

PRUDENT RATIONALIZABILITY IN GENERALIZED EXTENSIVE-FORM GAMES*

Aviad Heifetz[†] Martin Meier[‡] Burkhard C. Schipper[§]

March 30, 2011

Abstract

We define an extensive-form analogue of iterated admissibility, called Prudent Rationalizability (PR). In each round of the procedure, for each information set of a player a surviving strategy of hers is required to be rational vis-a-vis a belief system with a *full-support* belief on the opponents' previously surviving strategies that reach that information set. Somewhat surprisingly, prudent rationalizable strategies may not refine the set of Extensive-Form Rationalizable (EFR) strategies (Pearce 1984). However, we prove that the paths induced by PR strategy-profiles (weakly) refine the set of paths induced by EFR strategies.

PR applies also to generalized extensive-form games which model mutual unawareness of actions (Heifetz, Meier and Schipper, 2011a). We demonstrate the applicability of PR in the analysis of verifiable communication, and show that it yields the same, full information unraveling prediction as does the unique sequential equilibrium singled out by Milgrom and Roberts (1986); yet, we also show that under unawareness full unraveling might fail.

Keywords: Prudent rationalizability, caution, extensive-form rationalizability, extensive-form games, unawareness, verifiable communication.

JEL-Classifications: C70, C72, D80, D82.

*The material of this paper was previously part of our paper “Dynamic Unawareness and Rationalizable Behavior”. We are grateful to Pierpaolo Battigalli for numerous insightful comments and suggestions and to Elchanan Ben Porath for helpful discussions and exchange of ideas. We thank Paulo Borelli, Andy McLennan, Yoram Halevy as well as to seminar participants at Barcelona, Bocconi, Caltech, Maryland, Maastricht, Pittsburgh, USC, UC Davis, Tel Aviv, Stony Brook 2007, LOFT 2008, Games 2008 and NSF/NBER/CEME 2009. Aviad is grateful for financial support from the Open University of Israel's Research Fund grant no. 46106. Martin was supported by the Spanish Ministerio de Educación y Ciencia via a Ramon y Cajal Fellowship (IAE-CSIC) and a Research Grant (SEJ 2006-02079). Burkhard is grateful for financial support from the NSF SES-0647811.

[†]The Economics and Management Department, The Open University of Israel. Email: aviadhe@openu.ac.il

[‡]Institut für Höhere Studien, Wien, and Instituto de Análisis Económico - CSIC, Barcelona. Email: martin.meier@ihs.ac.at

[§]Department of Economics, University of California, Davis. Email: bcschipper@ucdavis.edu

1 Introduction

In normal-form games, iterated admissibility (*IA*) is a refinement of rationalizability. In the latter solution concept, at every round of elimination a player's strategy survives only if it is a best-reply to some belief over the opponents' strategies which survived the previous rounds, while in the former a strategy survives only if it is a best reply to such a *full-support* belief, which doesn't completely exclude any strategy of the other players that has not been thus far eliminated. In this paper we investigate the connection between the counterparts of these two notions in *dynamic games*.

The prominent analogue for rationalizability in dynamic games is Extensive-Form Rationalizability (*EFR*) (Pearce, 1984, Battigalli, 1997). *EFR* is particularly interesting because it may be used to refine the notion of sequential equilibrium (Pearce, 1984). Moreover, in generic perfect-information games *EFR* induces the unique backward-induction path (Reny, 1992, Battigalli, 1997, Robles, 2006, Perea, forthcoming), even though the *EFR* strategies may be distinct from the backward-induction strategies (Reny, 1992).

EFR is a notion that captures forward induction: at every information set, the active player looks for a best rationalization for the way this information set has been reached (in terms of her opponents' rationality, their belief in *their* opponents' rationality etc.), and replies optimally to a belief on these best-rationalizable strategies. In this paper we define *Prudent Rationalizability (PR)* in a likewise fashion by additionally requiring this belief to have *full-support* on the opponents' (recursively defined) prudently-best-rationalizable strategies. In Theorem 1 we prove that *PR* strategies exist in every dynamic game, including generalized extensive-form games (Heifetz, Meier and Schipper 2011a) which allow for the modeling of mutual unawareness of actions.

In normal-form games every *IA* strategy is also rationalizable. Somewhat surprisingly, we show that a similar inclusion does not obtain in dynamic games: in Section 4 we bring an example of a game in which a player's *PR* strategies is not a subset of her *EFR* strategies but is rather *disjoint* from it. Nevertheless, in Theorem 2 we prove that inclusion does obtain in terms of *outcomes*: The set of paths induced by *PR* strategy profiles is always contained in the set of paths induced by *EFR* strategy profiles.

In Section 5 we exemplify the attractiveness of *PR* in the Milgrom-Roberts (1986) model of verifiable communication. They proved that the model has a unique sequential equilibrium, and that in this unique equilibrium all the asymmetric information gets unraveled. We show that prudent rationalizability is sufficient to entail the same result

– the unique sequential equilibrium outcome is also the unique PR outcome.

Full unraveling of information is somewhat unrealistic, though. In Section 6 we show that when unawareness is introduced into the model, PR need not necessarily entail full information unraveling. We also analyze a sender-receiver game with unawareness introduced by Ozbay (2007), and show that PR delivers the same prediction as does his equilibrium notion which incorporates forward-induction reasoning.

These two applications demonstrate the attractiveness of PR in dynamic games with unawareness. Indeed, in many games with unawareness PR rules out implausible EFR strategies, with which a player makes an opponent aware of an action which the player would actually like the opponent to *avoid*, just because the player has a firm belief that the opponent wouldn't take it (even if the opponent is indifferent between the revealed action and another one, of which she was aware also before); PR rules out such imprudent behavior.

PR is equivalent to *iterated admissibility on the tree* as defined by Brandenburger and Friedenberg (2007). Iterated admissibility on the tree is analogous to iterated dominance conditional on normal-form information sets à la Shimoji and Watson (1994), in which strict dominance is replaced by weak dominance. Brandenburger and Friedenberg (2007) show that iterated admissibility on the tree is equivalent to IA of the strategic form of the game. IA does not require conditioning on normal-form information sets. In Heifetz, Meier and Schipper (2011b) we show, however, that such an analogous equivalence for dynamic games with unawareness still requires normal-form information sets since they encode also the awareness of players.

Some readers may be interested in PR independently of unawareness, and to this effect we first introduce the concept in Sections 2 and 3 within standard finite extensive-form games with perfect recall, finite horizon, and possibly simultaneous moves. In this standard setting we also demonstrate some of the refining power of prudent rationalizability in Sections 4 and 5. In Section 5, we apply prudent rationalizability to the Milgrom-Roberts (1986) communication game. Only in Section 6 do we introduce generalized extensive-form games that allow for unawareness and develop the general results on prudent rationalizability.

2 Extensive-form Games

We consider finite extensive-form games with finite horizon, perfect recall and possibly simultaneous moves (for standard properties see Dubey and Kaneko, 1984, and Osborne and Rubinstein, 1994). To fix notation, denote by I the finite set of players, by N the finite set of decision nodes, by I_n the active players at node n , by A_n^i the finite action set of player $i \in I_n$ (for $n \in N$), by C the chance nodes, and by Z the terminal nodes with a payoff vector $(p_i^z)_{i \in I} \in \mathbb{R}^I$ for the players for every $z \in Z$. The nodes $\bar{N} = N \cup C \cup Z$ shall constitute a tree. We denote by N_i the set of nodes in which player $i \in I$ is active. The information set of player i at node n is denoted by $\pi_i(n)$. Let H_i be the set of i 's information sets. For two information sets h_i, h'_i , we say that h_i precedes h'_i (or that h'_i succeeds h_i) if for every $n' \in h'_i$ there is a path n, \dots, n' such that $n \in h_i$. We denote it by $h_i \rightsquigarrow h'_i$. Standard properties on information sets imply that if $n', n'' \in h_i$ then $A_{n'}^i = A_{n''}^i$. Thus, if $n \in h_i$ we write also $A_{h_i}^i$ for A_n^i .

A (pure) strategy

$$s_i \in S_i \equiv \prod_{h_i \in H_i} A_{h_i}^i$$

for player i specifies an action of player i at each of her information sets $h_i \in H_i$. Denote by

$$S = \prod_{j \in I} S_j$$

the set of strategy profiles in the game.

If $s_i = (a_{h_i})_{h_i \in H_i} \in S_i$, we denote by

$$s_i(h_i) = a_{h_i}$$

the player's action at the information set h_i . If player i is active at node n , we say that at node n the strategy prescribes to her the action $s_i(\pi_i(n))$.

We say that a strategy profile $s \in S$ reaches the information set $h_i \in H_i$ if the players' actions and nature's moves (if there are any) lead to h_i with a positive probability. We say that the strategy $s_i \in S_i$ reaches the information set h_i if there is a strategy profile $s_{-i} \in S_{-i}$ of the other players such that the strategy profile (s_i, s_{-i}) reaches h_i . Otherwise, we say that the information set h_i is excluded by the strategy s_i . Similarly, we say that the strategy profile $s_{-i} \in S_{-i}$ reaches the information set h_i if there exists a strategy $s_i \in S_i$ such that the strategy profile (s_i, s_{-i}) reaches h_i . A strategy profile $(s_j)_{j \in I}$

reaches a node n if the players' actions $s_j (\pi_j (n'))_{j \in I}$ and nature's moves in the nodes n' lead to n with a positive probability. Since we consider only finite trees, $(s_j)_{j \in I}$ reaches an information set $h_i \in H_i$ if and if there is a node $n \in h_i$ such that $(s_j)_{j \in I}$ reaches n .

For an information set h_i , let $s_i/\tilde{s}_i^{h_i}$ denote the strategy that is obtained by replacing actions prescribed by s_i at the information set h_i and its successors by actions prescribed by \tilde{s}_i . The strategy $s_i/\tilde{s}_i^{h_i}$ is called an h_i -replacement of s_i .

The set of *behavioral* strategies is

$$\prod_{h_i \in H_i} \Delta(A_{h_i}).$$

3 Prudent Rationalizability

A *belief system* of player i

$$b_i = (b_i(h_i))_{h_i \in H_i} \in \prod_{h_i \in H_i} \Delta(S_{-i})$$

is a profile of beliefs - a belief $b_i(h_i) \in \Delta(S_{-i})$ on the other players' strategies, for each information set $h_i \in H_i$, with the following properties

- $b_i(h_i)$ reaches h_i , i.e. $b_i(h_i)$ assigns probability 1 to the set of strategy profiles of the other players that reach h_i .
- If h_i precedes h'_i ($h_i \rightsquigarrow h'_i$) then $b_i(h'_i)$ is derived from $b_i(h_i)$ by Bayes rule whenever possible.

Denote by B_i the set of player i 's belief systems.

For a belief system $b_i \in B_i$, a strategy $s_i \in S_i$ and an information set $h_i \in H_i$, define player i 's expected payoff **at** h_i to be the expected payoff for player i given $b_i(h_i)$, the actions prescribed by s_i at h_i and its successors, and conditional on the fact that h_i has been reached.¹

¹Even if this condition is counterfactual due to the fact that the strategy s_i does not reach h_i . The conditioning is thus on the event that nature's moves, if there are any, have led to the information set h_i , and assuming that player i 's past actions (in the information sets preceding h_i) have led to h_i *even if these actions are distinct than those prescribed by s_i* .

We say that with the belief system b_i and the strategy s_i player i is *rational* at the information set $h_i \in H_i$ if there exists no action $a'_{h_i} \in A_{h_i}$ such that only replacing the action $s_i(h_i)$ by a'_{h_i} results in a new strategy s'_i which yields player i a higher expected payoff at h_i given the belief $b_i(h_i)$ on the other players' strategies S_{-i} .

Definition 1 (Prudent rationalizability) *Let*

$$\bar{S}_i^0 = S_i$$

For $k \geq 1$ define inductively

$$\bar{B}_i^k = \left\{ b_i \in B_i : \begin{array}{l} \text{for every information set } h_i, \text{ if there exists some profile} \\ s_{-i} \in \bar{S}_{-i}^{k-1} = \prod_{j \neq i} \bar{S}_j^{k-1} \text{ of the other players' strategies} \\ \text{such that } s_{-i} \text{ reaches } h_i, \text{ then the support} \\ \text{of } b_i(h_i) \text{ is the set of strategy profiles } s_{-i} \in \bar{S}_{-i}^{k-1} \text{ that reach } h_i \end{array} \right\}$$

$$\bar{S}_i^k = \left\{ s_i \in \bar{S}_i^{k-1} : \begin{array}{l} \text{there exists } b_i \in \bar{B}_i^k \text{ such that for all } h_i \in H_i \text{ player } i \\ \text{is rational at } h_i \end{array} \right\}$$

The set of prudent rationalizable strategies of player i is

$$\bar{S}_i^\infty = \bigcap_{k=1}^{\infty} \bar{S}_i^k$$

At each level each player and each information set of the player, she forms full support beliefs over the opponents' strategies remaining from the previous level and reaching this information set. The player retains any of her previous level strategies for which there exists such a belief such that the strategy is rational at all information sets.

Theorem 1 below implies that every player's set of prudent rationalizable strategies is non-empty.

4 Extensive-Form Rationalizability versus Prudent Rationalizability

4.1 Extensive-Form Rationalizability

Pearce (1984) defined extensive-form (correlated) rationalizable strategies by a procedure of an iterative elimination of strategies. The inductive definition below is essentially the version in Battigalli (1997).²

Definition 2 (Extensive-form rationalizable strategies) *Define, inductively, the following sequence of belief systems and strategies of player i .*

$$B_i^1 = B_i$$

$$S_i^1 = \{s_i \in S_i: \text{there exists a belief system } b_i \in B_i^1 \text{ with which for every information set } h_i \in H_i \text{ player } i \text{ is rational at } h_i\}$$

⋮

$$B_i^k = \{b_i \in B_i^{k-1}: \text{for every information set } h_i, \text{ if there exists some profile of the other players' strategies } s_{-i} \in S_{-i}^{k-1} = \prod_{j \neq i} S_j^{k-1} \text{ such that } s_{-i} \text{ reaches } h_i, \text{ then } b_i(h_i) \text{ assigns probability 1 to } S_{-i}^{k-1}\}$$

$$S_i^k = \{s_i \in S_i: \text{there exists a belief system } b_i \in B_i^k \text{ with which for every information set } h_i \in H_i \text{ player } i \text{ is rational at } h_i\}$$

The set of player i 's **extensive-form rationalizable strategies** is


$$S_i^\infty = \bigcap_{k=1}^{\infty} S_i^k.$$

The set of extensive-form rationalizable strategies (Pearce 1984, Battigalli 1997) is nonempty.

²Battigalli's (1997) definition differs from Pearce's (1984) in at least two respects. First, Battigalli allows for correlated beliefs over opponents' strategies. Second, Battigalli's definition defines a procedure of an iterative elimination of beliefs. This allows for a more intuitive interpretation as a reasoning procedure. Nevertheless, Battigalli shows that when one allows for correlation in Pearce's original definition, then both procedures are equivalent. Definition 2 differs in one respect from Battigalli (1997) as he uses optimization over replacements of strategies whereas we optimize over actions. See Heifetz, Meier and Schipper (2011b) for further discussions and results.

Some normal-form games have Nash equilibria in weakly dominated strategies. Similarly, the game in Figure 1.³ is an example of a perfect-information extensive-form game with a subgame-perfect equilibrium involving a strategy which is *EFR* but not *PR*. To wit, this is the subgame-perfect equilibrium (af, dg) . The strategy af is *EFR* for player 1 – it is supported by the belief system which starts at the root with the belief that player 2 will play dg , revised completely at 1’s second decision node by the belief that 2 is playing cg ; the strategy dg is *EFR* for player 2 with the belief system with which player 2 is initially certain that 1 is playing af , revised completely at 2’s second decision node to the belief that 1 played ac . However, af is rational for player 1 at the root for no full-support belief on 2’s entire strategy set $\{dg, dh, cg, ch\}$, and hence af is not *PR*.


Figure 1:



4.2 PR refines EFR outcomes, but not EFR strategies

In normal-form games, iterated admissibility is a refinement of rationalizability. Somewhat surprisingly, in extensive-form games prudent rationalizability is *not* a refinement of extensive-form rationalizability, as the following example (Figure 2) demonstrates.

Figure 2:



³This example is from recent work by one of the authors with Ronen Gradwohl.

In this example, player 1 can guarantee herself the payoff 6 by choosing a and ending the game. If player 2 is called to play, should he believe that player 1 chose b or c ? If player 1 is certain that player 2 is rational, she is certain that player 2 will not choose f . Hence, if player 2 is certain that player 1 is certain that he (player 2) is rational, then at his information set player 2 is certain that player 1 chose c . The reason is that among player 1's actions leading to 2's information set, c is the only action which, assuming 2 believes c was chosen and that 2 is rational and will hence choose e , yields player 1 the payoff 6, which is just as high as the payoff she could guarantee herself with the outside option a . Hence (a, e) and (c, e) are the profiles of extensive-form (correlated) rationalizable strategies (as well as extensive-form rationalizable strategies) in this game.

The notion of prudence, in contrast, embodies the idea that being prudently rational, player 1 shouldn't rule out completely any of 2's possible choices, and hence that c is strictly inferior for player 1 relative to her outside option a . Hence, if 2's information set is ever reached, the only way for 2 to rationalize this is to believe that 1 chose b , based on a belief ascribing a high probability to the event that 2 will foolishly choose f . Player 2's best reply to b is d ; and player 1's best reply to d is a . Thus, the only profile of prudent rationalizable strategies in this game is (a, d) .

This example demonstrates that in dynamic interactions the notions of rationalization and prudence might involve a *tension*. Extensive-form rationalizability embodies a best-rationalization principle (Battigalli 1997, Battigalli and Siniscalchi 2002); it is driven by the assumption that in each of his information sets, a player assesses the other players' future behavior by attributing to them the 'highest' level of rationality and mutual certainty of rationality consistent with the fact that the information set has indeed been reached. But, with the additional criterion of 'prudence', what should a player believe about the behavior of his opponent if, as in the example, the opponent's only action which is compatible with common certainty of rationality is imprudent on the part of the opponent?

The definition of prudent rationalizability resolves this tension unequivocally in favor of the prudence consideration. It remains open whether and how a more balanced and elaborate definition could resolve the tension in less an extreme fashion. We plan to address this challenge in future work. However, any definition would have to cut the Gordian knot in the above example in one particular way, choosing either d or e , and indeed both potential resolutions are backed by sensible intuitions. This suggests that for dynamic interactions we need not necessarily expect one ultimate definition of

rationalizability taking into account both rationalization and prudence.


Remark 1 *The definition of prudent rationalizability employs extensive-form rationalizability. For standard extensive-form games, Brandenburger and Friedenberg (2007) studied the connection between extensive-form iteratively admissible strategies (defined on the basis of rationality rather than extensive-form rationality) and extensive-form rationalizability. They showed that under a “no relevant convexities” condition, extensive-form rationalizability and extensive-form iterated admissibility coincide. However, the example in Figure 2 does not satisfy this condition, and hence demonstrates that in general extensive-form iterated admissibility is not a refinement of extensive-form rationalizability.*

Nevertheless, as far as paths of play are concerned, in the above example the set of paths induced by prudent rationalizability (the path a) is a subset of the paths induced by extensive-form rationalizability (the paths a and (c, e)). This is an instance of a general phenomenon as we show in Theorem 2 below.

4.3 Forward Induction: the tension between EFR and PR

In Figure 2 we demonstrated the tension between the considerations of rationalization and prudence when a player tries to divine his opponent’s past actions. A related but distinct tension arises when a player tries to deduce the opponent’s future behavior from past actions of that opponent. Consider the following example in Figure 3.

Figure 3:



In this example, in is imprudent for player 1 (since by going out she can guarantee a payoff of 10, while by moving in she risks getting 0 if player 2 would rather foolishly choose r). This means that if player 1 does move in and player 2 gets to play, no prudent

strategy in \bar{S}_1^1 reaches 2's information set. Hence, the beliefs \bar{B}_2^2 of player 2 about player 1's future actions are not restricted. In particular, it contains beliefs by which if player 2 chooses m , player 1 will foolishly choose R (with a high probability). That's why both m and ℓ are prudent rationalizable for player 2.


However, it is not very sensible on the part of player 2 to believe that following m player 1 may choose R . After all, when player 2 has to move, player 1 has already proved to be imprudent, but not irrational. Indeed, player 1's rationalizable (though imprudent) strategy (in, L) yields her the payoff 10 in conjunction with 2's only (extensive-form) rationalizable strategy ℓ , as well as in conjunction with 2's prudent rationalizable strategy m ; and this payoff is the same as the payoff player 1 gets from her only prudent rationalizable strategy (out, L).

Thus, as long as player 1 has been rational (even if imprudent) thus far, it makes more sense for player 2 to believe that player 1 will continue to be rational (though possibly imprudent) in the future. Restricting player 2's beliefs according to this logic would cross out the nonsensical choice m .

Already Pearce (1984) was well aware of this tension, which motivated his definition of **cautious extensive-form rationalizability**. That definition involves refining the set of rationalizable strategies by another round of strategy elimination with full support beliefs about the other players' surviving strategies; and then repeating this entire procedure – the standard iterative elimination process as in the definition of rationalizability, followed by one round assuming full-support beliefs – ad infinitum. In the above example, cautious extensive-form rationalizability does indeed rule out the strategy m for player 2.

However, as Pearce (1984) himself admits, the definition of cautious extensive-form rationalizability is not really satisfactory, as the following simple example of his shows.

Figure 4:



In this example, the strategy d is irrational for player 2. Once d is crossed out, both a and b are extensive-form rationalizable for player 1, and are actually also cautious

extensive-form rationalizable. Notice that in contrast, b *does* get crossed out by *prudent rationalizability*, and the only prudent rationalizable strategy for player 1 is a .

4.4 Discussion of the PR definition

Definition 2 of extensive-form rationalizable strategies involves, as in Battigalli (1997), an iterative reduction procedure of belief systems (that is, by definition $B_i^k \subseteq B_i^{k-1}$), and this definition implies that strategies get iteratively eliminated ($S_i^k \subseteq S_i^{k-1}$); and the same is true also for extensive-form correlated rationalizable strategies – by definition $\hat{B}_i^k \subseteq \hat{B}_i^{k-1}$ and hence $\hat{S}_i^k \subseteq \hat{S}_i^{k-1}$. In contrast, the inductive definition of **prudent** rationalizable strategies involves an iterative elimination of strategies (that is, by definition $\bar{S}_i^k \subseteq \bar{S}_i^{k-1}$, in analogy with the original formulation of Pearce (1984) for extensive-form rationalizability by an iterative elimination procedure), but in the case of prudence it is **not** generally the case that $\bar{B}_i^k \subseteq \bar{B}_i^{k-1}$. Indeed, when $\bar{S}_{-i}^k \subsetneq \bar{S}_{-i}^{k-1}$:

- If the set of strategy profiles in \bar{S}_{-i}^k reaching some information set $h_i \in H_i$ is a proper, non-empty subset of the strategy profiles in \bar{S}_{-i}^{k-1} that reach h_i , then the support of each belief $\bar{b}_i^{k-1}(h_i)$ in each belief system $\bar{b}_i^{k-1} \in \bar{B}_i^{k-1}$ is strictly larger than the support of any belief $\bar{b}_i^k(h_i)$ for $\bar{b}_i^k \in \bar{B}_i^k$.
- For information sets h_i *not* reached by \bar{S}_{-i}^k , there is no restriction (beyond Bayes rule) on $\bar{b}_i^k(h_i)$ for $\bar{b}_i^k \in \bar{B}_i^k$. No such restriction is needed, because if we define

$$m_{h_i}^k = \max \{ m < k : \text{there exists } s_{-i} \in \bar{S}_{-i}^m \text{ that reaches } h_i \}$$

then for $s_i^k \in \bar{S}_i^k$ the restrictions on i 's actions $s_i^k(h_i)$ at h_i were already determined at stage $m_{h_i}^k$, since by definition $s_i^k \in \bar{S}_i^k \subseteq \bar{S}_i^{m_{h_i}^k}$.

Is it nevertheless feasible to define prudent rationalizability via a reduction process of belief systems? Asheim and Perea (2005) proposed to look at **systems of conditional lexicographic probabilities** – belief systems in which each belief at an information set is itself a **lexicographic probability system** (Blume, Brandenburger and Dekel 1991) about the other players' strategy profiles. Using belief systems which are conditional lexicographic probabilities we could, in the spirit of Stahl (1995), put forward an equivalent definition of prudent rationalizable strategies involving an iterative reduction procedure of belief systems rather than an iterative elimination procedure of strategies.

In each round of the procedure, the surviving belief systems would be those in which at each information set, ruled-out strategy profiles of the other players (i.e. strategy profiles outside $\bar{S}_{-i}^{m_i^k}$) would be deemed infinitely less likely than the surviving strategy profiles, but infinitely more likely than strategy profiles which had already been eliminated in previous rounds. We leave the precise formulation of such an equivalent definition to future work.

In their paper, Asheim and Perea (2005) proposed the notion of **quasi-perfect rationalizability**, which also involves the idea of cautious beliefs. Quasi-perfect rationalizability is distinct from our notion of prudent rationalizability. The difference is that with prudent rationalizability (as with extensive-form rationalizability), a player need not believe that another player's future behavior must be rationalizable to a higher order than that exhibited by that other player in the past; in contrast, with the quasi-perfect rationalizable strategies of Asheim and Perea (2005), a player should ascribe to her opponent the highest possible level of rationality in the future even if this opponent has already proved to be less rational in the past. That's why quasi-perfect rationalizability implies backward induction in generic perfect information games, while our prudent rationalizable strategies need not coincide with the backward induction strategies in such games (though they do generically lead to the backward induction path – the argument is the same as in Reny 1992 and Battigalli 1997, since in generic perfect information games prudent rationalizability coincides with extensive-form rationalizability in terms of realized paths).

5 An Application to Verifiable Information

In this section we provide an application of prudent rationalizability to the problem of relying on information of interested parties, which was introduced by Milgrom and Roberts (1986).

Consider a merchandise whose quality $q_i \in \{q_1, \dots, q_n\}$ is known to its seller, while a buyer knows only the prior probability distribution (p_1, \dots, p_n) of the qualities, where $p_i > 0$ for all $i = 1, \dots, n$. For each quality level q_i the seller is better off the larger the quantity that she sells, while the utility of the buyer from the merchandise is strictly concave in the quantity purchases with a single peak at $\beta(q_i)$. Furthermore,

$$\beta(q_1) < \dots < \beta(q_n).$$

Before sale takes place, the seller has the option of providing the buyer with a certified signal about the quality of her merchandise, proving to the seller that the quality is within some range $\{q_{\min}, \dots, q_{\max}\}$ containing the actual quality q_i .

Milgrom and Roberts (1986) proved that if the buyer's utility is strictly concave then there is a unique sequential equilibrium, in which when the quality is q_i the seller certifies to the buyer a range (possibly a singleton) $\{q_{\min}, \dots, q_{\max}\}$ in which $q_{\min} = q_i$, while the buyer is skeptical and always buys $\beta(q_{\min})$. Thus, in this unique sequential equilibrium the quality q_i is fully revealed to the seller, who buys the optimal quantity $\beta(q_i)$ for him.

We proceed with the caveat that the quantities which can be demanded by the buyer belong to a finite, fine grid (recall that, formally, in our formulation each player has finitely many available actions in each information set). For simplicity, we assume further that the quantities $\beta(q_i)$, $i = 1, \dots, n$ belong to this grid. For $1 \leq m < n$ we denote by $[\beta(q_m), \beta(q_n)]$ the set of quantities in this grid at least as large as $\beta(q_m)$ and no larger than $\beta(q_n)$.

Proposition 1 *The strategy to buy $\beta(q_{\min})$ when confronted with the certification that the quality is in the range $\{q_{\min}, \dots, q_{\max}\}$ is also the unique prudent rationalizable strategy for the buyer, and certifying some range $\{q_{\min}, \dots, q_{\max}\}$ in which $q_{\min} = q_i$ constitute the prudent rationalizable strategies of the seller.*

Thus, any profile of prudent rationalizable strategies in this game yields the full revelation outcome indicated by Milgrom and Roberts (1986).

Proof of Proposition 1. When the buyer is confronted with the certificate $\{q_n\}$, his unique level-1 (prudent) rationalizable action is to buy $\beta(q_n)$, while when he is confronted with some range $\{q_m, \dots, q_n\}$ all the quantities in the interval $[\beta(q_m), \beta(q_n)]$ are level-1 (prudent) rationalizable (because any posterior belief of the buyer about the qualities with support $\{q_m, \dots, q_n\}$ can be derived from a belief of the buyer that the seller provides the certificate $\{q_m, \dots, q_n\}$ with an appropriate probability r_i when the seller knows that the quality is $q_i \in \{q_m, \dots, q_n\}$.) Consequently, the only level-2 prudent rationalizable strategies of the seller are those in which she provides the certificate $\{q_n\}$ when the quality is q_n (because any other certificate that she can provide $\{q_m, \dots, q_n\}$ will yield an expected sale strictly smaller than $\beta(q_n)$ with a full support belief about the level-1 prudent rationalizable strategies of the buyer, that have actions in the range $[\beta(q_m), \beta(q_n)]$).

Assume, inductively, that we have already proved that in all the level- $(2k - 1)$ prudent

rationalizable strategies of the buyer, for every $i = 0, \dots, k - 1$ he buys the quantity $\beta(q_{n-i})$ when confronted with a certificate of the form $\{q_{n-i}, \dots, q_\ell\}$, and that in all the level- $2k$ prudent rationalizable strategies of the seller she indeed provides such a certificate when the quality is q_{n-i} . Then in all the level- $(2k + 1)$ (prudent) rationalizable strategies of the buyer he buys the quantity $\beta(q_{n-k})$ when confronted with a certificate of the form $\{q_{n-k}, \dots, q_\ell\}$ (because he believes that such a certificate could only be presented to him with the quality q_{n-k} , as by the induction hypothesis with each higher quality all the level- $2k$ prudent rationalizable strategies of the seller present a certificate where that higher value is the minimal value). Furthermore, when confronted with some range $\{q_m, \dots, q_{n-k}, \dots, q_\ell\}$ all the quantities in the interval $[\beta(q_m), \beta(q_{n-k})]$ are level- $(2k + 1)$ (prudent) rationalizable (because any posterior belief of the buyer about the qualities with support $\{q_m, \dots, q_{n-k}\}$ can be derived from a belief of the buyer on the level- $2k$ prudent rationalizable strategies of the seller in which the seller provides the certificate $\{q_m, \dots, q_\ell\}$ with an appropriate probability r_i when the seller knows that the quality is $q_i \in \{q_m, \dots, q_{n-k}\}$.)

Consequently, in all the level- $(2k + 2)$ prudent rationalizable strategies of the seller she provides the certificate $\{q_{n-k}, \dots, q_\ell\}$ when the quality is q_{n-k} (because any other certificate that she can provide $\{q_m, \dots, q_{n-k}, \dots, q_\ell\}$ will yield an expected sale strictly smaller than $\beta(q_{n-k})$ with a full support belief about the level- $(2k + 1)$ prudent rationalizable strategies of the buyer, that have actions in the range $[\beta(q_m), \beta(q_{n-k})]$).

Hence, the inductive claim obtains in particular for $k = n - 1$, concluding what we wanted to prove. \square

In fact, it is not difficult to see that the above argument does not depend on the assumption that the available certificates consist of ranges of qualities (containing the true quality). For the argument to hold it is enough to assume that for each quality level q_i one of the available certificates is the fully revealing certificate $\{q_i\}$.

Note that the result would not obtain when we employ extensive-form rationalizability rather than prudent rationalizability. The reason is that when the buyer is presented with a certificate $\{q_m, \dots, q_n\}$, then the buyer could optimistically believe that the seller's quality is q_n and buy a larger quantity than with a prudent full support belief.

6 The refining power of prudent rationalizability in generalized extensive-form games with unawareness

We now present two examples demonstrating the refining power of prudent rationalizability. The first example was originally analyzed (for the full awareness case) by Milgrom and Roberts (1986) using sequential equilibrium and in Section 5 using prudent rationalizability. A second example is due to Ozbay (2007). We will also show a general result on existence and the refinement power of prudent rationalizability.

Unawareness and mutual unawareness is modeled in a generalized extensive-form game by a family \mathbf{T} of trees (see Heifetz, Meier and Schipper 2011a for the formal definition and properties). Each tree $T \in \mathbf{T}$ represents a subjective view of a player (or a subjective view of a player about another player's subjective view, etc.) of the feasible paths; one of the trees represents the modeler's 'objective' view. The set of nodes that a player $i \in I$ considers as possible at a node n of a tree $T \in \mathbf{T}$ – her information set $h_i = \pi_i(n)$ there – may be a subset of nodes of a different tree $T' \in \mathbf{T}$, in case T' represents the set of feasible paths of which player i is aware at the node n . (In such case we write $T \succrightarrow T'$, and we denote by \leftrightarrow the transitive closure of \succrightarrow). As the game proceeds, each player's view of the game may evolve, and the trees T' in which the information sets are contained may be different at distinct nodes along a path. The players have perfect recall, and remember their past views and information sets as well as their respective chosen actions.

Player i 's strategy $s_i \in S_i$ defines her choice of action $s_i(h_i) \in A_{h_i}$ for all of her information sets $h_i \in H_i$ in all trees $T \in \mathbf{T}$. In particular, $s_i(\pi_i(n))$ defines i 's action at the node $n \in T$ even when $\pi_i(n)$ is a subset of a different tree $T' \in \mathbf{T}$. Thus, a profile of strategies $s = (s_i)$ defines a path in every tree $T \in \mathbf{T}$. Denote by T_{h_i} the tree containing the information set h_i .

For every tree $T \in \mathbf{T}$, the T -partial game is the partially ordered set of trees including T and all trees T' satisfying $T \leftrightarrow T'$. A T -partial game is a generalized game by itself. We denote by H_i^T the set of i 's information sets in the T -partial game. We denote by s_i^T the strategy in the T -partial game induced by s_i . If $R_i \subseteq S_i$ is a set of strategies of player i , denote by R_i^T the set of strategies induced by R_i in the T -partial game. The set of i 's strategies in the T -partial game is thus denoted by S_i^T .

A *belief system* of player i

$$b_i = (b_i(h_i))_{h_i \in H_i} \in \prod_{h_i \in H_i} \Delta(S_{-i}^{T_{h_i}})$$

is a profile of beliefs - a belief $b_i(h_i) \in \Delta(S_{-i}^{T_{h_i}})$ about the other players' strategies in the T_{h_i} -partial game, for each information set $h_i \in H_i$, with the following properties

- $b_i(h_i)$ reaches h_i , i.e. $b_i(h_i)$ assigns probability 1 to the set of strategy profiles of the other players that reach h_i .
- If h_i precedes h'_i ($h_i \rightsquigarrow h'_i$) then $b_i(h'_i)$ is derived from $b_i(h_i)$ by Bayes rule whenever possible.

Denote by B_i the set of player i 's belief systems.

For a belief system $b_i \in B_i$, a strategy $s_i \in S_i$ and an information set $h_i \in H_i$, define player i 's expected payoff **at** h_i to be the expected payoff for player i in T_{h_i} given $b_i(h_i)$, the actions prescribed by s_i at h_i and its successors, and assuming that h_i has been reached.

We say that with the belief system b_i and the strategy s_i player i is *rational* at the information set $h_i \in H_i$ if there exists no action $a'_{h_i} \in A_{h_i}$ such that only replacing the action $s_i(h_i)$ by a'_{h_i} results in a new strategy s'_i which yields player i a higher expected payoff at h_i given the belief $b_i(h_i)$ on the other players' strategies $S_{-i}^{T_{h_i}}$.

The definition of prudent rationalizability in generalized extensive-form games is almost identical to its definition in standard games. The only difference is the fact that a belief of a player at an information set h_i is about the opponents strategies $S_{-i}^{T_{h_i}}$ in the T_{h_i} -partial game:

Definition 3 (Prudent rationalizability in generalized extensive-form games)

Let

$$\bar{S}_i^0 = S_i$$

For $k \geq 1$ define inductively

$$\bar{B}_i^k = \left\{ b_i \in B_i : \begin{array}{l} \text{for every information set } h_i, \text{ if there exists some profile} \\ s_{-i} \in \bar{S}_{-i}^{k-1} = \prod_{j \neq i} \bar{S}_j^{k-1} \text{ of the other players' strategies} \\ \text{such that } s_{-i} \text{ reaches } h_i \text{ in the tree } T_{h_i}, \text{ then the support} \\ \text{of } b_i(h_i) \text{ is the set of strategy profiles } s_{-i} \in \bar{S}_{-i}^{k-1, T_{h_i}} \text{ that reach } h_i \end{array} \right\}$$

$$\bar{S}_i^k = \left\{ s_i \in \bar{S}_i^{k-1} : \begin{array}{l} \text{there exists } b_i \in \bar{B}_i^k \text{ such that for all } h_i \in H_i \text{ player } i \\ \text{is rational at } h_i \end{array} \right\}$$

The set of prudent rationalizable strategies of player i is

$$\bar{S}_i^\infty = \bigcap_{k=1}^{\infty} \bar{S}_i^k$$

The proof of the next result is contained in the appendix.

Theorem 1 *The set of player i 's prudent rationalizable strategies is non-empty.*

In Section 4, we showed that prudent rationalizability is *not a refinement of extensive-form rationalizable strategies*. However, we can show that it refines the set of extensive-form rationalizable paths.

Theorem 2 *The set of paths induced by profiles of prudent rationalizable strategies is a subset of the paths induced by profiles of extensive-form rationalizable strategies (or, equivalently, the paths induced by profiles of extensive-form correlated rationalizable strategies).*

The proof is contained in the appendix.

6.1 Milgrom-Roberts (1986) with unawareness

Assume now that there are several dimensions of quality along which such certifications could be provided. To fix ideas, consider two dimensions L, H and $^0, ^*$. The four combinations are

$$L^0, H^0, L^*, H^*.$$

So, for instance, in the state L^0 the available certificates are $\{L, H\} \times \{^0, *\}$, $\{L\} \times \{^0, *\}$, $\{L, H\} \times \{^0\}$ and $\{L\} \times \{^0\}$.

Assume further that

$$\beta(L^*) < \beta(L^0) < \beta(H^0) < \beta(H^*)$$

Since the singleton certificates

$$\{L\} \times \{^*\}, \{L\} \times \{^0\}, \{H\} \times \{^0\}, \{H\} \times \{^*\}$$

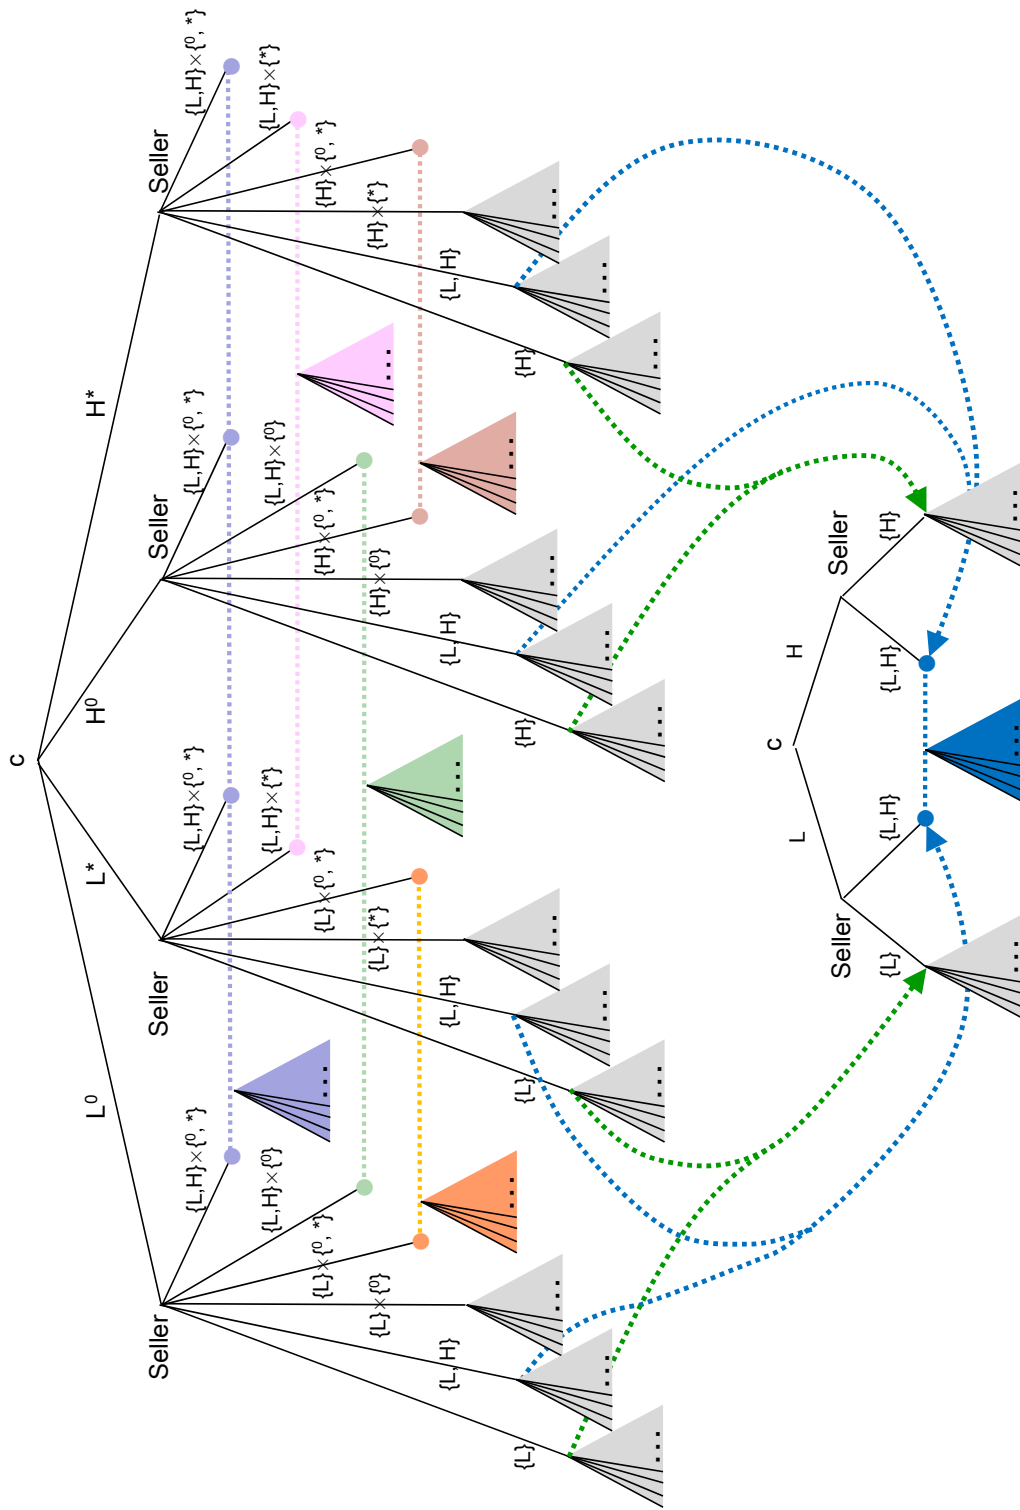
are available, the above argument obtains and full revelation takes place in any profile of prudent rationalizable strategies of the players.

Assume, however, that the buyer is initially aware only of the $\{L, H\}$ dimension and is unaware of the $\{^0, *\}$ dimension; he evaluates the merchandize as having the default quality L^0 when confronted with the certificate $\{L\}$, and similarly, with the certificate $\{H\}$ he evaluates the merchandize as having the default quality H^0 . Assume further that the seller knows this, and that by presenting the certificates $\{^*\}$, $\{^0\}$ or $\{^0, *\}$ the seller inter alia makes the buyer aware of the $\{^0, *\}$ dimension.

Intuitively, it is clear that the seller will want to make the buyer aware of this extra dimension when the quality is H^* , because this will lead the buyer to demand the high quantity $\beta(H^*)$. In contrast, when the actual quality is L^* , the seller will prefer not to present any certificate at all along the dimension $\{^0, *\}$: This way the buyer will remain unaware of this extra dimension, and will demand the quantity $\beta(L^0)$ (because unraveling and full revelation will occur only along the $\{L, H\}$ dimension); if the seller were to make the buyer aware of this extra dimension, the buyer would have demanded only $\beta(L^*) < \beta(L^0)$.

This strategic interaction is represented in the following generalized game form (Figure 5). Initially, nature, c , selects a state out of $\{L^0, L^*, H^0, H^*\}$. The seller observes the state of nature and chooses a certificate. Unless the seller presents a certificate involving the dimension $\{^0, *\}$, the buyer remains unaware of it. This is indicated by the intermitted arrows from nodes in the upper tree to nodes in the lower tree. E.g., if the seller selects the certificate $\{L\}$, then the buyer remains unaware of the $\{^0, *\}$ dimension and views the game as represented by the lower tree. In particular, his information set is a singleton containing the node after nature selects L and the seller reports $\{L\}$ in the lower tree. If the seller presents a certificate involving the $\{^0, *\}$ -dimension, then the

Figure 5:



buyer becomes aware of it and he conceives of the entire generalized game. For instance, if the seller selects the certificate $\{L, H\} \times \{^0, ^*\}$, then the buyer's information set is given by the upmost information set drawn as an intermitted line connecting four nodes.

We summarize the discussion in the following proposition.

Proposition 2 *In the verifiable information model in which the buyer is unaware of some dimension of the the good's quality, the seller may not fully reveal the quality in any prudent rationalizable outcome.*

This is in sharp contrast to the case with full awareness discussed in Section 5.

6.2 An example by Ozbay (2007)

To demonstrate the extra power of prudent rationalizability, consider the following example of dynamic interaction with unawareness, which is a variant of example 3 in Ozbay (2007). There are 3 states of nature, $\omega_1, \omega_2, \omega_3$. A chance move chooses one out of four potential distributions over the states of nature:

$$\begin{aligned} \delta_1 &= (1, 0, 0) \\ \delta_2 &= (0, 1, 0) \\ \delta_3 &= (0, 0, 1) \\ \delta_4 &= \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right) \end{aligned}$$

An Announcer gets to know the distribution (but not the realization of the state of nature). A Decision Maker (DM) is initially aware only of the state ω_1 (and hence the DM is certain that ω_1 will be realized with certainty). However, before the DM chooses what to do, the Announcer can choose to make the DM aware of either ω_2, ω_3 , none of them or both of them. Increased awareness makes the DM aware of the relevant marginals of the distributions. For instance, if the Announcer makes the DM aware of ω_2 , the DM becomes aware of the set of distributions

$$\begin{aligned} \delta_{1|\{\omega_1, \omega_2\}} &= (1, 0) \\ \delta_{2|\{\omega_1, \omega_2\}} &= (0, 1) \\ \delta_{4|\{\omega_1, \omega_2\}} &= \left(\frac{1}{2}, \frac{1}{2}\right) \end{aligned}$$

and also becomes certain that the Announcer knows which of these is the true distribution.⁴

Subsequently, the DM should choose one out of three possible actions – left, middle or right. The payoffs to the players as a function of the chosen action and the state of nature appear in the following table:

	left	middle	right
ω_1	3, 3	0, 0	2, 2
ω_2	0, 0	5, 5	2, 2
ω_3	2, 2	0, 0	2, 2

The game is thus described in Figure 6 in the following page.

It is obvious that if the Announcer announces nothing, and hence the DM is certain that ω_1 prevails, the DM will choose ‘left’.

What happens if the Announcer makes the DM aware of ω_2 ? The information set of the DM becomes

$$\left\{ \delta_{1|\{\omega_1, \omega_2\}}, \delta_{2|\{\omega_1, \omega_2\}}, \delta_{4|\{\omega_1, \omega_2\}} \right\}$$

The DM may then assign a high probability to $\delta_{1|\{\omega_1, \omega_2\}}$,⁵ and this will lead the DM to choose ‘left’. Hence, assuming such a belief by the DM, it is rationalizable for the Announcer to make the DM aware of ω_2 when the Announcer knows that the true distribution is δ_1 (i.e. when the Announcer knows that ω_1 will be realized with probability 1).


This is not very sensible, though. After all, the Announcer can ensure that the DM chooses ‘left’ by not announcing any new state. When the Announcer likes the DM to choose ‘left’, it makes no sense on the Announcer’s part to announce ω_2 and

⁴In Ozbay’s example and in what follows the DM’s beliefs about these marginal distributions will not be necessarily related to the prior probabilities with which the distributions were chosen by the chance move. That’s why we do not even bother to specify the probabilities with which the chance move chooses the different distributions.

Put differently, instead of describing this game by a partially ordered set of trees, one for each level of awareness as in Figure 6, we could have replaced each tree with an arborescence in which the initial chance move is erased. Allowing for arborescences instead of trees in the framework for dynamic unawareness is straightforward, but for the sake of clarity of the exposition we avoid this explicit generalization in the body of the paper.

⁵That is, the DM may assign a high probability to strategies of the Announcer by which the Announcer announces ω_2 (and cause the DM’s information set to become $\left\{ \delta_{1|\{\omega_1, \omega_2\}}, \delta_{2|\{\omega_1, \omega_2\}}, \delta_{4|\{\omega_1, \omega_2\}} \right\}$) when the Announcer has learned that the true distribution is δ_1 .

Figure 6:



thus face the risk that the DM assigns a low probability to $\delta_{1|\{\omega_1, \omega_2\}}$ and consequently choose ‘middle’. This idea is captured by Ozbay’s reasoning refinement to his awareness equilibrium notion⁶, as well as by prudent rationalizability:

Proposition 3 *The DM has a unique prudent rationalizable strategy. With this strategy the DM chooses ‘left’ when no new state is announced, ‘middle’ when only ω_2 is announced, ‘left’ when only ω_3 is announced, and ‘right’ when both ω_2, ω_3 are announced.*

Proof. \bar{B}_{DM}^1 contains belief systems in which in the information set $\{\delta_{1|\{\omega_1, \omega_2\}}, \delta_{2|\{\omega_1, \omega_2\}}, \delta_{4|\{\omega_1, \omega_2\}}\}$ (which follows the announcement of only ω_2 by the Announcer) the DM’s belief assigns high probabilities to $\delta_{2|\{\omega_1, \omega_2\}}, \delta_{4|\{\omega_1, \omega_2\}}$. The strategies in \bar{S}_{DM}^1 corresponding to these belief systems prescribe ‘middle’ to the DM in the information set $\{\delta_{1|\{\omega_1, \omega_2\}}, \delta_{2|\{\omega_1, \omega_2\}}, \delta_{4|\{\omega_1, \omega_2\}}\}$. The crucial point is that $\bar{B}_{Announcer}^2$ contains *only* belief systems that assign strictly positive probabilities to these strategies of the DM. Thus, with *any* belief system in $\bar{B}_{Announcer}^2$, it is sub-optimal for the Announcer to announce ω_2 in the announcer’s information set $\{\delta_1\}$, in which the Announcer is certain of ω_1 .⁷ Hence, $\bar{S}_{Announcer}^2$ does not contain strategies in which the Announcer announces just ω_2 when the announcer’s information set is $\{\delta_1\}$. We conclude that \bar{B}_{DM}^3 contains only belief systems in which the belief at the information set $\{\delta_{1|\{\omega_1, \omega_2\}}, \delta_{2|\{\omega_1, \omega_2\}}, \delta_{4|\{\omega_1, \omega_2\}}\}$ assigns probability zero to $\delta_{1|\{\omega_1, \omega_2\}}$. Hence, \bar{S}_{DM}^3 contains only strategies with which the DM chooses ‘middle’ at the information set $\{\delta_{1|\{\omega_1, \omega_2\}}, \delta_{2|\{\omega_1, \omega_2\}}, \delta_{4|\{\omega_1, \omega_2\}}\}$.

Furthermore, already \bar{S}_{DM}^1 contains only strategies with which the DM chooses ‘left’ at the information set $\{\delta_{1|\{\omega_1, \omega_3\}}, \delta_{3|\{\omega_1, \omega_3\}}, \delta_{4|\{\omega_1, \omega_3\}}\}$ (i.e. when the Announcer announces just the new state ω_3). This is because prudent rationalizability implies that all the belief systems in \bar{B}_{DM}^1 assign a positive probability to strategies of the Announcer with which the Announcer announces the new state ω_3 even when the Announcer’s information set (from the point of view of the DM!) is $\{\delta_{1|\{\omega_1, \omega_3\}}\}$ or $\{\delta_{4|\{\omega_1, \omega_3\}}\}$.

Also, \bar{B}_{DM}^1 contains belief systems in which the DM’s belief in the information set $\{\delta_1, \delta_2, \delta_3, \delta_4\}$ (when the Announcer announces both new states ω_2, ω_3) assigns high probability to δ_2 . The strategies in \bar{S}_{DM}^1 corresponding to these belief systems prescribe ‘middle’ to the DM in the information set $\{\delta_1, \delta_2, \delta_3, \delta_4\}$. Hence, $\bar{B}_{Announcer}^2$ contains *only* belief systems that assign strictly positive probabilities to these strategies of the DM. Thus,

⁶We believe that equilibrium notions are somewhat questionable in the context of unawareness, and hence our focus on rationalizability. See Heifetz, Meier and Schipper (2011a) for further discussions.

⁷Because according to every belief system in $\bar{B}_{Announcer}^2$, announcing just ω_2 will lead the DM with a positive probability to choose ‘middle’.

with *any* belief system in $\bar{B}_{Announcer}^2$, it is sub-optimal for the Announcer to announce both ω_2 and ω_3 in the announcer's information sets $\{\delta_1\}$ and $\{\delta_3\}$. Similarly, \bar{B}_{DM}^1 contains belief systems in which the DM's belief in the information set $\{\delta_1, \delta_2, \delta_3, \delta_4\}$ assigns high probability to δ_1 . The strategies in \bar{S}_{DM}^1 corresponding to these belief systems prescribe 'left' to the DM in the information set $\{\delta_1, \delta_2, \delta_3, \delta_4\}$. Hence, $\bar{B}_{Announcer}^2$ contains *only* belief systems that assign strictly positive probabilities to these strategies of the DM. Thus, with *any* belief system in $\bar{B}_{Announcer}^2$, it is sub-optimal for the Announcer to announce both ω_2 and ω_3 in the Announcer's information sets $\{\delta_1\}, \{\delta_2\}$ or $\{\delta_3\}$. We conclude that \bar{B}_{DM}^3 contains only belief systems in which the belief at the information set $\{\delta_1, \delta_2, \delta_3, \delta_4\}$ assigns probability zero to $\delta_1, \delta_2, \delta_3$. That is, \bar{B}_{DM}^3 contains only belief systems that assign probability 1 to δ_4 at the information set $\{\delta_1, \delta_2, \delta_3, \delta_4\}$. Hence, \bar{S}_{DM}^3 contains only strategies with which the DM chooses 'right' at the information set $\{\delta_1, \delta_2, \delta_3, \delta_4\}$.

We thus conclude that \bar{S}_{DM}^3 contains a unique strategy s_{DM}^* . This strategy prescribes the DM to choose 'left' in the information set $\{\delta_1|_{\{\omega_1\}}\}$ (i.e. when the Announcer does not announce any new state), to choose 'middle' in the information set $\{\delta_1|_{\{\omega_1, \omega_2\}}, \delta_2|_{\{\omega_1, \omega_2\}}, \delta_4|_{\{\omega_1, \omega_2\}}\}$ (i.e. when the Announcer announces just the new state ω_2), to choose 'left' in the information set $\{\delta_1|_{\{\omega_1, \omega_3\}}, \delta_3|_{\{\omega_1, \omega_3\}}, \delta_4|_{\{\omega_1, \omega_3\}}\}$ (i.e. when the Announcer announces just the new state ω_3) and to choose 'right' in the information set $\{\delta_1, \delta_2, \delta_3, \delta_4\}$ (i.e. when the Announcer announces both new states ω_2, ω_3).⁸ \square

A Proofs

A.1 Proof of Theorem 1

First, observe that $\bar{B}_i^k \neq \emptyset$ for every $k \geq 1$, because if an information set $h_i \in H_i$ is reached by some $s_{-i} \in \bar{S}_{-i}^{k-1}$, then s_{-i} reaches also all of i 's information sets that precede h_i in the tree T_{h_i} .

We proceed by induction. $\bar{S}_i^0 = S_i$ and hence non-empty. Notice also that for every $b_i \in \bar{B}_i^1$, a standard backward induction procedure on the arborescence of information sets H_i yields a strategy $s_i \in \bar{S}_i^1$ with which player i would be rational $\forall h_i \in H_i$ given b_i .

Suppose, inductively, we have already shown that $\forall i \in I \ \bar{S}_i^{k-1} \neq \emptyset$ (and hence that

⁸This is also the unique strategy of the DM which is part of an awareness equilibrium satisfying reasoning refinement in Ozbay (2007).

$\bar{S}_{-i}^{k-1} \neq \emptyset$), and also that for every $b_i \in \bar{B}_i^{k-1}$ there exists a strategy $s_i \in \bar{S}_i^{k-1}$ with which player i would be rational $\forall h_i \in H_i$ given b_i .

Let $b_i \in \bar{B}_i^k$. Let $\dot{H}_i \subseteq H_i$ be the set of i 's information sets not reached by any profile $s_{-i} \in \bar{S}_{-i}^{k-1}$ but reached by some profile $s_{-i} \in \bar{S}_{-i}^{k-2}$. If $\dot{H}_i \neq \emptyset$, for every $h_i \in \dot{H}_i$ with no predecessor in \dot{H}_i , modify (if necessary) $b_i(h_i)$ so as to have full support on the profiles in \bar{S}_{-i}^{k-2} that reach h_i , and in succeeding information sets modify b_i by Bayes rule whenever possible. Denote the modified belief system by \dot{b}_i . Then by construction also $\dot{b}_i \in \bar{B}_i^k$.

Consider a sequence of belief systems $b_{i,n} \in \bar{B}_i^{k-1}$ such that

$$\dot{b}_i = \left(\dot{b}_i(h'_i) \right)_{h'_i \in H_i} \equiv \left(\lim_{n \rightarrow \infty} b_{i,n}(h'_i) \right)_{h'_i \in H_i}$$

and given this sequence⁹ $b_{i,n} \in \bar{B}_i^{k-1}$ let $s_{i,n} \in \bar{S}_i^{k-1}$ be a corresponding sequence of strategies with the property that given $b_{i,n}$, it is the case that with the strategy $s_{i,n}$ player i would be rational at every $h_i \in H_i$. Since player i has finitely many strategies, some strategy s_i appears infinitely often in the sequence $s_{i,n}$. Since expected utility is linear in beliefs and hence continuous, also given \dot{b}_i it is the case that with the strategy s_i player i would be rational at every $h_i \in H_i$. Hence $s_i \in \bar{S}_i^k$ as well.

Now, since player i 's set of strategies S_i is finite and by definition $\bar{S}_i^{k+1} \subseteq \bar{S}_i^k$ for every $k \geq 1$, for some ℓ we eventually get $\bar{S}_i^\ell = \bar{S}_i^{\ell+1} \forall i \in I$ and hence $\bar{B}_i^{\ell+1} = \bar{B}_i^{\ell+2} \forall i \in I$. Inductively,

$$\emptyset \neq \bar{S}_i^\ell = \bar{S}_i^{\ell+1} = \bar{S}_i^{\ell+2} = \dots$$

and therefore

$$\bar{S}_i^\infty = \bigcap_{k=1}^{\infty} \bar{S}_i^k = \bar{S}_i^\ell \neq \emptyset$$

as required. □

⁹To construct such a sequence $b_{i,n} \in \bar{B}_i^{k-1}$, for every information set $h'_i \in H_i$ not reached by any $s_{-i} \in \bar{S}_{-i}^{k-1}$ define $b_{i,n}(h'_i) = \dot{b}_i(h'_i)$ for every $n \geq 1$; and for every $h'_i \in H_i$ with no predecessors but reached by some profile $s_{-i} \in \bar{S}_{-i}^{k-1}$ define $b_{i,n}(h'_i) \in \Delta(\bar{S}_{-i}^{k-1})$ to be any converging sequence of beliefs such that for every $n \geq 1$ the support of $b_{i,n}(h'_i)$ is the subset of profiles in \bar{S}_{-i}^{k-2} that reach h'_i , while $\lim_{n \rightarrow \infty} b_{i,n}(h'_i) = \dot{b}_i(h'_i)$. In succeeding information sets reached by some $s_i \in \bar{S}_i^{k-1}$ define $b_{i,n}(h'_i)$ by Bayes rule whenever possible.

A.2 Proof of Theorem 2

Denote by (a_i, h_i) the copy of the action a_i of player $i \in I$ whenever it appears in the information set h_i . For the purpose of this proof the word “action” will refer to a copy (a_i, h_i) of an action *at a given information set*.

Define a *menu* of a player to be a (possibly empty) subset of (the union of) her actions in her information sets.

Define a *menu profile* to be a profile of menus, one for each player, with the following property: For each information set h_i of player i , her menu in the menu profile contains at least one action in h_i if and only if that information set is reached by a sequence of actions of the players in the menu profile.

For a menu profile M , denote by M_i the menu of player i in M .

For a menu profile M , denote by $P^T(M)$ all the paths from the roots to leaves in the trees of the T -partial game that one can compose from actions in M and moves of nature (if there are any). Denote also by $P(M)$ the set of paths from roots to leaves in all the trees of the generalized games that one can compose from actions in M and moves of nature.

Now, every product of sets of strategies $R = \prod_{i \in I} R_i$ (where R_i is a subset of i 's strategies) induces a menu profile, in which player i 's menu is defined as follows. For each information set of the player:

- 1) If the information set is reached by some strategy profile in the set R , the player's menu contains all the actions ascribed in that information set by i 's strategies in R_i that reach the information set.
- 2) If the information set is not reached by any strategy profile in R , then player i 's menu contains no action of hers in that information set.

Intuitively, player i 's menu is mute about an information set if and only if that information set is excluded by the set of strategy profiles R (case 2); otherwise (case 1) the menu contains all the actions in that information set that appear in some strategy of hers in R_i that reaches that information set.

If M is the menu profile induced by R , then every strategy in R_i together with a belief about R_{-i} induce a belief β^T about the paths of actions in $P^T(M)$ for every tree T of the generalized game.

Next, denote by M^k the menu profile induced by $S^k = \prod_{i \in I} S_i^k$, the set of level

k extensive-form rationalizable strategy profiles; and denote by \bar{M}^k the menu profile induced by $\bar{S}^k = \prod_{i \in I} \bar{S}_i^k$, the set of level k prudent rationalizable strategy profiles.

Proposition 2 is implied by the following lemma:

Lemma 1 *For all $\ell \geq 0$, $\bar{M}^\ell \subseteq M^\ell$. In particular $\bar{M}^\infty \subseteq M^\infty$.*

Proof. The proof is by induction.

For $\ell = 0$ we have $M^0 = \bar{M}^0$, the menu profile which includes all actions at all the information sets of all the players.

Suppose the claim holds for $\ell \leq k$.

By the induction hypothesis $P(\bar{M}^\ell) \subseteq P(M^\ell)$ for every $\ell \leq k$.

We will now prove the claim for $\ell = k + 1$, i.e. that $\bar{M}_i^{k+1} \subseteq M_i^{k+1}$ for every player $i \in I$.

To this end we have to show that for every player $i \in I$, every $\bar{s}_i^{k+1} \in \bar{S}_i^{k+1}$, every information set $h_i \in H_i$ which is reached both by \bar{s}_i^{k+1} and by some strategy profile in \bar{S}_{-i}^{k+1} (meaning that $\bar{s}_i^{k+1}(h_i) \in \bar{M}_i^{k+1}$), it is the case that

- a) h_i is also reached by S^{k+1} , and
- b) $\bar{s}_i^{k+1}(h_i) \in M_i^{k+1}$ as well.

In fact, it is enough to show that b) holds. To see this, proceed inductively along each feasible path of the generalized game (in each of its trees). If player i is the first to play in this path (apart from nature, if there are nature moves in the path), and if h_i is the information set in which she makes this initial move, then condition a) automatically obtains for h_i , and we only need to prove b). Inductively, if we reach a node in the path which is not in $P(\bar{M}^{k+1})$, we have nothing to prove for this node's information set when considering this path.¹⁰ If all the nodes $n_1 \dots n_m$ in an initial segment of the path are on a path in $P(\bar{M}^{k+1})$ and we have already proved conditions a) and b) for all the information sets of these nodes, then it already follows that a) holds for the information set of the next node n_{m+1} in the path [because b) holds for the previous node n_m for the player (or players) active in n_m]. It thus remains to show b) for such an information set.

So we now proceed to prove b).

Suppose h_i is reached by \bar{S}_{-i}^{k+1} and by $\bar{s}_i^{k+1} \in \bar{S}_i^{k+1}$. Since by definition $\bar{S}_{-i}^{k+1} \subseteq \bar{S}_{-i}^k$,

¹⁰We may have to consider this information set again when we analyze another path passing through it.

we have $\bar{s}_i^{k+1} \in \bar{S}_i^k$ and hence $m_i^{k+1}(h_i) = k$. Consider a belief system $b_i \in \bar{B}_i^{k+1}$ with a full-support belief $b_i(h_i)$ on the strategy profiles \bar{S}_{-i}^k that reach h_i , and with which \bar{s}_i^{k+1} would be rational at h_i (i.e. player i cannot improve her expected payoff by changing \bar{s}_i^{k+1} only at h_i , from $\bar{s}_i^{k+1}(h_i)$ to some other action a'_{h_i} available there).

The strategy \bar{s}_i^{k+1} together with the belief $b_i(h_i)$ on the other players' strategies induce a full support belief β on the paths of actions in $P(\bar{M}^k)$ reaching h_i and along which player i uses the strategy \bar{s}_i^{k+1} . Since by the induction hypothesis $P(\bar{M}^k) \subseteq P(M^k)$, it follows that β is a belief on the paths of actions in $P(M^k)$ reaching h_i and along which player i uses the strategy \bar{s}_i^{k+1} .

Denote by $\bar{s}_i^{k+1}|a'_{h_i}$ the strategy one gets from \bar{s}_i^{k+1} by altering the action at the information set h_i from $\bar{s}_i^{k+1}(h_i)$ to a'_{h_i} . The altered strategy $\bar{s}_i^{k+1}|a'_{h_i}$ together with the belief $b_i(h_i)$ on the other players' strategies induce a full support belief β' on the paths of actions in $P(\bar{M}^k)$ reaching h_i and along which player i uses the strategy $\bar{s}_i^{k+1}|a'_{h_i}$.

The fact that \bar{s}_i^{k+1} extensive-form rational given the belief system b_i means that in particular at the information set h_i , with the belief $b_i(h_i)$ on the other players' strategies, the expected payoff to player i given β is not smaller than the expected payoff to player i given β' .

This yields the conclusion b) that we wanted, namely that $\bar{s}_i^{k+1}(h_i) \in M_i^{k+1}$. \square

References

- [1] Asheim, G.B. and A. Perea (2005), Sequential and quasi-perfect rationalizability in extensive games, *Games and Economic Behavior* **53**, 15-42.
- [2] Battigalli, P. (1997). On rationalizability in extensive games, *Journal of Economic Theory* **74**, 40-61.
- [3] Battigalli, P. and M. Siniscalchi (2002). Strong belief and forward induction reasoning, *Journal of Economic Theory* **106**, 356-391.
- [4] Blume, L., Brandenburger, A., and E. Dekel (1991). Lexicographic probabilities and choice under uncertainty, *Econometrica* **69**, 61-79.
- [5] Brandenburger, A. and A. Friedenberg (2007). The relationship between rationality on the matrix and the tree, mimeo.

- [6] Dubey, P. and M. Kaneko (1984). Information patterns and Nash equilibria in extensive games: I, *Mathematical Social Sciences* **8**, 111-139.
- [7] Heifetz, A., Meier, M. and B. C. Schipper (2011a). Dynamic unawareness and rationalizable behavior, mimeo.
- [8] Heifetz, A., Meier, M. and B. C. Schipper (2011b). Conditional dominance in dynamic games with unawareness, mimeo.
- [9] Milgrom, R. and J. Roberts (1986). Relying on the information of interested parties, *Rand Journal of Economics* **17**, 18-32.
- [10] Osborne, M. and A. Rubinstein (1994). *A course in game theory*, MIT Press.
- [11] Ozbay, E. (2007). Unawareness and strategic announcements in games with uncertainty, in: Samet, D. (ed.), Proceedings of the 11th conference on Theoretical Aspects of Rationality and Knowledge, Presses Universitaires de Louvain, pp. 231-238.
- [12] Pearce, D.G. (1984). Rationalizable strategic behavior and the problem of perfection, *Econometrica* **52**, 1029-1050.
- [13] Perea, A. (forthcoming). *Epistemic game theory: Reasoning and choice*, Cambridge University Press.
- [14] Reny, P. (1992). Backward induction, normal form perfection and explicable equilibria, *Econometrica* **60**, 627-649.
- [15] Robles, J. (2006). Order independence of conditional dominance, mimeo., Victoria University of Wellington.
- [16] Shimoji, M. and J. Watson (1998). Conditional dominance, rationalizability, and game forms, *Journal of Economic Theory* **83**, 161-195.
- [17] Stahl, D. (1995). Lexicographic rationalizability and iterated admissibility, *Economics Letters* **47**, 155-159.