
Reflective Reason and Equilibrium Refinements

If we allow that human life can be governed by reason, the possibility of life is annihilated.

Leo Tolstoy

If one weight is twice another, it will take half as long to fall over a given distance.

Aristotle, *On the Heavens*

In previous chapters, we have stressed the need for a social epistemology to account for the behavior of rational agents in complex social interactions. However, there are many relatively simple interactions in which we can use some form of reflective reason to infer how individuals will play. Since reflective reason is open to the players as well as ourselves, in such cases we expect Nash equilibria to result from play. However, in many cases there are a plethora of Nash equilibria, only some of which will be played by reasonable agents.

A *Nash equilibrium refinement* of an extensive form game is a criterion that applies to all Nash equilibria that are deemed reasonable, but fails to apply to other Nash equilibria that are deemed unreasonable, based on our informal understanding of how rational individuals might play the game. A voluminous literature has developed in search of an adequate equilibrium refinement criterion. A number of criteria have been proposed, including *subgame perfect*, *perfect*, *perfect Bayesian*, *sequential*, and *proper equilibrium* (Harsanyi 1967, Myerson 1978, Selten 1980, Kreps and Wilson 1982, Kohlberg and Mertens 1986), which introduce player error, model beliefs off the path of play, and investigate the limiting behavior of perturbed systems as deviations from equilibrium play go to zero.¹

I present a new refinement criterion that better captures our intuitions, and elucidates the criteria we use implicitly to judge a Nash equilibrium as

¹Distinct categories of equilibrium refinement for normal form games, not addressed in this paper, are focal point (Schelling 1960; Binmore and Samuelson 2006), and risk-dominance (Harsanyi and Selten 1988) criteria. The perfection and sequential criteria are virtually coextensive (Blume and Zame 1994), and extend the subgame perfection criterion.

reasonable or unreasonable. The criterion does not depend on counterfactual or disequilibrium beliefs, trembles, or limits of nearby games. I call this the *local best response* (LBR) criterion. The LBR criterion appears to render the traditional refinement criteria superfluous.

The traditional refinement criteria are all variants of subgame perfection, and hence suffer from the fact that there is generally no good reason for rational agents to choose subgame perfect strategies in situations where CKR cannot be assumed. The LBR criterion, by contrast, is a variant of *forward induction*, in which agents infer from the fact that a certain node in the game tree has been attained that certain future behaviors can be inferred from the fact that the other players are rational.

We assume a finite extensive form game \mathcal{G} of perfect recall, with players $i = 1, \dots, n$ and a finite pure strategy set S_i for each player i , so $S = S_1 \times \dots \times S_n$ is the set of pure strategy profiles for \mathcal{G} , with payoffs $\pi_i : S \rightarrow \mathbf{R}$. Let S_{-i} be the set of pure strategy profiles of players other than i , and let $\Delta^* S_{-i} = \prod_{j \neq i} \Delta S_j$ be the set of mixed strategies over S_{-i} . Let \mathcal{N} be the set of information sets of \mathcal{G} , and let \mathcal{N}_i be the information sets where player i chooses.

A *behavioral strategy* p at an information set v is a probability distribution over the actions A_v available at v . We say p is *part of* a strategy profile σ if p is the probability distribution over A_v induced by σ . We say a mixed strategy profile σ *reaches* an information set v if a path through the game tree that occurs with strictly positive probability, given σ , passes through a node of v .

For player i and $v \in \mathcal{N}_i$, we call $\phi^v \in \Delta^* S_{-i}$ a *conjecture* of i at v . If ϕ^v is a conjecture at $v \in \mathcal{N}_i$ and $j \neq i$, we write ϕ_μ^v for the marginal distribution of ϕ^v on $\mu \in \mathcal{N}_j$, so ϕ_μ^v is i 's conjecture at v of j 's behavioral strategy at μ .

Let N_v be the set of Nash equilibrium strategy profiles that reach information set v , and let \mathcal{N}^σ be the set of information sets reached when strategy profile σ is played. For $\tau \in N_v$, we write p_μ^τ for the behavioral strategy at μ (i.e., the probability distribution over the choices A_μ at μ) induced by τ . We say a set of conjectures $\{\phi^v | v \in \mathcal{N}\}$ *supports* a Nash equilibrium σ if, for any i and any $v \in \mathcal{N}^\sigma \cap \mathcal{N}_i$, σ_i is a best response to ϕ^v .


We say Nash equilibrium σ is an LBR equilibrium if there is a set of conjectures $\{\phi^v | v \in \mathcal{N}\}$ supporting σ with the following properties: (a) For each i , each $j \neq i$, each $v \in \mathcal{N}_i$, and each $\mu \in \mathcal{N}_j$, if $N_\mu \cap N_v \neq \emptyset$, then $\phi_\mu^v = p_\mu^\tau$ for some $\tau \in N_\mu \cap N_v$; and (b) If player i choosing at v has several

choices that lead to different information sets of the other players (we call such choices *decisive*), i chooses among those with the highest payoff.

We can state the first condition verbally as follows. An LBR equilibrium is a Nash equilibrium σ supported by a set of conjectures of players at each information set reached, given σ , where players are constrained to conjecture only behaviors on the part of other players that are part of a Nash equilibrium.

9.1 Perfect, Perfect Bayesian, and Sequential Equilibria

The traditional refinement criteria can be understood as reactions to a problem appearing in the game to the right. This game has two Nash equilibria, Ll and Rr . The former, however, includes an *incredible threat*, because if Bob is rational, when he is faced with a choice at v_B , he will surely choose r rather than l . If Alice believes Bob is rational, she will not find Ll plausible, and will play R , to which r is Bob's best response.



Selten (1975) treated this as a problem with subgame perfection (§5.2). He noted that if we assume that there is always a small positive probability that a player will make a wrong move and choose a random action at each information set, (called a “tremble”) then all nodes of the game tree must be visited with positive probability, and the non-subgame-perfect equilibria will disappear, while the subgame perfect equilibria will remain. Selten defines an equilibrium as *perfect* if remains a Nash equilibrium of the game as we let the probability ϵ of a tremble go to zero. Clearly, in the game above, Rr is the only perfect equilibrium.

One weakness of this solution is that the Ll equilibrium is unreasonable even if there is zero probability of a tremble. A more pertinent equilibrium refinement, the so-called *perfect Bayesian* equilibrium (Fudenberg and Tirole 1991), directly incorporates beliefs in the refinement. Let \mathcal{N} be the set of information sets of the game. A Nash equilibrium σ determines a behavioral strategy p_v for all $v \in \mathcal{N}$ (i.e., a probability distribution over the actions A_v at v). An *assessment* μ is defined to be a probability distribution over the nodes at each information set v . An assessment must be consistent with the behavior strategy $\{p_v | v \in \mathcal{N}\}$. Consistency means that if σ reaches $v \in \mathcal{N}$, and x is a node in v , then $\mu(x)$ must equal the probability of reaching x , given σ . On an information set not reached, given σ μ can be defined arbitrarily. We say σ is a perfect Bayesian equilibrium if there is a


consistent assessment μ such that, for every player i and every information set v where i chooses, p_v maximizes i 's payoff when played against p_{-v} , using the probability weights given by μ at v .²

This is a rather complicated definition, but the intuition is clear. A Nash equilibrium is perfect Bayesian if there are consistent beliefs rendering each player's choice at every information set payoff maximizing. Note that for the game above, the equilibrium Ll is not perfect Bayesian, because at v_B Bob's payoff from r is greater than his payoff from l .

Perhaps the most influential refinement criterion other than subgame perfect is the *sequential equilibrium* (Kreps and Wilson 1982). This criterion is a hybrid of perfect and perfect Bayesian. Rather than allowing arbitrary assessments off the path of play (i.e., where the Nash equilibrium does not reach with positive probability), the sequential equilibrium approach perturbs the strategy choices using some pattern of errors, and requires that the assessment at an information set off the path of play be the limit of assessment in the perturbed game as the error rate goes to zero. If the pattern of errors is chosen appropriately, Bayes rule plus σ uniquely determine a limit assessment μ at all nodes, which is the limit of consistent assessments as the error rate goes to zero. We say σ is a sequential equilibrium if there is a limit assessment μ such that, for every player i and every information set v where i chooses, p_v maximizes i 's payoff when played against p_{-v} , using the probability weights given by μ at v .

9.2 Incredible Threats

In the game to the right, first suppose $a = 3$. All Nash equilibria have the form (L, σ_B) for arbitrary mixed strategy σ_B for Bob. At v_A , any conjecture for Alice supports all Nash equilibria. Since no Nash equilibrium reaches v_B , there are no constraints on Alice's conjecture. At v_B , Bob's conjecture must put probability 1 on L , and any σ_B for Bob is a best response to this conjecture. Thus, all (L, σ_B) equilibria satisfy the LBR criterion. While only one of these equilibria is subgame perfect, none involves an incredible threat, and hence there is no reason a rational Bob would choose one strategy profile over another. This is why all satisfy the LBR criterion.



²We define p_{-v} as the behavioral strategies at all information sets other than v given by σ .

Now, suppose $a = 1$ in the figure. The Nash equilibria are now (R, r) and (L, σ_B) , where $\sigma_B(r) \leq 1/2$. Alice conjectures r for Bob, because this is the only strategy at v_B that is part of a Nash equilibrium. Because R is the only best response to r , the (L, σ_B) are not LBR equilibria. Bob must conjecture R for Alice, because this is her only choice in a Nash equilibrium that reaches v_B . Bob's best response is r . Thus (R, r) , the subgame perfect equilibrium, is the unique LBR equilibrium.

Note that this argument does not require any out-of-equilibrium belief or error analysis. Subgame perfection is assured by epistemic considerations alone; i.e., a Nash equilibrium in which Bob plays l with positive probability is an incredible threat.

One might argue that subgame perfection can be defended because there is, in fact, always a small probability that Alice will make a mistake and play R in the $a = 3$ case. However, why single out this possibility? There are many possible "imperfections" that are ignored in the passage from a real-world strategic interaction to the game depicted in above figure, and they may work in different directions. Singling out the possibility of an Alice error is thus arbitrary. For instance, suppose l is the default choice for Bob, in the sense that it costs him a small amount ϵ_d to decide to choose r over l , and suppose it costs Bob ϵ_B to observe Alice's behavior. The new decision tree is depicted in Figure 9.1.




Figure 9.1. Adding an Infinitesimal Decision Cost for Bob

In this new situation, Bob may choose not to observe Alice's choice (i), with payoffs as before, and with Bob choosing l by default. But, if Bob chooses to view (v), he pays inspection cost ϵ_B , observes Alice's choice and shifts to the non-default r when she accidentally plays R , at cost ϵ_d . If Alice plays R with probability ϵ_A , it is easy to show that Bob will choose to inspect only if $\epsilon_A \geq \epsilon_B / (1 - \epsilon_d)$.

The LBR criterion is thus the correct refinement criterion for this game. Standard refinements fail by rejecting non-subgame perfect equilibria

whether or not there is any rational reason to do so (e.g., that the equilibrium involves an incredible threat). The LBR gets to the heart of the matter, which is expressed by the argument that when there is an incredible threat, if Bob gets to choose, he will choose the strategy that gives him a higher payoff, and Alice knows this. Thus Alice maximizes by choosing *R*, not *L*. If there is no incredible threat, Bob can choose as he pleases.

In the remainder of this chapter, I compare the LBR criterion with traditional refinement criteria in a variety of typical game contexts. I address where and how the LBR differs from traditional refinements, and which criterion better conforms to our intuition of rational play. I include some cases where both perform equally well, for illustrative purposes. Mainly, however, I treat cases where the traditional criteria perform poorly and the LBR criterion performs well. I am aware of no cases where the traditional criteria perform better than the LBR criterion. Indeed, I know of no cases where the LBR criterion, possibly strengthened by other epistemic criteria, does not perform well, assuming our intuition is that a Nash equilibrium will be played. My choice of examples follows Vega-Redondo (2003). I have tested the LBR criterion for all of Vega-Redondo's examples, and many more, but present only a few of the more informative examples here.

The LBR shares with the traditional refinement criteria the presumption that a Nash equilibrium will be played, and indeed, in every example in this chapter, I would expect rational players to choose an LBR equilibrium (although this expectation is not backed by empirical evidence). In many games, however, such as the Rosenthal's centipede game (Rosenthal 1981), Basu's traveler's dilemma (Basu 1994), and Carlsson and van Damme's global games (Carlsson and van Damme 1993), both our intuition and the behavioral game theoretic evidence violate the presumption that rational agents play Nash equilibria. The LBR criterion does not apply to these games.

Many games have multiple LBR equilibria, only a strict subset of which would be played by rational players. Often epistemic criteria supplementary to the LBR criterion single out this subset. In this chapter, I use the Principle of Insufficient Reason and what I call the Principle of Honest Communication to this end.




Figure 9.2. Only the equilibrium BbV is reasonable, but there is a connected set of Nash equilibria, including the pure strategy AbU , all members of which are perfect Bayesian.

9.3 Unreasonable Perfect Bayesian Equilibria

Figure 9.2 depicts a game in which all Nash equilibria are subgame perfect and perfect Bayesian, but only one is reasonable, and this is the only equilibrium that satisfies the LBR. The game has two sets of equilibria. The first, \mathcal{A} , chooses A with probability 1 and $\sigma_B(b)\sigma_C(V) \leq 1/2$, which includes the pure strategy equilibrium AbU , where σ_A , σ_B and σ_C are mixed strategies of Alice, Bob, and Carole, respectively. The second is the strict Nash equilibrium BbV . Only the latter is a reasonable equilibrium in this case. Indeed, while all equilibria are subgame perfect because there are no proper subgames, and AbU is perfect Bayesian if Carole believes Bob chose a with probability at least $2/3$, it is not sequential, because if Bob actually gets to move, Bob chooses b with probability 1 because Carole chooses V with positive probability in the perturbed game.


The forward induction argument for the unreasonability of the \mathcal{A} equilibria is as follows. Alice can insure a payoff of 1 by playing A . The only way she can secure a higher payoff is by playing B and having Bob play b and Carole play V . Carole knows that if she gets to move, Alice must have chosen B , and because choosing b is the only way Bob can possibly secure a positive payoff, Bob must have chosen b , to which V is the unique best response. Thus, Alice deduces that if she chooses B , she will indeed secure the payoff 2. This leads to the equilibrium BbV .

To apply the LBR criterion, note that the only moves Bob and Carole use in a Nash equilibrium where they get to choose (i.e., that reaches one of

their information sets) are b and V , respectively. Thus, Alice must conjecture this, to which her best response is B . Bob conjectures V , so choose b , and Carole conjectures b , so chooses V . Therefore, only BbV is an LBR equilibrium.

9.4 The LBR Picks Out the Sequential Equilibrium

The figure to the right depicts another example where the LBR criterion rules out unreasonable equilibria that pass the subgame perfection and perfect Bayesian criteria, but sequentiality and LBR are equally successful in this case. In addition to the Nash equilibrium Ba , there is a set \mathcal{A} of equilibria in which Alice plays A with probability 1 and Bob plays b with probability $\geq 2/3$. The set \mathcal{A} are not sequential, but Ba is sequential. The LBR criterion requires that Alice conjecture that Bob plays a if he gets to choose, because this is Bob's only move in a Nash equilibrium that reaches his information set. Alice's only best response to this conjecture is B . Bob must conjecture B , because this is the only choice by Alice that is part of a Nash equilibrium and reaches his information set, and a is a best response to this conjecture. Thus, Ba is an LBR equilibrium, and the others are not.



9.5 Selten's Horse: Sequentiality and LBR Disagree

Selten's Horse is depicted in Figure 9.3. This game shows that sequentiality is neither strictly stronger than, nor strictly weaker than the LBR criterion, since the two criteria pick out distinct equilibria in this case.

There is a connected component \mathcal{M} of Nash equilibria given by

$$\mathcal{M} = \{(A, a, p_\lambda \lambda + (1 - p_\lambda) \rho) \mid 0 \leq p_\lambda \leq 1/3\},$$

where p_λ is the probability that Carole chooses λ , all of which of course have the same payoff (3, 3, 0). There is also a connected component \mathcal{N} of Nash equilibria given by

$$\mathcal{N} = \{(D, p_a a + (1 - p_a) d, \lambda) \mid 1/2 \leq p_a \leq 1\},$$

where p_a is the probability that Bob chooses a , all of which have the same payoff (4, 4, 4).




Figure 9.3. Selten's Horse

The \mathcal{M} equilibria are sequential, but the \mathcal{N} equilibria are not even perfect Bayesian, since if Bob were given a choice, his best response would be d , not a . Thus the standard refinement criteria select the \mathcal{M} equilibria as reasonable.

The only Nash equilibrium in which Carole gets to choose is in the set \mathcal{N} , where she plays λ . Hence, for the LBR criterion, Alice and Bob must conjecture that Carole chooses λ . Also a is the only choice by Bob that is part a Nash equilibrium that reaches his information set. Thus, Alice must conjecture that Bob play a and Carole play λ , so her best response is D . This generates the equilibrium $Da\lambda$. At Bob's information set, he must conjecture that Carole plays λ , so a is a best response. Thus, only the pure strategy $Da\lambda$ in the \mathcal{N} component satisfies the LBR criterion.

Selten's horse is thus a case where the LBR criterion chooses an equilibrium that is reasonable even though it is not even perfect Bayesian, while the standard refinement criteria choose an unreasonable equilibrium in \mathcal{M} . The \mathcal{M} equilibria are unreasonable because if Bob did get to choose, he would conjecture that Carole play λ , because that is her only move in a Nash equilibrium where she gets to move, and hence will violate the LBR condition of choosing an action that is part of a Nash equilibrium that reaches his choice node. However, if he is rational, he will violate the LBR stricture and play d , leading to the payoff $(5,5,0)$. If Alice conjectures that Bob will play this way, she will play a , and the outcome will be the non-Nash equilibrium $Aa\lambda$. Of course, Carole is capable of following this train of thought, and she might conjecture that Alice and Bob will play non-Nash strategies, in which case, she could be better off playing the non-Nash ρ herself. But of course, both Alice and Bob might realize that Carole might reason in this manner. And so on. In short, we have here a case where the sequential equilibria are all unreasonable, but there are non-Nash choices that are as reasonable as the Nash equilibrium singled out by the LBR.

9.6 The Spence Signaling Model




Figure 9.4. The Unreasonable Pooling Equilibrium is Rejected by the LBR

Figure 9.4 represents the famous Spence signaling model (Spence 1973). Alice is either a low quality worker (L) with probability $p = 1/3$ or a high quality worker (H) with probability $p = 2/3$. Only Alice knows her own quality. Bob is an employer who has two types of jobs to offer, one for an unskilled worker (U) and the other for a skilled worker (S). If Bob matches the quality of a hire with the skill of the job, his profit is 10; otherwise, his profit is 5. Alice can invest in education (Y) or not (N). Education does not enhance Alice's skill, but if Alice is low quality, it costs her 10 to be educated, while if she is high quality, it costs her nothing. Education is thus purely a signal, possibly indicating Alice's type. Finally, the skilled job pays 6 more than the unskilled job, the uneducated high quality worker earns 2 more than the uneducated low quality worker in the unskilled job, and the base pay for a low quality, uneducated worker in an unskilled job is 12. This gives the payoffs listed in Figure 9.4.

This model has a separating equilibrium in which Alice gets educated only if she is high quality, and Bob assigns educated workers to skilled jobs and uneducated workers to unskilled jobs. In this equilibrium, Bob's pay-off is 10 and Alice's payoff is 17.33 prior to finding out whether she is of low or high quality. Low quality workers earn 12 and high quality workers earn 20. There is also a pooling equilibrium in which Alice never gets an education and Bob assign all workers to skilled jobs. Indeed, any combination of strategies SS (assign all workers to skilled jobs) and SU (assign uneducated workers to skilled jobs and educated workers to unskilled jobs)

is a best response for Bob in the pooling equilibrium. In this equilibrium, Bob earns 8.33 and Alice earns 19.33. However, uneducated workers earn 18 in this equilibrium and skilled workers earn 20.

Both sets of equilibria are sequential. For the pooling equilibrium, consider a completely mixed strategy profile σ_n in which Bob chooses SS with probability $1 - 1/n$. For large n , Alice's best response is not to be educated, so the approximate probability of being at the top node a_t of Bob's right-hand information set v_r is approximately $1/3$. In the limit, as $n \rightarrow \infty$, the probability distribution over v_t computed by Bayesian updating approaches $(1/3, 2/3)$. Whatever the limiting distribution over the left-hand information set v_l (note that we can always assure that such a limiting distribution exists), we get a consistent assessment in which SS is Bob's best response. Hence the pooling equilibrium is sequential.

For the separating equilibrium, suppose Bob chooses a completely mixed strategy σ_n with the probability of US (allocate uneducated workers to unskilled jobs and educated workers to skilled jobs) equal to $1 - 1/n$. Alice's best response is NY (only high quality Alice gets educated), so Bayesian updating calculates probabilities at v_r as placing almost all weight on the top node, and at Bob's left-hand information set v_l , almost all weight is placed on the bottom node. In this limit, we have a consistent assessment in which Bob believes that only high quality workers get educated, and the separating equilibrium is Bob's best response given this belief.

Both Nash equilibria specify that Bob choose S at v_B^l , so Alice must conjecture this, and that Alice choose N at L , so Bob must conjecture this. It is easy to check that (NN, SS) and (NY, US) thus both satisfy the LBR criterion.

9.7 Irrelevant Node Additions

Kohlberg and Mertens (1986) use Figure 9.5 with $1 < x \leq 2$ to show that an irrelevant change in the game tree can alter the set of sequential equilibria. We use this game to show that the LBR criterion chooses the reasonable equilibrium in both panes, while the sequential criterion only does so if we add an "irrelevant" node, as in the right panel of Figure 9.5. The reasonable equilibrium in this case is ML , which is sequential. However, TR is also sequential in the left panel. To see this, let $\{\sigma_A(T), \sigma_A(B), \sigma_A(M)\} = \{1 - 10\epsilon, \epsilon, 9\epsilon\}$ and $\{\sigma_B(L), \sigma_B(R)\} = \{\epsilon, 1 - \epsilon\}$. These converge to T and R , respectively, and the conditional probability of being at the left node of





Figure 9.5. The only reasonable equilibrium for $1 < x \leq 2$ is ML , which is sequential and satisfies the LBR criterion. However, the unreasonable equilibrium TR is sequential in the left panel, but not the right, when an “irrelevant” node is added.

v_2 is 0.9, so Bob’s mixed strategy is ϵ distance from a best response. In the right panel, however, M strictly dominates B for Alice, so TR is no longer sequential.

To apply the LBR criterion, note that the only Nash equilibrium allowing Bob to choose is ML , which gives Alice payoff 3, as opposed to the payoff 2 from choosing T . Therefore, conjecturing this, Alice maximizes her payoff by allowing Bob to choose; i.e., ML is the only LBR equilibrium.

9.8 Improper Sequential Equilibria

Consider the game to the right. There are two Nash equilibria, Fb and Aa . These are also sequential equilibria, since if Alice intends F , but if the probability that B is chosen by mistake is much greater than the probability that A is chosen, then b is a best response. Conversely, if the probability of choosing A by mistake is much greater than the probability that B is chosen by mistake, then a is a best response. Since B is the more costly mistake for Alice, the proper equilibrium concept assumes it occurs very infrequently compared to the A mistake, Bob will play a when he gets to move, so Alice should chose A . Therefore, Aa is the only proper equilibrium, according to the Myerson (1978) criterion.



To see that Aa is the only LBR equilibrium of the game, note that the only Nash equilibrium that reaches Bob’s information set is Aa . The LBR

criterion therefore stipulates that Alice conjecture that Bob chooses a , wo which her best response is A .

The reader will note how simple and clear this justification is of Aa by comparison with the properness criterion, which requires an order of magnitude assumption concerning the rate at which trembles go to zero.

9.9 Second-order Forward Induction




Figure 9.6. Money Burning

Figure 9.6 depicts the famous Money Burning game analyzed by Ben-Porath and Dekel (1992), illustrating second order forward induction. By “burning money,” (to the amount of 1) Alice can ensure a payoff of 2, so if she does not burn money, she must expect a payoff of greater than 2. This induces her battle of the sexes partner to act favorably to Alice, giving her a payoff of 3.

The set of Nash equilibria can be described as follows. First, there is \mathcal{YB} in which Alice plays Yb (don’t burn money, choose favoring Bob), and Bob chooses any mixture of BB (play B no matter what Alice did) and SB (play S if Alice played X and play B if Alice played Y). This set represents the Pareto-optimal payoffs favoring Bob. The second is \mathcal{YS} , in which Alice plays Ys (don’t burn money, choose favoring Alice), and Bob chooses any mixture of BS (play B against X and S against Y) and SS (play S no matter what) This set represents the Pareto-optimal payoffs favoring Alice. The third is \mathcal{XS} , in which Alice plays Xs and Bob chooses any mixed strategy combination of SB and SS in which the former is played with probability $\geq 1/2$. This is the money-burning Pareto-inferior battle of the

sexes equilibrium favoring Bob. The fourth set is the set \mathcal{M} in which Alice plays Y and then $(1/4)b + (3/4)s$ and Bob chooses any mixed strategy that leads to the behavioral strategy $(3/4)B + (1/4)S$. Second order forward induction selects out $\mathcal{Y}S$.

All equilibria are sequential, and result from to different orders in eliminating weakly dominated strategies. The only Nash equilibrium where Bob chooses at v_{2l} involves choosing S there. Thus, Alice conjectures this, and knows that her best response compatible with v_{2l} is Xs , which gives her payoff 2. There are three sets of Nash equilibria where Alice chooses Y . In one, she chooses Yb and Bob chooses B , giving her payoff 2. In the second, she chooses Ys and Bob chooses S , giving her payoff 3. In the third, Alice and Bob play the battle of the sexes mixed strategy equilibrium, with payoff of $3/2$ for Alice. Each of these is compatible with a conjecture of Alice that Bob plays a Nash strategy. Hence, her highest payoff is with Ys . Because Y is decisive and includes a Nash equilibrium, where Alice plays Ys , with higher payoff for Alice than any Nash equilibrium using X , when Bob moves at v_{2r} , the LBR stipulates that he conjectures this, and hence his best response is S . Thus, Alice must conjecture that Bob plays S when she plays Y , to which Ys is the only best response. Thus, YSs is the only LBR equilibrium.

9.10 Beer and Quiche without the Intuitive Criterion

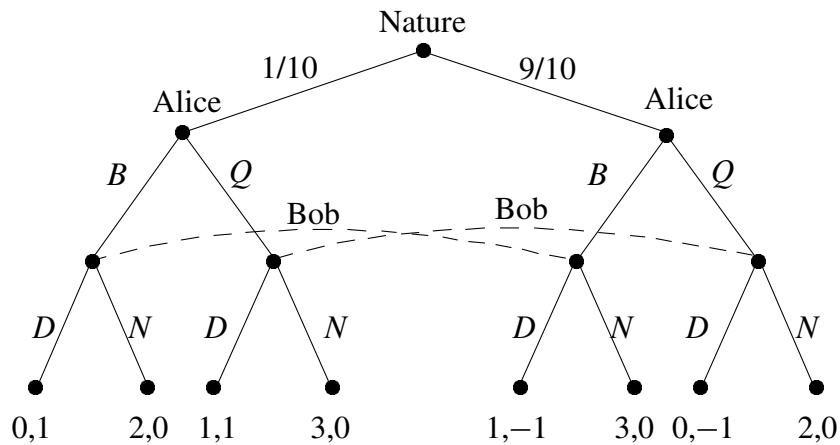


Figure 9.7. The Beer and Quiche Game

This is perhaps the most famous example arguing the importance of “beliefs” in equilibrium refinement, the Cho and Kreps (1987) Real Men Don’t

Eat Quiche game. It also illustrates the well-known *Intuitive Criterion*, which is complementary to sequentiality. However, the LBR criterion singles out the reasonable equilibrium without recourse to an additional criterion.

This game has two sets of (pooling) Nash equilibria. The first is \mathcal{Q} , in which Alice plays QQ (Q if wimp, Q if strong) and Bob uses any mixture of DN (play D against B and play N against Q) and NN (play N no matter what) that places at least weight $1/2$ on DN . The payoff to Alice is $21/10$. The second is \mathcal{B} , in which Alice plays BB (B if wimp, B if strong) and Bob uses any mixture of ND (play N against B and play D against Q) and NN that places at least weight $1/2$ on ND . The payoff to Alice is $29/10$.

The Cho and Kreps (1987) Intuitive Criterion notes that by playing QQ , Alice earns $21/10$, while by playing BB , she earns $29/10$. Therefore, a rational Alice will choose BB . To find the LBR equilibrium, we note that both the \mathcal{B} and \mathcal{Q} equilibria satisfy the first LBR conditions. Moreover, \mathcal{B} is decisive with respect to \mathcal{Q} and has higher payoff for Alice. Thus, the only LBR equilibria are the \mathcal{B} .

9.11 An Unreasonable Perfect Equilibrium




Figure 9.8. An Example where Perfection is Not Sensible, but the LBR criterion is


The game in Figure 9.8, taken from McLennan (1985), has a strict Nash equilibrium at RU and an interval of sequential equilibria \mathcal{L} of the form Lr and D with probability at least $3/4$. The \mathcal{L} equilibria are “unintuitive” for many reasons. Perhaps the simplest is forward induction. If Carole gets to move, either Alice played R or Bob played l . In the former case, Alice

must have expected Carole to play U , so that her payoff would be 3 instead of 2. If Bob moved, he must have moved l , in which case he also must expect Carole to move U , so that his payoff would be 2 instead of 1. Thus, U is the only reasonable move for Carole.

For the LBR criterion, note that \mathcal{L} is ruled out, since the only Nash equilibrium reaching Carole's information set uses U with probability one. The RU equilibrium, however, satisfies the LBR criteria.

9.12 The Principle of Insufficient Reason


The figure to the right, due to Jeffery Ely (private communication) depicts a game the only reasonable Nash equilibria of which form a connected component \mathcal{A} in which Alice plays A and Bob plays a with probability $\sigma_B(a) \in [1/3, 2/3]$, guaranteeing Alice a payoff of 0. There are two additional equilibria Ca and Bb , and there are states at Alice's information set that are consistent with each of these equilibria. At Bob's choice information set, all conjectures σ_A for Alice with $\sigma_A(A) < 1$ are permitted. Thus, all Nash equilibria are LBR equilibria. However, by the Principle of Insufficient Reason and the symmetry of the problem, if Bob gets to move, each of his moves is equally likely. Thus, it is only reasonable for Alice to assume $\sigma_B(a) = \sigma_B(b) = 1/2$, to which A is the only best response.



9.13 The Principle of Honest Communication

It is easy to check that the coordination game to the right has three Nash equilibria, each of which satisfies all traditional refinement criteria, and all are LBR equilibria as well. Clearly, however, only Ll is a reasonable equilibrium for rational players. One justification of this equilibrium is

that if we add a preliminary round of communication, and each player communicates a promise to make a particular move, and if each player believes the Principle of Honest Communication, according to which players keep their promises unless they can benefit by violating these promises and be-



ing believed, then each will promise to play L or l , and they will keep their promises.

9.14 Discussion

The LBR criterion is an improvement of refinement criteria, meant to render other standard refinement criteria superfluous. However, the LBR is not an impediment to standard game-theoretic reasoning. If a game has a unique Nash equilibrium, this will be both a subgame perfect and an LBR equilibrium. If a game has a non-subgame perfect equilibrium with an incredible threat, this equilibrium cannot be LBR because in an LBR equilibrium, players never conjecture incredible threats. In a repeated game, it is reasonable to argue, as does the LBR criterion, that players maximize at every information set that they reach. We do not need the extra baggage of the sequential equilibrium concept to justify this assumption.

The LBR supplies a rigorous and insightful equilibrium refinement criterion. Moreover, it clarifies the meaning of an 'intuitively reasonable' equilibrium as being one in which players conjecture that other players use actions that are part of Nash equilibria, and choose actions themselves that maximize their payoffs subject to such conjectures.