Extensive Games with Randomly Disturbed Payoffs: A New Approach to Equilibrium Refinements

V Bhaskar
University of Texas at Austin

Maxwell B Stinchcombe
University of Texas at Austin

August 2, 2024

Abstract

Refinements of Nash equilibrium hinge on the question: what inferences does a player draw about her opponent's future behavior or his type at an information set that has zero prior probability under the equilibrium. We address this question by adding, in the spirit of Harsanyi (*IJGT*, 1973), shocks to the payoffs of each player at every terminal node, that are independent across players and across nodes, that have a sufficiently large support, and have a continuous distribution. A behavior strategy profile b of the unperturbed game is *purifiable* if there exists some sequence of distributions, converging weakly to Dirac measures on 0, with a sequence of Bayes Nash equilibria whose aggregates converge to b. Strategy profile b is strongly purifiable if it is purifiable for *every* converging sequence of distributions. If the shocks are restricted to be also identically distributed for for each player (i.e. they are i.i.d), this yields the notions of *symmetric purification* and *symmetric strong purification.*

First we consider finite games of perfect information with generic payoffs, with a unique backwards induction (BI) strategy profile. If each player moves at most once along any path, then the backwards induction strategy profile is strongly purifiable, and no other strategy profile is purifiable. However, if a player moves more than once along some path, as in the centipede game, then there may exist purifiable Nash equilibria that are not subgame perfect. For example, in the perturbed centipede game, if player 1 does not play the backwards induction action at her initial node, then player 2 cannot conclude that player 1 will play her backward induction action with high probability at a subsequent node, even if the payoff shocks are independent and arbitrarily small. Furthermore, the backwards induction strategy profile is not strongly purifiable. However, every purifiable profile induces the backwards induction outcome.

We next consider signaling games. We focus on symmetric purification, and also assume that our sequence of distributions have thin tails. In the beer-quiche game, pooling on beer is symmetrically purifiable, while pooling on quiche is not. In general, symmetric purification is neither stronger nor weaker than D1. Also, symmetric purification cannot justify forward induction arguments.

Keywords: Equilibrium refinements, Purification, backward induction, signaling, D1, forward induction.

Preliminary and Incomplete!

1 Introduction

Refinements of Nash equilibrium hinge on the question: what inferences does a player draw about her opponent's future behavior or his type at an information set that has zero prior probability under the equilibrium. Classical approaches to this question assume that players "tremble" in their strategy choices (Selten (1975), Kreps and Wilson (1982), Myerson (1978), Kohlberg and Mertens (1986)). We address this question by adding, in the spirit of Harsanyi (1973), shocks to the payoffs of each player at every terminal node, that are independent across players and across nodes, that have a sufficiently large support, and have a continuous distributions. We assume that each player observes, at the outset, her vector of shock realizations but not those of any other player. A strategy for player i, σ_i, is a mapping from the space of possible payoff realizations to a (pure) behavior strategy in the unperturbed game. The *aggregate* corresponding to σ_i is the (mixed) behavior strategy $\tilde{b}_i$ that is obtained by integrating over i's payoff shocks. A behavior strategy profile b of the unperturbed game is *purifiable* if there exists some sequence of payoff shock distributions F^n, converging weakly to Dirac measures on 0, with a sequence of Bayes Nash equilibria σ^n whose aggregate profiles $\tilde{b}^n$ converge to b. Strategy profile b is strongly purifiable if it is purifiable for *every* converging sequence of distributions F^n. Our large support assumption implies at any information set, a player chooses all her actions with positive probability. Since the shocks have a continuous distribution, a player has a essentially unique best response to any profile of aggregates b_{-i} played by her opponents. Furthermore, her own aggregate is a continuous function of b_{-i}, ensuring the existence of a Bayes Nash equilibrium in any finite extensive form game. Standard arguments ensure the existence of a purifiable strategy profile (but not necessarily a strongly purifiable one).

Our analysis begins with finite games of perfect information, which can be solved by backwards induction. Assume generic payoffs, so that no player has equal payoffs at different terminal nodes. Such games have a unique backwards induction strategy profile, and the strategic form is dominance solvable. [1] If each player moves at most once along any path, we find that the backwards induction strategy profile is strongly purifiable, and no other strategy profile is purifiable. However, if a player moves more than once along some path, as in the centipede game, then there may exist purifiable Nash equilibria that are not subgame perfect. For example, in the perturbed centipede game, if player 1 does not play the backwards induction action at her initial node, then player 2 cannot conclude that player 1 will play

[1] That is, the backwards induction profile is the unique survivor also of the process of iteratively deleting weakly dominated strategies.

her backward induction action with high probability at a subsequent node, even if the payoff shocks are independent and arbitrarily small. Furthermore, there exist shock sequences such that the backwards induction strategy profile is not the limit of the equilibrium sequence, and hence it is not strongly purifiable. However, every purifiable profile induces the backwards induction outcome.

The negative result for the strong purifiability of the backward induction strategy profile leads us to consider the more restrictive notion of *symmetric purification*, where the shocks for a player are also identically distributed across terminal nodes. We can similarly define strong and weak versions of symmetric purification. We conjecture (but have not yet proved) that the the backwards induction strategy profile is uniquely symmetrically purifiable profile (and is hence strongly symmetrically purifiable). Generalizing the class of games to allow players to move simultaneously, we find that symmetric purification does not support forward induction arguments. That is, a subgame perfect equilibrium that does not satisfy forward induction (or the iterative elimination of weakly dominated strategies) is symmetrically purifiable. However, it is not strongly purifiable.

We next consider signaling games, and on focus on symmetric purification. Our cleanest results are obtained when we also assume that the distributions of payoff shocks have *thin tails*, so that larger shock realizationa become infinitely less likely than smaller shock realizations as the distribution converges. In the beer-quiche game, pooling on beer is symmetrically purifiable, while pooling on quiche is not. In general, symmetric purification is neither stronger nor weaker than D1 (Banks and Sobel (1987), Cho and Kreps (1987)). Indeed, it involves considerations that are very different from refinements in signaling games that place restrictions on beliefs at unreached information sets that are motivated by hypothetical speeches by the deviating type.

2 The model

Our set up covers finite extensive form games with perfect recall which we denote by Γ. Since the set up is standard, we will be economical in our formalization, and refer the reader to standard texts such as van Damme (1991) for more details. Let N denote the finite set of players, and X the finite collection of information sets, each of which contains finitely many nodes. The collection $\{X_i\}_{i \in N}$ is partition of X so that exactly one player moves at each information set. Let W denote the set of terminal nodes and $u_i : W \to \mathbb{R}$, the payoff function for player i, is defined for each $\imath \in N$. Let A_i denote the set of actions for player i.

This is partitioned into the sets $\{A_i(x)\}_{x \in X_i}$, so that $A_i(x)$ is the set of actions available at information set x. A behavior strategy for i, is a function $b_i : X_i \to \Delta(A_i)$, where at each x_i, $b_i(x)$ assigns positive probability only to the available actions. The set of behavior strategies for i is B_i, and $b = (b_i)_{i \in N}$ is a behavior strategy profile.

Now consider the perturbed game. The payoffs of each player i at each terminal node w are augmented by random shock z_{iw}. Each player observes his own vector of shock realizations at the beginning of the game and does not observe his opponents' shock realizations. Shocks for player i have support $Z_i \subset \mathbb{R}^{|W|}$. We assume that the shock distributions satisfy the following conditions:

- Independence, across players and terminal nodes.
- They have sufficiently large support, so that all actions have positive probability.
- They are continuous, i.e. have no mass points.

Let $F = (F_{iw})_{i \in \mathcal{I}, w \in W}$ denote the collection of distributions in the perturbed game, which we denote by $\tilde{\Gamma}(F)$ or $\tilde{\Gamma}$. A behavior strategy for player i in $\tilde{\Gamma}$ is $\sigma_i : X_i \times Z_i \to A_i$. Since the support of payoff shocks is large enough, this implies any σ_i that is not strictly dominated takes every available action with strictly positive probability at each decision node. σ_i defines an *aggregate*: a completely mixed behavior strategy $\tilde{b}_i \in B_i$. Our solution concept for $\tilde{\Gamma}$ is Bayes Nash equilibrium, which suffices since every decision node in X is reached with positive probability.

Let $\mathbf{0}$ denote the null-vector in $\mathbb{R}^{|W| \times |N|}$. The unperturbed game Γ corresponds to degenerate payoff shocks given by a Dirac measure on $\mathbf{0}$.

Definition 1. *A strategy profile b of the unperturbed game is:*

- *Purifiable if there exists some sequence of distributions F^n, converging weakly to the Dirac measure on $\mathbf{0}$, and a sequence of equilibria σ^n whose aggregates $\tilde{b}^n$ converge to b.*
- *Strongly purifiable if for every sequence of distributions F^n , that converge weakly to the Dirac measure on $\mathbf{0}$, there exists a sequence of equilibria σ^n whose aggregates $\tilde{b}^n$ converge to b.*

Existence of purifiable strategy profiles in general finite extensive form games is straightforward, and follows from the following proposition.

Theorem 2. *Let Γ be a finite extensive form game. For any sequence F^n of distributions that converge to a Dirac measure on $\mathbf{0}$, there exist a convergent subsequences (F^{n^k}, σ^{n^k}) such that the aggregates $\tilde{b}^{n^k}$ converge to b, where b is a Nash equilibrium of Γ.*

Proof. Fix payoff perturbations for player i, and let $\hat{\sigma}_i(b)$ denote (the essentially unique) best response to a behavior strategy profile b in Γ. This defines an aggregate best response map $\hat{b}(b)$, that is continuous since the shocks are atomless. Brouwer's fixed point theorem ensures the existence of a Bayes Nash equilibrium in any perturbed game. Since the space of behavior strategies is compact, the existence of a convergent subsequence of equilibria follows. Finally, the limit profile is a Nash equilibrium since the weak inequalities defining an equilibrium hold in the limit. ■

The existence of strongly purifiable equilibria is an open question. Recall that Harsanyi (1973) showed that in any strategic form game with generic payoffs, every Nash equilibrium is strongly purifiable. Of course, any non-trivial extensive form games gives rise to payoff ties in the strategic form, so that Harsanyi's theorem is not applicable.

3 Games of perfect information

We now specialize to finite games of perfect information, where each information set is a singleton decision node and a single player moves at each such node. That is X is the set of decision nodes. We will focus on games without any payoff ties. That is, if w and w' are distinct terminal nodes, then $u_i(w) \neq u_i(w')$ for every player $i \in N$. It is well known that in the absence of payoff ties, Γ has a unique backwards induction strategy profile b^*. Furthermore, the strategic form of the game is dominance solvable, so that the iterated elimination of weakly dominated strategies gives us b^* as the unique solution.

3.1 Simple games

Our first result is for simple games of perfect information where along any path of play, each player moves at most once.

Proposition 3. *Let Γ be a perfect information game where along any path of play every player moves at most once. The backwards induction strategy profile is strongly purifiable. Moreover, no other strategy profile is purifiable.*

Proof. The proof is by backwards induction. Fix a penultimate node x, and suppose that player i has to move at this node. Let w^* denote the terminal node induced by i's (unique) optimal action and let $W(x)$ denote the set of sub-optimal terminal nodes that are reachable from x. For $w \in W(x)$, let $\Delta(w) > 0$ be the payoff difference for player i between node action w_1 and node w. Since player i has not, by assumption, moved previously on the path to node x, the probability of this node being reached is independent of i's payoff shocks. Let $\tilde{F}_{iw}$ denote the distribution of the difference $z_{iw^*} - z_{iw}$. Thus the probability that i chooses the optimal action at node x is $\Pi_{w \in W(x)} \tilde{F}_{iw}(\Delta(w))$, which converges to 1 as F_i converges to zero.

Now consider any decision node x where player i has to move. Suppose, by the induction hypothesis, that at every successor node x', the backwards induction continuation path is played with a probability close to 1. It follows that player i strictly prefers the backwards induction continuation strategy, in the absence of any shocks to his payoffs. Thus when F_i is close to zero, she plays the action corresponding to the backwards induction action with a probability close to one. ■

Note that the proof is by backwards induction. Since any player moves at most once along a path of play, and since the payoff shocks are independent across players, no inferences regarding the behavior of other players can be drawn from an unlikely move by some player i. Along the backwards induction continuation path from any decision node, each player has strict preferences in the unperturbed game for complying with her backwards induction strategy given that others do so. Consequently, when payoff shocks are small, each player is very likely to continue with the backwards induction strategy.

The following example, and the one that follows, illustrate our result, and elucidates the differences between our assumptions and those of Dekel and Fudenberg (1990).

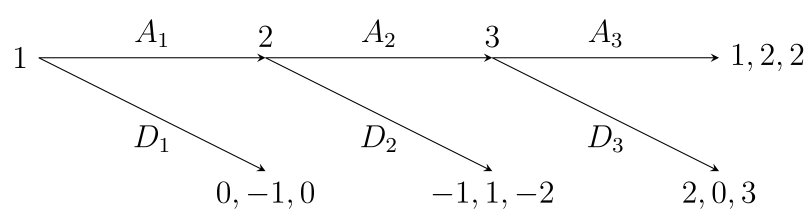

The backward induction profile is $(D_1; D_2; D_3)$, with the choices corresponding to 1, 2 and 3, respectively. In the perturbed game, player 3's decision node is reached with positive probability. Since the shocks of players 1 and 3 are independent, from player 2's point of view, the conditional distribution of player 3's shocks equals the ex ante distribution. Consequently, player 2 believes that player 3 will play D_3 with high probability, leading her to play D_2 with high probabililty. This in turn implies that 1 must play D_1 with high probability.

Dekel and Fudenberg (1990) analyze strategic form games and allow the shocks for different players to be correlated. They study strong convergence. This delivers $S^\infty W$, i.e. one round of elimination of weakly dominated strategies, followed by iterated strict dominance. This procedure eliminates A_3 for player 3, but no further elimination is possible. Our analysis differs in three dimensions. We rule out correlation of shocks, and we study payoff uncertainty in the extensive form, so that the dimensionality of the set of payoffs is lower than in the corresponding strategic form. This allows us to prove a stronger result even though we only require weak convergence of payoffs. Indeed, if we allowed uncertainty about the payoffs in the corresponding strategic form, we would not be able to get even one round of elimination of weakly dominated strateges, as the entry deterrence example shows.

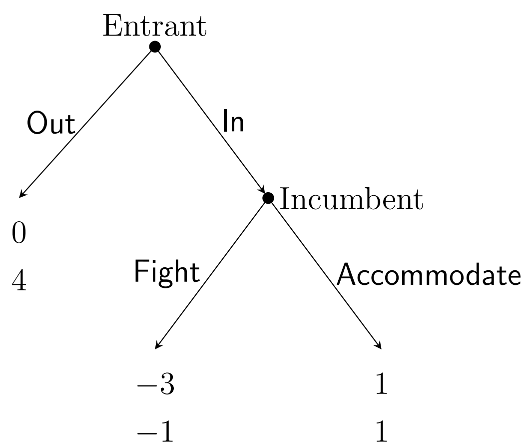

Proposition 3 implies that the backward induction strategy profile $(IN; A)$ is strongly and uniquely purifiable. Now consider the corresponding strategic form.

	A	F
IN	$1, 1$	$-3, -1$
OUT	$0, 4$	$0, 4$

Suppose that we perturb payoffs at each profile in $\{IN, OUT\} \times \{A, F\}$. Let Let $G = (G^1, G^2, G^3, G^4)$ denote the shock distributions for player 2, and let $\hat{G}$ denote the distribution of the difference in shock values between (OUT, A) and (OUT, F). Assume that $\hat{G}$ has median 0. Then the equilibrium $(OUT, p(F) = 0.5)$ is purifiable. Let When the entrant plays OUT with a probability close to 1, the incumbent plays F with a probability $\approx \hat{G}(0) = 0.5$. This implies that it is optimal for the entrant to stay OUT with high probability. In other words, we cannot eliminate weakly dominated strategies in non-generic strategic form if we only require weak convergence. However, since we perturb extensive form payoffs, the payoff of the incumbent when the entrant plays OUT does not depend on the incumbent's choice between A and F. This implies that the incumbent's choice between A and F only matters when the entrant chooses IN, implying that the incumbent chooses A with a probability close to one.

This example also suggests a route to studying non-generic extensive form games. We can partition terminal nodes W into sets of player-specific *outcomes*, so that the payoffs of i are generic on outcomes, but constant at same outcome. We shall explore this idea, which builds on ideas in Mailath, Samuelson, and Swinkels (1993), later.

3.2 Not so simple games

Things are more complicated when a player moves more than once along some path of play. To illustrate this, the minimal example is one where player 1 moves twice and player 2 moves once, as in the simplest version of the centipede game. Player 1 moves first choosing between D_1, yielding payoffs $(0, -1)$ and A_1. In the latter event, player 2 chooses between D_2, yielding $(-1, 1)$, and A_2. In the latter event, player 1 moves choosing between D_3 (yielding $(2, 0)$) and A_3, which yields $(1, 2)$.

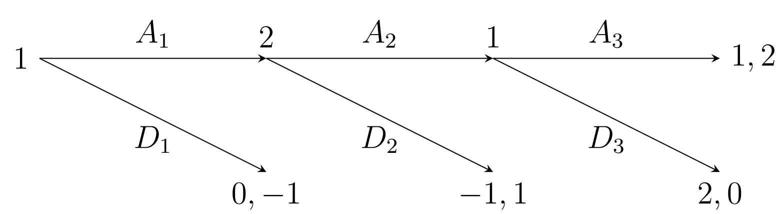

3.2.1 One-dimensional shocks

Let us first consider an illustrative example of payoff shocks where we shock the payoff of each player at a single terminal node. Player 1 has a shock ϵ_3, which has cdf F_3 at the terminal node that follows A_3. Player 2 has a shock η at the terminal node following D_2, with distribution G. If $\eta < -1$, D_2 is strictly dominated for player 2 and thus he plays A_2 with probability at least $G(-1)$. Thus, if ϵ_3 is large enough, 1 must play A_1, and continue with A_3. However, if A_1A_3 is played with positive probability and A_1D_3 is never played, then 2 will find it optimal to play A_2 for any $\eta < 1$, i.e. with high probability $G(1)$, which is greater than 2/3 if shocks to 2's payoff are small. But this implies that 1 will play A_1

for every shock value, i.e. even if $\epsilon_3 < 1$, so that she will continue with D_3 at her second decision node. We conclude therefore that in any equilibrium, 1 must play A_1 with positive probability when $\epsilon_3 < 1$, which also implies that she plays A_1 for sure if $\epsilon_3 > 1$. Let $\bar{\eta}$ be the value of player 2's shock where he is indifferent between D_2 and A_2, and let $\mu = \frac{\Pr(A_1A_3)}{\Pr(A)}$. Since player 1 must be indifferent between A_1 and D_1 when her shock value is low (and she therefore intends to play A_3), $G(\bar{\eta}) = \frac{2}{3}$, and $1 + \bar{\eta} = 2\mu$. Furthermore, as F_3 converges weakly to zero, $\Pr(D_1)$ converges to zero, and as G converges weakly to zero, $\bar{\eta}$ converges to zero and μ converges to 0.5. In other words, as the shocks vanish, we converge to the following behavior strategy profile $\hat{b}$:

- Player 1 plays D_1 with probability one at her first decision node, and randomizes with equal probability between her two actions at her second decision node.
- Player2 plays D_2 with probability 2/3 at his single decision node.

Observe that $\hat{b}$ is a Nash equilibrium that is not subgame perfect. Since every perturbed game in the sequence has a unique equilibrium, which converges to $\hat{b}$ we conclude that $\hat{b}$ is purifiable but that the backwards induction strategy profile is *not* purifiable for this sequence of shocks. Of course, since we have only shocked one payoff for each player, this example is merely suggestive.

3.2.2 Shocks at all terminal nodes

We now consider shocks to both at every terminal node. Our analysis of the centipede will assume throughout that the shocks for player 1 at the terminal node following A_3, with distribution F_3, along the sequence, are more dispersed, along the sequence, than the shocks at her other terminal nodes. Our first result is the following:

Proposition 4. *Fix a sequence of convergent shock distributions for player* 1 *at the terminal node following* A_3*. There exist convergent sequences of shocks for player 1 at the other nodes such that the non-backward induction strategy profile* $\hat{b}$ *is always a limit, regardless of the shock sequence for player 2.*

Proof. Observe that for any player, we can normalize the payoff shock at one of the terminal nodes to zero, by differencing it out. Accordingly, we normalize the shock value for player 1 at the terminal node following D_3 to zero, by differencing out this value from each terminal node shock value. With this normalization, we have independent shocks to player 1's payoff at the nodes resulting from D_1, D_2 and A_3. Let ϵ_1 and ϵ_2 denote the shock values and F_1 and F_2 the distributions for the first two shocks. As before, let ϵ_3 and F_3 denote values and distribution for the third shock, at A_3.

Let q denote the probability that 2 plays D_2. Player 1 will play D_3 at her second decision node if $\epsilon_3 \leq 1$, an event that has probability $F_3(1)$. Given this, she will play A_1 only if $\epsilon_2 - \frac{1}{q}\epsilon_1 \geq 3 - \frac{2}{q}$. Thus the probability that she plays strategy $A_1 D_3$ is

$$\Pr(A_1 D_3) = \Pr\left\{\epsilon_1 - q\epsilon_2 < 2 - 3q\right\} F_3(1). \tag{1}$$

Now suppose that $\epsilon_3 > 1$, so that player 1 will play A_3 if her second decision node is reached. She will play A_1 at her first decision node if $\epsilon_2 - \frac{1}{q}\epsilon_1 \geq 3 - \frac{1+\epsilon}{q}$. Thus the probability that she plays strategy $A_1 A_3$, conditional on ϵ_3, is

$$\Pr(A_1 A_3 | \epsilon_3) = \Pr\left\{\epsilon_1 - q\epsilon_2 < (1+\epsilon_3)(1-q) - q\right\}. \tag{2}$$

Thus the total probability that 1 plays $A_1 A_3$ is

$$\Pr(A_1 A_3) = \int_1^\infty \Pr(A_1 A_3 | \epsilon_3) F_3(\epsilon_3) d\epsilon_3. \tag{3}$$

Observe that both the above equations 1 and 3 describe player 1's best response to q, the probability that 2 plays D_2. Let $\Pr(A) := \Pr(A_1 D_3) + \Pr(A_1 A_3)$ be the probability that 1 plays action A.

Now consider player 2. We normalize his payoff shock after D_3 is played to zero and let η_1 denote his payoff shock after D_2 is played and η_2 the shock after A_3 is played.[2] Playing D_2 is optimal at player 2's single decision node if

$$(1+\eta_1) > \frac{\Pr(A_1 A_3)}{\Pr(A)}(2+\eta_2). \tag{4}$$

Thus the probability that player 2 chooses strategy D_2 is

$$q = \Pr\left\{(2+\eta_2)\Pr(A_1 A_3) > \Pr(A)(1+\eta_1)\right\}. \tag{5}$$

This defines player 2's best response to $\Pr(A_1 D_3), \Pr(A_1 A_3)$.

An equilibrium of the perturbed game consists of a triple of probabilities $(\Pr(A_1 A_3), \Pr(A_1 D_3), q)$

[2] The payoff shock after D_1 is played is irrelevant.

that solve the system of equations 1,3 and 5. Since 1 and 3 depend only on q and q depends only on $(\Pr(A_1A_3), \Pr(A_1D_3))$, we define the composition of best responses, $\tilde{q}(q)$. An equilibrium is a fixed point of the function $\tilde{q}$. However, our goal is to construct an equilibrium that converges to σ, the non backwards induction Nash equilibrium, which corresponds to $q = \frac{2}{3}$ and $\frac{\Pr(A_1D_3)}{\Pr(A)} = \frac{1}{2}$. Consequently we define a function $q^\dagger$, with domain and range $[\underline{q}, \bar{q}]$, where $\underline{q} < \frac{2}{3} < \bar{q}$.

$$q^\dagger(x) = \begin{cases} \max\{\tilde{q}(x), \underline{q}\} & \tilde{q}(x) \leq \bar{q} \\ \min\{\tilde{q}(x), \bar{q}\} & \tilde{q}(x) \geq \underline{q}. \end{cases} \tag{6}$$

Fix a distribution F_3 such that $F_3(1) >> 1 - F_3(1)$. At $\underline{q}$, $3 - \frac{2}{\underline{q}} < 0$. Thus the probability of the event defined in curly brackets in equation 1 tends to one as F_1 and F_2 converge weakly to zero, and $\Pr(A_1D_3) \to F_3(1)$. On the other hand, $\Pr(A_1A_3)$ is bounded above by $1 - F_3(1)$, so that $\frac{\Pr(A_1A_3)}{\Pr(A)} << \frac{1}{2}$. Consequently, $\tilde{q}(\underline{q})$ is close to one if the shocks to player 2's payoffs, G, are small enough, implying that $q^\dagger(\underline{q}) = \bar{q}$.

Now consider $\bar{q}$, where $3 - \frac{2}{\bar{q}} > 0$. The probability of the event defined in curly brackets in equation 1 tends to zero as F_1 and F_2 converge weakly to zero, and $\Pr(A_1D_3) \to 0$. However, $\Pr(A_1A_3)$ is bounded away from zero, for a fixed F_3. Pick a $\Delta > 0$ such that $3 - \frac{1+\Delta}{\bar{q}} > 0$. Thus $\Pr(A_1A_3|\epsilon_3) \to 1$ as F_1 and F_2 converge weakly to zero, so that $\Pr(A_1A_3)$ converges to $K \geq 1 - F_3(1 + \Delta)$. Consequently, $\tilde{q}(\bar{q})$ is close to zero if the shocks to player 2's payoffs, G, are small enough, implying that $q^\dagger(\bar{q}) = \underline{q}$).

Finally, since the shocks have a continuous distribution, without mass points, $\Pr(A_1A_3)$ and $\Pr(A_1D_3)$ are continuous in q,and the expression for q in 5 is continuous in $\Pr(A_1A_3)$ and $\Pr(A_1D_3)$. Thus $q^\dagger$ is a continuous function, and the intermediate value theorem ensures the existence of a fixed point in the interval $[\underline{q}, \bar{q}]$ that contains $\frac{2}{3}$. Since this interval can be taken to be arbitrarily small, and since we can now take a sequence of F_3 converging weakly to zero, this ensures that the equilibrium σ, where player 1 plays D_1 at her first information set and randomizes equi-probably at her second information set, while player 2 plays D_2 with probability $\frac{2}{3}$, is purifiable. ■

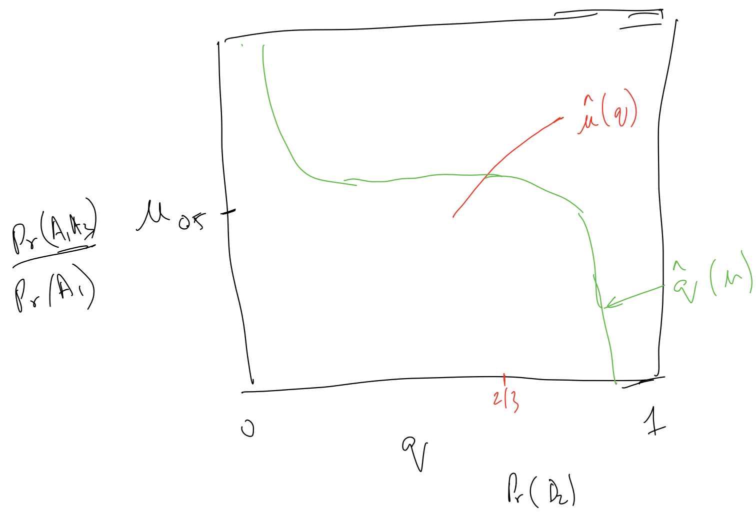

The above figure illustrates the argument. We depict q, the probability that player 2 plays D_2, on the horizontal axis, and μ, the conditional probability that player 1 chooses A_3 at her second decision node on the vertical axis. The green curve depicts player 2's best response $\hat{q}(\mu)$, which is a smooth strictly decreasing function close to a "step function" that is horizontal at $\mu = 0.5$ when the shocks to 2's payoffs are small. The red curve depicts $\hat{\mu}(q)$, that is derived from player 1's best response. When $q \leq 2/3$, most types of player 1 choose A_1, so that μ is small, being close to the unconditional probability that A_3 is better than D_3, $1 - F_3(1)$. When $q > 2/3$ but is close to it, the payoff from the strategy A_1D_3 is strictly negative in the unperturbed game, and thus the probability of this strategy can be made arbitrarily small by choosing F_1 and F_2 sufficiently close to 0. However, if player 1 has a large shock value of ϵ_3, say 2, she will still finds it profitable to choose A_1 for most (ϵ_1, ϵ_2) values, so that $\Pr(A_1A_3) \approx 1 - F_3(2)$. This implies that $\hat{\mu}(\bar{q}) > 0.5$, ensuring an intersection between the two best responses.

Our next proposition implies that the backwards induction strategy profile is not strongly purifiable. Recall that G denotes the distribution of shocks for player 2.

Proposition 5. *Fix shock sequence G converging weakly to zero. There exists a converging sequence F^n of player 1's shocks such that the backwards induction strategy profile is not a limit of any sequence of equilibria.*

Proof. Fix a distribution G for player 2'shocks. This defines $\bar{q} := \hat{q}(1)$, the probability that player 2 plays D_2 when $\mu = 1$. Since $\hat{q}$ is increasing in μ, $\bar{q}$ is an upper bound for q. Let $\underline{q} \in (2/3, \bar{q}]$. We show that for any $q \in [\underline{q}, \bar{q}]$, there exists a distribution of $\epsilon = (\epsilon_1, \epsilon_2, \epsilon_3)$ such that $\mu(q) = \frac{\Pr(A_1A_3)}{\Pr(A_1)} \approx 1$, so that 2's best response is $q \approx 0$.

From equation 1, $\Pr(A_1D_3) \rightarrow 0$ for any $q \in [\underline{q}, \bar{q}]$ as F^1 and F^2 converge weakly to 0. Choose $\bar{\epsilon} > \frac{\bar{q}}{1-\bar{q}} - 1$. For any $\epsilon_3 > \bar{\epsilon}$ and any $q \leq \bar{q}$, $\Pr(A_1A_3|\epsilon_3) \rightarrow 1$ as F_1 and F_2 converge weakly to 0, so that $\Pr(A_1A_3)$ converges to a value greater than $1 - F_3(\bar{\epsilon})$. Thus $\mu(q) \approx 1$, implying that 2's best response is $q \approx 0$. Thus there is no equilibrium with $q > \underline{q}$. ■

Proposition 6. *Fix any shock sequence $F^n = (F_1^n, F_2^n, F_n^3)$ converging weakly to zero. There exists a sequence of player 2 shocks G^n converging to zero, such that the backwards induction strategy profile is a limit of some sequence of equilibria of the games $\tilde{\Gamma}_n$.*

Proof. Consider the composition of best responses for player 2, $\tilde{q}(q)$. Shocks to player 2's payoffs implies $\hat{q}(\mu)$ is bounded above by $\bar{q} < 1$, where $\bar{q} := \tilde{q}(1)$ is close to one if G is close enough to zero. At $q = 1$, player 1's decision to play A_1 is independent of ϵ_3 and only depends upon (ϵ_1, ϵ_2). Thus, $\hat{\mu}(1) = 1 - F_3(1) \approx 1$. Let F be sufficiently close to zero such that $\hat{\mu}(\bar{q}) > 3/4$. By choosing G sufficiently concentrated, we can make $\hat{q}(\mu) > \bar{q} \forall \mu \geq 3/4$. Thus $\tilde{q}(\bar{q}) > \bar{q}$. Thus there must be a fixed point of the composition $\tilde{q}(q)$ in the interval $(\bar{q}, 1)$. Since $\bar{q}$ converges to 1 as G converges weakly, this suffices to prove the proposition. ■

The following two pictures illustrate the two cases. Player 2's best response is plotted in green, while that of player 1 is in red. in the first graph, player 2's shocks are relatively dispersed relative to player 1's shocks, and there is no fixed point close to the backwards induction profile. In the second figure. we have kept player 1's shocks unaltered but made those of two more concentrated, and this ensures an equilibrium close to the backwards induction strategy profile.

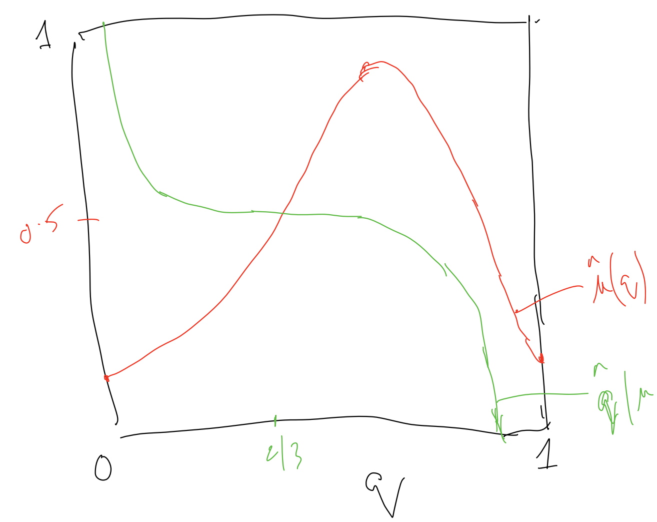
1
0.5
0
2/3
1
q
$\hat{\mu}(q)$
$\hat{q}_f/\mu$

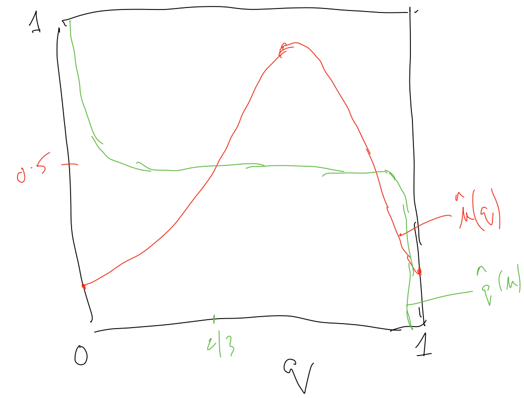

3.3 Stong purifiability of outcomes

Definition 7. *The outcome of a strategy profile is the distribution over terminal nodes induced by it. An outcome $\omega \in \Delta(W)$ is strongly purifiable if for every converging sequence F^n, there exists a sequence of behavior strategy profiles b^n whose outcomes converge to ω.*

Proposition 8. *Let Γ be a generic game of perfect information, without payoff ties. If b is a purifiable strategy profile, then it induces the backwards induction outcome. The backwards induction outcome is strongly purifiable.*

Proof.(Sketch) We will show that any path that is not the backward induction path has a probability that converges to zero as the shocks vanish. In the case of outcomes, we can base the argument on ex ante probabilities rather than conditional ones. We illustrate the argument in the context of the centipede game. In the perturbed game, the ex ante probability that the terminal node following A_3 is reached is small, being no greater than $1 - F_3(1)$. Thus the ex ante probability that player 2 plays A_2 must be small, since his ex ante expected payoff from playing A_2 is strictly less than the payoff from D_2, in the unperturbed game. This in turn implies that the ex ante probability that player 1 plays A_1 must be small. The general argument is similar: at any penultimate node, the probability of any choice other than the uniquely optimal one must have a small ex ante probability. The rest of the argument follows by backwards induction. ■

3.4 Symmetric purification

Our results imply that it if wants strong purifiabilty to deliver sequential rationality of strategy profiles, we must restrict the type of perturbations. We now investigate this by requiring that for any player, the shock distributions are the same at all terminal nodes. In addition to our already stated assumptions, we now invoke:

Assumption 9. *The shock distributions satisfy symmetry if every player i, $F_{iw} = F_{iw'}$ for every pair of terminal nodes w and w'.*

Definition 10. *A convergent sequence of distributions F^n has* thin tails *if* $\forall x > y > 0$

$$\lim_{n\to\infty} \frac{1 - F^n(x)}{1 - F^n(y)} = 0,$$

$$\lim_{n\to\infty} \frac{F^n(-x)}{F^n(-y)} = 0.$$

A sequence of mean-zero normal distributions with the variances converging to zero is thin tailed.

We will invoke these twin assumptions (symmetry and thin tails) together. This allows us to define *symmetric purification* and *strong symmetric purification*, by which we mean that the shock sequences satisfy both assumptions.

We now investigate the implications of symmetric purification in the context of the centipede game. It will be convenient to adopt a new normalization. For player 1, set the shock following D_2 to zero and let ϵ_1 be the shock at the terminal node following D_1, ϵ_3 the shock following A_3 and ϵ_4 the shock following D_3.

Conditioning on a realization of ϵ_3, the probability that A_1A_3 is played is

$$\Pr(A_1A_3|\epsilon_3) = F_1(-q + (1-q)(1+\epsilon_3))F_4(\epsilon_3 - 1). \tag{7}$$

Conditioning on a realization of ϵ_4, the probability that A_1D_3 is played is

$$\Pr(A_1 D_3|\epsilon_4) = F_1(-q + (1-q)(2+\epsilon_4))F_3(\epsilon_4 + 1). \quad (8)$$

Invoking the assumption that the distributions are the same,

$$\Pr(A_1 A_3) = \int F(-q + (1-q)(1+\epsilon))F(\epsilon - 1)f(\epsilon)d\epsilon. \quad (9)$$

$$\Pr(A_1 D_3) = \int F(-q + (1-q)(2+\epsilon))F(\epsilon + 1)f(\epsilon)d\epsilon. \quad (10)$$

Now for any ϵ, $F(-q+(1-q)(2+\epsilon) \geq F(-q+(1-q)(1+\epsilon)$. Furthermore, as F converges weakly to the Dirac measure on 0,

- For any $\epsilon \in (-1/2, 1/2)$, F$(\epsilon - 1) \to 0$ and F$(\epsilon + 1) \to 1$.
- $F(-1/2)$ and $1 - F(1/2) \to 0$.

We conjecture but have not yet proved the following.

Conjecture 11. *In any generic game of perfect information, the unique symmetrically purifiable strategy is the backwards induction strategy profile b^*. b^* is therefore strongly symmmetrically purifiable.*

3.5 Simultanenous moves

We now consider games where more than one player moves at the same time.[3] That is, at any decision node, more than one player can move, and players have perfect information on all past moves (including those by nature). Let us consider the following example.

[3]These are called games of perfect information by Osborne and Rubinstein (1994), but the terminology is not universal, and so we shall avoid it.

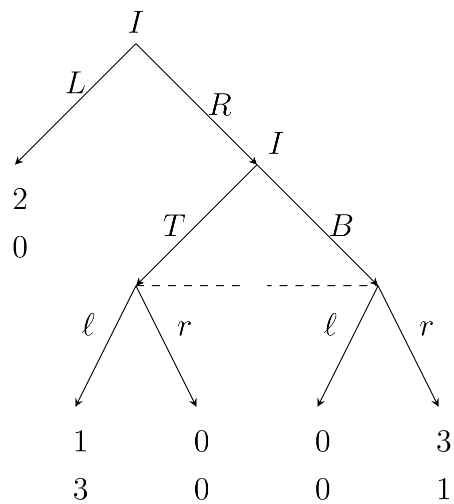

Observe that (LT, ℓ) and (RB, r) are subgame perfect equilibria. The profile (LT, ℓ) does not satisfy forward induction or iterative elimination of weakly dominated strategies.

Observe first that (LT, ℓ) is symmetrically purifiable. To see this, suppose that player 1 believes that that player 2 will play ℓ with probability one. Then, since the payoff loss in the base game when choosing RT is 1, while the loss from choosing RB is 2, player 1 is infinitely more likely to choose RT than RB, implying that player 2 will play ℓ with probability close to one. This ensures that there is a fixed point close to the profile (LT, ℓ).

However, the profile (LT, ℓ) is not strongly purifiable. Suppose that player 1's shocks ϵ_1 at (LT, r) are much more variable than the shocks ϵ_2 at (LT, ℓ). For example, let the shocks be normally distributed, with the standard deviations σ_1 and σ_2, with $\sigma_2 = M\sigma_1$, $M > 2$. Then as $\sigma_1 \to 0$, player 1 is infinitely more likely to choose RB rather than RT, so that player 2 must play r.

Finally, note that (RB, r) is strongly purifiable. No matter what the distributions of player 1's shocks, if they are small, there exists an equilibrium of the perturbed game where player 1 plays RB with probability close to one and player 2 plays r with probability close to one. Indeed, it seems that this is generally true: any equilibrium where every information set is reached is strongly purifiable.

Consider now an alternative extensive form representation of the above game, where both players move simultaneously, with action sets $\{L, T, B\}$ and $\{\ell, r\}$, and the payoffs after 1 plays L are $(2, 0)$, independent of player 2's choice. Assume that the payoff shocks after L are one-dimensional for each player, i.e. the shock realization is identical regardless of player 2's action. The analysis of the perturbed version of this game is identical to that of the previous extensive form, and thus the same results apply.

This example raises the following questions.

- In games of perfect information with only one player moving at a time, every subgame perfect equilibrium (backwards induction) outcome is always strongly purifiable. The forward induction example shows that this is not true for games with simultaneous moves.

- Strong purifiability of outcomes, in this example, corresponds to the outcome given by iterated elimination of weakly dominated strategies. This is also true in the perfect information case.

- Finally, is there a relation between strongly purifiable outcomes and strategically stable outcomes in (i.e. outcome in a strategically stable set of a generic game)? One conjecture is that the two are equivalent.

3.6 Strategy trembles vs payoff shocks

The main approaches to refining Nash equilibria rely on strategy trembles. Extensive trembling hand perfection is based on trembles that are independent across information sets. The same is true of sequential equilibrium, where beliefs are derived as the limit of those induced by independent trembles. Our analysis of the centipede game shows that payoff shocks that are independent across terminal nodes do not deliver trembling hand perfect/sequential equilibrium. In simple games, where along every path of play any player moves at most once, then payoff perturbations induce independent trembles. This would seem to ensure that any purifiable equilibrium is also trembling hand perfect, not just in games of perfect information, but other such games, such as signaling games.

Strategic stability relies on trembles to the normal form strategies, not to choices at each decision node. This allows for arbitrary correlation of choices across information sets. It also requires robustness (of a set of strategies) to all possible trembles. A natural conjecture is that any strategically stable outcome is also a strongly purifiable, since the possible strategy trembles that one can derive from independent payoff shocks is a subset of those that are possible when requiring strategically stability. This remains to be proven. Finally, strategic stability is a normal form concept. This is also the case for purification, since we obtain limits of Bayes Nash equilibria, without any direct reference to sequential rationality.

Our examples suggest that the relationship between strong purifiability and strategic stability is close:

- In the centipede game, we have shown that the backwards induction strategy profile $(D_1; D_2; D_3)$ and the profile $(D_1; 2/3 \times D_2, 1/3 \times A_2; 1/2 \times D_3, 1/2 \times A_3)$ are both purifiable, and neither is strongly purifiable. The strategically stable set in this game defines strategies only in the reduced normal form, and has $(D_1; D_2)$ and $(D_1; 2/3 \times D_2, 1/3 \times A_2)$ as its elements. In other words, the purifiable strategy profiles, when restricted to the reduced normal form, coincide with elements of the strategically stable sets.

- In the forward induction example, strategically stability and strong purification yield identical conclusions, regarding strategy profiles (and therefore, of outcomes).

Finally, the reader may wonder why we analyze extensive form games, since Bayes Nash equilibria in perturbed games can be defined without requiring sequential rationality explicitly. The reason is that the extensive form allows us to consider lower dimensional payoff perturbations. Essentially, when payoff assignments are identical at two strategy profiles (in the strategic form), the realization of payoff shocks is identical for all players. Of course, one could impose this assumption directly in the strategic form, and obtain identical conclusions. The weaker requirement that whenever the payoffs of a player at two strategy profiles (or two terminal nodes) are equal, the realized payoffs of that player are identical in the perturbed game, would allow us to also non-generic extensive form games, where a player obtains equal payoffs at two terminal nodes.

3.7 Games with payoff ties

We now consider games of perfect information with payoff ties. In the example below, players 1 and 2 are investors in an project. At date 0 player 1, the senior partner, must decide whether to retain the decision rights on the project or to transfer the rights to player 2, an option that is labelled delegation. At date 1, the person who has the decision rights must decide whether to continue with the project, C or liquidate it, L, and the project is profitable for both partners. Players only care about whether the project is continued or liquidated, and not per'se about the allocation of decision rights. This game has a continuum of subgame perfect equilibria, where player 1 delegates with a probability $p \in [0, 1]$, and the decision maker always chooses C. Each of these equilibria is trembling hand perfect. In particular, the equilibrium where player 1 delegates with probability 1 is uniquely selected when player 1 is more likely to tremble than player 2 at date 1.

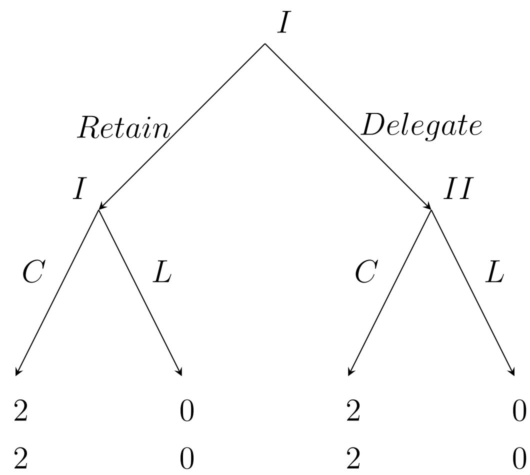

Observe that the set of terminal nodes can be partitioned into two subsets: one where the project is continued, which yields both players a payoff of 2, and the complement where it is liquidated, yielding payoff 0 to both. Let us add payoff shocks (ϵ_C, ϵ_L) for player 1 and (η_C, η_L) for player 2. We now argue that there is a unique purifiable equilibrium, where player 1 keeps the decision rights. In the perturbed game, when player 1 liquidates the project, it is in her interest to do so, since liquidation only occurs is $\epsilon_L \geq \epsilon_C + 2$. However, when player 2 liquidates the project, this decision is independent of player 1's shock value, and hence is always costly for player 1. In other words, purification selects quasi-perfect equilibria (van Damme (1984)) in this example, not trembling hand perfect ones.

Now consider a different example, with only a single player, who can decide at date 0 to give up the right to liquidate the project at date 1, so that the project must continue, or to retain this right.

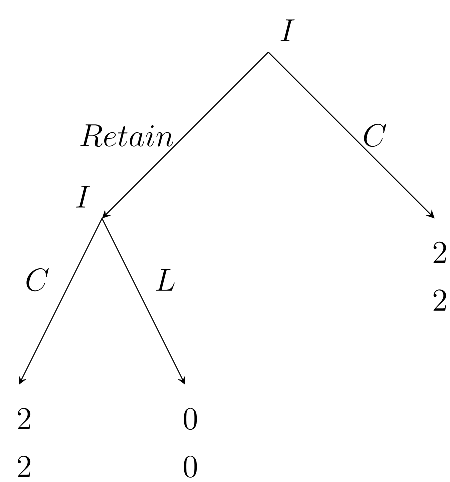

The unique trembling hand perfect equilibrium has the player giving up the right to liquidate, to protect herself against trembles by her future self. Purification is agnostic on this question since we assume identical payoff shock realizations at both terminal nodes that correspond to project continuation. If $\epsilon_C - \epsilon_L > -2$, the player is indifferent between giving up the right today or retaining it, since she will continue in both instances.

4 Symmetric purification in signaling games

Let $\Gamma =< \Theta, M, A, u, v >$ be a finite signaling game, where Θ is the set of sender types, M the set of messages, A the set of receiver actions and u and v are payoff functions for sender and receiver respectively. For simplicity, we assume that the set of available receiver actions A does not depend upon m, the message sent by the sender, an assumption that is without loss of generality. We will use s to denote the sender's strategy, and r to denote the receiver's strategy.

First, we define a partial assessment, $(s, \bar{r}; \bar{\mu})$. This consists of a complete strategy for the sender, $s : \Theta \to \Delta(M)$, and a partial strategy for the receiver, $\bar{r}$, that is defined with reference to s. Let $\bar{M}$ denote the set of messages that are sent with positive probability by some sender type, $\bar{M} := \{m \in M : \exists\theta \in \Theta : s(m|\theta) > 0\}$. A partial strategy for the receiver is a function $\bar{r} : \bar{M} \to \Delta(A)$. Finally, $\bar{\mu}(m)$ assigns, via Bayes rule, receiver beliefs $\mu(m) \in \Delta(\Theta)$ for any $m \in \bar{M}$. Observe that the pair $(s, \bar{r})$ suffices to define a payoff vector for each sender type, $u := (u^*_\theta)_{\theta\in\Theta}$, and a payoff for the receiver, v^*. We will restrict attention to partial strategy profiles which satisfy twin properties:

- For each $\theta \in \Theta$, $s(\theta)$ is optimal given $\bar{r}$ when the sender is restricted to messages in $\bar{M}$.

- $\bar{r}(m)$ is optimal for the receiver at every $m \in \bar{M}$, when beliefs are given by $\bar{\mu}$.

It remains to assign receiver beliefs $\mu(m) \in \Delta(\Theta)$ and receiver best responses $r(m)$ to messages in $\bar{M}^C$, the complement of $\bar{M}$, to complete the partial assessment $(s, \bar{r}, \bar{\mu})$. Our focus will be on purifiable sequential equilibria, when the shocks for the sender have the same distribution at each terminal node, and also have thin tails.

First, we set out a sufficient condition for purifiable sequential equilibria when the receiver is restricted to point beliefs regarding the type who sends a deviating message.

4.1 P-completion

Definition 12. *Let Δ^v denote the vertices of the set of $\Delta(\Theta)$. These correspond to beliefs that assign probability one to single type, and we will denote of a typical element of Δ^v, by θ (by which we mean the Dirac measure on θ). Assume that the receiver has a unique best response to each element of Δ^v, an assumption that is satisfied generically.*

A partial assessment $(s, \bar{r}; \bar{\mu})$ can be P-completed if for every $m \in \bar{M}^C$, there exists $\hat{\theta}(m) \in \Theta$, such that the receiver's best response to belief $\hat{\theta}(m) \in \Delta^v$, $BR(\hat{\theta}(m))$ is such that

- $\forall \theta \in \Theta, u(\theta, m, BR(\hat{\theta}(m))) < u^*_\theta$, *and*
- $\hat{\theta}(m)$ *is the unique value of* θ *that maximizes* $u(\theta, m, BR(\hat{\theta}(m))) - u^*_\theta$.

If a partial assessment can be P-completed, we complete it by assigning belief $\hat{\theta}(m)$ and receiver best response $r(m)$ to every $m \in \bar{M}^C$.[4] The resulting assessment will be called a P-equilibrium. The first condition in definition 12 ensures that the completion is a sequential equilibrium. The second part ensures purifiability, as we will show.

Proposition 13. *Let (s, r, μ) be a P-completion of a partial assessment. Assume that s is a pure strategy and that the payoff shocks of the sender have the same distribution at each terminal node, and have thin tails. Then for almost all payoff assignments to terminal nodes, the profile (s, r, μ) is purifiable.*

Proof. Since s is a pure strategy, this implies that for generic values of the receiver's payoffs and the prior, the receiver will have a unique strict best response at any $m \in \bar{M}$. Furthermore, for generic sender payoffs, each sender type is strictly worse off by choosing a message $m \in \bar{M}, m \neq s(\theta)$. Furthermore, by the first condition in the definition of P-completion (cf. definition 12), each type of sender has strict incentives to not deviate to any message in $\bar{M}^C$. The only difficulty pertains to the receiver having strict incentives at messages in $\bar{M}^C$. Under the two assumptions on payoff shocks (that the sender has identically distributed shocks at each terminal node, and that the shocks have thin tails), the limit beliefs at $m \in \bar{M}^C$ equal $\mu(m)$, i.e.they assign probability one to $\hat{\theta}(m)$. Thus both players have strict incentives at each information set, implying that the equilibrium is purifiable. ■

Remark 14. *The condition that s is a pure strategy is probably superfluous, and is made only in order to ensure that we have a sequentially strict equilibrium. If s is a mixed strategy,*

[4] There maybe multiple P-completions (or none).

then one would probably have to show that it regular, so that standard proofs of purification (Harsanyi (1973), Govindan, Reny, and Robson (2003)) can be used.

Now, let us consider non-degenerative beliefs which assign positive probability to multiple types after a deviating message.

4.2 M-completion

Definition 15. *A partial assessment* $(s, \bar{r}; \bar{\mu})$ *can be M-completed if for every* $m \in \bar{M}^C$*, there exists a subset of types* $T \subset \Theta$*, a belief* $\mu(m) \in \Delta(T)$*, and a receiver mixed best response to* $\mu(m)$*,* $r(m) \in \Delta(S)$*, where* $S \subset A$*, such that*

- $\forall \theta \in \Theta, u(\theta, m, r(m)) < u^*_\theta$,
- $T = argmax_\theta[u(\theta, m, r(m)) - u^*_\theta]$,
- $S = argmax_a[v(m, \mu(m), a)]$.

If a partial assessment can be M-completed, we complete it by assigning belief $\mu(m)$ and receiver best response $r(m)$ to every $m \in \bar{M}^C$. The resulting assessment will be called a M-equilibrium. The first condition in definition 15 ensures that the completion is a sequential equilibrium. The second part ensures purifiability.

Proposition 16. *Let* (s, r, μ) *be a M-completion of a partial assessment. Assume that* s *is a pure strategy and that the payoff shocks of the sender have the same distribution at each terminal node, and have thin tails. Then for almost all payoff assignments to terminal nodes, the profile* (s, r, μ) *is purifiable.*

Proof. To be completed. ■

Conjecture : if a partial assessment cannot be M-completed, then no completion of the assessment is purifiable.

4.3 Examples

4.3.1 Beer-Quiche and P-completion

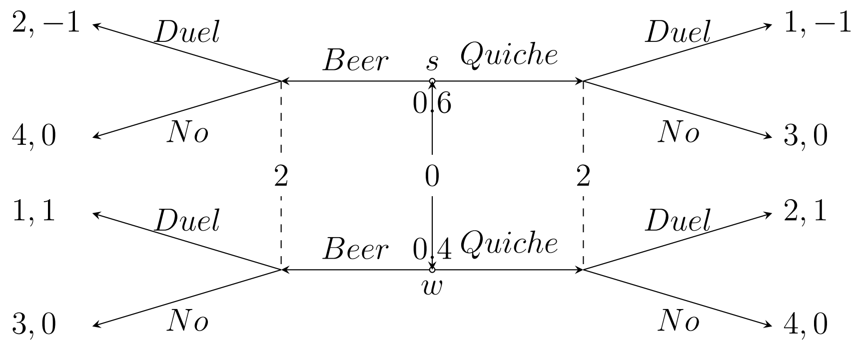

Our first example is the Beer-Quiche game from Cho and Kreps (1987). The partial assessment corresponding to the pooling on beer equilibrium is where both players choose Beer, and the receiver does not Duel. It can be P-completed as follows. The receiver believes that a deviant to Quiche is the weak type, and chooses Duel after Quiche. It is easy to verify that neither type wants to deviate to Quiche, and also that the type with the largest deviation incentive is the weak type. Thus the equilibrium is purifable.

Pooling on Quiche cannot be P-completed, nor can it be M-completed. For any response by the sender at Beer, the strong type has strictly smaller loss from playing Beer than the weak type, and hence is infinitely more likely to choose Beer in the perturbed game. Thus the limit beliefs assign probability one to strong at Beer, so that the unique best response is to not Duel.

4.3.2 Example of M-completion

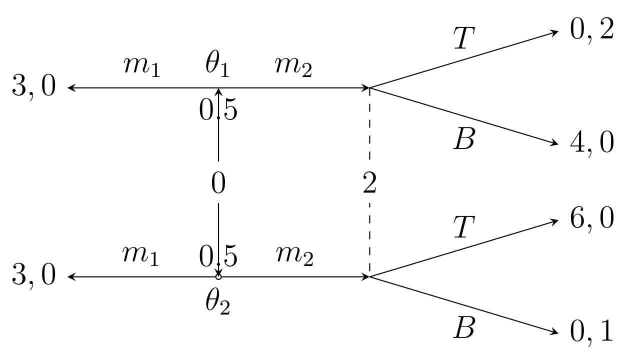

There is a unique sequential equilibrium outcome. The set of sequential equilibria has pooling on m_1. After the unsent message m_2, there is a unique belief, $\mu(\theta_1|m_2) = 1/3$, but multiple receiver best responses, $r(T) \in [1/4, 1/2]$. If we choose $r(T) = 2/5$, this is a M-completion, and hence purifiable.

4.3.3 NWBR, D1 & purification

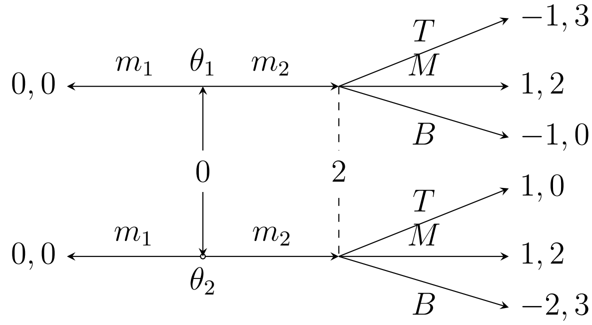

This example is from Cho and Kreps (1987). The sequential equilibrium with pooling on m_1, $\mu(\theta_2|m_2) = 1, r(m_2) = B$ satisfies D1. However, this violates NWBR. Cho and Kreps state: "we cannot suggest an intuitive inferential process" for NWBR in this game. However, symmetric purification justifies eliminating this equilibrium. If player 2 plays B with high probability, then type θ_1 has the smallest deviation loss. Thus θ_1 is infinitely more likely to play m_2 than θ_2 is in perturbed game.

4.3.4 Violates Equilibrium dominance but purifiable

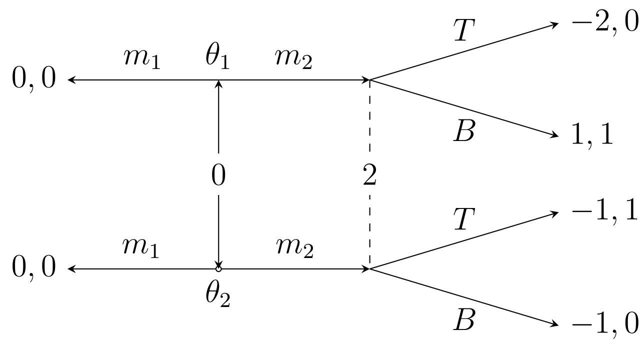

The sequential equilibrium with pooling on m_1, $\mu(\theta_2|m_2) = 1, r(m_2) = T$ is a P-completion. The payoff loss to deviating to m_2 is 1 for θ_2, and 2 for θ_1. Hence θ_2 is infinitely more likely to deviate to m_2 than θ_1 in perturbed game, and thus the equilibrium is symmetrically purifiable. This is even though m_2 is strictly dominated for θ_2 but not equilibrium dominated for θ_1.

4.4 Discussion

We briefly discuss the thin tails assumption and the logic of purification versus refinements based on belief restrictions.

4.4.1 Without thin tails

If we do not make the thin tails assumption, then we will not get extremal beliefs, and thus P-completion will not be the relevant notion. Let us consider the beer-quiche game, and assume symmetric shocks, but without thin tails. Suppose that the prior, $\pi := \pi(s)$, is large. Consider the partial assessment where both types pool on quiche. If the receiver observes Beer, then $\mu(s|B) > \pi$ and so the receiver will not duel. Thus there cannot be a purifiable equilibrium with pooling on Quiche. In other words, we do not need to assume thin tails to refine away pooling on beer.

Now consider pooling on Beer. If the receiver observes Quiche, then $\mu(s|Q) < \pi$, but this may be greater than 0.5. Thus, it may seem that an equilibrium with pooling on Beer may not also work. However, this conclusion is premature. In the sequence of perturbed game, we could have a sequence of equilibria where the receiver's beliefs converge to 0.5, and along the sequence, the behavior corresponds to partial separation. I.e. the weak type is almost indifferent between B and Q in the absence of payoff shocks, and thus plays Q with much higher probability than the strong type. While this is complicated to show directly, one can make a more general argument. First, for any sequence of payoff shocks, including symmetric ones without thin tails, one has a purifiable Nash equilibrium. Second, this Nash equilibrium must be a sequential equilibrium since we have a game in which each player moves at most once along any path of play. Since no assessment with pooling on quiche is purifiable, this ensures the existence of a purifiable assessment with pooling on beer (since these are the only two types of equilibria in the beer-quiche game).

4.4.2 Purification vs. Belief restriction based refinements

Belief restriction based refinements (equilibrium dominance, intuitive criterion, D1) based on an implicit speech by sender:I have deviated, and the only reason I have done so is because I think I can convince you to change your strategy. Ask yourself, which type could I be, given that I would benefit from you changing your strategy? This rationale is set out most clearly in Cho and Kreps (1987). Symmetric purification is resolutely equilibrium based. The receiver asks herself: given that I am playing my equilibrium strategy r, which type (or types) are most likely to have sent this unexpected message? That is, purification asks, is r self-enforcing given the answer to the above question? The question, which types benefit from receiver changing her strategy, plays no role in the argument. Similarly, forward induction is based on the same type of reasoning (see Kohlberg and Mertens (1986) or Kohlberg (1990). Similarly, symmetric purification does not support forward induction.

4.5 Strong purifiability

Let us now reconsider Beer-Quiche and focus on strong purifiability. The Quiche pooling equilibrium is clearly not strongly purifiable, since it is not symmetrically purifiable. So let us focus on the set of equilibria with pooling on beer. Since it is symmetrically purifiable, let us consider shock sequences where the shocks to the strong type are much more variable than those to the weak type. To keep things simple, we will perturb the payoff of each of the types of player 1 at a single node, that following (B, No), and not perturb player 2's payoffs. More specificlly, consider normally distributed where the weak type has a mean zero shock at (B, No) that is normally distributed with standard deviation τ, while the strong type's shock has standard deviation $K\tau$, $K > 3$. Consider the equilibrium where player 2 Duels with probability one at Q. As $\tau \to 0$, the limit belief assigns probability one to the deviant Quiche eater being strong, and thus this equilibrium is not purifiable. However, there is a purifible equilibrium for these shocks where $\mu(s|Q) = 0.5$ and $r(Duel|Q) = 1/2$. Let q denote the probability that player 2 duels at Q. It suffices to show that at any τ, there exists a $q \in (0.5, 1)$ such that the best responses of player 1 induce $\mu(s|Q) = 0.5$. Now when $q = 1$, we have already demonstrated that $\mu(s|Q) > 0.5$ when τ is small. When $q = 0.5$, the weak type is indifferent between B and Q in the absence of shocks, and thus plays Q with probability one-half, independent of τ, while the strong type strictly prefers B in the unperturbed game, and thus plays Q with probability less than one-half, which implies $\mu(s|Q) < 0.5$. By continuity, there must be a value of $q \in (0, 5, 1)$ that $\mu(s|Q) = 0.5$. Indeed, one can explicitly calculate q for this example. Denoting the proir on the strong type by π,

$$\pi\Phi(\frac{-1 - 2q}{K\sigma}) = (1 - \pi)\Phi(\frac{1 - 2q}{K\sigma}),$$

where Φ denotes the cdf of the standard normal distribution. This has a solution q in $(0.5, 1)$ that is independent of τ, and is decreasing in K.

This suggests that the *outcome* of the pooling on beer equilibrium is strongly purifiable. However, no equilibrium strategy is strongly purifiable.

Finally, when the shocks are much more variable for the strong type, pooling on quiche is not purifiable. However, when shocks are more variable for the weak type, then pooling on Quiche is purifiable.

References

Banks, J. S., and J. Sobel (1987): "Equilibrium Selection in Signalling Games," *Econometrica*, 55, 647–661.

Cho, I.-K., and D. Kreps (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(2), 179–221.

Dekel, E., and D. Fudenberg (1990): "Rational behavior with payoff uncertainty," *Journal of Economic Theory*, 52, 243–267.

Govindan, S., P. J. Reny, and A. J. Robson (2003): "A Short Proof of Harsanyi's Purification Theorem," *Games and Economic Behavior*, 45(2), 369–374.

Harsanyi, J. C. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory*, 2(1), 1–23.

Kohlberg, E. (1990): "Refinement of Nash Equilibrium: The Main Ideas," in *Game Theory and Applications*, ed. by T. Ichiishi, A. Neyman, and Y. Tauman, pp. 3–45. Academic Press, San Diego.

Kohlberg, E., and J.-F. Mertens (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54(5), 1003–1037.

Kreps, D., and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50(4), 863–894.

Mailath, G. J., L. Samuelson, and J. M. Swinkels (1993): "Extensive Form Reasoning in Normal Form Games," *Econometrica*, 61(2), 273–302.

Myerson, R. B. (1978): "Refinements of the Nash Equilibrium Concept," *International Journal of Game Theory*, 7, 73–80.

Osborne, M. J., and A. Rubinstein (1994): *A Course in Game Theory*. The MIT Press, Cambridge, MA.

Selten, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 22–55.

van Damme, E. (1984): "A Relation between Perfect Equilibria in Extensive Form Games and Proper Equilibria in Normal Form Games," *International Journal of Game Theory*, 13, 1–13.

——— (1991): *Stability and Perfection of Nash Equilibria.* Springer-Verlag, Berlin, second, revised and enlarged edn.