The Local Best Response Criterion: An Epistemic Approach to Equilibrium Refinement

Herbert Gintis*

May 2, 2008

Abstract

The standard refinement criteria for extensive form games, including subgame perfect, perfect, perfect Bayesian, sequential, and proper are too strong, because they reject important classes of unobjectionable Nash equilibria, and too weak, because they accept many objectionable Nash equilibria. This paper develops a new refinement criterion, based on epistemic game theory, that captures the concept of a Nash equilibrium that is plausible when players are rational. I call this the *local best response* (LBR) criterion. This criterion is conceptually simpler than the standard refinement criteria, because it does not depend on out-of-equilibrium, counterfactual, or passage to the limit arguments. The LBR is also informationally richer, because it clarifies the epistemic conditions that render a Nash equilibrium acceptable.

1 Introduction

A *Nash equilibrium refinement* of an extensive form game is a criterion that applies to all Nash equilibria that are deem "plausible," but fails to apply to other Nash equilibria that are deemed "implausible," based on our informal understanding of how "rational" individuals might play the game. A voluminous literature has developed in search of an acceptable equilibrium refinement criterion. A number of useful 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

*Santa Fe Institute and Central European University. I would like to thank Robert Aumann, Jeffrey Ely, Larry Samuelson, and Fernando Vega-Redondo for their help, and the European Science Foundation for financial support.

behavior of perturbed systems as deviations from equilibrium play go to zero.[1]

I present a new refinement criterion that is more accurate is reflecting our intuitions concerning the implications of rationality, and reveals the implicit criteria we use to judge a Nash equilibrium as plausible or implausible. Moreover, 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 asserts that the choices of the players in an extensive form game must be consistent with the assumption that for each information set ν that is reached in some Nash equilibrium, player i choosing at ν forms a conjecture ϕ^ν of other players' actions that consists of best responses in some Nash equilibrium that reaches ν, and chooses an action at ν that is a best response to such a conjecture. If player i choosing at ν 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 for i. The *LBR equilibria* of the game are the Nash equilibria in which every move by every player conforms to this criterion.

Most readers will find this informal description sufficient to begin to apply the LBR criterion to specific cases, and might want to skip directly to section 2. Formally, the LBR criterion is defined using epistemic game theory (Aumann and Brandenburger 1995). However, standard epistemic game theory applies only to normal form games, so I will define a distinct structure for extensive form games of perfect recall. Its relationship to the normal form epistemic game structure is layed out in the Appendix.

We assume a finite extensive form game $\mathcal{G}$ of perfect recall, with players $i = 1, \ldots, n$ and a finite pure strategy set S_i for each player i, so $S = S_1 \times \ldots \times S_n$ 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 ΔS_{-i} be the set of probability distributions over S_{-i}. Let $\mathcal{N}_i$ be the information sets where player i chooses. For player i and $\nu \in \mathcal{N}_i$, we call $\phi^\nu \in \Delta S_{-i}$ a *conjecture* of i at ν. If ϕ^ν is a conjecture at $\nu \in \mathcal{N}_i$ and $j \neq i$, we write $\phi^\nu_{\nu_j}$ for the marginal distribution of ϕ^ν on $\nu_j \in \mathcal{N}_j$, and we call $\phi^\nu_{\nu_j}$ i's conjecture at ν of j's behavioral strategy at ν_j.

We consider $\mathcal{G}$ an extensive form epistemic game with state space Ω, which consists of all states of the form (ν, a_ν, ϕ^ν), where $\nu \in \mathcal{N}_i$ for some i, $a_\nu \in A_\nu$, the set of choices available to i at ν, and $\phi^\nu \in \Delta S_{-i}$ is a conjecture that reaches ν with positive probability for some $s_i \in S_i$. We say an action $a_\nu \in A_\nu$ is a *best response* to ϕ^ν if there is a strategy $s_i \in S_i$ such that (s_i, ϕ^ν) reaches ν, s_i is a best response to ϕ^ν, and $s_i(\nu) = a_\nu$.

[1] Distinct 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.

Let N_ν be the set of Nash equilibria that reach information set ν with positive probability. We say a choice between two actions at information set ν is *decisive* if they actions lead to distinct information sets for players choosing after ν. We say a state (ν, a_ν, ϕ^ν) for $\nu \in \mathcal{N}_i$ is *admissible* if (a) a_ν is a best response to ϕ^ν; (b) for all $j \neq i$, $\mu \in \mathcal{N}_j$, ϕ^ν_μ is the marginal on A_μ of some $\sigma \in N_\mu$; and (c) if a choice between $a_1, a_2 \in A_\nu$ is decisive, and one Nash equilibrium using a_2 has a higher payoff than all Nash equilibria using a_1, then $a_\nu \neq a_1$, and the other players conjecture this. We say a Nash equilibrium $\sigma = (\sigma_1, \ldots, \sigma_n)$ is admissible at information set $\nu \in \mathcal{N}_i$ if either σ does not reach ν, or σ reaches ν with positive probability and there is some admissible state (ν, a_ν, ϕ^ν) such that $\sigma(a_\nu) > 0$. We say a Nash equilibrium is and *LBR equilibrium* if it is admissible at every information set.

2 Subgame Perfection

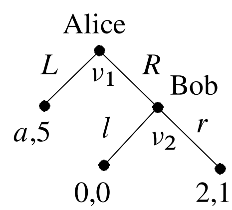

Figure 1: A Simple Game with ($a = 1$) and without ($a = 3$) an Incredible Threat

In Figure 1, first suppose $a = 3$. Using the informal definition, we see that if Bob gets to move, he chooses r, so Alice conjectures this. Her best response is L. All Nash equilibria (L, σ_2), where σ_2 is any behavioral strategy for Bob, are LBR equilibria because they conform to Alice's choice and never get to Bob's choice set.

Formally, Ω has four types of states:

$$\Omega = \{(\nu_1, L, \sigma_2), (\nu_1, R, \sigma_2), (\nu_2, l, \sigma_1), (\nu_2, r, \sigma_1)\},$$

where $\sigma_i \in \Delta S_i$, and $\sigma_1(R) > 0$. The states (ν_1, L, σ_2) are admissible because all σ_2 are part of a Nash equilibrium, and L is a best response to all σ_2. The states (ν_1, R, σ_2) are not admissible because R is not a best response to any σ_2. The states (ν_2, l, σ_1) and (ν_2, r, σ_1) are not admissible because there is no conjecture σ_1 of Bob for Alice that is part of a Nash equilibrium and satisfies $\sigma_1(R) > 0$.

All Nash equilibria (L, σ_2) are admissible at state (ν_1, L, σ_2), and do not reach states at ν_2. Hence all Nash equilibria are LBR equilibria. Note that only one of these LBR equilibria is subgame perfect. Nevertheless, Bob's behavior is rational in all LBR equilibria.

Now, suppose $a = 1$. The Nash equilibria are now (R, r) and (L, σ_2), where $\sigma_2(r) \leq 1/2$. Informally, Alice still conjectures r for Bob, but now R is the only best response, so the (L, σ_2) are not LBR equilibria. Bob must conjecture R for Alice, because this is the only move she makes in a Nash equilibrium that reaches Bob's information set. His best response is r. Thus (R, r) is the unique LBR equilibrium.

Formally, the state space has the same four types of states. Because the only Nash equilibrium to reach ν_2 has Bob play r, the only admissible conjecture for Alice at ν_1 is r, to which her best response is R. Thus, (ν_1, R, r) is the only admissible state at ν_1. Because the only Nash equilibrium that reaches ν_2 has Alice play R, the only admissible state at ν_2 is (ν_2, r, R). Thus, only (R, r) is a 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 Figure 1, 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, but by spending ϵ_B on a decision-making procedure, Bob can choose r when it is to his advantage to do so. The new decision tree is depicted in Figure 2.

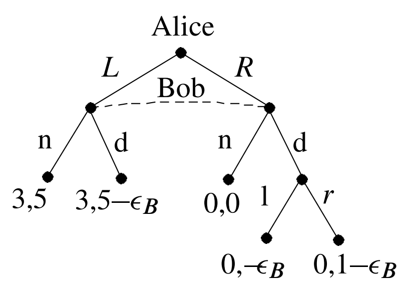

Figure 2: Addition Infinitesimal Decision Costs for Bob

In this new situation, Bob does not know Alice's choice, and if he chooses not to find out (n), the payoffs are as before, with Bob choosing l by default. But, if Bob pays inspection cost ϵ_B, he observes Alice's choice and shifts to the non-default r when she accidentally plays R. If Alice plays R with probability ϵ_A, it is easy to show that Bob will choose to inspect only if $\epsilon_A \geq \epsilon_B$.

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, or simply refuse to choose, permitting the outcome to be a byproduct of other processes in which he is engaged.

In the remainder of this paper, 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), which includes an extremely complete treatment of equilibrium refinement theory. 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 the paper 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 (Gintis 2009) 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 paper, I use the Principle of Insufficient Reason and what I call the Principle of Honest Communication to this end.

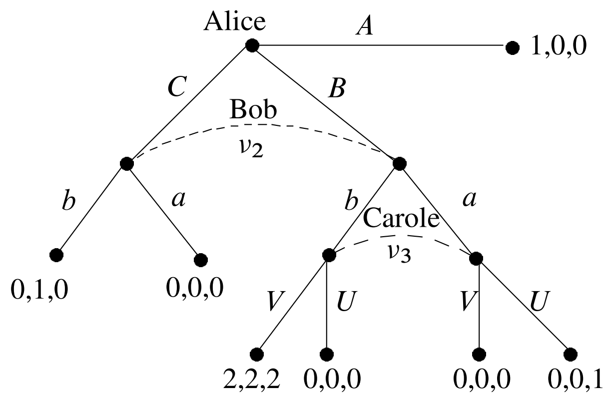

Figure 3: Only the equilibrium BbV is acceptable, but there is a connected set of Nash equilibria, including the pure strategy AbU, all members of which are perfect Bayesian.

3 The LBR Rejects Unacceptable Perfect Bayesian Equilibria

Figure 3 (Vega-Redondo 2003, p. 128) depicts a game in which all Nash equilibria are subgame perfect and perfect Bayesian, but only one is plausible, 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 an acceptable 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 informal argument begins by noting that the only moves Bob and Carole use in a Nash equilibrium where they get to choose are b and V, respectively. Thus, Alice conjectures this, to which her best response is B. Therefore, only BbV is an LBR equilibrium.

Formally, the state space for this problem has eight types of states: Let ν_A, ν_B and ν_C be the information sets for Alice, Bob, and Carole, respectively. The eight types are $(\nu_A, A, \sigma_B, \sigma_C)$, $(\nu_A, B, \sigma_B, \sigma_C)$, and $(\nu_A, C, \sigma_B, \sigma_C)$; $(\nu_B, a, \sigma_A, \sigma_C)$, and $(\nu_B, b, \sigma_A, \sigma_C)$, where $\sigma_A(A) < 1$; and $(\nu_C, U, \sigma_A, \sigma_B)$, $(\nu_C, V, \sigma_A, \sigma_B)$, where $\sigma_A(B) > 0$.

The only Nash equilibrium reaching ν_C has Carole choose V, so both Alice and Bob must conjecture that Carole play V in any admissible state. Similarly, the only Nash equilibrium reaching ν_B has Bob play b, so the only admissible states have Alice and Carole conjecture that Bob play b. The only admissible states at ν_C are thus (ν_C, V, σ_A, b) where $\sigma_A(B) > 0$. The only admissible state at ν_A is then (ν_A, B, b, V), since B is the only best response to (b, V). Thus the only admissible Nash equilibrium is BbV, which thus the only LBR equilibrium.

4 LBR Picks Out the Sequential Equilibrium

Figure 4 depicts another example where the LBR criterion rules out unacceptable 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 specifies that Bob play a if he gets to choose, and thus that Alice conjecture that Bob play a. Hence, Alice knows that she earns 2 by playing B, which is greater than the payoff to playing A. Alice thus chooses B, and Ba is the only LBR equilibrium.

Formally, using the above notation, there are five sets of states: (ν_A, A, σ_B), (ν_A, B, σ_B), and (ν_A, C, σ_B); and (ν_B, a, σ_A), and (ν_B, b, σ_A), where $\sigma_A(A) < 1$. In the Nash equilibria that reach Bob's information set, Bob chooses a. Thus, only (ν_B, a, σ_A) with $\sigma_A(A) < 1$ are admissible at Bob's information set. All states at Alice's information set must conjecture a and be a best response to a. Thus, the only admissible state at Alice's information set is (ν_A, B, a), so the only Nash equilibrium admissible at Alice's information set is Ba, which is thus the only LBR equilibrium.

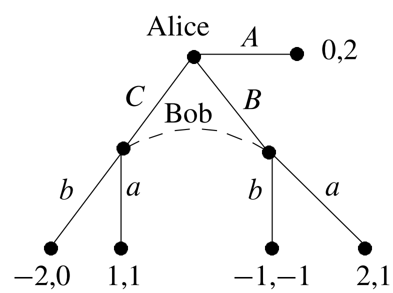

Figure 4: A Generalization of Subgame Perfection

5 Selten's Horse: Sequentiality and LBR Disagree

Selten's Horse is depicted in Figure 5. 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

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

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

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

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

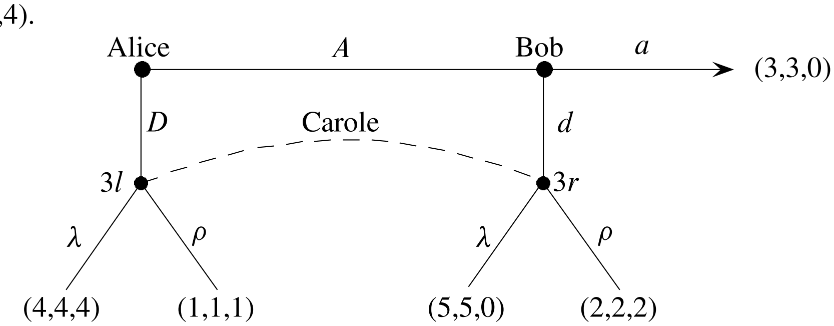

Figure 5: Selten's Horse

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 acceptable. By contrast, only the pure strategy $Da\lambda$ in the $\mathcal{N}$ component satisfies the LBR criterion, because the only Nash equilibrium in which Carole gets to choose is in the set $\mathcal{N}$, where she plays λ with probability 1. LBR also specifies that Bob choose a, because it is the only move that is part of a Nash equilibrium that reaches his choice node. LBR stipulates therefore that Alice conjecture that Bob play a and Carole play λ, so her best response is D. This generates the equilibrium $Da\lambda$.

Selten's horse is thus a case where the LBR criterion chooses an equilibrium that is acceptable even though it is not even perfect Bayesian, while the standard refinement criteria choose an unacceptable equilibrium in $\mathcal{M}$. The $\mathcal{M}$ equilibria are unacceptable 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 implausible, but there are non-Nash choices that are as plausible as the Nash equilibrium singled out by the LBR.

6 The Spence Signaling Model: LBR Correctly Rejects a Sequential Pooling Equilibrium

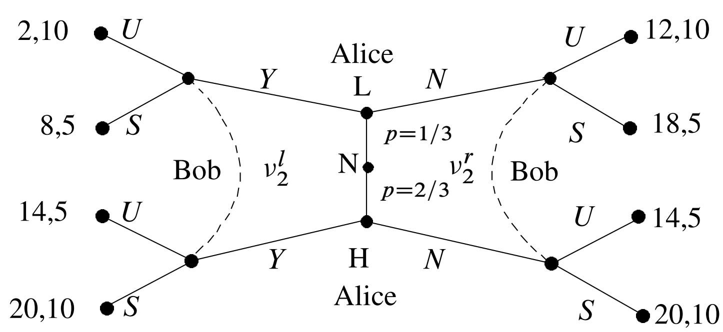

Figure 6: The Unacceptable Pooling Equilibrium is Rejected by the LBR

Figure 6 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 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 6.

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 payoff 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 ν_r is approximately 1/3. In the limit, as $n \to \infty$, the probability distribution over ν_t computed by Bayesian updating approaches (1/3,2/3). Whatever the limiting distribution over the left-hand information set ν_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 ν_r as placing almost all weight on the top node, and at Bob's left-hand information set ν_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.

There are eight types of states: $(H, Y, \sigma_2^l, \sigma_2^r)$, $(H, N, \sigma_2^l, \sigma_2^r)$, $(L, Y, \sigma_2^l, \sigma_2^r)$, and $(L, N, \sigma_2^l, \sigma_2^r)$, where σ_2^l and σ_2^r are Alice's conjectures for Bob at ν_2^l and ν_2^r, respectively; and Bob's states $(\nu_2^l, S, \sigma_1^L, \sigma_1^H)$, $(\nu_2^l, U, \sigma_1^L, \sigma_1^H)$, $(\nu_2^r, S, \sigma_1^L, \sigma_1^H)$, and $(\nu_2^r, U, \sigma_1^L, \sigma_1^H)$, where σ_1^H and σ_1^L are Bob's conjecture for Alice at nodes L and H, respectively. Both Nash equilibria specify that Bob choose S at ν_2^l, so Alice must conjecture this, and that Alice choose N at L, so Bob must conjecture this, so the only admissible states are (H, Y, S, σ_2^r),(H, N, S, S),(L, N, S, σ_2^r), $(\nu_2^l, S, Y, \sigma_1^H)$, $(\nu_2^r, S, N, \sigma_1^H)$ and $(\nu_2^r, U, N, \sigma_1^H)$.

The equilibrium (NN, SS) is admissible at (H, N, S, S) and (L, N, S, σ_2^r), and $(\nu_2^l, S, Y, \sigma_1^H)$ since it never reaches this information set, and at $(\nu_2^r, S, N, \sigma_1^H)$. Hence (NN, SS) is admissible. The Nash equilibrium (NY, US) is admissible at states (H, Y, S, σ_u^r) and (L, N, S, σ_2^r), and $(\nu_2^l, S, Y, \sigma_1^H)$ and (ν_2^r, U, Y, N). Hence (NY, NS) is admissible. Neither player has a decisive choice among Nash equilibria, so both the separating and pooling equilibria are LBR equilibria.

7 Sequentiality is Sensitive, and the LBR Insensitive, to Irrelevant Node Additions

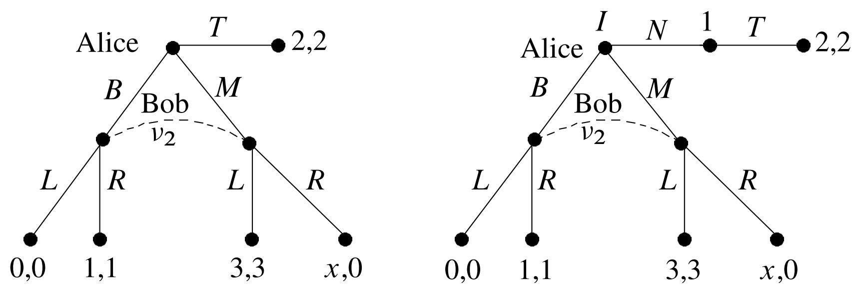

Figure 7: The only acceptable equilibrium for $1 < x \leq 2$ is ML, which is sequential and satisfies the LBR criterion. However, the unacceptable equilibrium TR is sequential in the left panel, but not the right, when an "irrelevant" node is added.

Kohlberg and Mertens (1986) use Figure 7 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 acceptable equilibrium in both panes, while the sequential criterion only does so if we add an "irrelevant" node, as in the right panel of Figure 7. The acceptable 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 ν_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.

8 The LBR Criterion Rejects Improper Sequential Equilibria

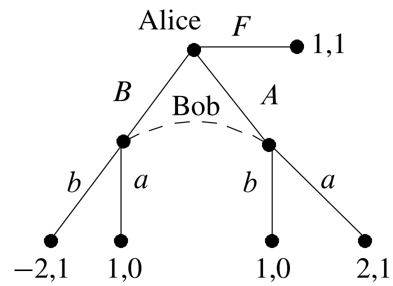

Figure 8: The LBR Criterion Selects Out Proper Equilibria from a Set of Sequential Equilibria

Consider the game in Figure 8. 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 Bob choose a at his information set, and that Alice conjecture this. With this conjecture, A has a higher payoff than F for Alice, who hence plays 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 The LBR Criterion Includes Second-order Forward Induction

Figure 9 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

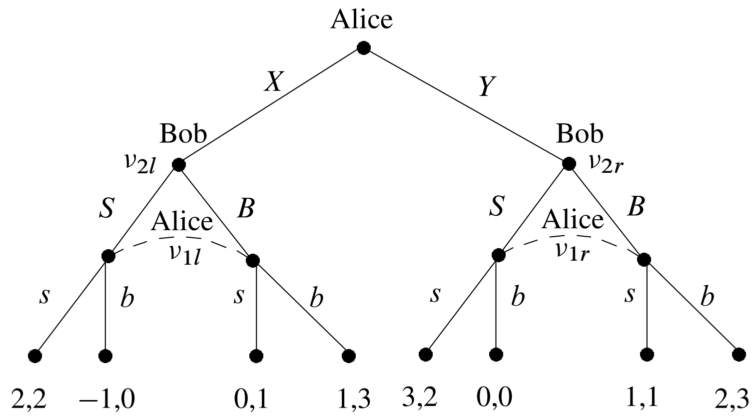

Figure 9: Money Burning

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 $\mathcal{YS}$.

All equilibria are sequential, and result from to different orders in eliminating weakly dominated strategies. The only Nash equilibrium where Bob chooses at ν_{2l} involves choosing S there. Thus, Alice conjectures this, and knows that her best response compatible with ν_{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 ν_{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.

10 Beer and Quiche without the Intuitive Criterion

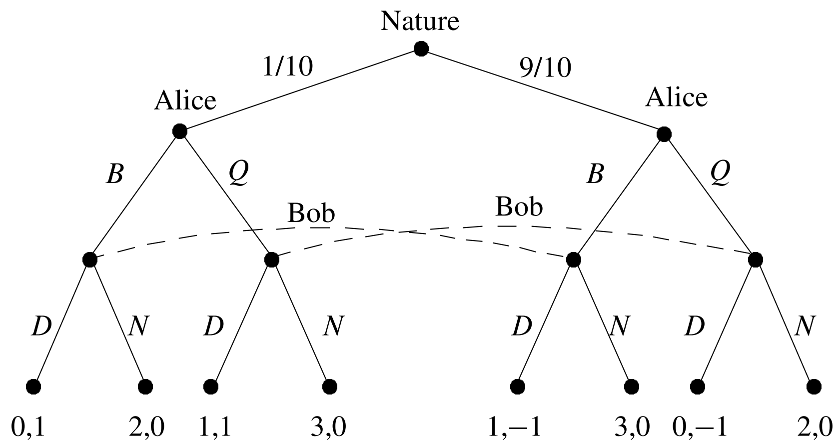

Figure 10: The Beer and Quiche Game

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 acceptable 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 whimp, 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 whimp, 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 the only admissible Nash equilibria are BB and QQ, and BB is decisive with respect to QQ and has higher payoff for Alice. The only admissible states at Bob's choice set, given BB, include Bob's conjectures of B for Alice, to which his best response is N. Thus, the LBR equilibria includ BB for Alice and ND with probability $\geq 1/2$ and NN otherwise, for Bob.

11 The LBR Criterion Selects Acceptable Equilibria in Cases where Sequentiality Fails

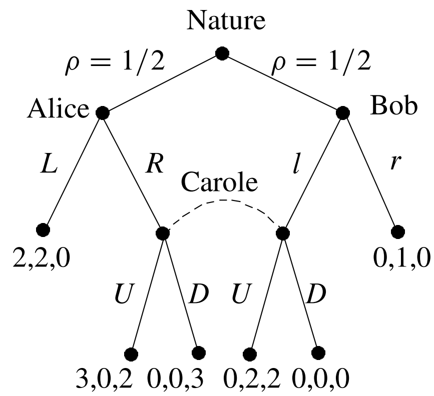

Figure 11: An Example where Perfection is Not Sensible, but the LBR criterion is

The game in Figure 11, taken from McLennan (1985), has a strict Nash equilibrium at RlU 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 plausible 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.

12 The Principle of Insufficient Reason

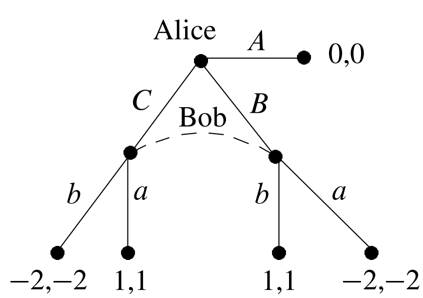

Figure 12: Using the Principle of Insufficient Reason

Figure 12, due to Jeffery Ely (private communication) depicts a game whose only acceptable Nash equilibria 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 admissible 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 admissible, so all states are admissible at this information set. 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.

13 The Principle of Honest Communication

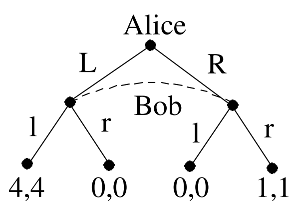

Figure 13: Using the Principle of Honest Communication

It is easy to check that the coordination game in Figure 13 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 plausible 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 being believed, then each will promise to play L or l, and they will keep their promises.

14 Discussion

The LBR criterion is a simplification and an improvement of refinement criteria, meant to render all other standard refinement criteria otiose. 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 plausible 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 acceptable" 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.

In a broader sense, this paper suggests the value of epistemic game theory in clarifying and resolving problems in classical game theory, in the spirit of Aumann and Dreze (2008). The attempt to extend the principles of Bayesian rationality directly to strategic interaction, which occupied game theorists in the 1980's and early 1990's (Kohlberg and Mertens 1986, Tan and da Costa Werlang 1988, Harsanyi and Selten 1988) failed for the simple reason that the epistemic assumptions of the Savage and related axiom systems are subjective, and hence cannot model intersubjectivity. The possible worlds framework corrects this weakness by assuming a common world Ω that is imperfectly known by different agents.

15 Appendix

A normal form *epistemic game* $\mathcal{G}$ consists of a normal form game with players $i = 1, \ldots, n$ and a finite pure strategy set S_i for each player i, so $S = S_1 \times \ldots \times S_n$ is the set of pure strategy profiles for $\mathcal{G}$, with payoffs $\pi_i : S \to \mathbf{R}$. In addition, $\mathcal{G}$ includes a finite *state space* Ω, elements of which are called *states* and subsets of which are called *events*. On the state space, for each player i, knowledge structure $(\mathbf{P}_i, \mathbf{K}_i, \mathcal{P}_i)$ is defined, and a subjective prior $p_i(\cdot;\omega)$ over the strategies S_{-i} for the other players that is a function of the state $\omega \in \Omega$.

The knowledge structure $(\mathbf{P}_i, \mathbf{K}_i, \mathcal{P}_i)$ has the following properties. We define $\mathcal{P}_i$ to be a partition of Ω. For state $\omega \in \Omega$, $\mathbf{P}_i\omega$ is the unique event $P \in \mathcal{P}_i$ such that $\omega \in P$. We interpret $\mathbf{P}_i\omega$ as the set of states that i considers possible when the actual state is ω. Because player i cannot distinguish among states in $\mathbf{P}_i\omega$, $p_i(\cdot;\omega) = p_i(\cdot;\omega')$ for all $\omega' \in \mathbf{P}_i\omega$; i.e., $p_i(\cdot;\omega)$ is measurable with respect to $\mathbf{P}_i$. We can thus write i's subjective prior as $p_i(\cdot;\mathbf{P}_i\omega)$. If E is an event, we define $\mathbf{K}_i E = \{\omega \in \Omega | \mathbf{P}_i\omega \subseteq E\}$, and we say that $\mathbf{K}_i E$ is the event that i *knows* the event E; i.e., $\mathbf{K}_i E$ is the set of states ω at which i knows that $\omega \in E$, which is equivalent to $\mathbf{P}_i\omega \subseteq E$. In particular, $\mathbf{P}_i\omega$ is the event that i considers ω possible.

If $\mathbf{K}_i E = E$, so E is the union of cells in the partition $\mathbf{P}_i$, we say E is *self-evident*.

The state space and knowledge structures are general features of the modal logic of knowledge (Hughes and Cresswell 1996, Aumann 1999), but are more closely specified in epistemic game theory, in which state $\omega \in \Omega$ specifies the strategy profile $s(\omega) = (s_1(\omega), \ldots, s_n(\omega)) \in S$ chosen by the players at ω, as well as their *conjectures* $\phi^\omega = (\phi_1^\omega, \ldots, \phi_n^\omega)$ concerning the choices of the other players, where $\phi_i^\omega \in \Delta S_{-i}$, the space of probability distributions on S_{-i}. Player i's conjecture is derived from i's subjective prior by $\phi_i^\omega(s_{-i}) = p_i(s_{-i}; \mathbf{P}_i\omega)$. A player i is *rational* at ω if $s_i(\omega)$ maximizes $\mathbf{E}[\pi_i(s_i, \phi_i^\omega)]$, which is defined by

$$\mathbf{E}[\pi_i(s_i, \phi_i^\omega)] = \sum_{s_{-i} \in S_{-i}} \phi_i^\omega(s_{-i})\pi_i(s_i, s_{-i}). \tag{1}$$

We extend the epistemic game structure to extensive form games of perfect recall as follows. Every state ω defines a unique path from the root node to a terminal node of the game. Let $\mathcal{N}_i$ be the information sets at which player i chooses. We say information set ν is *reached* in state ω if a node along this path is a member of ν. Then, given information set ν, we define E_ν to be the event that ν is reached. The epistemic assumptions of the extensive form game are that for each player i, and for all $\nu \in \mathcal{N}_i$, i considers all states in E_ν, and no others, possible, so $\mathbf{K}_i E_\nu = E_\nu$; i.e., E_ν is self-evident to the player choosing at ν.

By the perfect recall assumption, the tree property partially orders the information sets for the players, where $\nu \succ \mu$ if every node in μ has an ancestor in ν. This implies that $\nu \succ \mu \Rightarrow E_\mu \subseteq E_\nu$, and the inequality is strict if E_ν contains more than one node. We define i's subjective prior at $\nu \in \mathcal{N}_i$ as the conditional probability $\phi^\nu(s_{-i}) = p_i(s_{-i}; \mathbf{P}_i\omega | \omega \in E_\nu)$, which is well defined if $p_i(s_{-i}; \mathbf{P}_i\omega) \neq 0$ for some $\omega \in E_\nu$, because E_ν is the union of cells $\mathbf{P}_i\omega$, for $\omega \in E_\nu$. If $p_i(s_{-i}; \mathbf{P}_i\omega)$ are identially zero, we define $\phi^\nu(s_{-i})$ arbitrarily. We require that $p_i^\omega(s_{-i}; \mathbf{P}_i\omega) = 0$ if s_{-i} does not reach ν for some $s_i \in S_i$, so conjectures satisfy $\phi^\nu(s_{-i}) = 0$ if s_{-i} does not reach ν for some $s_i \in S_i$.

REFERENCES

Aumann, Robert J., "Interactive Epistemology I: Knowledge," *International Journal of Game Theory* 28 (1999):264–300.

— and Adam Brandenburger, "Epistemic Conditions for Nash Equilibrium," *Econometrica* 65,5 (September 1995):1161–80.

— and Jacques H. Dreze, "Rational Expectations in Games," *American Economic Review* 98,1 (2008):72–86.

Basu, Kaushik, "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory," *American Economic Review* 84,2 (May 1994):391–395.

Ben-Porath, Elchanan and Eddie Dekel, "Signaling Future Actions and the Potential for Self-sacrifice," *Journal of Economic Theory* 57 (1992):36–51.

Binmore, Kenneth G. and Larry Samuelson, "The Evolution of Focal Points," *Games and Economic Behavior* 55,1 (April 2006):21–42.

Blume, Lawrence E. and William R. Zame, "The Algebraic Geometry of Perfect and Sequential Equilibrium," *Econometrica* 62 (1994):783–794.

Carlsson, Hans and Eric van Damme, "Global Games and Equilibrium Selection," *Econometrica* 61,5 (September 1993):989–1018.

Cho, In-Koo and David M. Kreps, "Signalling Games and Stable Equilibria," *Quarterly Journal of Economics* 102,2 (May 1987):180–221.

Gintis, Herbert, *The Bounds of Reason: Game Theory for the Behavioral Sciences* (Princeton, NJ: Princeton University Press, 2009).

Harsanyi, John C., "Games with Incomplete Information Played by Bayesian Players, Parts I, II, and III," *Behavioral Science* 14 (1967):159–182, 320–334, 486–502.

— and Reinhard Selten, *A General Theory of Equilibrium Selection in Games* (Cambridge, MA: MIT Press, 1988).

Hughes, G. E. and M. J. Cresswell, *A New Introduction to Modal Logic* (London: Routledge, 1996).

Kohlberg, Elon and Jean-Franųis Mertens, "On the Strategic Stability of Equilibria," *Econometrica* 54,5 (September 1986):1003–1037.

Kreps, David M. and Robert Wilson, "Sequential Equilibria," *Econometrica* 50,4 (July 1982):863–894.

McLennan, Andrew, "Justifiable Beliefs in Sequential Equilibrium," *Econometrica* 53,4 (July 1985):889–904.

Myerson, Roger B., "Refinements of the Nash Equilibrium Concept," *International Journal of Game Theory* 7,2 (1978):73–80.

Rosenthal, Robert W., "Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox," *Journal of Economic Theory* 25 (1981):92–100.

Schelling, Thomas C., *The Strategy of Conflict* (Cambridge, MA: Harvard University Press, 1960).

Selten, Reinhard, "A Note on Evolutionarily Stable Strategies in Asymmetric Animal Conflicts," *Journal of Theoretical Biology* 84 (1980):93–101.

Spence, A. Michael, "Job Market Signaling," *Quarterly Journal of Economics* 90 (1973):225–243.

Tan, Tommy Chin-Chiu and Sergio Ribeiro da Costa Werlang, "The Bayesian Foundations of Solution Concepts of Games," *Journal of Economic Theory* 45 (1988):370–391.

Vega-Redondo, Fernando, *Economics and the Theory of Games* (Cambridge: Cambridge University Press, 2003).

c\Papers\LBR Criterion\Florence.tex May 2, 2008