

Incomplete Information Games with Ambiguity Averse Players*

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

This version: November 14, 2016

Abstract

We study incomplete information games of perfect recall 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 equilibrium concepts satisfying *sequential optimality* – each player’s strategy must be optimal at each stage given the strategies of the other players and the player’s conditional beliefs. We show that for the purpose of identifying strategy profiles that are part of a sequential optimum, it is without loss of generality to restrict attention to beliefs generated using a particular generalization of Bayesian updating. We also propose and analyze 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 GAMES 2016, FUR 2016, RUD 2015, SING11-GTM2015, the Paris conferences “Ambiguity in Games and Mechanisms”, “Time, Uncertainties and Strategies” and D-TEA 2014, and the workshop “Knightian Uncertainty in Strategic Interactions and Markets” at the ZiF, Center for Interdisciplinary Research at Bielefeld University for helpful comments. Klibanoff and Mukerji additionally thank the Robust Finance research group at the ZiF, Bielefeld University for helpful comments and hospitality during their time as visiting fellows in Spring/Summer 2015. This research was partially supported by grant No. 2014350 from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel to Hanany and Klibanoff.

[†]Faculty of Engineering, Tel Aviv University, Tel Aviv 69978, Israel. Email: hananye@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 multistage games with perfect recall, and on equilibrium notions capturing perfection analogous to those in standard theories for ambiguity neutral players, such as subgame perfect equilibrium, 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 parameters, a special case of which are 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. 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).

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

Furthermore, an important method facilitating analysis of dynamic games with standard preferences is the sufficiency of checking only one-stage deviations (as opposed to general deviations) when verifying optimality. 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 does not restrict player i 's beliefs at points where player i is forced to abandon his theory about what has happened in the game so far because of deviation(s) of other player(s). We propose a refinement of sequential optimality restricting such beliefs: sequential equilibrium with ambiguity (SEA). In addition to sequential optimality, SEA imposes a generalization of Kreps and Wilson's (1982) consistency condition from their definition of sequential equilibrium. We show that SEA exists for any finite multistage game with perfect recall and incomplete information, and for any specification of players' ambiguity aversion and initial beliefs. In Appendix B, we describe an alternative refinement of sequential optimality restricted to games with observed actions and private types for each player: perfect equilibrium with ambiguity (PEA). Instead of the consistency condition imposed by SEA, PEA requires 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 (i.e., only one action) available at that stage. PEA is shown to be a weaker refinement than SEA.

In Section 3, we provide several examples that apply our equilibrium notions. These examples are truly strategic in that they include player(s) reacting to the ambiguity aversion of others and changing play because of it. First, we present a game with a path that is played in an SEA given sufficient ambiguity aversion, but, given ambiguity neutrality, is not an ex-ante equilibrium (and thus also not sequentially optimal, an SEA, a PBE or a sequential equilibrium) when players hold common beliefs. Second, 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 Appendix B, we consider a game with a path that is played in a PEA given sufficient ambiguity aversion, but is never played in a PEA given ambiguity neutrality.

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 the strategic use of 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 an SEA. Even establishing that such behavior is sequentially optimal 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). Ellis (2016) examines dynamic consistency between the ex-ante and interim stages of these static games and shows that a version of the folk result that dynamic consistency plus consequentialism implies probabilistic beliefs in an individual choice context can be extended to this game setting. 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, including through an example, 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, finite multistage games with incomplete information and perfect recall where players have (weakly) ambiguity averse smooth ambiguity preferences. It is on this domain that we develop and apply our equilibrium concepts. Such games allow for both imperfectly observed actions and private observations as the game proceeds. Other than perfect recall and finiteness, the multistage structure (i.e., the assumption that all players move simultaneously at each point) is the additional potential limitation on the game forms we consider. While not entirely without loss of generality,

if one doesn't object to giving a player singleton action sets at stages where this player has no "real" move, the multistage assumption is not restrictive. Note that (finite) normal form games with incomplete information and (weakly) ambiguity averse smooth ambiguity preferences are the special case where there is a single stage (i.e., $T = 0$).

Definition 2.1 *A (finite) extensive-form multistage game with incomplete information and perfect recall and (weakly) ambiguity averse smooth ambiguity preferences, Γ , is a tuple $(N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$ where:*

- N is a finite set of players.
- H is a finite set of histories, each of which is a finite sequence of length $T + 2$ of the form $h = (h_{-1}, (h_{0,i})_{i \in N}, \dots, (h_{T,i})_{i \in N})$. For $0 \leq t \leq T + 1$, let $H^t \equiv \{h^t \equiv (h_{-1}, (h_{0,j})_{j \in N}, \dots, (h_{t-1,j})_{j \in N}) \mid h \in H\}$ be the set of partial histories up to (but not including) stage t . The set of all partial histories is $\mathcal{H} \equiv \{\emptyset\} \cup \bigcup_{0 \leq t \leq T+1} H^t$. Given a partial history $\eta \in \mathcal{H}$ and $-1 \leq t \leq |\eta| - 1$, η_t denotes the element of η indexed by t , $\eta_{t,i}$ is player i 's component of η_t and η^t is the partial history up to but not including η_t . For each $i \in N$, $0 \leq t \leq T$ and $h^t \in H^t$, $A_i(h^t) \equiv \{\hat{h}_{t,i} \mid \hat{h} \in H, \hat{h}^t = h^t\}$ is the set of actions available to player i at h^t . The set of initial moves or "types" is $\Theta \equiv H^0$.
- $\mathcal{I}_i \equiv \bigcup_{0 \leq t \leq T} \mathcal{I}_i^t$ are the information sets for player i , where each \mathcal{I}_i^t is a partition of H^t such that, for all $\eta^t, \hat{\eta}^t \in H^t$, $\hat{\eta}^t \in \mathcal{I}_i(\eta^t)$ implies $A_i(\eta^t) = A_i(\hat{\eta}^t)$ (where $\mathcal{I}_i(\eta^t)$ is the unique element of \mathcal{I}_i^t such that $\eta^t \in \mathcal{I}_i(\eta^t)$). For $0 \leq t \leq T$, $\eta^t \in H^t$ and $i \in N$, $R_i(\eta^t) \equiv ((\mathcal{I}_i(\eta^s), \eta_{s,i}^t)_{0 \leq s < t}, \mathcal{I}_i(\eta^t))$ is the ordered list of information sets i encounters and the actions i takes under partial history η^t . The game satisfies perfect recall in that, for each player i , $0 \leq t \leq T$ and $\eta^t, \hat{\eta}^t \in H^t$, $\mathcal{I}_i(\eta^t) = \mathcal{I}_i(\hat{\eta}^t)$ implies $R_i(\eta^t) = R_i(\hat{\eta}^t)$. Extend R_i and A_i to information sets in the natural way.
- μ_i is a probability over $\Delta(\Theta)$ having finite support such that $\sum_{\pi \in \Delta(\Theta)} \mu_i(\pi) \pi(\theta) > 0$ for all $i \in N$ and $\theta \in \Theta$, where $\Delta(\Theta)$ is the set of all probability measures over Θ .
- $u_i : H \rightarrow \mathbb{R}$ is the utility payoff of player i given the history.
- $\phi_i : \text{co}(u_i(H)) \rightarrow \mathbb{R}$ is a continuously differentiable, concave and strictly increasing function.

The first three bullet points in Definition 2.1 describe the game form. The only notable non-standard parts of this definition are the specification of preferences. In particular, ϕ_i and μ_i are part of the specification of smooth ambiguity preferences, with the degree of concavity

of ϕ_i reflecting ambiguity aversion and μ_i indicates the presence of ambiguity when it has multiple probability measures in its support.

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 $c_i : H \rightarrow Z$ 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 outcomes, $w_i : Z \rightarrow \mathbb{R}$. Assume that Z is big enough so that $u_i(H)$ is interior in $w_i(Z)$.²

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

Definition 2.2 *A (behavior) strategy for player i in a game Γ is a function σ_i such that $\sigma_i(I_i) \in \Delta(A_i(I_i))$ for each $I_i \in \mathcal{I}_i$.*

Let Σ_i denote the set of all strategies for player i . A strategy profile, $\sigma \equiv (\sigma_i)_{i \in N}$, is a strategy for each player. Given a strategy σ_i for player i , the continuation strategy at information set $I_i, \sigma_i^{I_i}$, is the restriction of σ_i to the information sets \hat{I}_i such that $I_i \in R_i(\hat{I}_i)$.

Given a strategy profile σ , history h and $0 \leq r \leq t \leq T + 1$, the probability of reaching h^t starting from h^r is $p_\sigma(h^t|h^r) \equiv \prod_{j \in N} \prod_{r \leq s < t} \sigma_j(I_j(h^s))(h_{s,j})$.³ It is useful to separate this probability into a part affected only by σ_i and a part affected only by σ_{-i} . These are $p_{i,\sigma_i}(h^t|h^r) \equiv \prod_{r \leq s < t} \sigma_i(I_i(h^s))(h_{s,i})$ and $p_{-i,\sigma_{-i}}(h^t|h^r) = \prod_{j \neq i} \prod_{r \leq s < t} \sigma_j(I_j(h^s))(h_{s,j})$ respectively, with $p_{i,\sigma_i}(h^t|h^r)p_{-i,\sigma_{-i}}(h^t|h^r) = p_\sigma(h^t|h^r)$. With this notation, we can now state formally the assumption that players ex-ante preferences over strategies are smooth ambiguity preferences (Klibanoff, Marinacci and Mukerji 2005) with the u_i, ϕ_i and μ_i as specified by the game.

Assumption 2.1 *Fix a game Γ . Ex-ante (i.e., given the empty partial history), each player i ranks strategy profiles σ according to*

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

²This will be convenient for some later optimality characterizations, the proofs of Theorem 2.3 and Lemma A.1 in particular.

³If $r = t$, so that the product is taken over an empty set, invoke the convention that a product over an empty set is 1.

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

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

$$V_i(\sigma^*) \geq V_i(\sigma'_i, \sigma^*_{-i})$$

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

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

We next turn to defining preferences beyond the ex-ante stage. In order to formulate interim beliefs, fixing a strategy profile, σ , it is useful to consider when one information set of player i is reachable from another without requiring a deviation from σ by players other than i .

Definition 2.4 For information set I_i , define $s(I_i)$ to be such that $I_i \in \mathcal{I}_i^{s(I_i)}$.

Definition 2.5 Information set I_i is reachable from information set \hat{I}_i given strategies σ_{-i} via partial history $h^t \in I_i$ if $h^{s(\hat{I}_i)} \in \hat{I}_i$ and $p_{-i, \sigma_{-i}}(h^t | h^{s(\hat{I}_i)}) > 0$. Say that I_i is reachable from \hat{I}_i given σ_{-i} whenever such a partial history exists.

It is possible, for some information set, that no partial histories make it reachable from the beginning of the game assuming other players follow σ_{-i} . At the first information set for which this happens (if it does), player i is forced to abandon his theory (i.e., the part of σ_{-i}) concerning the behavior of i 's opponents up to this stage and adopt some new theory. Call this information set \hat{I}_i . At the first subsequent information set not reachable from \hat{I}_i assuming other players follow σ_{-i} , player i would again be forced to abandon his theory, now consisting of some interim beliefs and the part of σ_{-i} specifying behavior following \hat{I}_i up to this stage. It will be useful in what follows to keep track of the most recent information set where i was forced to adopt a new theory. The function f_i , defined recursively below, records this for each information set of player i , where we adopt the convention that if i 's theory has never yet been falsified then f_i records player i 's initial information set.

Notation 2.1 For information set $I_i \not\subseteq \Theta$, define I_i^{-1} to be the information set immediately preceding I_i in $R_i(I_i)$.

Definition 2.6 Given information set I_i and strategies σ_{-i} , define $f_i(I_i)$ as follows: For $I_i \subseteq \Theta$, $f_i(I_i) = I_i$; for $I_i \not\subseteq \Theta$, if I_i is reachable from $f_i(I_i^{-1})$ given σ_{-i} , $f_i(I_i) = f_i(I_i^{-1})$, otherwise $f_i(I_i) = I_i$. Define $m_i(I_i) = s(f_i(I_i))$.

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

Definition 2.7 $I_{i,\sigma_{-i}} \equiv \{h^t \in I_i \mid I_i \text{ is reachable from } f_i(I_i) \text{ given } \sigma_{-i} \text{ via } h^t\}$ and $f_i^{\sigma_{-i}}(I_i) \equiv \{h^{m_i(I_i)} \in f_i(I_i) \mid I_i \text{ is reachable from } f_i(I_i) \text{ given } \sigma_{-i} \text{ via } h^{s(I_i)}\}$.

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

Definition 2.8 An interim belief for player i in a game Γ given information set I_i and strategy profile σ is a finite support probability measure ν_{i,I_i} over $\Delta(H^{m_i(I_i)})$ such that

$$\nu_{i,I_i}(\{\pi \mid \pi(f_i^{\sigma_{-i}}(I_i)) > 0\}) = 1. \quad (2.2)$$

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

Since the space $\Delta(H^{m_i(I_i)})$ over which interim beliefs are defined depends on $f_i(I_i)$, which in turn depends on the strategy profile σ , it will be useful, when defining our generalization of Kreps and Wilson's (1982) consistency condition in Section 2.1, to adapt an interim belief system so that it becomes a system of beliefs defined over $\Delta(I_i)$ for each i and I_i , a structure that is independent of the strategy profile. This is done via the mapping in the following definition, the idea of which is to condition each measure in the support of the interim belief ν_{i,I_i} on both the information set I_i and the strategies of i 's opponents, σ_{-i} (note that $\text{supp}()$ denotes the support of the measure in its argument):

Definition 2.9 For a strategy profile σ and an interim belief system ν , say that $\hat{\nu} \equiv (\hat{\nu}_{i,I_i})_{i \in N, I_i \in \mathcal{I}_i}$ is adapted from ν using σ if, for each player i , information set $I_i \in \mathcal{I}_i$ and $\pi \in \Delta(I_i)$, $\hat{\nu}_{i,I_i}(\pi) = \sum_{\hat{\pi} \in \text{supp } \nu_{i,I_i} \mid \hat{\pi}_{I_i, \sigma_{-i}} = \pi} \nu_{i,I_i}(\hat{\pi})$.

Fundamental to our equilibrium notion will be sequential optimality. It requires that each player plays optimally at each information set given the strategies of the others. This optimality is required even when the information set is before-the-fact viewed as unreachable according to the given strategy profile combined with the beliefs of the player. In order to describe optimality for i given I_i , we need to write i 's conditional preferences. These make use of interim beliefs.

Assumption 2.2 *Fix a game Γ , a strategy profile σ and an interim belief system ν . Any player i at information set I_i ranks strategies $\hat{\sigma}_i$ according to*

$$V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i}) \equiv \sum_{\pi \in \Delta(H^{m_i(I_i)} | \pi(f_i^{\sigma_{-i}}(I_i))) > 0} \phi_i \left(\sum_{h|h^t \in I_i} u_i(h) p_{(\hat{\sigma}_i, \sigma_{-i})}(h|h^t) \pi_{I_i, \sigma_{-i}}(h^t) \right) \nu_{i,I_i}(\pi), \quad (2.3)$$

where $t = s(I_i)$ and all the conditionals $\pi_{I_i, \sigma_{-i}} \in \Delta(I_i)$ are related to π via the following Bayes' formula:

For $I_i = f_i(I_i)$,

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

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

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

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 I_i and σ (i.e., μ_i is replaced by ν_{i,I_i}), (2) each π is conditioned on I_i, σ_{-i} , which is the subset of I_i reachable under σ_{-i} from $f_i(I_i)$, the most recent information set where i was forced to adopt a new theory (or from i 's initial information set), and (3) the probabilities of reaching various histories according to the strategy profile are now calculated starting from I_i rather than from the beginning of the game. The possibly multiple conditionals, $\pi_{I_i, \sigma_{-i}}$, may differ only in the relative weights assigned to elements of I_i, σ_{-i} . This, together with multiplication by $p_{-i, \sigma_{-i}}(h|h^t)$, reflects that there is no ambiguity about the part of σ_{-i} specifying behavior from $f_i(I_i)$ onward (i.e., the part of σ_{-i} that has not yet been contradicted). This shows that the ambiguity possibly impacting interim preferences concerns θ and/or opponents' actions prior to $f_i(I_i)$

that are not uniquely identified given $I_{i,\sigma_{-i}}$. Each conditional, $\pi_{I_{i,\sigma_{-i}}}$, is calculated using Bayes' formula (see the Remark below). In contrast, no restriction has been placed on the ν_{i,I_i} other than (2.2).

Remark 2.1 To see that the formula for $\pi_{I_{i,\sigma_{-i}}}$ in (2.5) is Bayesian conditioning, note that, because $p_{i,\sigma_i}(h^t|h^{t-1})$ is the same for all $h^t \in I_{i,\sigma_{-i}}$,

$$\frac{p_{-i,\sigma_{-i}}(h^t|h^{t-1})\pi_{I_{i,\sigma_{-i}}}^{-1}(h^{t-1})}{\sum_{\hat{h}|\hat{h}^t \in I_{i,\sigma_{-i}}} p_{-i,\sigma_{-i}}(\hat{h}^t|\hat{h}^{t-1})\pi_{I_{i,\sigma_{-i}}}^{-1}(\hat{h}^{t-1})} = \frac{p_{\sigma}(h^t|h^{t-1})\pi_{I_{i,\sigma_{-i}}}^{-1}(h^{t-1})}{\sum_{\hat{h}|\hat{h}^t \in I_{i,\sigma_{-i}}} p_{\sigma}(\hat{h}^t|\hat{h}^{t-1})\pi_{I_{i,\sigma_{-i}}}^{-1}(\hat{h}^{t-1})}$$

if $h^t \in I_{i,\sigma_{-i}}$ and $p_{i,\sigma_i}(h^t|h^{t-1}) > 0$.

An equivalent formulation applies Bayesian conditioning in a ‘‘all-at-once’’ manner from $f_i(I_i)$ to the present:

$$\pi_{I_{i,\sigma_{-i}}}(h^t) = \frac{p_{-i,\sigma_{-i}}(h^t|h^{m_i(I_i)})\pi_{f_i(I_i)_{i,\sigma_{-i}}}(h^{m_i(I_i)})}{\sum_{\hat{h}|\hat{h}^t \in I_{i,\sigma_{-i}}} p_{-i,\sigma_{-i}}(\hat{h}^t|\hat{h}^{m_i(I_i)})\pi_{f_i(I_i)_{i,\sigma_{-i}}}(\hat{h}^{m_i(I_i)})} \text{ if } h^t \in I_{i,\sigma_{-i}} \text{ and } 0 \text{ otherwise.} \quad (2.6)$$

Using these preferences, we may now define sequential optimality:

Definition 2.10 Fix a game Γ . A pair (σ, ν) consisting of a strategy profile and interim belief system is sequentially optimal if, for all players i and all information sets I_i ,

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

and

$$V_{i,I_i}(\sigma) \geq V_{i,I_i}(\sigma'_i, \sigma_{-i}) \quad (2.8)$$

for all $\sigma'_i \in \Sigma_i$, where the V_i and V_{i,I_i} are as specified in (2.1) and (2.3).

Note that since $V_{i,I_i}(\tilde{\sigma}_i, \sigma_{-i}) = V_{i,I_i}(\hat{\sigma}_i, \sigma_{-i})$ if $\tilde{\sigma}_i^{I_i} = \hat{\sigma}_i^{I_i}$, requiring the inequalities for the V_{i,I_i} to hold as i changes only her continuation strategy given I_i would result in an equivalent definition. A strategy profile σ is said to be *sequentially optimal* whenever there exists an interim belief system ν such that (σ, ν) is sequentially optimal.

Assuming a common μ , sequential optimality implies subgame perfection adapted to allow for smooth ambiguity preferences. The only proper subgames start from partial histories η where $I_i(\eta) = \{\eta\}$ for all players i (i.e., all uncertainty (if any) about the past has been

resolved). For any such proper subgame, (2.8) ensures that the continuation strategy profile derived from σ forms a Nash equilibrium of the subgame. For the overall game, (2.7) ensures σ is an ex-ante equilibrium.

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

Definition 2.11 *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 information set I_i , letting $t = s(I_i)$: If $I_i \subseteq \Theta$, for all π such that $\pi(I_i) > 0$,*

$$\nu_{i,I_i}(\pi) \propto \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi_{I_i, \sigma_{-i}}(h^t) \right)} \pi(I_i) \mu_i(\pi); \quad (2.9)$$

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

$$\begin{aligned} \nu_{i,I_i}(\pi) \propto & \frac{\phi'_i \left(\sum_{h|h^{t-1} \in I_i^{-1}} u_i(h) p_\sigma(h|h^{t-1}) \pi_{I_i, \sigma_{-i}}^{-1}(h^{t-1}) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi_{I_i, \sigma_{-i}}(h^t) \right)} \\ & \cdot \left(\sum_{h|h^t \in I_i, \sigma_{-i}} p_{-i, \sigma_{-i}}(h^t|h^{t-1}) \pi_{I_i, \sigma_{-i}}^{-1}(h^{t-1}) \right) \nu_{i,I_i^{-1}}(\pi). \end{aligned} \quad (2.10)$$

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 becomes standard Bayesian updating of μ_i and the ν_{i,I_i} , applied whenever possible. More generally, the ϕ'_i ratio terms, which reflect changes in the motive to hedge against ambiguity (see Hanany and Klibanoff 2009 and Baliga, Hanany and Klibanoff 2013), are the only difference from Bayesian updating. These changes can be motivated via dynamic consistency. For ambiguity averse preferences, Bayesian updating does not ensure dynamic consistency. The smooth rule

is dynamically consistent for all ambiguity averse smooth ambiguity preferences (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, ν) is sequentially optimal. Then, there exists an interim belief system $\hat{\nu}$ satisfying the smooth rule using σ as the ex-ante equilibrium such that $(\sigma, \hat{\nu})$ 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. The version of perfect Bayesian equilibrium (PBE) in, for example, Gibbons (1992) imposes sequential optimality (defined using only expected utility preferences) and also that beliefs are related via Bayesian updating wherever possible. From our theorem, it follows that in the expected utility case, sequential optimality alone (i.e., without additionally requiring Bayesian updating) identifies the same set of strategy profiles as sequential optimality plus Bayesian updating. This Bayesian version of the result was 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 information sets. In particular, all ex-ante equilibria for which all information sets are on-path can be supported as sequentially optimal.

Theorem 2.2 *Fix a game Γ . Suppose σ is an ex-ante equilibrium and, for each player i , each of i 's information sets is reachable from some information set in \mathcal{I}_i^0 given σ_{-i} (i.e., $m_i(I_i) = 0$ for all I_i). Then, there exists an interim belief system ν such that (σ, ν) is sequentially optimal.*

We next provide conditions under which (σ, ν) is sequentially optimal if, for each player i and information set I_i , there are no deviations by i at I_i alone that are desirable according to i 's preferences at I_i . These “one-stage” deviations are typically a small fraction of the deviations available to players. Formally, the absence of these profitable one-stage deviations is the following property:

Definition 2.12 *The pair (σ, ν) satisfies the no profitable one-stage deviation property if for each player i and each information set I_i , $V_{i,I_i}(\sigma) \geq V_{i,I_i}(\sigma'_i, \sigma_{-i})$ for all σ'_i agreeing with σ_i everywhere except possibly at I_i .*

To describe when checking one-stage deviations is sufficient, it is helpful to define smooth rule updating given a strategy profile σ , even when σ is not necessarily an ex-ante equilibrium.

Definition 2.13 *An interim belief system ν satisfies extended smooth rule updating using σ as the strategy profile if ν satisfies the conditions in Definition 2.11 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. As a result, fixing a strategy profile σ , at any information set one may check for any player i whether σ_i is an optimal strategy for i for the remainder of the game using a “folding back” algorithm. It works as follows – for each information set for player i at the final stage, calculate an optimal mixture over the actions she has available at that stage given the information set, the strategies of the other players and i 's beliefs at that information set. Then, holding these optimal mixtures fixed, repeat this process for information sets one stage earlier in the game. Continue backwards iteratively. The only thing that (in common with the standard approach under ambiguity neutrality) cannot be calculated via folding back are the beliefs at each information set. These may be determined by updating according to the smooth rule (which reduces to Bayes' rule under ambiguity neutrality). Recall that smooth rule updating is without loss of generality for the purposes of identifying sequentially optimal strategies.

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

2.1 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 where player i is forced to abandon his theory about what has happened in the game so far because of deviation(s) of other player(s). Formally, these are information sets $I_i \not\subseteq \Theta$ such that $f_i(I_i) = I_i$. The condition extends Kreps and Wilson's (1982) *consistency* condition that they use in defining sequential equilibrium. We extend consistency in order to accommodate ambiguity aversion by replacing Bayes' rule in their definition with the (extended) smooth rule (and adapting from ν using σ – see Definition 2.9). 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 2.14 *Fix a game Γ . A pair (σ, ν) consisting of a strategy profile and interim belief system satisfies smooth rule consistency if ν is determined by extended smooth rule updating using σ as the strategy profile, and there exists a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$, with $\lim_{k \rightarrow \infty} \sigma^k = \sigma$, such that $\hat{\nu} = \lim_{k \rightarrow \infty} \hat{\nu}^k$, where $\hat{\nu}$ is adapted from ν using σ , each $\hat{\nu}^k$ is adapted from ν^k using σ^k , and each ν^k is the interim belief system determined by extended smooth rule updating using σ^k as the strategy profile.*

Definition 2.15 *A sequential equilibrium with ambiguity (SEA) of a game Γ is a pair (σ, ν) consisting of a strategy profile and interim belief system such that (σ, ν) is sequentially optimal and satisfies smooth rule consistency.*

By definition, any SEA is sequentially optimal. In general, a sequential optimum might not be an SEA. However, if all information sets 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 for each player i , each of i 's information sets is reachable from some information set in \mathcal{I}_i^0 given σ_{-i} (i.e., $m_i(I_i) = 0$ for all I_i). Then, there exists an interim belief system $\hat{\nu}$ such that $(\sigma, \hat{\nu})$ is an SEA.*

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 may use 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 equilibrium strategies.

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

2.2 Comparative statics in ambiguity aversion

In this section, we explore the extent to which changes in ambiguity aversion affect equilibrium play.

Definition 2.16 *For a game $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$, let $E_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ be the set of all ex-ante equilibria of the game $\hat{\Gamma} = (N, H, (\mathcal{I}_i)_{i \in N}, (\hat{\mu}_i)_{i \in N}, (u_i, \hat{\phi}_i)_{i \in N})$ differing from Γ only in ambiguity aversions and beliefs. Let $Q_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ be the analogous set of sequentially optimal strategy profiles and $S_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ be the analogous set of SEA strategy profiles.*

Notation 2.2 *Denote the identity function by ι .*

We start with the simplest and most direct comparative static question: Do changes in ambiguity aversion affect the set of equilibrium strategy profiles (and play paths) for fixed beliefs? The answer is yes they can. The following result demonstrates this and also shows that the change need not be a simple expansion or contraction of the equilibrium set as ambiguity aversion increases.

Theorem 2.6 *There exists a game Γ and $(\hat{\phi}_i)_{i \in N}$ such that $E_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) \cap E_\Gamma((\mu_i)_{i \in N}, (\iota)_{i \in N}) = \emptyset$.*

Examination of the proof of Theorem 2.6 shows that, fixing beliefs, not only are the equilibrium strategies distinct under ambiguity aversion compared to ambiguity neutrality, but it can also be that the strategies under ambiguity aversion generate paths of play that do not occur in equilibrium under ambiguity neutrality. An analogue of Theorem 2.6 is true for sequential optima, SEA and any other refinement of ex-ante equilibria as well, as they are all ex-ante equilibria. Thus, with fixed beliefs, change in ambiguity aversion can impact the set of equilibrium strategies and realized play.

The result above requires at least one player to perceive some ambiguity. As intuition would suggest, in games where all beliefs μ_i are degenerate so that there is no ambiguity, ambiguity aversion does not affect the set of equilibria:

Theorem 2.7 Fix a game Γ with degenerate μ_i for each player. For all $(\hat{\phi}_i)_{i \in N}$, $E_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) = E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$, $Q_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) = Q_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and $S_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) = S_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$.

Further examination of the proof of Theorem 2.6 shows that ambiguity aversion continues to affect the equilibrium set even if we impose *common beliefs* (i.e., $\mu_i = \mu$ for all players i). We next explore introducing ambiguity aversion under common beliefs while dropping the assumption that these beliefs are fixed when we change ambiguity aversion. We ask whether ambiguity aversion plus the assumption of common beliefs has equilibrium implications that are different from ambiguity neutrality plus the assumption of common beliefs. The answer again is yes. In fact, under the assumption of common beliefs, we show that ambiguity aversion always weakly expands the set of equilibria compared to ambiguity neutrality and may do so strictly:

Theorem 2.8 There exists a game Γ and $(\hat{\phi}_i)_{i \in N}$ such that $\bigcup_{\hat{\mu}} S_\Gamma((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supset \bigcup_{\hat{\mu}} E_\Gamma((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ and some of the new equilibrium strategies induce new paths of play. For all games Γ and $(\hat{\phi}_i)_{i \in N}$, $\bigcup_{\hat{\mu}} E_\Gamma((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} E_\Gamma((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$, $\bigcup_{\hat{\mu}} Q_\Gamma((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} Q_\Gamma((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ and $\bigcup_{\hat{\mu}} S_\Gamma((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} S_\Gamma((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$.

Thus, under an assumption of common beliefs, ambiguity aversion may generate new equilibrium behavior (and new paths of play) and does not eliminate possible equilibria compared to ambiguity neutrality.

We now turn to the level of greatest generality by allowing complete freedom to change beliefs (possibly different for each player) when changing ambiguity aversion. In such a circumstance, are all equilibria under ambiguity neutrality still equilibria under any specified ambiguity aversion? The answer is yes, because one can always specify degenerate $\hat{\mu}_i$ with support consisting of the reduced measure $\sum_{\pi \in \Delta(\Theta)} \mu_i(\pi) \pi(\theta)$ and, for $I_i \not\subseteq \Theta$ such that $f_i(I_i) = I_i$, degenerate interim beliefs $\hat{\nu}_{i, I_i}$ with support equal to $\sum_{\pi \in \Delta(\Theta)} \nu_{i, I_i}(\pi) \pi(\theta)$. Less obviously, a similar result in the opposite direction holds as well. Putting the two directions together yields the next result of this section: All equilibrium strategy profiles (ex-ante or sequentially optimal or SEA) of a game are still equilibrium profiles when ambiguity aversion(s) change (possibly, but not necessarily, to ambiguity neutrality) for some modified beliefs $\hat{\mu}_i$. Our constructive proof shows that the beliefs $\hat{\mu}_i$ and interim beliefs $\hat{\nu}_{i, I_i}$ that work for a given equilibrium profile σ are related to the beliefs μ_i and ν_{i, I_i} in the game with the original ambiguity aversion(s) by the formulae $\hat{\mu}_i \propto \frac{\phi'_i(\sum_{\hat{h}} u_i(\hat{h}) p_\sigma(\hat{h} | \hat{h}^0) \pi(\hat{h}^0))}{\hat{\phi}'_i(\sum_{\hat{h}} u_i(\hat{h}) p_\sigma(\hat{h} | \hat{h}^0) \pi(\hat{h}^0))} \mu_i$ and $\hat{\nu}_{i, I_i} \propto \frac{\phi'_i(\sum_{\hat{h} | \hat{h}^s(I_i) \in I_i} u_i(\hat{h}) p_\sigma(\hat{h} | \hat{h}^s(I_i)) \pi_{I_i, \sigma_{-i}}(\hat{h}^s(I_i)))}{\hat{\phi}'_i(\sum_{\hat{h} | \hat{h}^s(I_i) \in I_i} u_i(\hat{h}) p_\sigma(\hat{h} | \hat{h}^s(I_i)) \pi_{I_i, \sigma_{-i}}(\hat{h}^s(I_i)))} \nu_{i, I_i}$ where the ϕ_i are the original and $\hat{\phi}_i$ the new specifications of ambiguity aversions. Formally the result is:

Theorem 2.9 Fix a game Γ . For all $(\hat{\phi}_i)_{i \in N}$, $\bigcup_{(\hat{\mu}_i)_{i \in N}} E_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) = \bigcup_{(\hat{\mu}_i)_{i \in N}} E_\Gamma((\hat{\mu}_i)_{i \in N}, (\phi_i)_{i \in N})$, $\bigcup_{(\hat{\mu}_i)_{i \in N}} Q_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) = \bigcup_{(\hat{\mu}_i)_{i \in N}} Q_\Gamma((\hat{\mu}_i)_{i \in N}, (\phi_i)_{i \in N})$ and $\bigcup_{(\hat{\mu}_i)_{i \in N}} S_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) = \bigcup_{(\hat{\mu}_i)_{i \in N}} S_\Gamma((\hat{\mu}_i)_{i \in N}, (\phi_i)_{i \in N})$.

Theorem 2.9 shows that for ex-ante equilibrium, sequential optimum and SEA, when ranging over all possible beliefs μ_i for each player, there is no change in the set of (ex-ante or sequentially optimal or SEA) equilibrium strategy profiles as ambiguity aversion changes. For a result that in individual decision problems, under standard assumptions (including reduction, broad framing, statewise dominance and expected utility evaluation of objective lotteries), all observed behavior optimal according to ambiguity averse preferences is also optimal for *some* subjective expected utility preferences, see e.g., Kuzmics (2015). Bade (2016) independently shows a result very similar to invariance of the set of ex-ante equilibrium to ambiguity aversion. Our construction relies on changing the beliefs of different players in different ways. In light of our earlier ambiguity aversion comparative static for the case of common beliefs (Theorem 2.8), the equivalence in Theorem 2.9 depends crucially on diversity of beliefs among the players in $\hat{\Gamma}$. If $\hat{\mu}_i$ is required to be the same for all i , the theorem fails to hold (even if μ_i is the same for all i as well).

The result in Theorem 2.9 is reminiscent of a result in Battigalli et al. (2015a, p. 667) showing that when mixed strategies are allowed (as here), the set of their smooth self-confirming equilibria does not change as ambiguity aversion changes. In contrast, the result they emphasize (their Theorem 1 together with an example of strict inclusion), in which they show that the set of self-confirming equilibria increases as ambiguity aversion increases and that this increase can be strict, relies crucially on limiting attention to pure strategies (both in terms of the equilibrium profile and in terms of the deviations against which optimality is checked). If we were also to limit attention to pure strategies in both these respects, would an analogous comparative static to that in Battigalli et al. (2015a) apply to the sets of strategy profiles that are ex-ante equilibria or sequentially optimal? The answer is yes, as we now show. To this end modify the equilibrium set notation to restrict attention to pure strategies as follows.

Definition 2.17 For a game $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$, let $\tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ be the set of all ex-ante equilibria with respect to pure strategies of a game $\hat{\Gamma} = (N, H, (\mathcal{I}_i)_{i \in N}, (\hat{\mu}_i)_{i \in N}, (u_i, \hat{\phi}_i)_{i \in N})$ differing from Γ only in ambiguity aversions and beliefs. Let $\tilde{Q}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ be the analogous set of sequentially optimal strategy profiles with respect to pure strategies.

Theorem 2.10 Fix a game Γ . For all $(\hat{\phi}_i)_{i \in N}$ such that, for each i , $\hat{\phi}_i$ is at least as concave as ϕ_i , $\tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ and $\tilde{Q}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{Q}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

To complete the argument, we observe that the set of strategy profiles that are ex-ante equilibria (or sequentially optimal) with respect to pure strategies can strictly increase:

Theorem 2.11 There exists a game Γ and $(\hat{\phi}_i)_{i \in N}$ such that for each i , $\hat{\phi}_i$ is at least as concave as ϕ_i , $\bigcup_{(\mu_i)_{i \in N}} \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subset \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ and $\bigcup_{(\mu_i)_{i \in N}} \tilde{Q}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subset \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{Q}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

2.3 Robustness

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

We propose definitions of robustness to ambiguity aversion and belief robustness and show that robustness to ambiguity aversion implies increased belief robustness.

Definition 2.18 For a game $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$, an equilibrium $\sigma \in E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ is ex-ante robust to increased ambiguity aversion if it remains an equilibrium whenever, for each i , ϕ_i is replaced by an at least as concave $\hat{\phi}_i$, i.e. $\sigma \in E_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

Definition 2.19 For a game $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$, consider an equilibrium $\sigma \in E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$. Ambiguity aversion makes σ ex-ante more belief robust if there exist, for each i , $\hat{\phi}_i$ strictly more concave than ϕ_i such that, for each $(\hat{\mu}_i)_{i \in N}$ with the same supports as the $(\mu_i)_{i \in N}$, $\sigma \in E_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ whenever $\sigma \in E_\Gamma((\hat{\mu}_i)_{i \in N}, (\phi_i)_{i \in N})$. It makes σ ex-ante strictly more belief robust if, in addition, there exist $(\tilde{\mu}_i)_{i \in N}$ with the same supports as the $(\mu_i)_{i \in N}$ such that $\sigma \in E_\Gamma((\tilde{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ but $\sigma \notin E_\Gamma((\tilde{\mu}_i)_{i \in N}, (\phi_i)_{i \in N})$.

Sequentially optimal robust to increased ambiguity aversion and Sequentially optimal more belief robust are defined analogously, replacing E_Γ with Q_Γ in the above two definitions.

SEA robust to increased ambiguity aversion and SEA (strictly) more belief robust are defined by, in addition to replacing E_Γ with S_Γ in each definition, requiring that there exists a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$, such that $\lim_{k \rightarrow \infty} \sigma^k = \sigma$, with respect to which smooth rule consistency for the assumed SEAs are simultaneously satisfied.

Theorem 2.12 Fix a game $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i, \phi_i)_{i \in N})$. If an equilibrium σ is ex-ante (resp. sequentially optimal or SEA) robust to increased ambiguity aversion and, for

each player i , $\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0)$ has a unique minimizing π over the support of μ_i , then ambiguity aversion makes σ ex-ante (resp. sequentially optimal or SEA) more belief robust.

If, additionally, there exist $(\tilde{\mu}_i)_{i \in N}$, with the same supports as the $(\mu_i)_{i \in N}$, such that σ is not an equilibrium given $(\tilde{\mu}_i)_{i \in N}$, and, for each player i , $\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0)$ can be strictly ordered across the π in the support of μ_i , then ambiguity aversion makes σ ex-ante (resp. sequentially optimal or SEA) strictly more belief robust.

Note that without the unique minimizer condition the first part of the theorem would be false. Similarly, even with the unique minimizer condition, without the strict ordering condition the second part of the theorem would be false.

The proof of Theorem 2.12 makes use of the following first-order characterization of optimality of σ_i :

The requirement that σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i$ and $\hat{\mu}_i$ is equivalent to (see (A.4))

$$\sigma_i \in \arg \max_{\sigma'_i} \sum_{\pi \in \Pi_i} \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi(h^0) \right) \hat{\phi}'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h|h^0) \pi(h^0) \right) \hat{\mu}_i(\pi). \quad (2.11)$$

Similarly, the requirement that σ_i is an interim best response to σ_{-i} for player i at information set I_i given $\hat{\phi}_i$ and $\hat{\nu}_{i, I_i}$ is equivalent to (see (A.3))

$$\sigma_i \in \arg \max_{\sigma'_i} \sum_{\pi | \pi(J_i^{\sigma_{-i}}(I_i)) > 0} \left(\sum_{h | h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^{s(I_i)}) \pi_{I_i, \sigma_{-i}}(h^{s(I_i)}) \right) \cdot \hat{\phi}'_i \left(\sum_{h | h^{s(I_i)} \in I_i} u_i(h) p_\sigma(h|h^{s(I_i)}) \pi_{I_i, \sigma_{-i}}(h^{s(I_i)}) \right) \hat{\nu}_{i, I_i}(\pi). \quad (2.12)$$

The proof applies (2.11) and (2.12) using the following particularly convenient set of $\hat{\phi}_i$ strictly more concave than ϕ_i :

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

Definition 2.20 For any $l \geq 1$ such that e_i^{l+1} exists, let $\hat{\phi}_i^l \equiv \psi_i^l \circ \phi_i$, where ψ_i^l is defined by

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

with $b > 1$.

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

3 Examples

In this section, we present 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 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 Example 1: New Strategic Behavior in Equilibrium: 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, it is determined whether player 1 is of type I or type II and 1 observes her own type. Players 2 and 3 have only one type, so there is complete information about them. The payoff triples in Figure 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

⁴Note that to eliminate any possible effects of varying players' risk aversion, think of the payoffs being




Figure 3.1: SEA Example

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

As the proof of Proposition 3.2 mentions, the example μ is chosen for simplicity, and degeneracy of the measures in its support is not necessary for the result.

3.2 Example 2: Limit Pricing under Ambiguity

In this section we use a parametric class of games based on the Milgrom and Roberts (1982) limit pricing entry model and show that ambiguity aversion may make limit pricing behavior more robust. These games have an SEA involving limit pricing and non-trivial updating on the equilibrium path that departs from Bayes rule. 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. 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).^{5,6} 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

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.

⁵The 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.

⁶The 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.

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 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 (2.7) 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 deterring 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

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

This is equivalent to

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

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

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

This is equivalent to

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

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 $\sum_{\pi} \mu(\pi)\pi(\tau) > 0$ for $\tau \in \{L, M, H\}$. Denote the entrant's Cournot profit net of entry costs when facing an incumbent of type τ by $w_{\tau} \equiv b(\frac{a+c_{\tau}-2c_E}{3b})^2 - K$. 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 (2.7) of sequential optimality if and only if the four inequalities (3.1), (3.2), (3.4) and (3.5) are satisfied.

The following proposition provides sufficient conditions for the existence of an 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 3.3 *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 SEA 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 Use of Ambiguity

At first glance, since the only source of on-path ambiguity in our framework is ambiguity about types, one might think that this is too restrictive to allow for a player to perceive ambiguity about the strategies of others. However, since strategies are type-contingent, ambiguity about types directly translates into ambiguity about players' actions. Thus, even though we examine equilibrium concepts where players take as given the behavior strategies of their opponents, there is still ambiguity about other players' actions in equilibrium.

In this section, we go further and show that our framework can also address the strategic use of ambiguity – intentionally choosing strategies that will be perceived as ambiguous by others. 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: since players' payoff functions may depend on types, generating strategic ambiguity through conditioning play on types 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 types 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 ambiguous strategies and normal form games by Riedel and Sass (2013).

Formally, consider placing the following structure on the finite set of types: $\Theta \equiv \Theta^{\mathcal{U}} \times \Theta^{\mathcal{A}}$ (with generic element $\theta \equiv (\theta^{\mathcal{U}}, \theta^{\mathcal{A}})$) and the utility payoff function for each player i may depend on types only through the $\Theta^{\mathcal{U}}$ component of the type space, i.e., for all $h, \hat{h} \in H$ and $\theta, \hat{\theta} \in \Theta$, if $h^0 = \theta$, $\hat{h}^0 = \hat{\theta}$, $\theta^{\mathcal{U}} = \hat{\theta}^{\mathcal{U}}$ and $h_{t,i} = \hat{h}_{t,i}$ for all $i \in N$ and all $0 \leq t \leq T + 1$, then $u_i(h) = u_i(\hat{h})$.

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

Definition 4.1 *A strategy for player i , σ_i , is an Ellsberg strategy if $\sigma_i(I_i) \neq \sigma_i(\hat{I}_i)$ for some $I_i, \hat{I}_i \in \mathcal{I}_i$ such that*

$$\{(\eta_{-1}^{\mathcal{U}}, (\eta_{0,j})_{j \in N}, \dots, (\eta_{s(I_i)-1,j})_{j \in N}) \mid \eta \in I_i\} = \{(\eta_{-1}^{\mathcal{U}}, (\eta_{0,j})_{j \in N}, \dots, (\eta_{s(\hat{I}_i)-1,j})_{j \in N}) \mid \eta \in \hat{I}_i\}.$$

Notice that if μ makes $\Theta^{\mathcal{A}}$ ambiguous given some $\theta^{\mathcal{U}} \in \Theta^{\mathcal{U}}$ then an Ellsberg strategy potentially allows player i to create ambiguity about his strategy even fixing the payoff relevant component of the type. Because ignoring $\theta^{\mathcal{A}}$ is always an option, if, in a sequential optimum/SEA, a player uses such an Ellsberg strategy it must be the case that she views choosing to create this strategic ambiguity (and to 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.⁷

4.1 Example 3: Strategic Use of Ambiguity: Ambiguous Cheap Talk

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, $a1$ and $a2$. 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

⁷In 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.

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 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.6, 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,I_i} as in (2.1) and (2.3). For each player i and information set $I_i \in \mathcal{I}_i^T$, let

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

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

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

Finally, let

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

(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 ,

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

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

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

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

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

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.

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 PBE and sequential equilibrium. Thus, it is both important and 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




Figure 5.1: Violation of sequential optimality

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 2.10 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 2.10 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. BADE, Dynamic Semi-Consistency, mimeo. (2016).
- [8] S. BALIGA, E. HANANY, P. KLIBANOFF, Polarization and Ambiguity, *American Economic Review* **103(7)** (2013), 3071-3083.
- [9] P. BATTIGALLI, S. CERREIA VIUGLIO, F. MACCHERONI, M. MARINACCI, Self-confirming Equilibrium and Model Uncertainty, *American Economic Review* **105(2)** (2015a), 646-77.

- [10] P. BATTIGALLI , S. CERREIA VIOGLIO, F. MACCHERONI, M. MARINACCI, Selfconfirming Equilibrium and Sequential Rationality with Ambiguity Averse Agents, mimeo. (2015b).
- [11] A.L. BODOH-CREED, Ambiguous beliefs and mechanism design, *Games and Economic Behavior* **75(2)** (2012), 518-537.
- [12] S. BOSE, A. DARIPA, A dynamic mechanism and surplus extraction under ambiguity, *Journal of Economic Theory* **144(5)** (2009), 2084-2114.
- [13] S. BOSE, E. OZDENOREN, A. PAPE, Optimal Auctions with Ambiguity, *Theoretical Economics* **1(4)** (2006), 411-438.
- [14] S. BOSE, L. RENO, Mechanism Design With Ambiguous Communication Devices, *Econometrica* **82** (2014), 1853-1872.
- [15] Y. CHEN, P. KATUSCAK, E. OZDENOREN, Sealed Bid Auctions with Ambiguity: Theory and Experiments, *Journal of Economic Theory*, **136(1)** (2007), 513-535.
- [16] A. DOMINIAK, M. S. LEE, Coherent Dempster-Shafer Equilibrium and Ambiguous Signals, mimeo. Virginia Tech (2015).
- [17] J. DOW, S.R.C WERLANG, Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction, *Journal of Economic Theory*, **64(2)** (1994), 305-324.
- [18] 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.
- [19] J. EICHBERGER, D. KELSEY, Non-Additive Beliefs and Strategic Equilibria, *Games and Economic Behavior* **30(2)** (2000), 183-215.
- [20] J. EICHBERGER, D. KELSEY, Sequential Two-Player Games with Ambiguity, *International Economic Review* 45 (2004), 1229-1261.
- [21] A. ELLIS, On Dynamic Consistency in Ambiguous Games, mimeo., (2016).
- [22] D. FUDENBERG, J. TIROLE, *Game Theory* (1991a), MIT Press.
- [23] D. FUDENBERG, J. TIROLE, Perfect Bayesian Equilibrium and Sequential Equilibrium, *Journal of Economic Theory*, **53** (1991b), 236-260.

- [24] 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.
- [25] R. GIBBONS, *Game Theory for Applied Economists* (1992), Princeton University Press.
- [26] 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.
- [27] I. GILBOA, D. SCHMEIDLER, Maxmin Expected Utility with a Non-Unique Prior, *Journal of Mathematical Economics*, **18** (1989), 141-153.
- [28] 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.
- [29] J. GREENBERG, The Right to Remain Silent, *Theory and Decision*, **48** (2000), 193-204.
- [30] E. HANANY, P. KLIBANOFF, Updating Preferences with Multiple Priors, *Theoretical Economics* **2** (2007), 261-298.
- [31] E. HANANY, P. KLIBANOFF, Updating Ambiguity Averse Preferences, *The BE Journal of Theoretical Economics (advances)* **9(1)** (2009), article 37.
- [32] A. KAJII, T. UI, Incomplete Information Games with Multiple Priors, *Japanese Economic Review* **56(3)** (2005), 332-351.
- [33] C. KELLNER, Tournaments as a Response to Ambiguity Aversion in Incentive Contracts, *Journal of Economic Theory*, **159(A)** (2015), 627-655.
- [34] C. KELLNER, M. LE QUEMENT, Modes of ambiguous communication, mimeo. University of Bonn (2013).
- [35] C. KELLNER, M. LE QUEMENT, Endogenous ambiguity in cheap talk, mimeo. University of Bonn (2015).
- [36] P. KLIBANOFF, Uncertainty, decision, and normal form games, mimeo. Northwestern University (1996).
- [37] P. KLIBANOFF, M. MARINACCI, S. MUKERJI, A Smooth Model of Decision Making under Ambiguity, *Econometrica* **73** (2005), 1849-1892.

- [38] D.M. KREPS, R. WILSON, Sequential Equilibria, *Econometrica* **50(4)** (1982), 863-894.
- [39] C. KUZMICS, Abraham Wald's complete class theorem and Knightian uncertainty, mimeo. Bielefeld University (2015).
- [40] K.C. LO, Equilibrium in beliefs under uncertainty, *Journal of Economic Theory* **71(2)** (1996), 443-484.
- [41] K.C. LO, Extensive Form Games with Uncertainty Averse Players, *Games and Economic Behavior* **28** (1999), 256-270.
- [42] G. LOPOMO, L. RIGOTTI, C. SHANNON, Uncertainty in Mechanism Design, mimeo. (2010).
- [43] M. MARINACCI, Ambiguous Games, *Games and Economic Behavior* **31(2)** (2000), 191-219.
- [44] P. MILGROM, J. ROBERTS, Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis, *Econometrica* **50(2)** (1982), 443-459.
- [45] I. MOURAVIEV, F. RIEDEL, L. SASS, Kuhn's Theorem for Extensive Form Ellsberg Games, mimeo. (2015).
- [46] S. MUKERJI, H.S. SHIN, Equilibrium departures from common knowledge in games with non-additive expected utility, *Advances in Theoretical Economics* **2(1)** (2002).
- [47] 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.
- [48] R.B. MYERSON, *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA (1991).
- [49] E. OZDENOREN, D. LEVIN, Auctions with Uncertain Numbers of Bidders, *Journal of Economic Theory*, **118(2)**, (2004) 229-251.
- [50] F. RIEDEL, L. SASS, Ellsberg games, *Theory and Decision* **76** (2013), 469-509.
- [51] A. SALO, M. WEBER, Ambiguity Aversion in First Price Sealed Bid Auctions, *Journal of Risk and Uncertainty* **11** (1995), 123-137.
- [52] M. SHIMOJI, J. WATSON, Conditional Dominance, Rationalizability, and Game Forms, *Journal of Economic Theory* **83** (1998), 161-195.

- [53] R.H. STROTZ, Myopia and Inconsistency in Dynamic Utility Maximization, *The Review of Economic Studies* **23(3)** (1955-56), 165-180.
- [54] A. DI TILLIO, N. KOS, M. MESSNER, The Design of Ambiguous Mechanisms, mimeo. Bocconi University (2012).
- [55] A. WOLITZKY, Bilateral Trading with Maxmin Agents, mimeo. Stanford University (2013).
- [56] 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 , a (σ, ν) such that σ is an ex-ante equilibrium, a player i and an information set I_i . If ν_{i, I_i} is derived from $\nu_{i, f_i(I_i)}$ (or, if $s(I_i) = 0$, from μ_i) via the smooth rule using σ as the ex-ante equilibrium and, for all $\sigma'_i \in \Sigma_i$,*

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

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

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

It is useful for the proof of this result (as well as that of Theorem 2.3) to refer to a player's "local ambiguity neutral measure" at an information set. Given (σ, ν) , for any player i , let $q^{\sigma, i}(h)$ denote i 's ex-ante σ -local measure, defined for each $h \in H$ by

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

additionally, for any information set I_i , let $q^{(\sigma,\nu),i,I_i}$ denote i 's (σ, ν) - local measure given I_i , defined for each $h \in H$ such that $h^{s(I_i)} \in I_i$ by, letting $t = s(I_i)$,

$$q^{(\sigma,\nu),i,I_i}(h) \equiv \sum_{\pi \in \Delta(H^{m_i(I_i)}) | \pi(f_i^{\sigma-i}(I_i)) > 0} \phi'_i \left(\sum_{\hat{h} | \hat{h}^{s(I_i)} \in I_i} u_i(\hat{h}) p_\sigma(\hat{h} | \hat{h}^{s(I_i)}) \pi_{I_i, \sigma_{-i}}(\hat{h}^{s(I_i)}) \right) \cdot p_{-i, \sigma_{-i}}(h | h^{s(I_i)}) \pi_{I_i, \sigma_{-i}}(h^{s(I_i)}) \nu_{i, I_i}(\pi). \quad (\text{A.2})$$

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

$$\sum_{\bar{h} | \bar{h}^{s(I_i)} \in I_i} u_i(\bar{h}) p_{i, \sigma'_i}(\bar{h} | \bar{h}^{s(I_i)}) q^{(\sigma,\nu),i,I_i}(\bar{h}), \quad (\text{A.3})$$

where $q^{(\sigma,\nu),i,I_i}$ is i 's (σ, ν) - local measure given I_i (defined in (A.2)), (respectively, $\sigma'_i = \sigma_i$ maximizes

$$\sum_{\bar{h}} u_i(\bar{h}) p_{i, \sigma'_i}(\bar{h} | \bar{h}^0) q^{\sigma,i}(\bar{h}), \quad (\text{A.4})$$

where $q^{\sigma,i}$ is i 's ex-ante σ - local measure (defined in (A.1)).

We want to show that $V_{i, I_i}(\sigma) \geq V_{i, I_i}(\sigma'_i, \sigma_{-i})$ for all σ'_i . By the above, it is sufficient to show that $\sigma'_i = \sigma_i$ maximizes (A.3).

Consider the case where $s(I_i) > 0$ (the case where $s(I_i) = 0$ is similar, using (A.4) instead of (A.5), and is omitted). By assumption in the statement of the lemma, $V_{i, f_i(I_i)}(\sigma) \geq V_{i, f_i(I_i)}(\sigma'_i, \sigma_{-i})$ for all $\sigma'_i \in \Sigma_i$. As in (A.3), this is equivalent to the condition that $\sigma'_i = \sigma_i$ maximizes

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

Notice that, since i 's strategy is a function only of i 's information sets and $R_i(h^{s(I_i)}) = R_i(I_i)$ for any h such that $h^{s(I_i)} \in I_i$, $p_{i, \sigma'_i}(h^{s(I_i)} | h^{m_i(I_i)})$ is the same for any such h . Thus, the objective function in (A.5) can be equivalently written as

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

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

$$\sum_{\bar{h} | \bar{h}^{s(I_i)} \in I_i} u_i(\bar{h}) p_{i, \sigma'_i}(\bar{h} | \bar{h}^{s(I_i)}) q^{(\sigma, \nu), i, f_i(I_i)}(\bar{h}) \quad (\text{A.6})$$

is affected by the specification of σ'_i from I_i onward and no other part of σ'_i affects (A.6). Therefore (A.5) implies that σ_i maximizes (A.6). For that to imply σ_i maximizes (A.3), it is sufficient to show that $q^{(\sigma, \nu), i, I_i}(\bar{h}) \propto q^{(\sigma, \nu), i, f_i(I_i)}(\bar{h})$ holds for $\{\bar{h} \mid \bar{h}^{s(I_i)} \in I_i\}$. This proportionality may be shown by using the local measure definition (A.2), applying the smooth rule iteratively to substitute for $\nu_{i, I_i}(\pi)$ for all $\pi \in \Delta(H^{m_i(I_i)})$ such that $\pi(f_i^{\sigma^{-i}}(I_i)) > 0$ (as $\nu_{i, I_i}(\pi) = 0$ for other π) and then using the definitions of $\pi_{I_i, \sigma_{-i}}$ and $\pi_{f_i(I_i), \sigma_{-i}}$ and cancelling terms. ■

Proof of Theorem 2.1. We show that $(\sigma^P, \hat{\nu}^P)$, where, for all i, I_i , $\hat{\nu}_{i, I_i}^P = \nu_{i, I_i}^P$ whenever $m_i(I_i) = s(I_i) > 0$, and where, everywhere else, $\hat{\nu}_{i, I_i}^P$ 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, I_i where $m_i(I_i) = s(I_i) > 0$. However, from the definition of the smooth rule (Definition 2.11), observe that it is exactly for i, I_i with $m_i(I_i) = s(I_i) > 0$ 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 $V_{i, I_i}(\sigma^P) \geq V_{i, I_i}(\sigma'_i, \sigma_{-i}^P)$ for all $\sigma'_i \in \Sigma_i$, observe that (a) for i, I_i such that $m_i(I_i) = s(I_i) > 0$, 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 , (2.7) in the definition of sequential optimality is satisfied. Choose a ν satisfying the smooth rule using σ as the ex-ante equilibrium. Since $m_i(I_i) = 0$ for all i and I_i , either $I_i \subseteq \Theta$ or $f_i(I_i) \neq I_i$, implying that ν is pinned down completely by the smooth rule. By Lemma A.1, (2.8) in the definition of sequential optimality is satisfied for all i and I_i . ■

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 optimality of σ_i according to V_{i, I_i} for all $I_i \in \mathcal{I}_i^T$. Next we proceed by induction on the stage t . Fix any t such that $0 < t \leq T$, and suppose that, for each player i , σ_i is optimal according to V_{i, I_i} for all $I_i \in \mathcal{I}_i^t$. We claim that, for each player i , σ_i is optimal according to V_{i, I_i} for all $I_i \in \mathcal{I}_i^{t-1}$. The argument for this is as follows: Fix a player i and $I_i \in \mathcal{I}_i^{t-1}$. Consider any strategy σ'_i

for player i . For any $J_i \in \mathcal{I}_i^t$, the optimality of σ_i according to V_{i,J_i} implies (see (A.3))

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

Since ν satisfies extended smooth rule updating using σ as the strategy profile, for all such J_i that are reachable from I_i given σ_{-i} , $q^{(\sigma,\nu),i,J_i}(h) \propto q^{(\sigma,\nu),i,I_i}(h)$ holds for $\{h \mid h^t \in J_i\}$ (see the argument for this near the end of the proof of Lemma A.1). Also, since i 's strategy is a function only of i 's information sets and $h^t \in J_i$ implies $h^{t-1} \in I_i$, $p_{i,\sigma'_i}(h|h^t)$ is the same for all h such that $h^t \in J_i$. After substituting in (A.7) for $q^{(\sigma,\nu),i,J_i}$, cancelling the constant of proportionality and multiplying by $p_{i,\sigma'_i}(h^t|h^{t-1})$, (A.7) becomes

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

Notice that for any $J_i \in \mathcal{I}_i^t$ not reachable from I_i given σ_{-i} , it holds that, for all h with $h^t \in J_i$, $q^{(\sigma,\nu),i,I_i}(h) = 0$ because $p_{-i,\sigma_{-i}}(h|h^{t-1}) = 0$. Thus, summing (A.8) for all $J_i \in \mathcal{I}_i^t$ reachable from I_i given σ_{-i} is the same as summing for all $J_i \in \mathcal{I}_i^t$ such that $J_i^{-1} = I_i$, yielding:

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

The no profitable one-stage deviation property implies σ_i is optimal according to V_{i,I_i} among all strategies deviating only at I_i . By the optimality representation invoked in (A.7) applied to I_i and restricted to such deviations,

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

Combining (A.10) and (A.9) implies

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

Since (A.11) holds for any σ'_i , it is the same as (A.7) with $t-1$ in the role of t and I_i in the role of J_i . Therefore σ_i is optimal according to V_{i,I_i} . Since this conclusion holds for any $I_i \in \mathcal{I}_i^{t-1}$, the induction step is completed. It follows that (σ, ν) satisfies the interim optimality conditions (2.8) in the definition of sequentially optimal.

It remains to show that σ also satisfies the ex-ante optimality conditions (2.7). Since ν satisfies extended smooth rule updating using σ as the strategy profile, for all $I_i \subseteq \Theta_i$, $q^{(\sigma,\nu),i,I_i}(h) \propto q^{\sigma,i}(h)$ holds for $\{h \mid h^0 \in I_i\}$. Using this to substitute for $q^{(\sigma,\nu),i,I_i}$ in (A.11) with $t=1$, cancelling the constant of proportionality and summing for all I_i , yields:

$$\begin{aligned} & \sum_h u_i(h) p_{i,\sigma_i}(h|h^0) q^{\sigma,i}(h) \\ & \geq \sum_h u_i(h) p_{i,\sigma'_i}(h|h^0) q^{\sigma,i}(h). \end{aligned} \quad (\text{A.12})$$

Since (A.12) holds for any σ'_i , σ satisfies (A.4) which is equivalent to the ex-ante optimality condition (2.7). ■

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 adapted from the interim belief systems determined by extended smooth rule updating using the strategies in the sequence as the strategy profile. By continuity of the extended smooth rule in the strategy profile at on-path information sets, this sequence has a limit, which is an adaptation using σ from the interim belief system $\tilde{\nu}$. By construction, $(\sigma, \tilde{\nu})$ satisfies smooth rule consistency, thus it is an SEA. ■

Proof of Theorem 2.5. Fix a sequence $\varepsilon^k = (\varepsilon_I^k)_{I \in \cup_{i \in N} \mathcal{I}_i}$ of strictly positive vectors of dimension $|\cup_{i \in N} \mathcal{I}_i|$, converging in the sup-norm to 0 and such that $\varepsilon_{I_i}^k \leq \frac{1}{|A_i(I_i)|}$ for all players i and information sets I_i . For any k , let Γ^k be the restriction of the game Γ defined such that the set of feasible strategy profiles is the set of all completely mixed σ^k satisfying $\sigma_i^k(I_i)(a_i) \geq \varepsilon_{I_i}^k$ for all i , I_i and actions $a_i \in A_i(I_i)$. Consider the agent normal form G^k of the game Γ^k (see e.g., Myerson, 1991, p.61). Since the payoff functions are concave and the set of strategies of each player in G^k is non-empty, compact and convex,

G^k has an ex-ante equilibrium by Glicksberg (1952). Let $\hat{\sigma}^k$ be the strategy profile in the game Γ^k corresponding to this equilibrium. Then $\hat{\sigma}^k$ is an ex-ante equilibrium of Γ^k . By Theorem 2.2, since all information sets are on the equilibrium path, there exists an interim belief system ν^k such that $(\hat{\sigma}^k, \nu^k)$ is sequentially optimal. By Theorem 2.1, there exists an interim belief system $\hat{\nu}^k$ satisfying the smooth rule using $\hat{\sigma}^k$ as the ex-ante equilibrium such that $(\hat{\sigma}^k, \hat{\nu}^k)$ is a sequential optimum of Γ^k . By compactness of the set of strategy profiles, the sequence $\hat{\sigma}^k$ has a convergent sub-sequence, the limit of which is denoted by $\hat{\sigma}$. By continuity in the strategy profile of the extended smooth rule formula and compactness of the set of interim belief systems, an associated sub-sequence of adaptations from $\hat{\nu}^k$ using $\hat{\sigma}^k$ converges to a limit, which itself is adapted using $\hat{\sigma}$ from some interim belief system. Denote this interim belief system by ν . By continuity of the payoff functions, $\hat{\sigma}$ is an ex-ante equilibrium of Γ . Given any information set I_i and continuation strategy $\tilde{\sigma}_i^{I_i}$ of player i in Γ , let $\tilde{\sigma}_i^{k, I_i}$ be a feasible strategy in Γ^k for this player that is closest (in the sup-norm) to $\tilde{\sigma}_i^{I_i}$. Since, by sequential optimality of $(\hat{\sigma}^k, \hat{\nu}^k)$ for each k , $\hat{\sigma}_i^{k, I_i}$ is weakly better than $\tilde{\sigma}_i^{k, I_i}$ for player i given belief $\hat{\nu}_{i, I_i}^k$, and since, along the sub-sequence, $\tilde{\sigma}_i^{k, I_i}$ converges to $\tilde{\sigma}_i^{I_i}$ and the adaptation from $\hat{\nu}_{i, I_i}^k$ using $\hat{\sigma}_i^{k, I_i}$ converges to the adaptation from ν_{i, I_i} using $\hat{\sigma}_i^{I_i}$, continuity of the payoff functions implies that $\hat{\sigma}_i^{I_i}$ is weakly better than $\tilde{\sigma}_i^{I_i}$ for this player given belief ν_{i, I_i} . Therefore $(\hat{\sigma}, \nu)$ satisfies sequential optimality. Finally, observe that $(\hat{\sigma}, \nu)$ satisfies smooth rule consistency (since it is explicitly constructed as the limit of an appropriate sequence). Therefore $(\hat{\sigma}, \nu)$ is an SEA of Γ . ■

Proof of Corollary 2.1. It is enough to show that the no profitable one-stage deviation property and smooth rule consistency imply (σ, ν) is sequentially optimal. This follows directly from Definition 2.14 and Theorem 2.3. ■

Proof of Theorem 2.6. Modify Example 3.1 by removing the action Q for Player 1. For each player i , let $\mu_i = \mu$ where μ puts probability $\frac{1}{2}$ on π_0 and $\frac{1}{2}$ on π_1 , where $\pi_0(I) = 1$ and $\pi_1(I) = 0$, and let $\hat{\phi}_i(x) = \phi(x) \equiv -e^{-x}$. With these preferences, the unique ex-ante equilibrium has player 2 play U with probability λ^* and player 3 play R with probability λ^* , where $\lambda^* = 1 - \frac{2}{5} \ln(3/2)$. In contrast, if $\phi(x) \equiv \nu$, using the same μ , then the unique ex-ante equilibrium has player 2 playing D with probability 1 and player 3 play L with probability 1. ■

Proof of Theorem 2.7. With degenerate μ_i , since there is no ambiguity, ex-ante preferences are independent of ϕ_i . This establishes $E_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) = E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$. To establish that $Q_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) = Q_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$, suppose that $(\sigma, \nu) \in Q_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ where ν is derived according to the smooth rule and apply the argument in the latter part of the proof of Theorem 2.9 to construct $\hat{\nu}$ such that $(\sigma, \hat{\nu}) \in Q_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. Note that at all on path information sets, since the μ_i are degenerate, the smooth rule gives ν_{i, I_i} that

are degenerate as well, and the construction has $\hat{\nu}_{i,I_i} = \nu_{i,I_i}$ so optimality on path is trivially preserved. The purpose of using the construction from the other proof is to ensure that off-path, where ν_{i,I_i} are unconstrained by μ_i and so may be non-degenerate, optimality can be maintained under $\hat{\phi}_i$ given appropriate choice of interim beliefs. Finally, for SEA, smooth rule consistency implies that even off-path ν_{i,I_i} are degenerate, so that interim optimality is trivially preserved under $\hat{\phi}_i$. ■

Proof of Theorem 2.8. That $\bigcup_{\hat{\mu}} E_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} E_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ follows by considering only degenerate beliefs on the left-hand side and choosing them to have the same reduced measure as the right-hand side beliefs. $\bigcup_{\hat{\mu}} Q_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} Q_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ follows using the same construction and additionally taking the left-hand side interim beliefs at each information set to be degenerate with the same reduced measure as the right-hand side interim beliefs at the corresponding information set and noting that this preserves optimality at each information set. $\bigcup_{\hat{\mu}} S_{\Gamma}((\hat{\mu})_{i \in N}, (\hat{\phi}_i)_{i \in N}) \supseteq \bigcup_{\hat{\mu}} S_{\Gamma}((\hat{\mu})_{i \in N}, (\iota)_{i \in N})$ follows using the same construction as for sequential optima, observing that the left-hand side degenerate beliefs satisfy smooth rule consistency since the right-hand side beliefs do so. Section 3.1 provides an example where the inclusion is strict and the new strategies generate new paths of play. ■

Proof of Theorem 2.9. Fix Γ and let $\sigma \in E_{\Gamma}((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$. Ex-ante equilibrium is equivalent to ex-ante optimality for all players i of σ_i according to i 's preferences given σ_{-i} . This ex-ante optimality is equivalent to (see (A.4)) $\sigma'_i = \sigma_i$ maximizing

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

with respect to σ'_i , where $q^{\sigma, i}(h)$ is i 's ex-ante σ -local measure as in (A.1). Let $\hat{\mu}_i$ be the probability measure such that

$$\hat{\mu}_i(\pi) \propto \frac{\phi'_i(\sum_h u_i(h) p_{\sigma}(h|h^0) \pi(h^0))}{\hat{\phi}'_i(\sum_h u_i(h) p_{\sigma}(h|h^0) \pi(h^0))} \mu_i(\pi). \quad (\text{A.13})$$

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

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

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

Turn now to sequentially optimal strategy profiles. Suppose $\sigma \in Q_{\Gamma}((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (σ, ν) is sequentially optimal for Γ . Let $\hat{\mu}_i$ be defined as in (A.13) and for each I_i , $\hat{\nu}_{i,I_i}$ be the probability measure such that

$$\hat{\nu}_{i,I_i}(\pi) \propto \frac{\phi'_i \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{\sigma}(h|h^{s(I_i)}) \pi_{I_i,\sigma_{-i}}(h^{s(I_i)}) \right)}{\hat{\phi}'_i \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{\sigma}(h|h^{s(I_i)}) \pi_{I_i,\sigma_{-i}}(h^{s(I_i)}) \right)} \nu_{i,I_i}(\pi). \quad (\text{A.14})$$

By the argument in the ex-ante equilibrium part of this proof, $\sigma \in E_{\Gamma}((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. Using $\hat{\phi}_i$ and $\hat{\nu}_{i,I_i}$, interim optimality for player i at I_i as a function of σ'_i is equivalent (see (A.3) and (A.2)) to $\sigma'_i = \sigma_i$ maximizing

$$\begin{aligned} & \sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{i,\sigma'_i}(h|h^{s(I_i)}) \sum_{\pi|\pi(f_i^{\sigma_{-i}}(I_i)) > 0} \hat{\phi}'_i \left(\sum_{\hat{h}|\hat{h}^{s(I_i)} \in I_i} u_i(\hat{h}) p_{\sigma}(\hat{h}|\hat{h}^{s(I_i)}) \pi_{I_i,\sigma_{-i}}(\hat{h}^{s(I_i)}) \right) \\ & \cdot p_{-i,\sigma_{-i}}(h|h^{s(I_i)}) \pi_{I_i,\sigma_{-i}}(h^{s(I_i)}) \hat{\nu}_{i,I_i}(\pi). \end{aligned} \quad (\text{A.15})$$

Since (A.15) is proportional to

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

optimality of σ_i at such an I_i in game Γ implies (see (A.3)) the same in game $\hat{\Gamma}$. This is true for each player i . Thus, $(\sigma, \hat{\nu})$ is sequentially optimal in $\hat{\Gamma}$.

We now extend the argument to SEA. Suppose $\sigma \in S_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (σ, ν) is an SEA for Γ . As above, let $\hat{\mu}_i$ be as in (A.13) and for each I_i , $\hat{\nu}_{i, I_i}$ be defined as in (A.14). By our previous arguments, $(\sigma, \hat{\nu})$ is sequentially optimal in $\hat{\Gamma}$. Since ν is determined by extended smooth rule updating using σ as the strategy profile, it follows from inspection of (A.14) and the formulation of the smooth rule in (A.26) and (2.9) that $\hat{\nu}$ satisfies extended smooth rule updating using σ as the strategy profile. It remains to show that $(\sigma, \hat{\nu})$ satisfies the remaining parts of smooth rule consistency. Since (σ, ν) satisfies smooth rule consistency, there exists a sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$, with $\lim_{k \rightarrow \infty} \sigma^k = \sigma$, such that $\hat{\nu} = \lim_{k \rightarrow \infty} \hat{\nu}^k$, where $\hat{\nu}^k$ is adapted from ν using σ , each $\hat{\nu}^k$ is adapted from ν^k using σ^k , and each ν^k is the interim belief system determined by extended smooth rule updating using σ^k as the strategy profile. Fix such sequences and i, I_i . In what follows, when we include π in an expression, the expression is understood to apply for all $\pi \in \Delta(\Theta)$ such that $\pi(\{h^0 \mid h^{s(I_i)} \in I_i\}) > 0$. Since σ^k is completely mixed, $\{h^0 \mid h^{s(I_i)} \in I_i\} = f_i^{\sigma^k}(I_i)$ and the extended smooth rule has bite at each information set and so, applying the ‘‘all-at-once’’ formulation of the smooth rule in (A.26) and the initial smooth rule step (2.9) from μ_i , for each k ,

$$\nu_{i, I_i}^k(\pi) \propto \frac{\phi_i' \left(\sum_{h \in H} u_i(h) p_{\sigma^k}(h \mid h^0) \pi(h^0) \right)}{\phi_i' \left(\sum_{h \mid h^{s(I_i)} \in I_i} u_i(h) p_{\sigma^k}(h \mid h^{s(I_i)}) \pi_{I_i, \sigma_{-i}^k}(h^{s(I_i)}) \right)} \left(\sum_{h \mid h^{s(I_i)} \in I_i, \sigma_{-i}^k} p_{-i, \sigma_{-i}^k}(h^{s(I_i)} \mid h^0) \pi(h^0) \right) \mu_i(\pi). \quad (\text{A.16})$$

Using (2.6) for I_{i, σ_{-i}^k} and $f_i(I_i)_{i, \sigma_{-i}^k}$, where $f_i(I_i)$ and $m_i(I_i)$ are determined from σ , and observing that $I_{i, \sigma_{-i}^k} = I_i$ and $f_i(I_i)_{i, \sigma_{-i}^k} = f_i(I_i)$ since σ^k is completely mixed,

$$\pi_{I_{i, \sigma_{-i}^k}}(h^{s(I_i)}) = \frac{p_{-i, \sigma_{-i}^k}(h^{s(I_i)} \mid h^{m_i(I_i)}) \pi_{f_i(I_i)_{i, \sigma_{-i}^k}}(h^{m_i(I_i)})}{\sum_{\hat{h} \mid \hat{h}^{s(I_i)} \in I_i} p_{-i, \sigma_{-i}^k}(\hat{h}^{s(I_i)} \mid \hat{h}^{m_i(I_i)}) \pi_{f_i(I_i)_{i, \sigma_{-i}^k}}(\hat{h}^{m_i(I_i)})} \text{ for all } h \text{ such that } h^{s(I_i)} \in I_i$$

and

$$\pi_{f_i(I_i)_{i, \sigma_{-i}^k}}(h^{m_i(I_i)}) = \frac{p_{-i, \sigma_{-i}^k}(h^{m_i(I_i)} \mid h^0) \pi(h^0)}{\sum_{\hat{h} \mid \hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} \mid \hat{h}^0) \pi(\hat{h}^0)} \text{ for all } h \text{ such that } h^{m_i(I_i)} \in f_i(I_i).$$

By Definition 2.9, $\lim_{k \rightarrow \infty} \hat{\nu}^k$ exists only if $\lim_{k \rightarrow \infty} \pi_{I_i, \sigma_{-i}^k}$ and $\lim_{k \rightarrow \infty} \pi_{f_i(I_i)_{i, \sigma_{-i}^k}}$ exist for $\pi \in \text{supp } \nu_{i, I_i}^k$. Since $\lim_{k \rightarrow \infty} \hat{\nu}^k$ indeed exists, the existence of $\lim_{k \rightarrow \infty} \pi_{I_i, \sigma_{-i}^k}$ and $\lim_{k \rightarrow \infty} \pi_{f_i(I_i)_{i, \sigma_{-i}^k}}$

implies $\bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0) \equiv \lim_{k \rightarrow \infty} \frac{p_{-i, \sigma_{-i}^k}(h^{m_i(I_i)} | h^0)}{\sum_{\hat{h} | \hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} | \hat{h}^0)}$ exists for all h such that $h^{m_i(I_i)} \in f_i(I_i)$. Define $\bar{\pi}^\pi \in \Delta(H^{m_i(I_i)})$ such that $\bar{\pi}^\pi(h^{m_i(I_i)}) = \bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0) \pi(h^0)$ for all h such that $h^{m_i(I_i)} \in f_i(I_i)$ and $\bar{\pi}^\pi(h^{m_i(I_i)})$ is constant for all other h . It follows that $\lim_{k \rightarrow \infty} \pi_{f_i(I_i)_{i, \sigma_{-i}^k}} = \bar{\pi}_{f_i(I_i)_{i, \sigma_{-i}^k}}^\pi$ and $\lim_{k \rightarrow \infty} \pi_{I_i, \sigma_{-i}^k} = \bar{\pi}_{I_i, \sigma_{-i}^k}^\pi$. Define an interim belief $\bar{\nu}_{i, I_i}$ such that for each π ,

$$\bar{\nu}_{i, I_i}(\bar{\pi}^\pi) \propto \frac{\phi'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h | h^0) \pi(h^0) \right)}{\phi'_i \left(\sum_{h | h^{s(I_i)} \in I_i} u_i(h) p_\sigma(h | h^{s(I_i)}) \bar{\pi}_{I_i, \sigma_{-i}}^\pi(h^{s(I_i)}) \right)} \left(\sum_{h | h^{s(I_i)} \in I_i, \sigma_{-i}} p_{-i, \sigma_{-i}}(h^{s(I_i)} | h^{m_i(I_i)}) \bar{\pi}^\pi(h^{m_i(I_i)}) \right) \mu_i(\pi) \quad (\text{A.17})$$

Dividing (A.16) by $\sum_{\hat{h} | \hat{h}^{m_i(I_i)} \in f_i(I_i)} p_{-i, \sigma_{-i}^k}(\hat{h}^{m_i(I_i)} | \hat{h}^0)$, which is constant with respect to π , $\lim_{k \rightarrow \infty} \hat{\nu}_{i, I_i}^k$ equals the adaptation from $\bar{\nu}_{i, I_i}$ using σ . Since, by inspection of (A.17) and the formulation of the smooth rule in (A.26) and (2.9), the interim belief system $\bar{\nu}$ satisfies extended smooth rule updating using σ as the strategy profile, $(\sigma, \bar{\nu})$ satisfies smooth rule consistency in Γ .

Apply the same construction to $\hat{\Gamma}$ by defining, for each I_i , $\hat{\nu}_{i, I_i}^k$ as in (A.16) and $\bar{\hat{\nu}}_{i, I_i}$ as in (A.17), with, in each case, $\hat{\phi}_i$ and $\hat{\mu}_i$ replacing ϕ_i and μ_i , respectively. Then, by the same arguments as above, with the exception that the existence of $\bar{p}_{-i, \sigma_{-i}}$ need not be re-established, $(\sigma, \bar{\hat{\nu}})$ satisfies smooth rule consistency in $\hat{\Gamma}$. Since the adaptations from ν and $\bar{\nu}$ using σ have each been established to be equal to $\hat{\nu}$, by using (A.13) and (A.14) and observing that the expectation terms in (A.14) are constant with respect to the summation in Definition 2.9, one can establish that the adaptations from $\hat{\nu}$ and $\bar{\hat{\nu}}$ using σ must be equal to each other. It follows that $(\sigma, \hat{\nu})$ satisfies smooth rule consistency in $\hat{\Gamma}$. Therefore $(\sigma, \hat{\nu})$ is an SEA of $\hat{\Gamma}$.

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

Proof of Theorem 2.10. Fix a game Γ . Suppose $\varsigma \in \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$, for each i , $\hat{\phi}_i = \chi_i(\phi_i)$ for some increasing, differentiable and concave χ_i (note that differentiability of χ_i is implied by the continuous differentiability of $\hat{\phi}_i$ in the class of games considered in this

paper) and $\hat{\mu}_i$ is the probability measure such that

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

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

$$\sum_\pi \phi_i \left(\sum_h u_i(h)p_\varsigma(h|h^0)\pi(h^0) \right) \mu_i(\pi) \geq \sum_\pi \phi_i \left(\sum_h u_i(h)p_{(\varsigma'_i, \varsigma_{-i})}(h|h^0)\pi(h^0) \right) \mu_i(\pi). \quad (\text{A.18})$$

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

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

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

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

where the last inequality follows from A.18. Since this is true for each i and each ς'_i , $\varsigma \in \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$. This shows $\tilde{E}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N}) \subseteq \bigcup_{(\hat{\mu}_i)_{i \in N}} \tilde{E}_\Gamma((\hat{\mu}_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$.

Turn now to the part of the theorem about sequentially optimal strategy profiles. Suppose $\varsigma \in \tilde{Q}_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ and ν is an interim belief system such that (ς, ν) is sequentially optimal for Γ with respect to pure strategies. Further suppose that for each i , $\hat{\phi}_i = \chi_i(\phi_i)$ for some increasing, differentiable and concave χ_i (note that differentiability of χ_i is implied by the continuous differentiability of $\hat{\phi}_i$ in the class of games considered in this paper), $\hat{\mu}_i$ is the probability measure such that

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

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

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

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

$$\begin{aligned} & \sum_{\pi} \phi_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_\varsigma(h|h^s(I_i)) \pi_{I_i, \varsigma_{-i}}(h^s(I_i)) \right) \nu_{i,I_i}(\pi) \\ & \geq \sum_{\pi} \phi_i \left(\sum_{h|h^s(I_i) \in I_i} u_i(h) p_{(\varsigma'_i, \varsigma_{-i})}(h|h^s(I_i)) \pi_{I_i, \varsigma_{-i}}(h^s(I_i)) \right) \nu_{i,I_i}(\pi). \end{aligned} \quad (\text{A.19})$$

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

Proof of Theorem 2.11. Consider the following modification of the game in Figure B.1: Remove player 1 and remove the action F for player 2, so that the initial stage has player 2 choosing between R and TB (\equiv an action that ends the game and results in payoffs as in T followed by B in the original game) without being informed of the type. Under ambiguity neutrality for players 2 and 3, $\bigcup_{(\mu_i)_{i \in N}} \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\iota)_{i \in N}) = \{(TB, (S, S), H), (TB, (S, S), G), (TB, (C, S), H), (TB, (C, S), G), (TB, (S, C), H), (TB, (S, C), G), (TB, (C, C), H), (TB, (C, C), G)\}$. To see this, first note that if $\sum_{\pi} \pi(I) \mu_2(\pi) \in (0, \frac{2}{5})$, then all the pure profiles where 2 plays TB are ex-ante equilibria under ambiguity neutrality. Second, any pure profile where 2 plays R cannot be an ex-ante equilibrium under ambiguity neutrality. Observe that 3 plays C following R (under either type) only if 2 plays H , 2 can play H rather than G on path if and only if $2 \geq 6 \sum_{\pi} \pi(I) \mu_2(\pi)$, and 2 can play R followed by H rather than TB if and only if $p(C)2 \geq 4(1 - \sum_{\pi} \pi(I) \mu_2(\pi))$ where $0 \leq p(C) \leq 1$ is 2's reduced probability that the type is such that 3 plays C . Since $\sum_{\pi} \pi(I) \mu_2(\pi)$ cannot be simultaneously $\leq \frac{1}{3}$ and $\geq \frac{1}{2}$, 2 cannot play R in pure strategy equilibrium under ambiguity neutrality. Similarly, under ambiguity neutrality the sequentially optimal pure profiles are $\bigcup_{(\mu_i)_{i \in N}} \tilde{Q}_\Gamma((\mu_i)_{i \in N}, (\iota)_{i \in N}) = \{(TB, (S, S), G), (TB, (C, C), H)\}$. To see this, note that optimality at the point where 3 is given the move requires each type of 3 to play C if and only if 2 plays H , thus eliminating all the other profiles in $\bigcup_{(\mu_i)_{i \in N}} \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\iota)_{i \in N})$. Furthermore, the two listed profiles are supported by beliefs for player 2 such that $\sum_{\pi} \pi(I) \mu_2(\pi) \in (0, \frac{2}{5})$ and $\sum_{\pi} \pi(I) \nu_{2, \{(I, R, C), (II, R, C)\}}(\pi) = \frac{1}{3}$, for example.

In light of Theorem 2.10, we know that if we change the ϕ to be more ambiguity averse (more concave) these equilibrium sets cannot shrink. It remains to show an example in which new elements appear. To this end, suppose $\hat{\phi}_2(x) \equiv -e^{-2x}$ and $\mu_2(\pi_1) = \mu_2(\pi_2) = \frac{1}{2}$, where $\pi_1(I) = \frac{1}{4}$ and $\pi_2(I) = \frac{3}{4}$. Furthermore suppose $\nu_{2,\{(I,R,C),(II,R,C)\}}(\pi_1) = \nu_{2,\{(I,R,C),(II,R,C)\}}(\pi_2) = \frac{1}{2}$ (as the smooth rule delivers in this example). Calculation then shows that $\hat{\phi}_2(2) > \frac{1}{2}(\hat{\phi}_2(3) + \hat{\phi}_2(1))$ and $\hat{\phi}_2(2) > \frac{1}{2}(\hat{\phi}_2(\frac{3}{2}) + \hat{\phi}_2(\frac{9}{2}))$, implying $(R, (C, C), H) \in \tilde{E}_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N}) \cap \tilde{Q}_\Gamma((\mu_i)_{i \in N}, (\hat{\phi}_i)_{i \in N})$ since 2 is worse off deviating either to TB or to R followed by G .

■

Proof of Theorem 2.12. Fixing an equilibrium $\sigma \in E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ (resp. Q_Γ with associated interim belief system ν or S_Γ with associated ν and sequence of completely mixed strategy profiles $\{\sigma^k\}_{k=1}^\infty$) and a player i , say that ambiguity aversion makes σ_i *ex-ante* (resp. *sequentially optimal or SEA*) *more belief robust* if there exists a $\hat{\phi}_i$ strictly more concave than ϕ_i such that, for each $\hat{\mu}_i$ with the same support, Π_i , as μ_i , σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i$ and $\hat{\mu}_i$ (resp. additionally interim best response to σ_{-i} at any information set I_i given $\hat{\phi}_i$ and ν_{i,I_i} or (for SEA) that plus satisfying the part for player i of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$) whenever it is an ex-ante best response to σ_{-i} given ϕ_i and $\hat{\mu}_i$ (resp. additionally interim best response to σ_{-i} at any information set I_i given ϕ_i and $\hat{\nu}_{i,I_i}$ or (for SEA) that plus satisfying the part for player i of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$). Similarly, say that ambiguity aversion makes σ_i *ex-ante* (resp. *sequentially optimal or SEA*) *strictly more belief robust* if, in addition, there exists $\tilde{\mu}_i$ with support Π_i such that σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i$ and $\tilde{\mu}_i$ (resp. additionally interim best response to σ_{-i} at any information set I_i given $\hat{\phi}_i$ and ν_{i,I_i} or (for SEA) that plus satisfying the part for player i of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$) but is not an ex-ante best response to σ_{-i} given ϕ_i and $\tilde{\mu}_i$ (resp., for all interim beliefs, fails either the ex-ante or interim best response to σ_{-i} property given ϕ_i and $\tilde{\mu}_i$, or (for SEA) fails the same property for all interim beliefs satisfying smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$). To prove that ambiguity aversion makes σ ex-ante (resp. sequentially optimal or SEA) (strictly) more belief robust, it is sufficient to show, for each player i , that ambiguity aversion makes σ_i ex-ante (resp. sequentially optimal or SEA) more belief robust (and that there is some player i for whom ambiguity aversion makes σ_i ex-ante (resp. sequentially optimal or SEA) strictly more belief robust). The argument is the same for each player, so for the remainder of the proof fix a player i .

We begin by proving that ambiguity aversion makes σ_i ex-ante more belief robust under the assumptions of the second part of the theorem. Observe that any strategies σ'_i that are weakly worse than σ_i (in terms of ex-ante expected utility, $\sum_{h \in H} u_i(h) p(\sigma'_i, \sigma_{-i})(h|h^0) \pi(h^0)$) for all $\pi \in \Pi_i$ can never interfere with optimality of σ_i and will thus, without loss of generality,

be ignored whenever making statements about strategies other than σ_i in what follows. For each l , denote by π_i^l the unique $\pi \in \Pi_i$ under which σ_i gives ex-ante expected utility e_i^l . The desired conclusion follows immediately if σ_i is weakly best for each $\pi \in \Pi_i$, as then σ_i is an ex-ante best response to σ_{-i} given any $\bar{\phi}_i$ and any $\bar{\mu}_i$. If $|\Pi_i| = 1$, σ_i an ex-ante best response to σ_{-i} given ϕ_i and μ_i immediately implies that σ_i is weakly best for each $\pi \in \Pi_i$. For the remainder of the argument, therefore, suppose that $|\Pi_i| > 1$ and there exists a strategy $\hat{\sigma}_i$ with strictly higher ex-ante expected utility than σ_i under π_i^l for some $1 \leq l \leq |\Pi_i|$.

We next show that all strategies must give weakly lower ex-ante expected utility than σ_i under π_i^1 . To see this, suppose, to the contrary, there exists a strategy $\hat{\sigma}_i$ with strictly higher ex-ante expected utility than σ_i under π_i^1 . Since σ is ex-ante robust to increased ambiguity aversion, σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^1$ and μ_i (and in particular is at least as good as $\hat{\sigma}_i$). Using (2.11) and applying the definition of $\hat{\phi}_i^1$ (Definition 2.20), this implies

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

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

Let \hat{M}_i denote the set of probabilities $\hat{\mu}_i$ with support contained in Π_i for which (2.11) holds with ϕ_i replacing $\hat{\phi}_i$. Let $m < |\Pi_i|$ be the smallest number l for which there exists a strategy that gives a strictly higher ex-ante expected utility than σ_i under π_i^{l+1} . By the previous paragraph, $m \geq 1$. Using (2.11) with ϕ_i replacing $\hat{\phi}_i$ and $\hat{\mu}_i \in \hat{M}_i$ together with the fact that all strategies give weakly lower ex-ante expected utility than σ_i under π_i^l for all $1 \leq l \leq m$, conclude that

$$\begin{aligned} \sigma_i \in \arg \max_{\sigma'_i} & \sum_{\pi \in \Pi_i} \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi(h^0) \right) \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi(h^0) \right) \hat{\mu}_i(\pi) \quad (\text{A.20}) \\ & + (b-1) \sum_{l=1}^m \left(\sum_{h \in H} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^0) \pi_i^l(h^0) \right) \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^l(h^0) \right) \hat{\mu}_i(\pi_i^l). \end{aligned}$$

Applying the definition of $\hat{\phi}_i^m$, we see that (A.20) is equivalent to (2.11) with $\hat{\phi}_i^m$ replacing $\hat{\phi}_i$. This shows that ambiguity aversion makes σ_i ex-ante more belief robust.

To complete the second part of the theorem, it remains to show the strict part of the result. By Lemma A.2, there exists some $b > 1$ and $\tilde{\mu}_i \notin \hat{M}_i$ with support Π_i such that σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. This optimality together with (2.11), where $\hat{\phi}_i^m$ with this b replaces $\hat{\phi}_i$ and for all $\hat{\mu}_i \in \hat{M}_i$, as was established above, implies that ambiguity aversion makes σ_i ex-ante strictly more belief robust.

Consider now sequential optimality in the second part of the theorem. Assume for the remainder of the argument that $|\Pi_i| > 1$, as otherwise the result follows immediately because $\hat{\mu}_i = \mu_i$. Consider m as defined above (or $m = 1$ if σ_i is ex-ante weakly best for each $\pi \in \Pi_i$). Since σ is sequentially optimal robust to increased ambiguity aversion, σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i^m$ and μ_i , and for any information set I_i there exists an interim belief ν_{i,I_i} such that σ_i is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ and ν_{i,I_i} . Consider any $\hat{\mu}_i$ with the same supports as the μ_i for which σ_i is an ex-ante best response to σ_{-i} given ϕ_i and $\hat{\mu}_i$. By the ex-ante equilibrium argument above, σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i^m$ and $\hat{\mu}_i$. Given $\hat{\phi}_i^m$ and $\hat{\mu}_i$, derive $\hat{\nu}_{i,I_i}$ for information sets I_i reachable from some information set in \mathcal{I}_i^0 given σ_{-i} via the smooth rule using σ as the ex-ante equilibrium. By Theorem 2.1, limiting attention to such interim beliefs is without loss of generality. By Lemma A.1, σ_i an ex-ante best response to σ_{-i} given $\hat{\phi}_i^m$ and $\hat{\mu}_i$ implies σ_i is an interim best response to σ_{-i} at these information sets given $\hat{\phi}_i^m$ and $\hat{\nu}_{i,I_i}$. Extend $\hat{\nu}$ by setting $\hat{\nu}_{i,I_i} = \nu_{i,I_i}$ elsewhere. Thus σ_i is an interim best response to σ_{-i} also at these remaining information sets given $\hat{\phi}_i^m$ and $\hat{\nu}_{i,I_i}$, as this fact is not affected by the shift from μ_i to $\hat{\mu}_i$. This shows that ambiguity aversion makes σ_i sequentially optimal more belief robust. It remains to show the strict part of the result. By the argument above for the ex-ante equilibrium case using Lemma A.2 to produce $\tilde{\mu}_i$, σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. By an analogous construction as for $\hat{\nu}_{i,I_i}$, construct interim beliefs $\tilde{\nu}_{i,I_i}$ for each I_i given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$ such that σ_i is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ and $\tilde{\nu}_{i,I_i}$. Thus ambiguity aversion makes σ_i sequentially optimal strictly more belief robust.

Finally turn to SEA under the assumptions of the second part of the theorem. Assume that $|\Pi_i| > 1$, as otherwise the result follows immediately because $\hat{\mu}_i = \mu_i$. For establishing the existence of interim beliefs so that both interim best response and i 's part of smooth rule consistency are satisfied, the arguments in the SEA portion of Theorem 2.9 concerning the properties of interim beliefs constructed using equation (A.17, p. 48) will prove particularly helpful. Given any $\hat{\phi}_i$, for each I_i , construct an interim belief, $\hat{\nu}_{i,I_i}$, to equal $\bar{\nu}_{i,I_i}$ as defined in (A.17) except with $\hat{\mu}_i$ replacing μ_i and $\hat{\phi}_i$ replacing ϕ_i . Then σ together with these interim beliefs for player i satisfies player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ given $\hat{\mu}_i$ and $\hat{\phi}_i$. Showing that σ_i is an interim best response to σ_{-i} for player i at I_i given $\hat{\phi}_i$ and

$\hat{\nu}_{i,I_i}$, i.e., that (2.12) holds, is equivalent to showing that

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

since each element in the support of $\hat{\nu}_{i,I_i}$ is $\bar{\pi}^\pi$ for some π in the support of $\hat{\mu}_i$ such that $\pi(\{h^0 | h^{s(I_i)} \in I_i\}) > 0$. Substituting for $\bar{\pi}_{I_i, \sigma_{-i}}^\pi$ using (2.6, p. 10) and for $\hat{\nu}_{i,I_i}$ and $\bar{\pi}^\pi$ using the (A.17) construction, using the notation $\bar{p}_{-i, \sigma_{-i}}(h^{s(I_i)} | h^0) \equiv p_{-i, \sigma_{-i}}(h^{s(I_i)} | h^{m_i(I_i)}) \bar{p}_{-i, \sigma_{-i}}(h^{m_i(I_i)} | h^0)$ for any h such that $h^{s(I_i)} \in I_i$, noticing that $\pi(\{h^0 | h^{s(I_i)} \in I_i\}) = 0$ implies $\pi(h^0) = 0$ for all h such that $h^{s(I_i)} \in I_i$, so that we may add terms multiplied by $\pi(h^0)$ to the sum in the objective function without changing its value, and, finally, simplifying, yields

$$\sigma_i \in \arg \max_{\sigma'_i} \sum_{\pi \in \Pi_i} \left(\sum_{h | h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h | h^{s(I_i)}) \bar{p}_{-i, \sigma_{-i}}(h^{s(I_i)} | h^0) \pi(h^0) \right) \quad (\text{A.21}) \\ \cdot \hat{\phi}'_i \left(\sum_{h \in H} u_i(h) p_\sigma(h | h^0) \pi(h^0) \right) \hat{\mu}_i(\pi).$$

Since σ is SEA robust to increased ambiguity aversion, σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i^1$ and μ_i , and for any I_i there exists an interim belief ν_{i,I_i} , constructed as was $\hat{\nu}_{i,I_i}$ at the beginning of the SEA part of the proof except now using $\hat{\phi}_i^1$ and μ_i , such that σ_i is an interim best response to σ_{-i} given $\hat{\phi}_i^1$ and ν_{i,I_i} . From the definition of $\hat{\phi}_i^1$ and the assumption that π_i^1 is well-defined, (A.21) with $\hat{\mu}_i = \mu_i$ and $\hat{\phi}_i = \hat{\phi}_i^1$ with b large enough (i.e., $\hat{\phi}_i^1$ sufficiently concave) implies that, for each I_i , $\sigma'_i = \sigma_i$ must maximize the interim expected utility under π_i^1 ,

$$\sum_{h | h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h | h^{s(I_i)}) \bar{p}_{-i, \sigma_{-i}}(h^{s(I_i)} | h^0) \pi_i^1(h^0), \quad (\text{A.22})$$

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

For each I_i , let \hat{M}_{i,I_i} denote the set of probabilities $\hat{\mu}_i$ with support contained in Π_i for which (A.21) holds with ϕ_i replacing $\hat{\phi}_i$. Let $1 \leq m < |\Pi_i|$ be the smallest number l for which there exists a strategy that, under π_i^{l+1} , gives either a strictly higher ex-ante expected utility than σ_i or a strictly higher interim expected utility than σ_i at some I_i (or let $m = 1$

if such l does not exist). Consider any $\hat{\mu}_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ with support Π_i . By the argument above for ex-ante equilibrium, σ_i is an ex-ante best response to σ_{-i} given $\hat{\phi}_i^m$ and $\hat{\mu}_i$. Analogous to the construction of $\hat{\nu}_{i,I_i}$ at the beginning of the SEA part of the proof, for each I_i construct $\hat{\nu}_{i,I_i}^m$ such that σ_i is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ and $\hat{\nu}_{i,I_i}^m$. Substituting $\hat{\phi}_i^m$ for $\hat{\phi}_i$, (A.21) is equivalent to

$$\begin{aligned} \sigma_i \in \arg \max_{\sigma'_i} & \sum_{\pi \in \Pi_i} \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^{s(I_i)}) \bar{p}_{-i, \sigma_{-i}}(h^{s(I_i)}|h^0) \pi(h^0) \right) \\ & \cdot \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi(h^0) \right) \hat{\mu}_i(\pi) \\ & + (b-1) \sum_{l=1}^m \left(\sum_{h|h^{s(I_i)} \in I_i} u_i(h) p_{(\sigma'_i, \sigma_{-i})}(h|h^{s(I_i)}) \bar{p}_{-i, \sigma_{-i}}(h^{s(I_i)}|h^0) \pi_i^l(h^0) \right) \\ & \cdot \phi'_i \left(\sum_{h \in H} u_i(h) p_{\sigma}(h|h^0) \pi_i^l(h^0) \right) \hat{\mu}_i(\pi_i^l). \end{aligned} \quad (\text{A.23})$$

That (A.23) holds follows from the fact that $\sigma'_i = \sigma_i$ maximizes each of the two terms – the first by the fact that $\hat{\mu}_i \in \bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i}$, and the second since, by definition of m , under π_i^l for all $1 \leq l \leq m$, all strategies give weakly lower ex-ante expected utility than σ_i and weakly lower interim expected utility at each I_i than σ_i . Thus σ_i is an interim best response to σ_{-i} for player i at information set I_i given $\hat{\phi}_i^m$ and $\hat{\nu}_{i,I_i}^m$. Furthermore, by construction, σ together with interim beliefs $\hat{\nu}_{i,I_i}^m$ for player i satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^{\infty}$ given $\hat{\phi}_i^m$ and $\hat{\mu}_i$. Therefore ambiguity aversion makes σ_i SEA more belief robust.

It remains to show the strict part of the SEA result. The conclusion follows immediately when no l exists for defining m , as then σ_i is an ex-ante best response to σ_{-i} for player i given any $\bar{\phi}_i$ and any $\bar{\mu}_i$ and is an interim best response to σ_{-i} given any $\bar{\phi}_i$ and any $\bar{\nu}_{i,I_i}$ for each I_i , and σ together with any interim beliefs $\bar{\nu}_{i,I_i}$ for player i , constructed analogously to the construction of $\hat{\nu}_{i,I_i}$ at the beginning of the SEA part of the proof, satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^{\infty}$. If there exists a l for defining m , then by Lemma A.3 there exists some $b > 1$ and $\tilde{\mu}_i \notin \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ with support Π_i and interim beliefs $\tilde{\nu}_{i,I_i}$ for each I_i such that σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$ and is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ and $\tilde{\nu}_{i,I_i}$ for each I_i , and σ together with the interim beliefs $\tilde{\nu}_{i,I_i}$ for player i satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^{\infty}$ given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. These facts together with the SEA more belief robustness results already proven above given $\hat{\phi}_i^m$ with this b and for all $\hat{\mu}_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ with support Π_i , imply that ambiguity aversion makes σ SEA strictly more belief robust.

We conclude the proof by showing that ambiguity aversion makes σ_i ex-ante (resp. sequentially optimal or SEA) more belief robust under the assumptions of the first part of the theorem. Under these assumptions, π_i^1 is well defined but π_i^l for $l > 1$ may not be. If $|\Pi_i| = 1$, then $\hat{M}_i = \{\mu_i\}$ and thus $\sigma \in E_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ (resp. $\sigma \in Q_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$ or $\sigma \in S_\Gamma((\mu_i)_{i \in N}, (\phi_i)_{i \in N})$) immediately implies that ambiguity aversion makes σ_i ex-ante (resp. sequentially optimal or SEA) more belief robust. Otherwise, $|\Pi_i| > 1$ and the same conclusion follows from the argument above for ex-ante (resp. sequentially optimal or SEA) in the second part of the theorem, with the exception that $\hat{\phi}_i^1$ should be used in place of $\hat{\phi}_i^m$ everywhere and excluding entirely the arguments used only for the strictly more belief robust part of the conclusions in the second part of the theorem.

Lemma A.2 *Fix a player i . If, additionally to the assumptions in the second part of Theorem 2.12, there exists $1 \leq m < |\Pi_i|$ such that m is the smallest number l for which there exists a strategy that gives a strictly higher ex-ante expected utility than σ_i under π_i^{l+1} , then there exists some $b > 1$ and $\tilde{\mu}_i \notin \hat{M}_i$ with support Π_i such that σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$.*

■

Proof of Lemma. Under the assumptions of the lemma, the definition of m implies that each σ'_i gives a weakly lower ex-ante expected utility than σ_i under each π_i^l for $l = 1, \dots, m$. For each $\hat{\sigma}_i$ that gives a strictly higher ex-ante expected utility than σ_i under π_i^{m+1} , $\hat{\sigma}_i$ gives a strictly lower ex-ante expected utility than σ_i for some π_i^l for $l = 1, \dots, m$. This is necessary to avoid $\hat{\sigma}_i$ being used together with $\hat{\phi}_i^{m+1}$ and μ_i to generate a contradiction to σ_i being ex-ante robust to increased ambiguity aversion. Consider the following function of σ'_i for $\lambda \in [0, 1]$,

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

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

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

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

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

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

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

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

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

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

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

One can verify that for

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

the version of the objective function in (2.11) given ϕ_i and $\check{\mu}_i$ is, up to normalization, equal to the version of the objective function in (2.11) given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. Thus σ_i being an ex-ante best response to σ_{-i} for player i given ϕ_i and $\check{\mu}_i$ is equivalent to σ_i being an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ (with this b) and $\tilde{\mu}_i$. The former holds since $\check{\mu}_i \in \hat{M}_i$, and thus we have established the latter. This proves that this b and $\check{\mu}_i$ satisfy the conclusion of the lemma. ■

Lemma A.3 *Fix a player i . If, additionally to the assumptions in the second part of The-*

orem 2.12, there exists $1 \leq m < |\Pi_i|$ such that m is the smallest number l for which there exists a strategy that, under π_i^{l+1} , gives either a strictly higher ex-ante expected utility than σ_i or a strictly higher interim expected utility than σ_i at some I_i , then there exists some $b > 1$ and $\tilde{\mu}_i \notin \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ with support Π_i and interim beliefs $\tilde{\nu}_{i,I_i}$ for each I_i such that σ_i is an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$ and is an interim best response to σ_{-i} given $\hat{\phi}_i^m$ and $\tilde{\nu}_{i,I_i}$ for each I_i , and σ together with the interim beliefs $\tilde{\nu}_{i,I_i}$ for player i satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$.

Proof. Under the assumptions of the lemma, the definition of m implies that, under each π_i^l for $l = 1, \dots, m$, each σ'_i gives a weakly lower ex-ante expected utility than σ_i and a weakly lower interim expected utility than σ_i at each I_i . Define $\xi(\sigma'_i; \lambda)$ as in the proof of A.2. For any $\hat{\sigma}_i$ that gives a strictly higher ex-ante expected utility than σ_i under π_i^{m+1} , there exists a $\underline{\lambda}(\hat{\sigma}_i)$ as defined in the proof of Lemma A.2. Consider any I_i for which there exists some $\hat{\sigma}_i$ that gives a strictly higher interim expected utility than σ_i at I_i under π_i^{m+1} . Any such $\hat{\sigma}_i$ gives a strictly lower interim expected utility than σ_i at I_i for some π_i^l for $l = 1, \dots, m$. This is necessary to avoid I_i and $\hat{\sigma}_i$ being used together with $\hat{\phi}_i^{m+1}$ and μ_i to generate a contradiction to σ_i being SEA robust to increased ambiguity aversion. Consider the function $\xi_{i,I_i}(\sigma'_i; \lambda)$ of σ'_i for $\lambda \in [0, 1]$, defined similarly to $\xi(\sigma'_i; \lambda)$ with interim expected utilities replacing ex-ante expected utilities. Observe that $\sigma_i \notin \arg \max_{\sigma'_i} \xi_{i,I_i}(\sigma'_i; 0)$ while $\{\sigma_i\} = \arg \max_{\sigma'_i} \xi_{i,I_i}(\sigma'_i; 1)$. Thus for each $\hat{\sigma}_i$ that gives a strictly higher interim expected utility than σ_i at I_i under π_i^{m+1} , there exists some $\underline{\lambda}_{i,I_i}(\hat{\sigma}_i) \in (0, 1)$ such that, for $\lambda \in [0, 1]$, $\lambda \geq \underline{\lambda}_{i,I_i}(\hat{\sigma}_i)$ if and only if

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

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

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

Since $\mu_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ and $\hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$ is closed and convex, there exists $\alpha \in (0, 1)$ and $\lambda \in (0, \underline{\lambda})$ such that the $\tilde{\mu}_i$ defined in the proof of Lemma A.2 satisfies $\tilde{\mu}_i \notin \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i,I_i} \right)$. By construction, $\tilde{\mu}_i$ has the same support as μ_i . By the properties

of $\underline{\lambda}$ and convexity of \hat{M}_i , for any $\gamma \in [0, 1]$, the $\check{\mu}_i$ defined in the proof of Lemma A.2 satisfies $\check{\mu}_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i, I_i} \right)$. One can verify that for the b defined in the proof of Lemma A.2, the version of the objective functions in (2.11) and (A.21) given ϕ_i and $\check{\mu}_i$ are, up to normalization, correspondingly equal to the versions of the objective functions in (2.11) and (A.21) given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. Analogous to the construction of $\hat{\nu}_{i, I_i}$ at the beginning of the SEA part of the proof of Theorem 2.12, for each I_i construct interim beliefs $\check{\nu}_{i, I_i}$ given $\hat{\phi}_i^m$ with this b and $\tilde{\mu}_i$ and interim beliefs $\check{\nu}_{i, I_i}$ given ϕ_i and $\check{\mu}_i$. Thus σ_i being an ex-ante best response to σ_{-i} for player i given ϕ_i and $\check{\mu}_i$ and an interim best response to σ_{-i} given ϕ_i and $\check{\nu}_{i, I_i}$ for each I_i is equivalent to σ_i being an ex-ante best response to σ_{-i} for player i given $\hat{\phi}_i^m$ with this b and $\tilde{\mu}_i$ and an interim best response to σ_{-i} given $\hat{\phi}_i^m$ and $\hat{\nu}_{i, I_i}^m$ for each I_i . The former holds since $\check{\mu}_i \in \hat{M}_i \cap \left(\bigcap_{I_i \in \mathcal{I}_i} \hat{M}_{i, I_i} \right)$, and thus we have established the latter. By construction, σ together with interim beliefs $\check{\nu}_{i, I_i}^m$ for player i satisfy player i 's part of smooth rule consistency using $\{\sigma^k\}_{k=1}^\infty$ given $\hat{\phi}_i^m$ and $\tilde{\mu}_i$. This proves that this b and $\tilde{\mu}_i$ satisfy the conclusion of the lemma. ■

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

$$p_I \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(\pi(I)) \geq \frac{3}{2} p_{II} \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(1 - \pi(I)). \quad (\text{A.24})$$

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

$$p_I \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(\pi(I)) \leq \frac{2}{3} p_{II} \sum_{\pi \in \Delta(\Theta)} \mu(\pi)(1 - \pi(I)). \quad (\text{A.25})$$

Since (A.24) and (A.25) 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 $\pi_0(I) = 1$ and $\pi_1(I) = 0$.⁸ Let $\phi(x) \equiv -e^{-x}$.⁹ 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 $\nu_{1,(\Pi, \eta)}^k(\pi_0) = 1$ and $\nu_{1,(\Pi, \eta)}^k(\pi_1) = 1$ for all partial

histories η . Furthermore, $\nu_{2,(\theta_2)}^k(\pi_0) = \frac{\frac{1}{2} \frac{\phi'((1-\frac{1}{2k})\lambda^*)}{\phi'((1-\frac{1}{2k})\lambda^*)}}{\frac{1}{2} \frac{\phi'((1-\frac{1}{2k})\lambda^*)}{\phi'(\lambda^*)} + \frac{1}{2} \frac{\phi'((1-\frac{1}{2k})(\lambda^* + \frac{5}{2}(1-\lambda^*)))}{\phi'((1-\frac{1}{2k})(\lambda^* + \frac{5}{2}(1-\lambda^*)))}} = \frac{1}{2}$, $\nu_{3,(\theta_3)}^k(\pi_0) = \frac{1}{2}$, $\nu_{2,(\theta_2, P)}^k(\pi_0) = \frac{\frac{1}{2}(1-\frac{1}{2k}) \frac{\phi'((1-\frac{1}{2k})\lambda^*)}{\phi'(\lambda^*)}}{\frac{1}{2}(1-\frac{1}{2k}) \frac{\phi'((1-\frac{1}{2k})\lambda^*)}{\phi'(\lambda^*)} + \frac{1}{2}(1-\frac{1}{2k}) \frac{\phi'((1-\frac{1}{2k})(\lambda^* + \frac{5}{2}(1-\lambda^*)))}{\phi'((\lambda^* + \frac{5}{2}(1-\lambda^*)))}} = \frac{1}{1+e^{-\frac{(5(1-\lambda^*))}{4k}}}$, 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 $\lim_{k \rightarrow \infty} \nu_{2,(\theta_2, P)}^k(\pi_0) =$

⁸The 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.

⁹Any more concave ϕ will also work, as will any ϕ more concave than $-e^{-\alpha x}$ for $\alpha = \frac{-4(\ln(2/3))}{5(2-\sqrt{2})} \approx 0.554$.

$\lim_{k \rightarrow \infty} \nu_{3,(\theta_3,P)}^k(\pi_0) = \lim_{k \rightarrow \infty} \nu_{2,(\theta_2)}^k(\pi_0) = \lim_{k \rightarrow \infty} \nu_{3,(\theta_3)}^k(\pi_0) = \frac{1}{2}$, 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 and, by construction, satisfies smooth rule consistency. It is therefore an SEA. ■

Rather than the “one-step ahead” formulation of smooth rule updating in (2.10) for an information set from an immediately prior one, one could alternatively (and equivalently) write the smooth rule as updating $\nu_{i,f_i(I_i)}$ to ν_{i,I_i} all at once:

For each information set I_i , for all π such that $\pi(f_i^{\sigma^{-i}}(I_i)) > 0$,

$$\begin{aligned} \nu_{i,I_i}(\pi) \propto & \frac{\phi'_i \left(\sum_{h|h^{m_i(I_i)} \in f_i(I_i)} u_i(h) p_\sigma(h|h^{m_i(I_i)}) \pi_{f_i(I_i), \sigma^{-i}}(h^{m_i(I_i)}) \right)}{\phi'_i \left(\sum_{h|h^t \in I_i} u_i(h) p_\sigma(h|h^t) \pi_{I_i, \sigma^{-i}}(h^t) \right)} \\ & \cdot \left(\sum_{h|h^t \in I_i, \sigma^{-i}} p_{-i, \sigma^{-i}}(h^t|h^{m_i(I_i)}) \pi_{f_i(I_i), \sigma^{-i}}(h^{m_i(I_i)}) \right) \nu_{i,f_i(I_i)}(\pi). \end{aligned} \quad (\text{A.26})$$

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

Proof of Proposition 3.3. 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.4, 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 (2.7) 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 $\gamma_{\tau,q}^k > 0$ is the probability that type τ of the incumbent chooses first period quantity q , $\lambda_q^k > 0$ is the probability that the entrant enters after observing quantity q , $\delta_{\tau,(q,enter,r)}^k > 0$ and $\delta_{(q,enter,r)}^k > 0$ are the probabilities of second period quantity r being chosen by, respectively, type τ of the incumbent and the entrant, after observing first period quantity q followed by entry, and $\delta_{\tau,(q,no\ entry,r)}^k > 0$ is the probability of second period quantity r being chosen by type τ of the incumbent after observing first period quantity q followed by no entry.

Specifically, let $\gamma_{\tau,q}^k \equiv \frac{\beta_{\tau,q}^k}{\sum_{\hat{q}} \beta_{\tau,\hat{q}}^k}$ for $k = 1, 2, \dots$, where $\beta_{\tau,q}^k$ is defined by

τ	q			
	$q = q_H$	$q = q_L$	$q_H \neq q < q_L$	$q > q_L$
L	1	k^2	1	k
M	1	k^2	1	1
H	k^2	1	k	1

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

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

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, $\lim_{k \rightarrow \infty} \pi_q^k(\tau)$, 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 $q = q_L$,

$$\lim_{k \rightarrow \infty} \pi_{q_L}^k(M) = \frac{\pi(M)}{\pi(L) + \pi(M)} \text{ and } \lim_{k \rightarrow \infty} \pi_{q_L}^k(L) = \frac{\pi(L)}{\pi(L) + \pi(M)}. \quad (\text{A.27})$$

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

$$\nu_{E,q_L}^k(\pi) \propto \frac{\phi' \left(\sum_{\tau} \pi(\tau) \sum_{q,x,y} \gamma_{\tau,q}^k \lambda_q^k \delta_{\tau,(q,enter,x)}^k \delta_{(q,enter,y)}^k w(x,y) \right)}{\phi' \left(\sum_{\tau} \pi_{q_L}^k(\tau) \sum_{x,y} \lambda_{q_L}^k \delta_{\tau,(q_L,enter,x)}^k \delta_{(q_L,enter,y)}^k w(x,y) \right)} \left(\sum_{\tau} \gamma_{\tau,q_L}^k \pi(\tau) \right) \mu(\pi),$$

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.27) yields,

$$\nu_{E,q_L}(\pi) \propto \frac{\phi'(\pi(H)w_H)}{\phi'(0)} \left(\frac{\pi(L) + \pi(M)}{2} \right) \mu(\pi), \quad (\text{A.28})$$

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 (2.8) 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:

$$\sum_{\pi | \pi(H) < 1} \nu_{E,q_L}(\pi) \frac{1}{1 - \pi(H)} (\pi(L)w_L + \pi(M)w_M) \phi'(0) \leq 0. \quad (\text{A.29})$$

Substituting (A.28) into (A.29) yields that not entering remaining optimal is equivalent to (3.5). Therefore (σ^*, ν) satisfies the interim optimality conditions (2.8) of sequential optimality.

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

Lemma A.4 *Under the assumptions of Proposition 3.3 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 proposition. We show that (3.5) is satisfied for concave enough ϕ . If $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M \leq 0\}) = 1$ then (3.5) is trivially satisfied for any ϕ . For the remainder of the proof, therefore, suppose that $\mu(\{\pi \mid \pi(L)w_L + \pi(M)w_M > 0\}) > 0$. Let $\Pi^- \equiv \{\pi \mid \pi(L)w_L + \pi(M)w_M < 0\}$, $\Pi^+ \equiv \{\pi \mid \pi(L)w_L + \pi(M)w_M > 0\}$, $N \equiv \sum_{\pi \in \Pi^-} \mu(\pi)(\pi(L)w_L + \pi(M)w_M)$, and $P \equiv \sum_{\pi \in \Pi^+} \mu(\pi)(\pi(L)w_L + \pi(M)w_M)$. Let $\pi^- \in \arg \max_{\pi \in \Pi^-} \pi(H)$ and $\pi^+ \in \arg \min_{\pi \in \Pi^+} \pi(H)$. The left-hand side of (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.5, $\frac{\pi^-(L)}{\pi^-(M)} > \frac{\pi^+(L)}{\pi^+(M)}$ 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.5 *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 $a1$ plays h after message m_l , $l = 1, 2$, and similarly let δ_l be the corresponding probabilities for $a2$. 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 $\pi_k(II B) + \pi_k(II R)$ 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

$$\begin{aligned} & \frac{1}{2} \sum_{k=1}^2 \phi_{a1}(\pi_k(I)[r_I(1 + \delta_1)\gamma_1 + (1 - r_I)(1 + \delta_2)\gamma_2] \\ & \quad + \pi_k(II)[r_{II}((1 + \delta_1)\gamma_1 + 5(1 - \gamma_1)) + (1 - r_{II})((1 + \delta_2)\gamma_2 + 5(1 - \gamma_2))]) \end{aligned} \quad (\text{A.30})$$

and $a2$ chooses δ_1, δ_2 to maximize

$$\begin{aligned} & \frac{1}{2} \sum_{k=1}^2 \phi_{a2}(\pi_k(I)[r_I((1 + \gamma_1)\delta_1 + 5(1 - \delta_1)) + (1 - r_I)((1 + \gamma_2)\delta_2 + 5(1 - \delta_2))] \\ & \quad + \pi_k(II)[r_{II}(1 + \gamma_1)\delta_1 + (1 - r_{II})(1 + \gamma_2)\delta_2]). \end{aligned} \quad (\text{A.31})$$

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 a_2 , as can be seen by the fact that the partial derivative of (A.31) with respect to δ_1 (resp. δ_2) evaluated at those strategies is

$$\frac{1}{2}\left(4 - 5 \sum_{k=1}^2 \pi_k(I)\right)\phi'_{a_2}(2) = -\frac{3}{8}\phi'_{a_2}(2) < 0.$$

Similarly, one can show that h_1 is not a best response for a_1 .

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 a_2 is that the partial derivatives of (A.31) 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.31) 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 a_1 because the partial derivative of (A.30) with respect to γ_1 evaluated at these strategies using the values for the π_k is,

$$\frac{3}{8}(2r_I - 1)\phi'_{a_1}\left(\frac{15r_I + 25}{20}\right) + \frac{1}{5}(r_I - 6)\phi'_{a_1}\left(\frac{4r_I + 36}{20}\right) < 0.$$

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 $\phi_{a_1}(x) = -e^{-11x}$ yields that the left-hand side is negative.

Case 4: When $r_{II} = 0$ and $0 < r_I \leq 1$, (A.31) 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 a_1 because the partial derivative of (A.30) with respect to γ_2 evaluated at these strategies using the values for the π_k is,

$$\frac{3}{4}\left(\frac{1}{2} - r_I\right)\phi'_{a_1}\left(2 - \frac{3}{4}r_I\right) + \left(-\frac{1}{5}r_I - 1\right)\phi'_{a_1}\left(2 - \frac{r_I}{5}\right) < 0.$$

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 $\phi_{a_1}(x) = -e^{-11x}$ yields that the left-hand side is negative. ■

Lemma A.6 *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 $\frac{h'(f(x))}{h'(f(y))} \geq 1$. ■

B Appendix: Perfect Equilibrium with Ambiguity

In this section we limit our analysis to multistage 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. 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 her type. There are no private observations as the game proceeds.

Definition B.1 *A (finite) extensive-form multistage 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.*

In such games, information sets all take the form $I_i = \{(\theta, h^t) \mid \theta_i = \tau_i\}$, where τ_i is player i 's component of the type and h^t is the partial history of players' actions up to but not including stage t . So in this section, for example, instead of writing ν_{i, I_i} , we write ν_{i, τ_i, h^t} .

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.

Notation B.1 $\Theta_{i,\tau_i,h^t} \equiv \{\theta \in \Theta \mid \theta_i = \tau_i \text{ and } p_{-i,\sigma_{-i},\theta}(h^t \mid h^{m_i(\{(\theta,h^t) \mid \theta_i = \tau_i\})}) > 0\}$.

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 B.2 *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 B.3 *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*

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

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

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

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 B.4 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 $\nu_{i, \tau_i, h}^P m_i(\{(\theta, h^{t-1}) | \theta_i = \tau_i\})$.

In the case where players are ambiguity neutral, Definition B.4 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 B.5 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 B.1 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 .

Proof of Corollary B.1. We show that $(\sigma^P, \hat{\nu}^P)$, where, for all i, τ_i , $\hat{\nu}_{i, \tau_i, h^t}^P = \nu_{i, \tau_i, h^t}^P$ whenever $(t > 0$ and $m_i(\{(\theta, h^t) | \theta_i = \tau_i\}) = t)$, and where, everywhere else, $\hat{\nu}_{i, \tau_i, h^t}^P$ 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. ■

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

Theorem B.1 *Fix an SEA, (σ^S, ν^S) , of a game Γ . (σ^S, ν^S) is also a PEA of Γ .*

Proof of Theorem B.1. That (σ^S, ν^S) satisfies sequential optimality follows directly from the definition of a SEA. By Lemma B.1, 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 . ■

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.

Lemma B.1 *Any $(\hat{\sigma}, \hat{\nu})$ satisfying smooth rule consistency also naturally extends updating.*

Proof of Lemma B.1. 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 (2.10) simplifies to

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

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} , $p_{-i,\sigma_{-i},\theta}(h^t|h^{t-1}) = \prod_{k \neq i,j} \sigma_k(h^{t-1}, \theta_k)(h_{t-1,k})$. 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,

$$\sum_{\pi \text{ s.t. } \nu_{i,\tau_i,h^{t-1}}(\pi) > 0 \text{ and } \pi(\Theta_{i,\tau_i,h^t}) > 0} \pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) \nu_{i,\tau_i,h^t}(\pi) \quad (\text{B.1})$$

$$= A \sum_{\pi \text{ s.t. } \nu_{i,\tau_i,h^{t-1}}(\pi) > 0 \text{ and } \pi(\Theta_{i,\tau_i,h^t}) > 0} \pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) \nu_{i,\tau_i,h^{t-1}}(\pi) \cdot \left(\sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t}} \left(\prod_{k \neq i,j} \sigma_k(h^{t-1}, \hat{\theta}_k) (h_{t-1,k}) \right) \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \right). \quad (\text{B.2})$$

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

$$\begin{aligned} & \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^t}}(\{\theta_j\} \times \Theta_{-j}) \nu_{i,\tau_i,h^t}(\pi) \quad (\text{B.3}) \\ &= A \sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t} \cap (\{\theta_j\} \times \Theta_{-j})} \left(\prod_{k \neq i,j} \sigma_k(h^{t-1}, \hat{\theta}_k) (h_{t-1,k}) \right) \\ & \cdot \sum_{\pi \text{ s.t. } \nu_{i,\tau_i,h^{t-1}}(\pi) > 0 \text{ and } \pi(\Theta_{i,\tau_i,h^t}) > 0} \pi_{\Theta_{i,\tau_i,h^{t-1}}}(\hat{\theta}) \nu_{i,\tau_i,h^{t-1}}(\pi). \end{aligned}$$

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 (B.3) yields,

$$\begin{aligned}
& \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^t}(\{\theta_j\} \times \Theta_{-j})} \nu_{i,\tau_i,h^t}(\pi) \tag{B.4} \\
&= A \sum_{\hat{\theta} \in \Theta_{i,\tau_i,h^t} \cap (\{\theta_j\} \times \Theta_{-j})} \left(\prod_{k \neq i,j} \sigma_k(h^{t-1}, \hat{\theta}_k)(h_{t-1,k}) \right) \\
&\quad \cdot \left(\prod_{k \neq i,j} \sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\hat{\theta}_k\} \times \Theta_{-k})} \nu_{i,\tau_i,h^{t-1}}(\pi) \right) \left(\sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\theta_j\} \times \Theta_{-j})} \nu_{i,\tau_i,h^{t-1}}(\pi) \right) \\
&= B \left(\sum_{\pi \in \Delta(\Theta)} \pi_{\Theta_{i,\tau_i,h^{t-1}}(\{\theta_j\} \times \Theta_{-j})} \nu_{i,\tau_i,h^{t-1}}(\pi) \right)
\end{aligned}$$

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 B.4) 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. ■

B.0.3 Example 4: 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 a PEA when player 2 is sufficiently ambiguity averse, but never occurs as part of a PEA when player 2 has expected utility preferences (and is thus ambiguity neutral). Furthermore, under the PEA we construct, 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 B.1.

There are three players: 1, 2 and 3. First, it is determined whether player 3 is of type I or type II and 3 observes his own type. Players 1 and 2 do not observe the type. The payoff triples in Figure B.1 describe vNM utility payoffs given players’ actions and players’ types (i.e., (u_1, u_2, u_3) means that player i receives u_i). 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.¹⁰

¹⁰Note that to eliminate any possible effects of varying 2’s risk attitude, think of the payoffs of player 2




Figure B.1: 3-Player Game

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 a PEA?

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

Proposition B.1 *There exist ϕ_2 and μ_2 (e.g., $\phi_2(x) = -e^{-2x}$ and $\mu_2(\frac{3}{4}) = \mu_2(\frac{1}{4}) = \frac{1}{2}$ where $\frac{3}{4}$ and $\frac{1}{4}$ denote probabilities of type I) such that C is played on the equilibrium path with probability 1 as part of a PEA.*

Proof of Proposition B.1. 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 the distribution used to determine player 3's type is given by $\hat{p} \geq 1/2$ or $\check{p} = 1 - \hat{p}$. Specifically, let $\mu_2(\hat{p}) = \mu_2(\check{p}) = 1/2$. We now show that $(R, F, (C, C), \lambda^*H + (1 - \lambda^*)G)$ is a PEA given $\hat{p} = \frac{3}{4}$ and $\phi_2(x) = -e^{-2x}$ and defining λ^* by

$$\lambda^* = \arg \max_{\lambda \in [0,1]} \frac{1}{2} \phi_2(2\lambda + 6(1 - \lambda)\hat{p}) + \frac{1}{2} \phi_2(2\lambda + 6(1 - \lambda)\check{p}). \quad (\text{B.5})$$

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

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.

of ϕ_3 and μ_3) if and only if $\lambda^* \geq 1/2$. First-order conditions for an interior λ^* are given by

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

Notice that the left-hand side of (B.6) 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 (B.6) is non-negative at $\lambda^* = 1/2$ i.e.,

$$(1 - 3\hat{p})\phi'_2(1 + 3\hat{p}) + (3\hat{p} - 2)\phi'_2(1 + 3(1 - \hat{p})) \geq 0. \quad (\text{B.7})$$

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 updating and naturally extends updating and verify that interim optimality is satisfied. The smooth rule using σ (in particular, applying the formula in (A.26)) has bite for player 2 only following the play of R and C :

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

Thus,

$$\nu_{2,(R,C)}\left(\frac{3}{4}\right) = \frac{1}{2} = \nu_{2,(R,C)}\left(\frac{1}{4}\right).$$

By the definition of λ^* (via (B.5)), since $\nu_{2,(R,C)}\left(\frac{3}{4}\right) = \frac{1}{2}$, player 2 is indeed best responding after 1 plays R and 3 plays C .

Since the smooth rule imposes no restrictions immediately following the off-path play of T , suppose player 2 maintains beliefs μ_2 following T (so that $\nu_{2,T}(p) = \frac{1}{2}$). This specification of beliefs satisfies naturally extends updating. (Note that even if 2's updating did change his beliefs after T , since 2 *has* costly ambiguity exposure under the given strategies prior to 1 playing T ,

$$\phi'_2(2\lambda^* + 6(1 - \lambda^*)\hat{p}) \neq \phi'_2(2\lambda^* + 6(1 - \lambda^*)(1 - \hat{p})),$$

naturally extending updating has no bite here.) Then 2's best response following T is found by solving

$$\gamma^* = \arg \max_{\gamma \in [0,1]} \frac{1}{2}\phi(2\gamma + 4(1 - \gamma)(1 - \hat{p})) + \frac{1}{2}\phi(2\gamma + 4(1 - \gamma)(1 - \check{p})).$$

Differentiating, since ϕ is increasing and concave and $\hat{p} > 1/2$, we have

$$\begin{aligned} & \frac{1}{2}\phi'(2\gamma + 4(1-\gamma)(1-\hat{p}))(2-4(1-\hat{p})) + \frac{1}{2}\phi'(2\gamma + 4(1-\gamma)\hat{p})(2-4\hat{p}) \\ \geq & \frac{1}{2}\phi'(2\gamma + 4(1-\gamma)\hat{p})(2-4(1-\hat{p}) + 2-4\hat{p}) \\ = & 0. \end{aligned}$$

Thus $\gamma^* = 1$ and, under these beliefs, F is indeed the best response of player 2 if player 1 plays T .

Putting everything together, $(R, F, (C, C), \lambda^*H + (1-\lambda^*)G)$ is sequentially optimal with respect to beliefs satisfying naturally extends updating. Therefore all the conditions for a PEA are satisfied. ■

The proof of Proposition B.1 relies on ambiguity aversion on the part of player 2. Our next result shows that this is essential:

Proposition B.2 *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.*

Proof of Proposition B.2. 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 B.1, we showed that no sequential optimum 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 (2.7) of sequential optimality (and thus PEA) when 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 (2.7) places no restrictions on 2's play following T . However, any beliefs for player 2 satisfying the naturally extends updating condition

of PEA (Definition B.4 in Appendix B) following T must continue to place weight γ on 3 being of type I because 3 has only one action 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, so that C is never played on the equilibrium path. ■

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

In light of Theorem 2.9 and the result that C can't be played as part of a PEA or SEA under ambiguity neutrality, it follows that (1) the PEA strategy profile we identified in Proposition B.1 cannot be an SEA strategy profile, and (2) ambiguity aversion can strictly expand PEA play compared to ambiguity neutrality.