \cdot \sum_{\pi \text{ s.t. } \nu_{i,\tau_i,h^{t-1}}(\pi) > 0 \text{ and } \pi(\Theta_{i,\tau_i,h^t}) > 0} \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \nu_{i,\tau_i,h^{t-1}}(\pi). \end{aligned}$$

By (2.2) applied to $\nu_{i,\tau_i,h^{t-1}}$, since $\hat{\theta} \in \Theta_{i,\tau_i,h^t} \cap (\{\theta_j\} \times \Theta_{-j})$ (so that if $\pi(\Theta_{i,\tau_i,h^t}) = 0$ then $\pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) = 0$),

$$\sum_{\pi \text{ s.t. } \nu_{i,\tau_i,h^{t-1}}(\pi) > 0 \text{ and } \pi(\Theta_{i,\tau_i,h^t}) > 0} \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \nu_{i,\tau_i,h^{t-1}}(\pi) = \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \nu_{i,\tau_i,h^{t-1}}(\pi).$$

Since $\nu_{i,\tau_i,h^{t-1}}$ satisfies the condition that the reduced measure $\sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\hat{\theta})} \nu_{i,\tau_i,h^{t-1}}(\pi) \in \Delta(\Theta)$ is a product measure, for $\hat{\theta} \in \Theta_{i,\tau_i,h^t} \cap (\{\theta_j\} \times \Theta_{-j})$

$$\begin{aligned}
& \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\hat{\theta})} \nu_{i,\tau_i,h^{t-1}}(\pi) \\
&= \prod_{k \in N} \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\hat{\theta}_k\} \times \Theta_{-k})} \nu_{i,\tau_i,h^{t-1}}(\pi) \\
&= \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\theta_j\} \times \Theta_{-j})} \nu_{i,\tau_i,h^{t-1}}(\pi) \left(\prod_{k \neq i,j} \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\hat{\theta}_k\} \times \Theta_{-k})} \nu_{i,\tau_i,h^{t-1}}(\pi) \right) \\
&= \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\theta_j\} \times \Theta_{-j})} \nu_{i,\tau_i,h^{t-1}}(\pi) \left(\prod_{k \neq i,j} \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\hat{\theta}_k\} \times \Theta_{-k})} \nu_{i,\tau_i,h^{t-1}}(\pi) \right).
\end{aligned}$$

Substituting into (A.14) yields,

$$\begin{aligned}
& \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^t}(\{\theta_j\} \times \Theta_{-j})} \nu_{i,\tau_i,h^t}(\pi) \tag{A.15} \\
&= A \sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t} \cap (\{\theta_j\} \times \Theta_{-j})} \left(\prod_{k \neq i,j} \sigma_k(h^{t-1}, \hat{\theta}_k)(h_{t-1,k}) \right) \\
&\quad \cdot \left(\prod_{k \neq i,j} \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\hat{\theta}_k\} \times \Theta_{-k})} \nu_{i,\tau_i,h^{t-1}}(\pi) \right) \left(\sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\theta_j\} \times \Theta_{-j})} \nu_{i,\tau_i,h^{t-1}}(\pi) \right) \\
&= B \left(\sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\theta_j\} \times \Theta_{-j})} \nu_{i,\tau_i,h^{t-1}}(\pi) \right)
\end{aligned}$$

where

$$B = A \sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t} \cap (\{\theta_j\} \times \Theta_{-j})} \left(\prod_{k \neq i,j} \left(\left(\sigma_k(h^{t-1}, \hat{\theta}_k)(h_{t-1,k}) \right) \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\hat{\theta}_k\} \times \Theta_{-k})} \nu_{i,\tau_i,h^{t-1}}(\pi) \right) \right),$$

showing that i 's marginal on player j 's type at partial history h^t remains the same as at h^{t-1} .

Since all the operations in the above argument are continuous in σ , it is also true that to ensure that i 's marginal on player j 's type at partial history h^t remains approximately the same as at h^{t-1} it is sufficient to know that conditions (b) and (d) in the definition of naturally extends updating (Definition 2.13) hold to within a given approximation (condition (c) always holds exactly, as it is part of the structure of the game).

By smooth rule consistency of $(\hat{\sigma}, \hat{\nu})$, there exists a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$, with $\lim_{k \rightarrow \infty} \sigma^k = \hat{\sigma}$, such that $\hat{\nu} = \lim_{k \rightarrow \infty} \nu^k$, where ν^k is determined by extended smooth rule updating using σ^k as the strategy profile. Fix any i, τ_i, h^t such that conditions (a)-(d) in the definition of naturally extends updating are satisfied for $(\hat{\sigma}, \hat{\nu})$. Then, by the argument above applied to (σ^k, ν^k) sufficiently far along the sequence, $(\hat{\sigma}, \hat{\nu})$ naturally extends updating. ■

Proof of Theorem 2.6. That (σ^S, ν^S) satisfies sequential optimality follows directly from the definition of a SEA. By Lemma A.2, and the fact that (σ^S, ν^S) satisfies smooth rule consistency, (σ^S, ν^S) naturally extends updating and is thus a PEA of Γ . ■

Proof of Theorem 2.7. Fix Γ, σ , and h as in the statement of the theorem. Ex-ante equilibrium, since only player j has non-trivial choices, is equivalent to ex-ante optimality of σ_j according to j 's preferences. This ex-ante optimality implies (see Hanany and Klibanoff 2009, Lemma A.1)

$$\begin{aligned} & \sum_{\bar{\theta} \in \Theta} \sum_{\bar{h} \in H} u_j(\bar{h}, \bar{\theta}) p_{j, \sigma, \bar{\theta}}(\bar{h} | \bar{h}^0) q^{\sigma, j}(\bar{h}, \bar{\theta}) \\ & \geq \sum_{\bar{\theta} \in \Theta} \sum_{\bar{h} \in H} u_j(\bar{h}, \bar{\theta}) p_{j, (\sigma'_j, \sigma_{-j}), \bar{\theta}}(\bar{h} | \bar{h}^0) q^{\sigma, j}(\bar{h}, \bar{\theta}) \end{aligned}$$

for all σ'_j , where $q^{\sigma, j}(\bar{h}, \bar{\theta})$ is the ‘‘local ambiguity neutral measure’’ as in (A.1):

$$\begin{aligned} q^{\sigma, j}(\bar{h}, \bar{\theta}) \equiv & \sum_{\pi \in \Delta(\Theta)} \phi'_j \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H} u_j(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | \hat{h}^0) \pi(\hat{\theta}) \right) \\ & \cdot p_{-j, \sigma, \bar{\theta}}(\bar{h} | \bar{h}^0) \pi(\bar{\theta}) \mu_j(\pi). \end{aligned}$$

Consider any $\hat{\mu}_j$ such that, for all θ and some $b > 0$,

$$\sum_{\pi} \pi(\theta) \hat{\mu}_j(\pi) = b q^{\sigma, j}(h, \theta).$$

Let $\hat{\Gamma}$ be the game identical to Γ except that ϕ_i is the identity for all i and player j has belief $\hat{\mu}_j$. Since, by construction, σ is ex-ante optimal for player j given beliefs $\hat{\mu}_j$, and no other players have non-trivial choices, σ is an ex-ante equilibrium of $\hat{\Gamma}$ such that $\sum_{\pi} \sum_{\theta} p_{\sigma, \theta}(h | h^0) \pi(\theta) \hat{\mu}_j(\pi) > 0$ for the history h .

We complete the proof by showing that σ is part of an SEA of $\hat{\Gamma}$. Consider a sequence of completely mixed strategy profiles σ^k with limit σ . Let $\hat{\nu}^k$ be the associated sequence of interim belief systems defined by Bayesian updating of $\hat{\mu}_j$ given σ^k . Let $\hat{\nu} \equiv \lim_{k \rightarrow \infty} \hat{\nu}^k$. Note

that $(\sigma, \hat{\nu})$ satisfies smooth rule consistency by construction, since smooth rule updating is Bayesian updating when players are ambiguity neutral. By Lemma 2.1, $\hat{\nu}$ is consistent with Bayesian updating using σ as the strategy profile. Because only j has non-trivial choices, (1) Bayesian updating uniquely pins down j 's interim beliefs after every partial history, and (2) the beliefs and interim beliefs of players other than j are irrelevant for checking sequential optimality. Therefore, since j is ambiguity neutral (subjective expected utility) in game $\hat{\Gamma}$ and Bayesian updating preserves ex-ante optimality under ambiguity neutrality, $(\sigma, \hat{\nu})$ is sequentially optimal and is therefore an SEA of $\hat{\Gamma}$. ■

Proof of Proposition 3.1. Observe that player 1 is willing ex-ante to play P with positive probability if and only if, after the play of P , (U, R) will be played with probability at least $\frac{1}{2}$. Suppose there is an ex-ante equilibrium, σ , in which P is played with positive probability. Let p_I and p_{II} denote the probabilities according to σ that types I and II, respectively, of player 1 play P . Then player 2 finds it optimal to play U with positive probability if and only if

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

which is equivalent to

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

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

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

which is equivalent to

$$p_I \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(\pi(I)) \leq \frac{2}{3} p_{II} \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(1 - \pi(I)). \quad (\text{A.17})$$

Since (A.16) and (A.17) 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. ■

Proof of Proposition 3.2. Let μ put probability $\frac{1}{2}$ on π_0 and $\frac{1}{2}$ on π_1 , where $\pi_0(\text{I}) = 1$ and $\pi_1(\text{I}) = 0$.¹⁰ Let $\phi(x) \equiv -e^{-x}$.¹¹ 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 . The associated interim belief system, ν^k , is derived from σ^k according to the extended smooth rule and we define ν as $\lim_{k \rightarrow \infty} \nu^k$. Recall that only player 1 has more than one possible type. Thus $\nu_{1, \text{I}, \eta}^k(\pi_0) = 1$ and $\nu_{1, \text{II}, \eta}^k(\pi_1) = 1$ for all partial

histories η . Furthermore, $\nu_{2, \theta_2, \emptyset}^k(\pi_0) = \frac{\frac{1}{2} \frac{\phi'((1-\frac{1}{2k})\lambda^*)}{\phi'((1-\frac{1}{2k})\lambda^*)}}{\frac{1}{2} \frac{\phi'((1-\frac{1}{2k})\lambda^*)}{\phi'((1-\frac{1}{2k})\lambda^*)} + \frac{1}{2} \frac{\phi'((1-\frac{1}{2k})(\lambda^* + \frac{5}{2}(1-\lambda^*)))}{\phi'((1-\frac{1}{2k})(\lambda^* + \frac{5}{2}(1-\lambda^*)))}} = \frac{1}{2}$, $\nu_{3, \theta_3, \emptyset}^k(\pi_0) =$

$\frac{1}{2}$, $\nu_{2, \theta_2, P}^k(\pi_0) = \frac{\frac{1}{2}(1-\frac{1}{2k}) \frac{\phi'((1-\frac{1}{2k})\lambda^*)}{\phi'(\lambda^*)}}{\frac{1}{2}(1-\frac{1}{2k}) \frac{\phi'((1-\frac{1}{2k})\lambda^*)}{\phi'(\lambda^*)} + \frac{1}{2}(1-\frac{1}{2k}) \frac{\phi'((1-\frac{1}{2k})(\lambda^* + \frac{5}{2}(1-\lambda^*)))}{\phi'((\lambda^* + \frac{5}{2}(1-\lambda^*)))}} = \frac{1}{1+e^{-\frac{5(1-\lambda^*)}{4k}}}$, and $\nu_{3, \theta_3, P}^k(\pi_0) =$

$\frac{\frac{1}{2}(1-\frac{1}{2k}) \frac{\phi'((1-\frac{1}{2k})(\lambda^* + \frac{5}{2}(1-\lambda^*)))}{\phi'(\lambda^* + \frac{5}{2}(1-\lambda^*))}}{\frac{1}{2}(1-\frac{1}{2k}) \frac{\phi'((1-\frac{1}{2k})(\lambda^* + \frac{5}{2}(1-\lambda^*)))}{\phi'(\lambda^* + \frac{5}{2}(1-\lambda^*))} + \frac{1}{2}(1-\frac{1}{2k}) \frac{\phi'((1-\frac{1}{2k})\lambda^*)}{\phi'(\lambda^*)}} = \frac{1}{1+e^{-\frac{5(1-\lambda^*)}{4k}}}$. Since $\lim_{k \rightarrow \infty} \nu_{2, \theta_2, P}^k(\pi_0) =$

$\lim_{k \rightarrow \infty} \nu_{3, \theta_3, P}^k(\pi_0) = \lim_{k \rightarrow \infty} \nu_{2, \theta_2, \emptyset}^k(\pi_0) = \lim_{k \rightarrow \infty} \nu_{3, \theta_3, \emptyset}^k(\pi_0) = \frac{1}{2}$, σ remains optimal for players 2 and 3 following the play of P given ν . (The beliefs at other partial histories can also be calculated using the smooth rule formula, but they will not matter for sequential optimality because the relevant player has no moves left.) Thus, (σ, ν) is sequentially optimal

¹⁰The 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.

¹¹Any more concave ϕ will also work, as will any ϕ more concave than $-e^{-\alpha x}$ for $\alpha = \frac{-4(\ln(2/3))}{5(2-\sqrt{2})} \approx 0.554$.

and, by construction, satisfies smooth rule consistency. It is therefore an SEA. ■

Rather than the “one-step ahead” formulation of smooth rule updating in (2.8) for a partial history from an immediately prior one, one could alternatively (and equivalently) write the smooth rule as updating $\nu_{i,\tau_i,h^{m_i}(h^t)}$ to ν_{i,τ_i,h^t} all at once:

For each partial history h^t , for all $\pi \in \Delta(\Theta)$ such that $\pi(\Theta_{i,\tau_i,h^t}) > 0$,

$$\nu_{i,\tau_i,h^t}(\pi) \propto \frac{\phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H | \hat{h}^{m_i}(h^t) = h^{m_i}(h^t)} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | h^{m_i}(h^t)) \pi_{\Theta_{i,\tau_i,h^{m_i}(h^t)}}(\hat{\theta}) \right)}{\phi'_i \left(\sum_{\hat{\theta} \in \Theta} \sum_{\hat{h} \in H | \hat{h}^t = h^t} u_i(\hat{h}, \hat{\theta}) p_{\sigma, \hat{\theta}}(\hat{h} | h^t) \pi_{\Theta_{i,\tau_i,h^t}}(\hat{\theta}) \right)} \cdot \left(\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} p_{-i,\sigma,\hat{\theta}}(h^t | h^{m_i}(h^t)) \pi_{\Theta_{i,\tau_i,h^{m_i}(h^t)}}(\hat{\theta}) \right) \nu_{i,\tau_i,h^{m_i}(h^t)}(\pi). \quad (\text{A.18})$$

This “all-at-once” formula is convenient for the proof of the next result.

Proof of Proposition 3.3. Since type uncertainty is only about player 3 and there are only two possible types, represent probabilities over the types by the probability, p , of Type I. Think of player 2 being subjectively uncertain whether nature used $\hat{p} \geq 1/2$ or $\check{p} = 1 - \hat{p}$ to determine player 3’s type. Specifically, let $\mu_2(\hat{p}) = \mu_2(\check{p}) = 1/2$. Also assume ϕ_2 is smooth and strictly concave.

We now show that $(R, F, (C, C), \lambda^*H + (1 - \lambda^*)G)$ is an SEA given $\hat{p} = \frac{3}{4}$ and $\phi_2(x) = -e^{-2x}$ and defining λ^* by

$$\lambda^* = \arg \max_{\lambda \in [0,1]} \frac{1}{2} \phi_2(2\lambda + 6(1 - \lambda)\hat{p}) + \frac{1}{2} \phi_2(2\lambda + 6(1 - \lambda)\check{p}). \quad (\text{A.19})$$

First, we show that the strategy profile is an ex-ante equilibrium of the game. Player 1 is best responding (for any specification of ϕ_1 and μ_1) because he gets a payoff of 1 on path and would get less than 1 by deviating since 2 plays F following T . Player 2 is best responding on-path by the definition of λ^* . Player 3 has (C, C) as a best response (for any specification of ϕ_3 and μ_3) if and only if $\lambda^* \geq 1/2$. First-order conditions for an interior λ^* are given by

$$(1 - 3\hat{p})\phi'_2(2\lambda^* + 6(1 - \lambda^*)\hat{p}) + (3\hat{p} - 2)\phi'_2(2\lambda^* + 6(1 - \lambda^*)(1 - \hat{p})) = 0. \quad (\text{A.20})$$

Notice that the left-hand side of (A.20) is always negative for $\lambda^* = 1$. By concavity therefore, $\lambda^* < 1$. Observe that $\lambda^* \geq 1/2$ if and only if the left-hand side of (A.20) is non-negative at $\lambda^* = 1/2$ i.e.,

$$(1 - 3\hat{p})\phi'_2(1 + 3\hat{p}) + (3\hat{p} - 2)\phi'_2(1 + 3(1 - \hat{p})) \geq 0. \quad (\text{A.21})$$

Observe that this is satisfied for our choice of \hat{p} and ϕ_2 . In particular $\lambda^* = 1 - \frac{\ln(5)}{6}$. Thus 3 is best responding and the strategy profile is an ex-ante equilibrium.

Before specifying interim beliefs observe that, since there is no type uncertainty about players 1 or 2, player 3's beliefs are trivial. Also, since player 1's payoffs do not depend on 3's type or actions, 1's best response to the strategies of the others is independent of his beliefs. Thus, the only important beliefs to specify are those of player 2. We construct such beliefs to satisfy smooth rule consistency. Consider the following sequence of strategy profiles with limit $(R, F, (C, C), \lambda^*H + (1 - \lambda^*)G)$:

$$\sigma^k \equiv \left(\frac{k}{k+1}R + \frac{1}{k+1}T, \frac{k}{k+1}F + \frac{1}{k+1}B, \left(\frac{k}{k+1}C + \frac{1}{k+1}S, \frac{k}{k+1}C + \frac{1}{k+1}S \right), \lambda^*H + (1 - \lambda^*)G \right).$$

Applying the extended smooth rule using σ^k (in particular, applying the formula in (A.18)),

$$\nu_{2,(R,C)}^k(p) \propto \frac{\phi'_i \left(\frac{p(\frac{1}{k+1}(\frac{k}{k+1}2 + \frac{1}{k+1}0) + \frac{k}{k+1}(\frac{1}{k+1}0 + \frac{k}{k+1}(2\lambda^* + 6(1 - \lambda^*)))}{+(1-p)(\frac{1}{k+1}(\frac{k}{k+1}2 + \frac{1}{k+1}4) + \frac{k}{k+1}(\frac{1}{k+1}0 + \frac{k}{k+1}(2\lambda^* + 0(1 - \lambda^*)))} \right)}{\phi'_i(p(2\lambda^* + 6(1 - \lambda^*)) + (1-p)(2\lambda^* + 0(1 - \lambda^*)))} \left(\frac{k}{k+1} \right)^2 \frac{1}{2}$$

Thus,

$$\nu_{2,(R,C)}^k\left(\frac{3}{4}\right) = \frac{(3 + \lambda^*)(-2 - 4k + k^2(-9 + 5\lambda^*))}{2(-30 + 8k(-3 + \lambda^*) + 14\lambda^* + k^2(-27 + 6\lambda^* + 5(\lambda^*)^2))}$$

and so

$$\nu_{2,(R,C)}\left(\frac{3}{4}\right) = \lim_{k \rightarrow \infty} \nu_{2,(R,C)}^k\left(\frac{3}{4}\right) = \frac{1}{2} = \lim_{k \rightarrow \infty} \nu_{2,(R,C)}^k\left(\frac{1}{4}\right) = \nu_{2,(R,C)}\left(\frac{1}{4}\right).$$

Furthermore,

$$\nu_{2,T}^k(p) \propto \frac{\phi'_i \left(\frac{p(\frac{1}{k+1}(\frac{k}{k+1}2 + \frac{1}{k+1}0) + \frac{k}{k+1}(\frac{1}{k+1}0 + \frac{k}{k+1}(2\lambda^* + 6(1 - \lambda^*)))}{+(1-p)(\frac{1}{k+1}(\frac{k}{k+1}2 + \frac{1}{k+1}4) + \frac{k}{k+1}(\frac{1}{k+1}0 + \frac{k}{k+1}(2\lambda^* + 0(1 - \lambda^*)))} \right)}{\phi'_i \left(p(\frac{k}{k+1}2 + \frac{1}{k+1}0) + (1-p)(\frac{k}{k+1}2 + \frac{1}{k+1}4) \right)} \frac{1}{k+1} \frac{1}{2}$$

Thus,

$$\nu_{2,T}^k\left(\frac{3}{4}\right) = \frac{(3 + 2k)(-2 - 4k + k^2(-9 + 5\lambda^*))}{2(-6 - 16k + k^2(-23 + 7\lambda^*) + 4k^3(-3 + \lambda^*))}$$

and so

$$\nu_{2,T}\left(\frac{3}{4}\right) = \lim_{k \rightarrow \infty} \nu_{2,T}^k\left(\frac{3}{4}\right) = \frac{9 - 5\lambda^*}{12 - 4\lambda^*} = \frac{24 + 5\ln(5)}{48 + 4\ln(5)} \approx 0.59.$$

We now verify that 2 is playing optimally with respect to these interim beliefs. By the definition of λ^* (via (A.19)), since $\nu_{2,(R,C)}(\frac{3}{4}) = \frac{1}{2}$, player 2 is indeed best responding after 1 plays R and 3 plays C . If player 2 is given the move after 1 plays T , 2 playing F if a best response if and only if

$$1 \in \arg \max_{\gamma \in [0,1]} \nu_{2,T}(\frac{3}{4})\phi_2(2\gamma + 4(1-\gamma)\frac{1}{4}) + (1 - \nu_{2,T}(\frac{3}{4}))\phi_2(2\gamma + 4(1-\gamma)\frac{3}{4}).$$

This may be verified for $\nu_{2,T}(\frac{3}{4}) = \frac{24+5\ln(5)}{48+4\ln(5)}$ and the ϕ_2 assumed.

Putting everything together, $(R, F, (C, C), \lambda^*H + (1 - \lambda^*)G)$ is sequentially optimal with respect to beliefs satisfying smooth rule consistency. Therefore all the conditions for an SEA are satisfied. ■

Proof of Proposition 3.4. Suppose player 2 is ambiguity neutral (without loss of generality, take ϕ_2 to be the identity). Let γ be player 2's initial reduced probability that 3 is of type I. For C to be played on the equilibrium path, player 1 must play R with positive probability, which can be a best response if and only if player 1's expected payoff following T is less than or equal to 1, the sure payoff after R . This is possible if and only if 2's strategy plays F with probability at least $\frac{5}{6}$ following T . If T is played with positive probability in equilibrium, then 2 playing F with probability at least $\frac{5}{6}$ following T is optimal for 2 if and only if $\gamma \geq 1/2$. In the explanation before Proposition 3.3, we showed that no PEA can have only type I of player 3 play C with positive probability on the equilibrium path. Suppose type II of player 3 plays C with positive probability on path. Optimality for 3 implies this can be true only if 2 plays H with probability weakly higher than G . But then type I of player 3 finds it strictly optimal to play C with probability 1. Note however that in this case 2 strictly prefers G over H , making C strictly worse than S for both types of player 3. It follows that playing C with positive probability on the equilibrium path cannot satisfy condition (2.6) of sequential optimality (and thus PEA) where T is played with positive probability.

It remains to consider the case where 1 plays R with probability 1. Then T is now an off equilibrium path action and thus condition (2.6) places no restrictions on 2's play following T . However, the naturally extends updating condition requires that 2's updated beliefs following T must continue to place weight γ on 3 being of type I because 3 has no move at that stage, 2's marginal over 3's type is a product measure, and there is no costly ambiguity exposure (since ϕ' is constant) for 2. From sequential optimality, it then follows that 2's best response to T is B whenever $\gamma < 1/2$, which contradicts the optimality of 1 playing R . Now suppose $\gamma \geq 1/2$. The same argument as used above after establishing that $\gamma \geq 1/2$ shows that C cannot be played with positive probability on the equilibrium path. In sum, when

player 2 is ambiguity neutral, in any PEA if $\gamma < 1/2$ then 1 plays T and 3 never is given the move, while if $\gamma \geq 1/2$ then 3 never plays C if given the move. ■

Proof of Proposition 3.5. Consider the following limit pricing strategy profile, σ^* : in the first period, 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. If entry occurs, the firms play the complete information Cournot quantities in the second period. If no entry occurs, the incumbent plays its monopoly quantity in the second period.

By Lemma A.3, under the assumptions of the Theorem there exists a $\hat{\phi}$ such that if the entrant's ϕ is at least as concave as $\hat{\phi}$, then (3.5) is satisfied. By the arguments in the text discussing this example, the assumptions of the Theorem together with (3.5) are sufficient for σ^* to satisfy inequality (2.6) of sequential optimality.

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

t	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, $\delta_{t,(q,enter,r)}^k$ converge to 1 as $k \rightarrow \infty$ when r is the Cournot quantity for type t 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 $\delta_{t,(q,no\ entry,r)}^k$ converge to 1 as $k \rightarrow \infty$ when r is the monopoly quantity for type t and converge to 0 otherwise. Note that σ^k converges to σ^* . For all π in the support of the ex-ante belief μ , let $\pi_q^k(t) \equiv \frac{\gamma_{t,q}^k \pi(t)}{\sum_{\hat{t}} \gamma_{\hat{t},q}^k \pi(\hat{t})}$ denote the conditional of π after observing q

in the first period under σ^k . Observe that

$$\lim_{k \rightarrow \infty} \pi_q^k(t) = \lim_{k \rightarrow \infty} \frac{\frac{\beta_{t,q}^k}{\sum_{\hat{q}} \beta_{t,\hat{q}}^k} \pi(t)}{\sum_{\hat{t}} \frac{\beta_{\hat{t},q}^k}{\sum_{\hat{q}} \beta_{\hat{t},\hat{q}}^k} \pi(\hat{t})}.$$

Thus, for $q \neq q_L$,

$$\lim_{k \rightarrow \infty} \pi_q^k(H) = 1 \text{ for all } q < q_L,$$

and

$$\lim_{k \rightarrow \infty} \pi_q^k(L) = 1, \text{ for all } q > q_L.$$

This yields that, for $q \neq q_L$, all measures in the support of $\nu_{E,q} \equiv \lim_{k \rightarrow \infty} \nu_{E,q}^k$ have the same conditional, $\lim_{k \rightarrow \infty} \pi_q^k(t)$, which is the point mass on H if $q < q_L$ and the point mass on L if $q > q_L$, so that the limit beliefs for the entrant constructed from the sequence σ^k according to the extended smooth rule are degenerate after observing anything except q_L . For $q = q_L$,

$$\lim_{k \rightarrow \infty} \pi_{q_L}^k(M) = \frac{\pi(M)}{\pi(L) + \pi(M)} \text{ and } \lim_{k \rightarrow \infty} \pi_{q_L}^k(L) = \frac{\pi(L)}{\pi(L) + \pi(M)}. \quad (\text{A.22})$$

The corresponding sequence of entrant's beliefs after observing q_L , $\nu_{E,q_L}^k(\pi)$ is defined via the extended smooth rule by

$$\nu_{E,q_L}^k(\pi) \propto \frac{\phi' \left(\sum_t \pi(t) \sum_{q,x,y} \gamma_{t,q}^k \lambda_q^k \delta_{t,(q,enter,x)}^k \delta_{(q,enter,y)}^k w(x,y) \right)}{\phi' \left(\sum_t \pi_{q_L}^k(t) \sum_{x,y} \lambda_{q_L}^k \delta_{t,(q_L,enter,x)}^k \delta_{(q_L,enter,y)}^k w(x,y) \right)} \left(\sum_t \gamma_{t,q_L}^k \pi(t) \right) \mu(\pi),$$

where $w(x,y) \equiv (a - b(x+y) - c_E)y - K$ is the entrant's Cournot profit net of entry costs when the incumbent produces x and the entrant produces y . Taking limits and applying (A.22) yields,

$$\nu_{E,q_L}(\pi) \propto \frac{\phi'(\pi(H)w_H)}{\phi'(0)} \left(\frac{\pi(L) + \pi(M)}{2} \right) \mu(\pi), \quad (\text{A.23})$$

for all π such that $\pi(L) + \pi(M) > 0$, recalling that w_H is the entrant's Cournot profit net of entry costs when facing an incumbent of type H . For all other π , $\nu_{E,q_L}(\pi) = 0$. Notice that $\nu_{E,q}$ is the only non-trivial part of the interim beliefs given σ^* : the incumbent is always completely informed after the ex-ante stage and the entrant becomes fully informed after its entry decision. As constructed, therefore, (σ^*, ν) satisfies smooth rule consistency.

The final step in the proof is to verify that (σ^*, ν) satisfies the interim optimality conditions (2.7) of sequential optimality. By construction, the Cournot strategies in the last stage

given entry are optimal with respect to the complete information degenerate beliefs. The fact that $w_L < 0$ plus (3.4) implies that it is optimal for the entrant to stay out if it believes the incumbent is type L and to enter if it believes the incumbent is type H . Therefore, given the constructed $\nu_{E,q}$, the play specified for the entrant by σ^* is indeed interim optimal following $q \neq q_L$. It remains to focus on the path where q_L is observed in the first period. σ^* says for the entrant not to enter. This being optimal from an interim perspective is equivalent to the following:

$$\sum_{\pi | \pi(H) < 1} \nu_{E,q_L}(\pi) \frac{1}{1 - \pi(H)} (\pi(L)w_L + \pi(M)w_M) \phi'(0) \leq 0. \quad (\text{A.24})$$

Substituting (A.23) into (A.24) yields that not entering remaining optimal is equivalent to (3.5). Therefore (σ^*, ν) satisfies the interim optimality conditions (2.7) of sequential optimality.

Having shown (σ^*, ν) is sequentially optimal and satisfies smooth rule consistency, it is therefore an SEA. ■

Lemma A.3 *Under the assumptions of Proposition 3.5 there exists an $\alpha > 0$ such that if ϕ is at least as concave as $-e^{-\alpha x}$ then (3.5) is satisfied.*

Proof. Assume the conditions of the theorem. We show that (3.5) is satisfied for concave enough ϕ . If $\mu(\{\pi | \pi(L)w_L + \pi(M)w_M \leq 0\}) = 1$ then (3.5) is trivially satisfied for any ϕ . For the remainder of the proof, therefore, suppose that $\mu(\{\pi | \pi(L)w_L + \pi(M)w_M > 0\}) > 0$. Let $\Pi^- \equiv \{\pi | \pi(L)w_L + \pi(M)w_M < 0\}$, $\Pi^+ \equiv \{\pi | \pi(L)w_L + \pi(M)w_M > 0\}$, $N \equiv \sum_{\pi \in \Pi^-} \mu(\pi)(\pi(L)w_L + \pi(M)w_M)$, 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 (3.5) 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 (3.5). 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.4, $\frac{\pi^-(L)}{\pi^-(M)} > \frac{\pi^+(L)}{\pi^+(M)}$ 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 (3.5) 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 (3.5). ■

Lemma A.4 *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. ■

Proof of Proposition 4.1. Since the proposed strategies involve the play of both messages, there are no off path actions before the last stage, thus, by Theorem 2.4, it is sufficient to verify that the proposed strategies are sequentially optimal. Furthermore, by Theorem 2.2, this is equivalent to verifying that these strategies form an ex-ante equilibrium, which is established in the rest of this proof.

P 's strategy is an ex-ante best response because it leads to payoff 2 for all types, which is the highest feasible payoff for this player. Let γ_l be the probability with which a_1 plays h after message m_l , $l = 1, 2$, and similarly let δ_l be the corresponding probabilities for a_2 . The

proposed strategies correspond to $\gamma_1 = \gamma_2 = \delta_2 = 1$ and $\delta_1 = 0$. We now verify that these are ex-ante best responses. Denoting $\pi_k(IIB) + \pi_k(IIR)$ by $\pi_k(II)$, given the strategies of the others, $a1$ maximizes

$$\frac{1}{2} \sum_{k=1}^2 \phi_{a1} (\pi_k(IB)\gamma_1 + 2\pi_k(IR)\gamma_2 + \pi_k(II)[2\gamma_2 + 5(1 - \gamma_2)]).$$

Since this function is strictly increasing in γ_1 , it is clearly maximized at $\gamma_1 = 1$. The first derivative with respect to γ_2 evaluated at $\gamma_1 = \gamma_2 = 1$ is

$$\begin{aligned} & \frac{1}{2} \sum_{k=1}^2 [2\pi_k(IR) - 3\pi_k(II)] \phi'_{a1} (2 - \pi_k(IB)) \\ &= \frac{11}{8} e^{-11 \cdot \frac{39}{20}} \left(e^{-11(\frac{7}{4} - \frac{39}{20})} - \frac{42}{5} \right) > 0, \end{aligned}$$

where the last equality uses $\phi_{a1}(x) = -e^{-11x}$ and the values of the π_k . Thus, by concavity in γ_2 , the maximum is attained at $\gamma_1 = \gamma_2 = 1$. Similarly, given the strategies of the others, $a2$ maximizes

$$\frac{1}{2} \sum_{k=1}^2 \phi_{a2} (\pi_k(IB)[2\delta_1 + 5(1 - \delta_1)] + \pi_k(IR)[2\delta_2 + 5(1 - \delta_2)] + 2\pi_k(II)\delta_2).$$

Since this function is strictly decreasing in δ_1 , it is clearly maximized at $\delta_1 = 0$. The first derivative with respect to δ_2 evaluated at $\delta_1 = 0$ and $\delta_2 = 1$ is

$$\begin{aligned} & \frac{1}{2} \sum_{k=1}^2 [-3\pi_k(IR) + 2\pi_k(II)] \phi'_{a2} (3\pi_k(IB) + 2) \\ &= -\frac{1}{2} \phi'_{a2} \left(\frac{11}{4} \right) + \frac{23}{40} \phi'_{a2} \left(\frac{43}{20} \right) \geq \frac{3}{40} \phi'_{a2} \left(\frac{11}{4} \right) > 0, \end{aligned}$$

where the last equality uses the values of the π_k . Since ϕ_{a2} is weakly concave, the problem is weakly concave in δ_2 , thus the maximum is attained at $\delta_1 = 0$ and $\delta_2 = 1$. ■

Proof of Proposition 4.2.

By definition, a non-Ellsberg strategy for P conditions only on the payoff relevant part of the type, I and II . Denote P 's probability of playing m_1 conditional on the payoff relevant part of the type by r_I and r_{II} . Let γ_l be the probability with which $a1$ plays h after message m_l , $l = 1, 2$, and similarly let δ_l be the corresponding probabilities for $a2$. Given r_I and r_{II} ,

$a1$ chooses γ_1, γ_2 to maximize

$$\frac{1}{2} \sum_{k=1}^2 \phi_{a1}(\pi_k(I)[r_I(1 + \delta_1)\gamma_1 + (1 - r_I)(1 + \delta_2)\gamma_2] + \pi_k(II)[r_{II}((1 + \delta_1)\gamma_1 + 5(1 - \gamma_1)) + (1 - r_{II})((1 + \delta_2)\gamma_2 + 5(1 - \gamma_2))]) \quad (\text{A.25})$$

and $a2$ chooses δ_1, δ_2 to maximize

$$\frac{1}{2} \sum_{k=1}^2 \phi_{a2}(\pi_k(I)[r_I((1 + \gamma_1)\delta_1 + 5(1 - \delta_1)) + (1 - r_I)((1 + \gamma_2)\delta_2 + 5(1 - \delta_2))] + \pi_k(II)[r_{II}(1 + \gamma_1)\delta_1 + (1 - r_{II})(1 + \gamma_2)\delta_2]). \quad (\text{A.26})$$

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

Case 1: When $r_I = r_{II} = 1$ (resp. $r_I = r_{II} = 0$) so that only one message is sent, for P to always receive the maximal payoff of 2 it is necessary that the agents play h_1, h_2 with probability 1 after this message, i.e. $\gamma_1 = \delta_1 = 1$ (resp. $\gamma_2 = \delta_2 = 1$). But h_2 is not a best response for $a2$, as can be seen by the fact that the partial derivative of (A.26) with respect to δ_1 (resp. δ_2) evaluated at those strategies is

$$\frac{1}{2} (4 - 5 \sum_{k=1}^2 \pi_k(I)) \phi'_{a2}(2) = -\frac{3}{8} \phi'_{a2}(2) < 0.$$

Similarly, one can show that h_1 is not a best response for $a1$.

Case 2: When $0 < r_{II} < 1$, since type II sends both messages with positive probability, it is necessary that h_1, h_2 are 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 $a2$ is that the partial derivatives of (A.26) with respect to δ_1, δ_2 are non-negative at $\gamma_1 = \gamma_2 = \delta_1 = \delta_2 = 1$. This is, respectively, equivalent to $14r_{II} \geq 19r_I$ and $14(1 - r_{II}) \geq 19(1 - r_I)$, which implies $14 \geq 19$, a contradiction.

Case 3: When $r_{II} = 1$ and $0 \leq r_I < 1$, (A.26) is strictly decreasing in δ_2 , thus the maximum is attained at $\delta_2 = 0$. For the principal to always receive the maximal payoff of 2, it is necessary that $\gamma_1 = \gamma_2 = \delta_1 = 1$. However, this is not a best response for $a1$ because the partial derivative of (A.25) with respect to γ_1 evaluated at these strategies using the values for the π_k is,

$$\frac{3}{8} (2r_I - 1) \phi'_{a1}\left(\frac{15r_I + 25}{20}\right) + \frac{1}{5} (r_I - 6) \phi'_{a1}\left(\frac{4r_I + 36}{20}\right) < 0.$$

To see this inequality, note that the second term is always negative, the first term is non-positive if $0 \leq r_I \leq \frac{1}{2}$, and, when $\frac{1}{2} < r_I < 1$, substituting $\phi_{a1}(x) = -e^{-11x}$ yields that the

left-hand side is negative.

Case 4: When $r_{II} = 0$ and $0 < r_I \leq 1$, (A.26) is strictly decreasing in δ_1 , thus the maximum is attained at $\delta_1 = 0$. For the principal to always receive the maximal payoff of 2, it is necessary that $\gamma_1 = \gamma_2 = \delta_2 = 1$. However, this is not a best response for $a1$ because the partial derivative of (A.25) with respect to γ_2 evaluated at these strategies using the values for the π_k is,

$$\frac{3}{4}\left(\frac{1}{2} - r_I\right)\phi'_{a1}\left(2 - \frac{3}{4}r_I\right) + \left(-\frac{1}{5}r_I - 1\right)\phi'_{a1}\left(2 - \frac{r_I}{5}\right) < 0.$$

To see this inequality, note that the second term is always negative, the first term is non-positive for $\frac{1}{2} \leq r_I \leq 1$, and, when $0 < r_I < \frac{1}{2}$, substituting $\phi_{a1}(x) = -e^{-11x}$ yields that the left-hand side is negative. ■

Lemma A.5 *Let f and g be continuously differentiable, concave and strictly increasing functions mapping reals to reals such that g is at least as concave as f and let $x < y$. Then $g'(x)/g'(y) \geq f'(x)/f'(y)$.*

Proof. By definition of at least as concave as, for all a in the domain of f , $g(a) = h(f(a))$ for some function h that is concave and strictly increasing on the range of f . Thus, $g'(a) = h'(f(a))f'(a)$. Therefore

$$\frac{g'(x)}{g'(y)} = \frac{h'(f(x))f'(x)}{h'(f(y))f'(y)} \geq \frac{f'(x)}{f'(y)}$$

where the inequality follows because concavity of h and $f(x) < f(y)$ implies $\frac{h'(f(x))}{h'(f(y))} \geq 1$. ■