Covering Your Posterior:

Teaching Signaling Games Using Classroom Experiments

by Theodore L. Turocy
Department of Economics
Texas A&M University
College Station TX 77843
Email: `turocy@econmail.tamu.edu`

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. A classroom experiment is a more valuable asset to an instructor if it reliably produces results consistent with theoretical predictions. An experiment that presents a concrete counterexample in which the theory does not appear to work undermines, in the eyes of a student, the credibility of the theory as 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 [3] 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 [1]. 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 [5], 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 probability 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). In these sessions, there were eight informed and eight uninformed players. Each player remained in the same role throughout the session. Each period, the players were randomly and anonymously rematched.

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.

A total of four class sections participated in these classroom experiments in an introductory course on game theory and strategy. At the end of each period, after all players had made their choices, information about the overall play in the period was posted on a projector at the front of the room.

I employed two treatments of the presentation of this public information. 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 [6] 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.

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 [5]. 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. Aces are the "high" card and kings the "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.

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 [4] (section 33.5). This suggests that some part 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 sub-

---

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

stantially, 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.

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 [5] 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 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 [4], 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, using abstract labels for the states and actions, is shown in Figure 4. For ease of exposition, this game can be presented as the "beer-quiche" game (Cho and Kreps [2]). The informed player may be "strong" (type A) or "weak" (type B). After his type is realized, he may choose to order beer ($P1$) or quiche ($P2$). After observing the order, the uninformed player may choose to flee ($R1$) or fight ($R2$).[2] This game has two pure-strategy equilibria, both of which are pooling. In one equilibrium, both types of informed player always choose beer ($P1$), and the uninformed player flees ($R1$) if the informed player chooses beer, but would fight ($R2$) if the informed player were to choose quiche. In the other equilibrium, both types of informed player always choose quiche ($P2$). The uninformed player flees ($R1$) if the informed player chooses quiche, but would fight ($R2$) if the informed player were to choose beer. Only the first equilibrium satisfies the Intuitive Criterion of Cho and Kreps. Note that the payoff structure for the uninformed player is very simple. The best reply is to flee ($R1$) if the probability the informed player is strong (type A) is greater than one-half, and to fight ($R2$) if the probability is less than one-half.

Figure 6 plots the posterior frequencies $P(A|P1)$ and $P(A|P2)$ 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 $P1$; in the other, there were only two instances of a type B player choosing $P2$ in that same span. In contrast, the sessions

---

[2] In classroom experiments, I prefer to use the abstract labeling rather than the beer-quiche, because some students appear to get utility out of choosing to pick a fight, even if it's clearly not in their best interest to do so in payoff terms.

using treatment COUNT did not converge, and behavior in those sessions can be best described as separating, even though that is not an equilibrium phenomenon. Furthermore, in COUNT some type A players chose the unintuitive signal $P2$, whereas all type A players in POST chose the signal $P1$ which is consistent with the equilibrium satisfying the Intuitive Criterion.

Based on debriefing students, 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 5 shows the table with which Veconlab presents the game.[3] 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 play $P1$ with A and $P2$ with B. Comparing 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 seeing $P1$ versus $P2$. 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.

Whatever reason is behind the initial separating behavior, the calculation of the posteriors in treatment POST helps students identify that type B players could be better off if they chose $P1$. If a small number of type B players experiment with $P1$, the posterior frequency $P(A|P1)$ will still be close to 1; therefore, the type B players choosing $P1$ will "get away with it." Seeing this posterior frequency, other informed players who originally did not consider playing $P1$ when they were of type B now will see it as a viable option. Therefore, students who do not initially understand the benefits of choosing $P1$ when they

---

[3] Veconlab refers to the informed player as the "proposer" and the uninformed player as the "responder."

are type B can learn from other students who do make the realization. This behavior is self-reinforcing, because it is an equilibrium for both types to play $P1$. Simply having the public information that some player of type B chose $P1$ is not enough to initiate this dynamic. The computation of the posterior shows that most players who choose $P1$ are of type A, and therefore the uninformed players still prefer to choose $R1$.

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 $P1$ in state A. Therefore, any time $P2$ was chosen, it was chosen by a type B player. Even after play has converged to pooling, from time to time a type B player might experiment with choosing $P2$. This is because he thinks he might have a chance to earn \$2.80, if the uninformed player were to respond with $R1$, as opposed to the \$2.00 he expects to earn by choosing $P1$. However, the posterior data indicate to the uninformed player that, historically, only players of type B chose $P2$; therefore, it is unlikely that the informed player will get away with playing $P2$. The fact that some type B players experiment in this way can motivate discussion of off-equilibrium beliefs and refinements, since the chain of reasoning behind a type B player possibly choosing $P2$ is exactly 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 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. Elizabeth Watson was instrumental in polishing and clarifying the presentation. Thanks to Charles Holt for providing the poker data presented in HOLT [4] 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

[1] Colin F. Camerer. *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton University Press, Princeton NJ, 2003.

[2] I. K. Cho and D. M. Kreps. Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102:179–221, 1987.

[3] Daniel Friedman. Equilibrium in Evolutionary Games: Some Experimental Results. *The Economic Journal*, 106:1–25, 1996.

[4] Charles A. Holt. *Markets, Games, and Strategic Behavior*. Pearson Addison Wesley, Boston MA, 2007.

[5] David Reiley, Michael B. Urbancic, and Mark Walker. Stripped-down Poker: A Classroom Game to Illustrate Equilibrium Bluffing. *Journal of Economic Education*. Forthcoming.

[6] Cedric A. B. Smith. Consistency in statistical inference and decision. *Journal of the Royal Statistical Society, Series B*, 23:1–25, 1961.

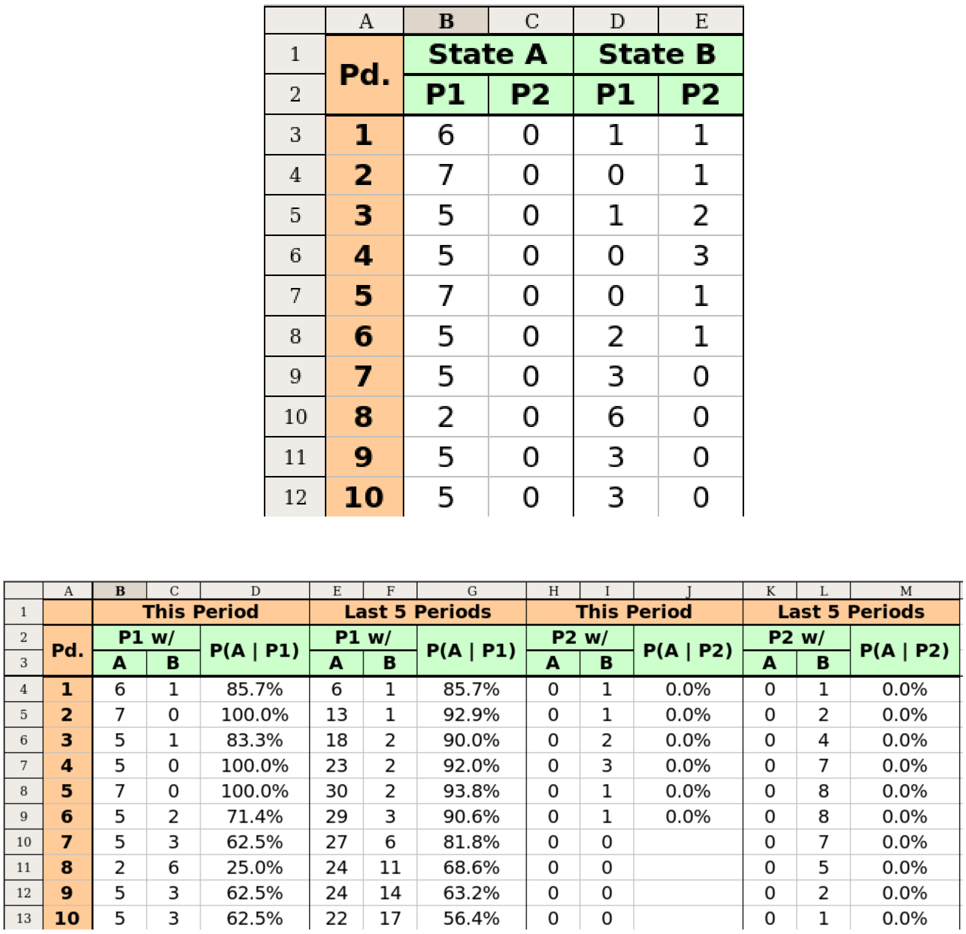

	A	B	C	D	E
1	**Pd.**	**State A**		**State B**	
2		**P1**	**P2**	**P1**	**P2**
3	**1**	6	0	1	1
4	**2**	7	0	0	1
5	**3**	5	0	1	2
6	**4**	5	0	0	3
7	**5**	7	0	0	1
8	**6**	5	0	2	1
9	**7**	5	0	3	0
10	**8**	2	0	6	0
11	**9**	5	0	3	0
12	**10**	5	0	3	0

	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.

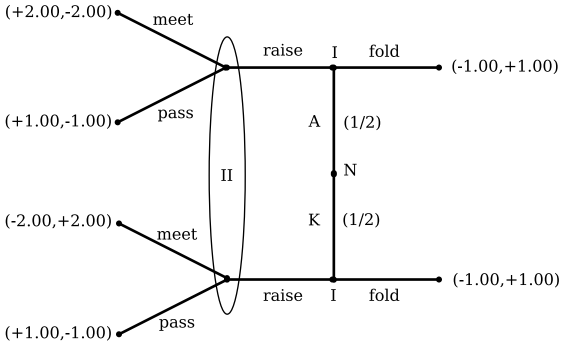

Figure 2. Extensive game for “stripped-down poker.”

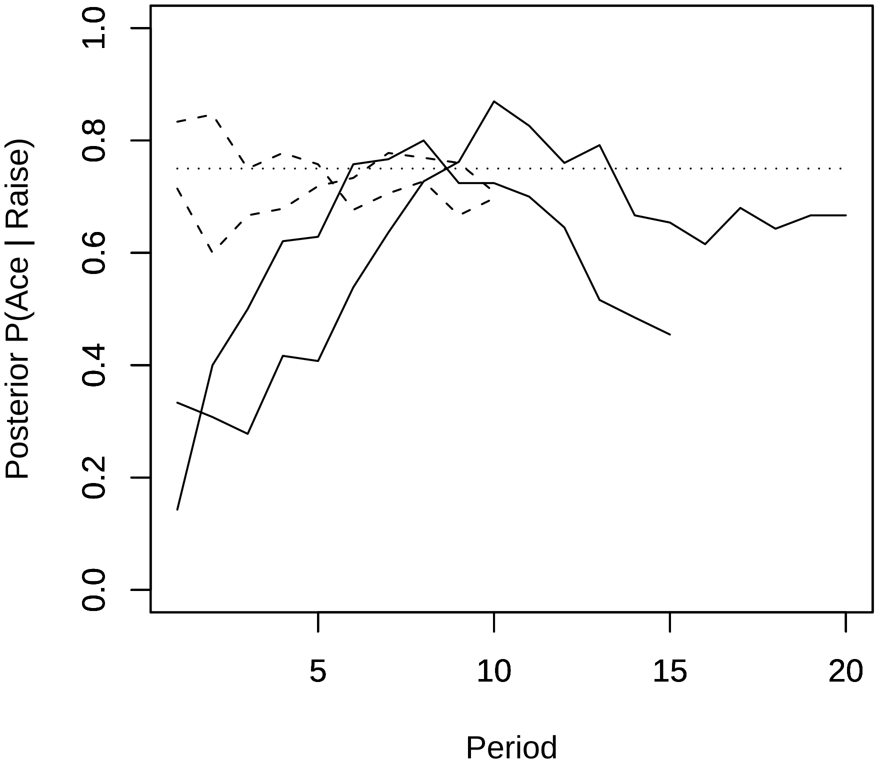

Figure 3. The posterior frequency P(Ace | Raise) in stripped-down poker. The solid lines represent sessions using treatment COUNT, the dashed lines sessions using treatment POST, and the dotted line the beliefs in equilibrium.

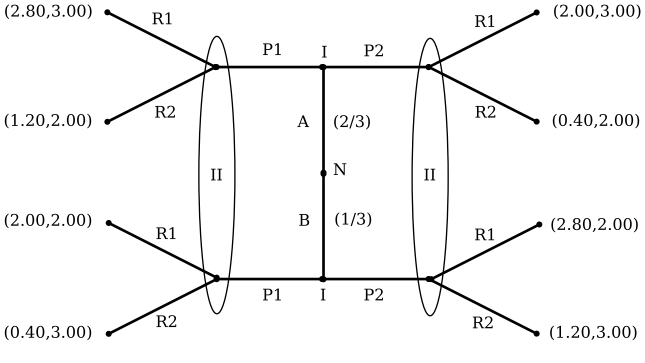

Figure 4. A signaling game with only pooling equilibria.

Payoffs: Proposer, Responder

	R1	R2
P1 (A)	$2.80, $3.00	$1.20, $2.00
P2 (A)	$2.00, $3.00	$0.40, $2.00
P1 (B)	$2.00, $2.00	$0.40, $3.00
P2 (B)	$2.80, $2.00	$1.20, $3.00

Figure 5. Tabular presentation of signaling game in Veconlab.

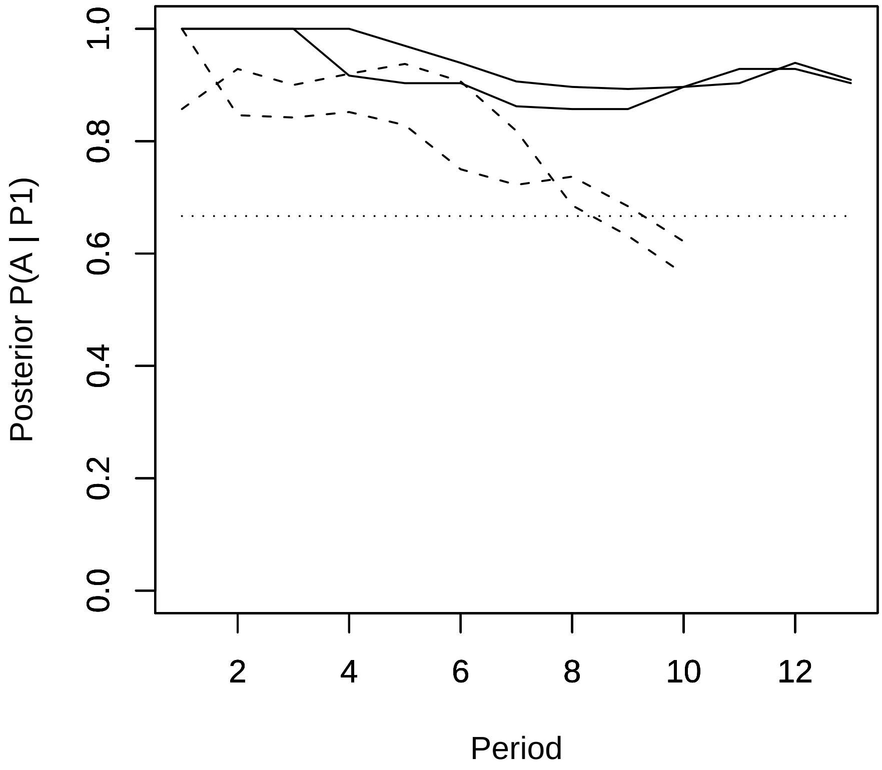

Figure 6. The posterior frequency $P(A|P1)$ for the game in Figure 4. The solid lines represent sessions using treatment COUNT, the dashed lines sessions using treatment POST. The dotted line indicates the beliefs in the intuitive pooling equilibrium.