Covering Your Posterior:

Teaching Signaling Games Using Classroom Experiments

by Theodore L. Turocy

Assistant Professor, Department of Economics

Texas A&M University

College Station TX 77843

Phone: 979.862.8082; Fax 979.847.8757

Email: turocy@econmail.tamu.edu

Covering Your Posterior:

Teaching Signaling Games Using Classroom Experiments

Abstract

This paper describes a protocol for classroom experiments for courses which introduce undergraduates to signaling games. Signaling games are conceptually difficult because, when analyzing the game, students are not naturally inclined to think in probabilistic, Bayesian terms. The experimental design explicitly presents the posterior frequencies of the unobserved events. The protocol's emphasis on the posterior enhances convergence to the equilibrium prediction, relative to a treatment in which posterior frequencies are not explicitly computed. This convergence reinforces the development of the theory in subsequent lecture periods.

Keywords: Signaling, Bayesian updating, classroom experiments.

1 Introduction

The incorporation of game theory into modern economic analysis was facilitated in part by advancements in tools for tractably representing private information. Successful integration of these applications into the undergraduate curriculum requires that students develop intuition for strategic thinking in games with asymmetric information. Classroom experiments engage students, in the role of active participants inside the game, to build this intuition prior to their undertaking the theoretical analysis as outside observers.

In this paper I describe an approach for introducing signaling games to students who are learning about games of asymmetric information for the first time. In signaling games, the posterior beliefs held by the uninformed player are central to determining whether separating or pooling equilibria exist. My experience has been that many students struggle with the logic of equilibrium in these games because they do not naturally think in terms of the posterior probabilities. Because of this, convergence to equilibrium in classroom experiments designed to demonstrate signaling games is unreliable in the absence of additional guidance for students during the play of the game. Experimental designs which enhance convergence to equilibrium play offer students a concrete example in which the theory is shown to be a useful analytical tool.

To simulate one-shot interactions while giving an opportunity to learn by experience, in this design the same signaling game is played repeatedly with random, anonymous matching each period. In some laboratory games, publicly providing aggregate information about the play of the game enhances the speed of convergence; see Friedman (1996) for one study where the effect of public information is systematically investigated in simultaneous-move games. In the signaling experiments reported in this paper, aggregate public information is in itself not enough to provide the guidance needed for convergence to a Nash equilibrium. Organizing the information to calculate explicitly the posterior frequencies of the state, conditional on the informed players' behavior, results in closer conformance with equilibrium.

The existing literature on research experiments with signaling games has not been concerned with the systematic manipulation of the presentation of public information. Much of this literature focuses on testing the predictions of refinements of Bayes-Nash equilibrium; see, for example, the survey of Camerer (2003). The research literature implicitly asks whether subjects will analyze the game in terms of Bayesian reasoning. In the classroom, the goal is to teach students to think in these terms. Presenting students with the posterior frequencies suggests to them that these frequencies are important data in formulating good strategy. Students who employ these data in making their decisions in the experiment learn that thinking in terms of these posterior frequencies results in better outcomes for them in the game.

Most undergraduate students have at least a passing familiarity with poker games. To ease the transition from games of perfect information, asymmetric information is introduced with a classroom experiment in which students play "stripped-down poker." The game is that described in Reiley et al (forthcoming), except students play both the informed and uninformed roles. Even novice poker players quickly realize that unpredictability and randomization are important elements of good strategy. Students in the uninformed role in this game instinctively ask themselves a question in probabilistic terms: What are the chances he's bluffing, given his betting behavior?

The paper is organized as follows. Section 2 describes the general procedure I use for these classroom games. Section 3 describes the one-card poker game used to introduce asymmetric information, and shows how the informational treatment helps to regulate bluffing to near equilibrium levels. Section 4 shows how presenting the posterior probabilities enhances convergence in a signaling game with pure-strategy pooling equilibria. Section 5 summarizes some concluding thoughts and observations.

2 General Protocol

The sessions reported here were implemented using the signaling game in Veconlab (http://veconlab.econ.virginia.edu/admin.htm). A total of four class sections participated in these classroom experiments in an introductory course on game theory and strategy. In each session there were eight informed and eight uninformed players. Each player remained in the same role, informed or uninformed, throughout the session. Each period, the players were randomly and anonymously rematched, and new realizations for the private information were drawn. At the end of each period, after all players had made their choices, information about the aggregate play in the period was posted on a projector screen at the front of the room.

Two classroom experiments were used to motivate the discussion of asymmetric information and signaling games. The first was a one-card poker game in which the unique equilibrium involves randomization; this game will be described in detail in Section 3. The second was a standard signaling game with pure-strategy pooling equilibria, which is covered in Section 4.

The treatment variable in the design is the organization of the public information on the projector screen. The two panels in Figure 1 present the same data series under each of the two treatments. Here, the states observed by the informed player are labeled A and B, and the signals which could be chosen by the informed player are $P1$ and $P2$. In the top panel, treatment COUNT, the public information simply reports the number of informed players who chose each signal in each state. In the bottom panel, treatment POST, the public information is organized according to the number of informed players in each state who chose each signal. Treatment POST augments the presentation by calculating the posterior frequencies of each state conditional on each signal. Since these posterior frequencies may vary significantly from period to period depending on the realizations of the underlying state, treatment POST also presents the posteriors aggregated over five-period intervals.

In this course, students play for points toward their final grade based upon their performance in the classroom experiments. In addition, at the end of the semester, two students from each section are chosen at random to have an opportunity to win a cash prize (on the order of US$50), where their total earnings for all sessions over the course translate into a probability of winning the prize. This is similar to the method suggested by Smith (1961) for inducing risk neutrality. Here it is not done to control for risk attitudes, but rather to allow a significant cash prize to be offered. Therefore, students have some incentives to take the payoffs in the game seriously.

3 Playing "Stripped-Down Poker"

Games of asymmetric information are introduced in the course with a classroom experiment in which students participate in the "stripped-down poker" game described by Reiley et al (forthcoming). At the beginning of the game, both players place an ante of $1.00 in the pot. The informed player receives a card drawn randomly from a deck consisting of an equal number of aces and kings. An ace is a "high" card, and a king a "low" card. After seeing the card, the informed player can raise, placing another $1.00 in the pot, or fold, ending the game and conceding the pot to the uninformed player. If the informed player raises, then the uninformed player may meet the raise, also placing an additional $1.00 in the pot, or pass, conceding the pot to the informed player. If the informed player raises and the uninformed player meets that raise, the outcome depends on the card that was drawn. The pot goes to the informed player if the card is an ace, and to the uninformed player if it is a king. The extensive game representation appears in Figure 2.

The uninformed player's choice of whether to meet a raise is governed by the probability he places on the event that the informed player holds an ace. If this probability is greater than $\frac{3}{4}$, then the uninformed player strictly prefers to pass; if it is less than $\frac{3}{4}$, the informed player strictly prefers to meet. In the unique equilibrium of this game, the informed player should behave in such a way that the uninformed player is exactly indifferent between his actions. This occurs when the uninformed player assesses exactly a $\frac{3}{4}$ chance that the card is an ace. To accomplish this, the informed player should always raise after drawing an ace, since that is his strictly dominant action, and should raise with probability $\frac{1}{3}$ after drawing a king. Since the informed player is randomizing after drawing a king, it must be that he is indifferent, in that situation, between raising and folding; to accomplish this, the uninformed player should meet a raise with probability $\frac{2}{3}$.

Ideally, a classroom experiment in which this game is played would teach students why this is the equilibrium. If bluffing occurs too frequently and the uninformed players recognize this, some of them will increase the frequency with which they meet. As the number of bluffs met goes up, the informed players more often suffer big losses, and are disciplined to cut back on their bluffing. However, overbluffing is often a persistent feature when students play this game; see, for example, Holt (2007, section 33.5). This suggests that some part of this cycle does not occur naturally.

The COUNT treatment presents all the information necessary to identify whether overbluffing is occuring in the population of informed players. However, it does not suggest how to combine the information to make the correct posterior inference. Treatment POST explicitly calculates this posterior. Calculating and posting this frequency is intended to signal to the students that this calculation is the right way to process the information. If students successfully incorporate this information in their decision-making, then $P(\text{ace} \mid \text{raise})$ should be closer to the equilibrium prediction of $\frac{3}{4}$ in this treatment.

Figure 3 presents the posterior frequency $P(\text{ace} \mid \text{raise})$ over time in the four sessions.[1] These are aggregated over five-period intervals to match the information shown to the students in treatment POST. The time series plotted with solid lines are from the sessions using treatment COUNT, and the dashed lines represent treatment POST. For reference, the horizontal dotted line indicates the equilibrium value of $\frac{3}{4}$. Both sessions in treatment POST were always close to the equilibrium prediction. In contrast, due to persistent overbluffing, the posterior frequencies in treatment COUNT were generally, and often substantially, below the equilibrium value. This is consistent with the hypothesis that students recognize the usefulness of the posterior when it is presented, but do not construct it on their own.

One function of the calculation of the empirical posterior frequency is that it makes it easy for the uninformed players to best respond to recent play; therefore, uninformed players are more likely to identify and discipline overbluffing. Figure 4 shows the proportion of raises met in both treatments, again over five-period intervals. In treatment COUNT, the proportion is lower than the equilibrium proportion, even though the best response to the high levels of bluffing observed in those treatments would be to meet all raises. The proportion of raises met in treatment POST is higher, slightly exceeding the equilibrium proportion. Overbluffing in treatment COUNT persists because it goes largely

1. The time series are of different lengths due to the constraints of completing the session within a lecture period.

unpunished. The time path of the meet probabilities under COUNT may reflect the unwillingness of some uninformed players to meet raises after being "burned" by meeting a raise and losing $2 when the card turns out to have been an ace. The calculation of the posterior in POST makes the cost-benefit analysis of meeting raises more transparent.[2]

Thus, in treatment POST, students learn to bluff in approximately the right proportions, solving the "problem" of overbluffing which often occurs in this game. Reiley et al (forthcoming) take a different approach to demonstrating the correct bluffing frequency. They implement this game with the instructor in the role of the informed player against a student volunteer. The instructor plays the optimal mixed strategy using a system for randomizing, such as looking at the second hand on his watch, which is not apparent to the student. When the instructor is "programmed" to play the minimax strategy, the classroom experiment teaches only part of the logic behind equilibrium. The cycle of mutual best responding is short-circuited; in fact, the instructor's commitment to the minimax strategy ensures that the student volunteer's expected earnings do not depend at all on how often he meets a raise.

The posterior frequency is more robust to idiosyncratic play because of the use of random rematching each period. Randomization thus needs to occur only at the population level rather than the individual. The posterior beliefs of an uninformed player can be generated either by playing the same randomizing opponent over and over, or by playing randomly-matched opponents one time each. The same posterior probability of $\frac{3}{4}$ is important in either case. I point out to students that this is the same whether considering playing the game over and over with a buddy, or playing one-off games with several different opponents encountered randomly on a poker website. If the game is played in fixed

2. Statistics on uninformed player behavior are not reported to the students in either treatment. This choice was motivated by the space constraints of constructing the spreadsheet by hand on the fly on the projector. A more refined interface screen without these layout constraints should also be able to report on uninformed player behavior as well, though the results reported here suggest this is of less importance than the calculation of the posterior frequencies.

pairs, a student may not learn the correct posterior probability if he is matched with another student who does not react to the strategic incentives in the game. For example, in the session reported in Holt (2007), which used fixed pairs, one informed player chose to raise in all 20 periods, even though he had the low card more than half the time.

4 A signaling game with pooling equilibria

After the lecture unit analyzing the poker game is complete, the study of topics in asymmetric information continues with a classroom experiment featuring a standard signaling game with pooling equilibria. The extensive form of the game is shown in Figure 5. For ease of exposition here, I present the game using the "beer-quiche" story of Cho and Kreps (1987).[3] The informed player may be "strong," with probability $\frac{2}{3}$ or "weak," with probability $\frac{1}{3}$. Knowing his own type, he goes to a tavern, where he may choose to order beer or quiche. Another patron in the tavern observes what is ordered, but does not know the type of the informed player. This second player then chooses whether to flee or to pick a fight. Note that the payoff structure for the uninformed player in this game is simple; his best reply is to flee if the probability the informed player is strong is greater than one-half, and to fight if the probability is less than one-half.

This game has two pure-strategy equilibria, both of which are pooling. In one equilibrium, the informed player always chooses beer; the uninformed player flees when the informed player chooses beer, but would fight if the informed player were to choose quiche. This equilibrum satisfies the Intuitive Criterion of Cho and Kreps. In this equilibrium, the

3. In classroom experiments, I prefer to use Veconlab's abstract labeling, with types A and B, informed player actions $P1$ and $P2$, and uninformed player actions $R1$ and $R2$. The beer-quiche story is attractive because its concreteness helps communicate the strategic structure of the game concisely. In my experience, though, some students get some extra enjoyment out of "picking a fight," even when the in-game incentives indicate it is not an optimal response.

strong type obtains his most preferred outcome; however, the weak type does not. Therefore, were the uninformed player to observe an order of quiche, it would be reasonable, according to the Intuitive Criterion, for him to believe that the quiche orderer must be weak. In the second pure-strategy equilibrium, the informed player always chooses quiche. The uninformed player flees when the informed player chooses quiche, but would fight if the informed player were to choose beer. This equilibrium requires that the uninformed player believes that, were the informed player to order beer, it must be more likely that the informed player is weak; otherwise, the uninformed player would not want to fight him. However, in this equilibrium, deviations by the weak type could only lower his payoff, making those beliefs unreasonable by the Intuitive Criterion.

Figure 6 plots the posterior frequency $P(\text{strong}|\text{beer})$ for each of the four sessions. As in Figure 3, these are reported over five-period intervals. In both sessions using treatment POST, play converged to the intuitive pooling equilibrium. In the last four periods of one session, all informed players chose beer; in the other, there were only two instances of a weak player choosing quiche in that same span. Therefore, the posterior frequencies are close to the prior probability of $\frac{2}{3}$. Behavior in the sessions using COUNT is best described as separating, even though that is not an equilibrium phenomenon. The posterior frequency $P(\text{strong}|\text{beer})$ is close to 1, reflecting that the weak types do not recognize that switching to beer would be profitable.

Based on classroom discussion, the initial condition of separating behavior in all sessions is explained by the students' interpretation of the presentation of the game by the Veconlab software. Figure 7 shows the table with which Veconlab presents the game.[4] In a game theory course, students are trained to identify dominant and dominated strategies in normal form games. For a player choosing the row, this involves comparing pairs of payoffs vertically. Applying this technique, many students initially think that the informed player has a dominant strategy to choose beer when strong and quiche when weak. Com-

4. Veconlab refers to the informed player as the "proposer" and the uninformed player as the "responder."

paring payoffs vertically, \$2.80 is bigger than \$2.00, and \$1.20 is bigger than \$0.40. However, this is not a simultaneous-move game; the uninformed player may choose to act differently after beer versus quiche. This observation illustrates the importance of emphasizing the sequential nature of the game. By using a tabular format similar to the ones used for simultaneous-move games, students are tacitly encouraged to use the wrong tools to analyze the game.

In this game, the uninformed players almost always choose to flee when beer is chosen, and fight when quiche is chosen. The dynamics are driven primarily by the behavior of the informed players. Given the initial separating behavior, the calculation of the posteriors in treatment POST helps students identify that weak players could be better off if they chose beer. If a small number of weak players experiment by choosing beer, the posterior frequency $P(\text{strong}|\text{beer})$ will still be close to 1; therefore, the weak players choosing beer will "get away with it." Seeing this posterior frequency, other informed players who originally did not consider choosing beer when they were weak now will see it as a viable option. Therefore, students who do not initially understand the benefits of choosing beer when weak can learn from other students who do make the realization. This behavior is self-reinforcing, because it is an equilibrium for both strong and weak types to choose beer. Simply having the public information that some weak player chose beer is not enough to initiate this dynamic. The computation of the posterior shows that most players who choose beer are strong, and therefore the uninformed players still prefer to flee.

Because the experiment does not begin in the pooling equilibrium, the experimental data can be used to motivate refinement concepts relating to beliefs off the equilibrium path. In the case of the sessions using treatment POST, the informed players always chose beer when strong. Therefore, any time quiche was chosen, it was chosen by a weak player. Even after play has converged to pooling, from time to time a weak player might experiment by choosing quiche. This is because he thinks he might have a chance to earn \$2.80, if the uninformed player were to respond by fleeing, as opposed to the $2.00 he expects to earn by choosing beer. However, the posterior data indicate to the uninformed player that, historically, only players who are weak chose quiche; therefore, it is unlikely that the informed player will get away with choosing quiche. The observation that some weak players experiment in this way, while strong players never do, can motivate discussion of off-equilibrium-path beliefs and refinements, since the observation parallels the reasoning behind the Intuitive Criterion refinement.

5 Conclusion

When introducing signaling games to students, classroom experiments are useful because they allow the students to participate actively in a specific realization of an environment with asymmetric information. In a well-designed classroom experiment, students without any prior knowledge of the game should be able to reach equilibrium on their own, through a combination of introspection and observation. In signaling games, focusing the students on the posterior probabilities creates a laboratory environment which encourages the process of discovering the equilibrium. When this occurs, the experiment lays the foundation for the subsequent theoretical analysis.

In a signaling game, information is communicated to the uninformed player if the informed player acts differently depending on the realization of the private information. The inclination of many students is to decompose the game by states, discarding the prior probability and ignoring the uninformed player's informational constraint. Placing the posterior frequency calculation in a central role shows students how to integrate the information learned from the informed players' behavior. The probabilities manipulated in the theoretical discussion can then be related directly to the frequencies observed in the laboratory.

The design of these classroom experiments is intended to convey the significance of the posterior probability calculation. A treatment intermediate between COUNT and POST would present the aggregrate behavior counts in the order used by POST, which makes calculating the posterior convenient, while omitting the explicit calculation. A question for future study is whether students, after having participated in one signaling experiment with the posterior computed for them, will continue to compute the posterior on their own in a subsequent session in which the posterior frequencies are not presented.

The Intuitive Criterion is generally the first refinement concept introduced to students in the study of signaling games. Research has shown that equilibria which do not satisfy the Intuitive Criterion are nevertheless sometimes selected. Brandts and Holt (1993) indicate that it is important that the experience gained in out-of-equilibrium play must be consistent with the out-of-equilibrium beliefs used to support selection of an equilibrium by the Intuitive Criterion. This is indeed what occurs in the design presented in this paper. In more advanced courses, the beer-quiche game with calculated posterior frequencies can be used to introduce Bayesian reasoning and to motivate the formulation of beliefs off the equilibrium path. A critical evaluation of the Intuitive Criterion, and refinements more generally, can then be undertaken by using the same protocol while playing games similar to those studied by Brandts and Holt.

The ultimate goal of augmenting lecture with classroom experiments is improving student achievement. Since adopting treatment POST to introduce signaling, scores have increased on the end-of-unit quiz, in which students are asked to find an equilibrium in a signaling game they have not seen before. The signaling quiz, which had been one of the lowest-scoring quizzes, has now become one of the highest in the course. The quiz includes a short answer portion in which students are asked to justify their answer. Students who participated in sessions using treatment POST more often articulate their explanation of the uninformed player's decision in probabilistic terms.

Acknowledgments

This paper was improved by helpful comments from Brit Grosskopf, Rajiv Sarin, and seminar participants at Texas A&M University, as well as the suggestions of three anonymous referees. Elizabeth Watson was instrumental in polishing and clarifying the presentation. Thanks to Charles Holt for providing the poker data presented in Holt (2007) for comparison, and for creating and maintaining the Veconlab site; and to John Van Huyck for making the Economic Research Laboratory at Texas A&M available for classroom games. The data for the sessions reported is available online at http://econweb.tamu.edu/turocy/papers/sigteach.html. All errors and shortcomings belong to the author alone.

Bibliography

Brandts, J. and C. A. Holt (1993). Adjustment patterns and equilibrium selection in experimental signaling games. *International Journal of Game Theory 22*, 279-302.

Camerer, C. F. (2003). *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton NJ: Princeton University Press.

Cho, I. K. and D. M. Kreps (1987). Signaling games and stable equilibria. *Quarterly Journal of Economics 102*, 179–221.

Friedman, D. (1996). Equilibrium in evolutionary games: Some experimental results. *The Economic Journal 106*, 1–25.

Holt, C. A. (2007). *Markets, Games, and Strategic Behavior*. Boston MA: Pearson Addison Wesley.

Reiley, D., M. B. Urbancic, and M. Walker. Stripped-down poker: A classroom game to illustrate equilibrium bluffing. *Journal of Economic Education*. Forthcoming.

Smith, C. A. B. (1961). Consistency in statistical inference and decision. *Journal of the Royal Statistical Society, Series B 23*, 1–25.

	A	B	C	D	E
1	**Pd.**	**State A**		**State B**	
2		**P1**	**P2**	**P1**	**P2**
3	**1**	6	0	1	1
4	**2**	7	0	0	1
5	**3**	5	0	1	2
6	**4**	5	0	0	3
7	**5**	7	0	0	1
8	**6**	5	0	2	1
9	**7**	5	0	3	0
10	**8**	2	0	6	0
11	**9**	5	0	3	0
12	**10**	5	0	3	0

	A	B	C	D	E	F	G	H	I	J	K	L	M
1		**This Period**			**Last 5 Periods**			**This Period**			**Last 5 Periods**		
2	**Pd.**	**P1 w/**		**P(A | P1)**	**P1 w/**		**P(A | P1)**	**P2 w/**		**P(A | P2)**	**P2 w/**		**P(A | P2)**
3		**A**	**B**		**A**	**B**		**A**	**B**		**A**	**B**	
4	**1**	6	1	85.7%	6	1	85.7%	0	1	0.0%	0	1	0.0%
5	**2**	7	0	100.0%	13	1	92.9%	0	1	0.0%	0	2	0.0%
6	**3**	5	1	83.3%	18	2	90.0%	0	2	0.0%	0	4	0.0%
7	**4**	5	0	100.0%	23	2	92.0%	0	3	0.0%	0	7	0.0%
8	**5**	7	0	100.0%	30	2	93.8%	0	1	0.0%	0	8	0.0%
9	**6**	5	2	71.4%	29	3	90.6%	0	1	0.0%	0	8	0.0%
10	**7**	5	3	62.5%	27	6	81.8%	0	0		0	7	0.0%
11	**8**	2	6	25.0%	24	11	68.6%	0	0		0	5	0.0%
12	**9**	5	3	62.5%	24	14	63.2%	0	0		0	2	0.0%
13	**10**	5	3	62.5%	22	17	56.4%	0	0		0	1	0.0%

Figure 1. Screenshots of the spreadsheets used to present the public information. Top panel: treatment COUNT. Bottom panel: treatment POST.

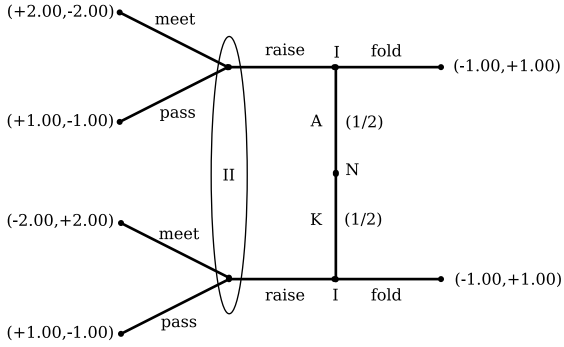

Figure 2. Extensive game for "stripped-down poker."

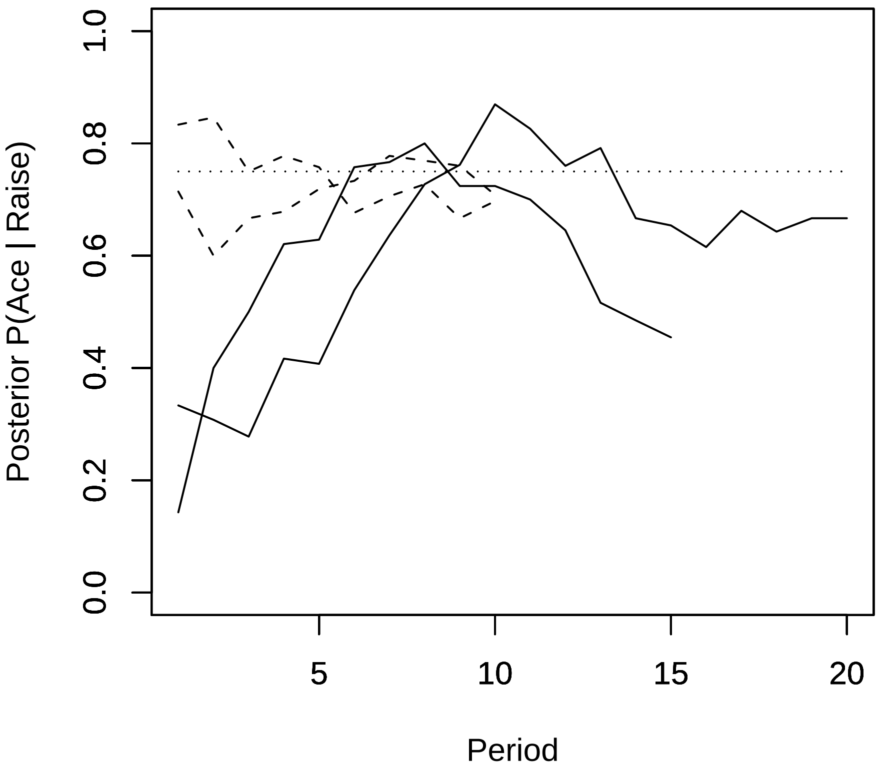

Figure 3. The posterior frequency P(Ace | Raise) in stripped-down poker. The solid lines represent sessions using treatment COUNT, the dashed lines sessions using treatment POST, and the dotted line the beliefs in equilibrium.

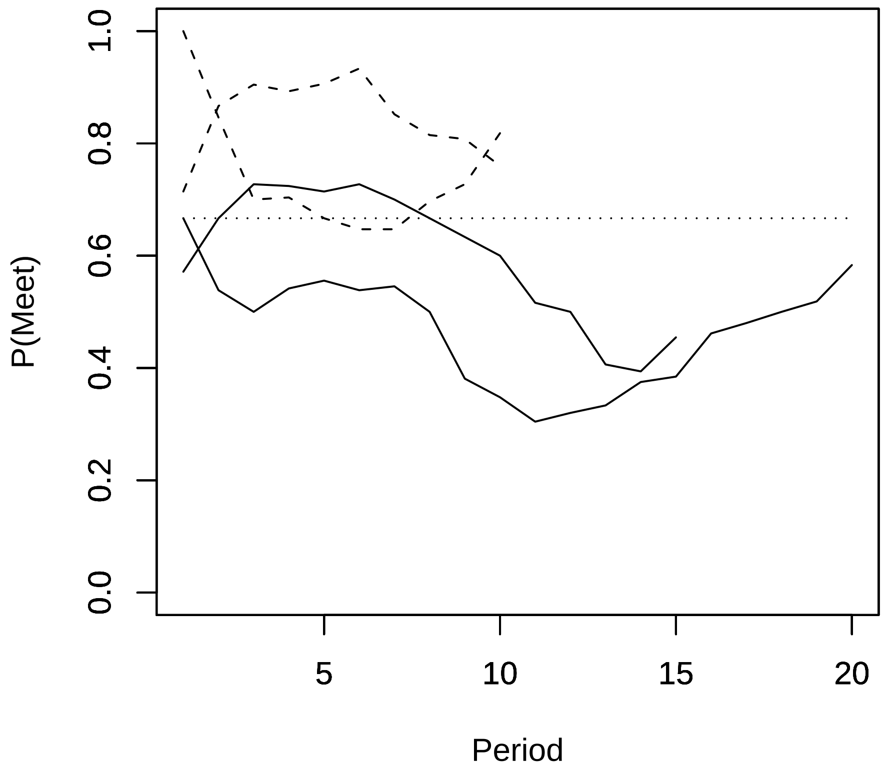

Figure 4. Proportion of raises met in stripped-down poker. The solid lines represent sessions using treatment COUNT, the dashed lines sessions using treatment POST, and the dotted line the equilibrium proportion.

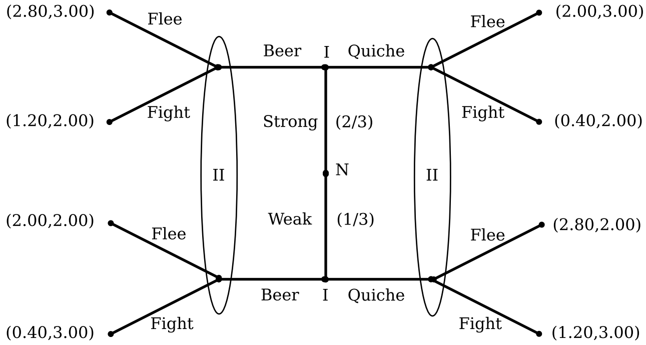

Figure 5. A signaling game with only pooling equilibria.

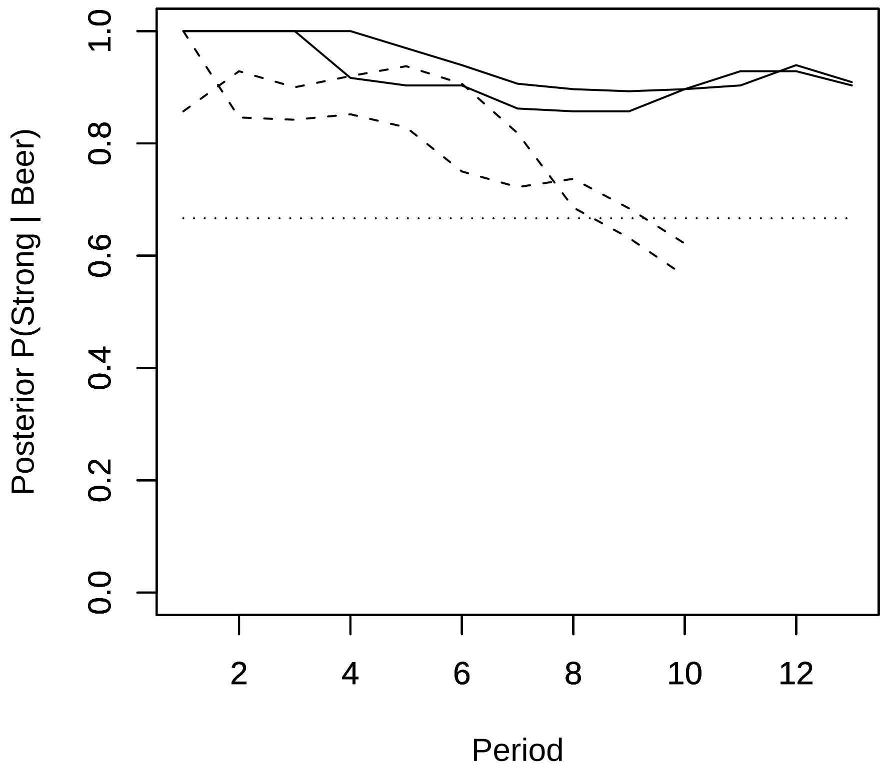

Figure 6. The posterior frequency $P(\text{Strong}|\text{Beer})$ for the game in Figure 5. The solid lines represent sessions using treatment COUNT, the dashed lines sessions using treatment POST. The dotted line indicates the beliefs in the intuitive pooling equilibrium.

Payoffs: Proposer, Responder

	Flee	Fight
Beer (Strong)	$2.80, $3.00	$1.20, $2.00
Quiche (Strong)	$2.00, $3.00	$0.40, $2.00
Beer (Weak)	$2.00, $2.00	$0.40, $3.00
Quiche (Weak)	$2.80, $2.00	$1.20, $3.00

Figure 7. Tabular presentation of signaling game in Veconlab, with "beer-quiche" terminology.