

Meaningful learning and transfer of learning in games played repeatedly without feedback

Scott Rick
The Wharton School
University of Pennsylvania

Roberto A. Weber
Department of Social & Decision Sciences
Carnegie Mellon University

April 22, 2009*

Abstract:

Psychologists have long recognized two kinds of learning: one that is relatively shallow and domain-specific; and another that is deeper, producing generalizable insights that transfer across domains. The game theory literature has only recently considered this distinction, and the conditions that stimulate the latter kind of “meaningful” learning in games are still unclear. Three experiments demonstrate that meaningful learning – the acquisition of the principle of iterated dominance – occurs in the absence of any feedback. We also demonstrate that meaningful learning transfers to new but strategically similar games, and that such transfer does not occur when prior games are played with feedback. The effects of withholding feedback are similar to, and substitutable with, those produced by requiring players to “self-explain” (provide written explanations for behavior), a method commonly employed in psychology to increase deliberation. This similarity suggests that withholding feedback encourages players to think more deeply about the game.

* We thank the Center for Organizational Innovation & Learning at Carnegie Mellon for funding. Rick was supported by an NSF Graduate Research Fellowship. We also thank three anonymous reviewers, the associate editor, Colin Camerer, George Loewenstein, Ted O’Donoghue, and participants at several seminars and conferences for helpful comments and suggestions. We thank Daniela Lejtneker, Patricia Zapater-roig, and several other members of the Wharton Behavioral Lab for their assistance in running Experiment 3.

Considerable research in economics attempts to understand how people learn in strategic environments. Many experimental studies on games demonstrate that players do not initially play equilibrium strategies, but that with repetition their behavior converges towards equilibrium. Several models attempt to provide a theoretical basis for this regularity (e.g., Cheung and Friedman, 1998; Fudenberg and Levine, 1998; Erev and Roth, 1998; see Chapter 6 of Camerer (2003) for a review). While these models vary in the details of how they assume learning occurs, most share the assumption that learning operates by players observing how well different strategies perform – either by playing those strategies, observing others playing them, or observing (foregone) outcomes produced by unselected strategies – and then adjusting their subsequent behavior in the direction of better-performing strategies. Thus, most learning models in economics focus on understanding how players gradually figure out what strategy produces the highest payoffs in a specific game, a process best described as *strategy learning*.

Economists have devoted considerably less attention to understanding a distinct process that might be called *meaningful learning*, whereby individuals come to obtain meaningful cognitive representations of higher-order concepts, rules, and relationships that can be transferred to novel domains. While a small amount of empirical and theoretical work in economics has uncovered some conditions under which higher-order concepts learned in one context transfer to new contexts (Rankin et al., 2000; Stahl 2000,a,b; Haruvy and Stahl 2008; Cooper and Kagel 2003, 2005, 2008), the factors that stimulate meaningful learning in games are still not well understood within economics.

In this paper, we attempt to advance economists' understanding of meaningful learning by importing relevant insights from psychology. Although economists have

only recently begun to study more than one type of learning, psychologists have for decades recognized a distinction between two kinds of learning, based on the way in which people learn, the kind of knowledge produced by learning, and the ability of individuals to transfer what they learn to new domains. This distinction is important because it highlights significant differences in the depth of what is learned and the ability of individuals to generalize their learning to new contexts.

One type of learning studied within psychology, variously referred to as “implicit,” “procedural,” or “unconscious” learning (Holyoak and Spellman, 1993), is an unconscious process that yields knowledge that is usually neither accessible to cognition nor verbalizable (Reber, 1967, 1989; Mandler, 2004). It is demonstrated, for instance, by showing that subjects exposed to massive amounts of information demonstrate improved performance in pattern matching, but that such improvement exceeds their ability to articulate or generalize their knowledge (Berry and Broadbent, 1984; Nissen and Bullemer, 1987; Hayes and Broadbent, 1988). A key property of this kind of learning is that it operates through perceptual and associative processes, rather than through cognition, and therefore fails to produce cognitive or conceptual representations of what is learned (Mandler, 2004). An important consequence of the absence of such meaningful representation is that what is learned through implicit learning cannot be consciously manipulated or transferred to new domains (Holyoak and Spellman, 1993).

The other type of learning, commonly referred to as “explicit,” “declarative,” or “conscious” learning (Holyoak and Spellman, 1993), is a process through which individuals come to obtain meaningful cognitive representations of underlying concepts, rules, and relationships. Unlike the knowledge acquired via implicit learning, the

knowledge acquired via explicit learning is consciously accessible, generalizable, and verbalizable. Moreover, explicit learning involves cognition, the evaluation of hypotheses, and often results in the development of improved general problem-solving ability (Hayes and Broadbent, 1988; Mandler, 2004). Thus, a key property that distinguishes explicit from implicit learning is that the former is less context-dependent and generates knowledge that can transfer to novel situations.

We propose that the strategy learning commonly observed in games more closely resembles implicit learning than explicit learning. The relationship between strategy learning as studied by game theorists and implicit learning as studied by psychologists is perhaps best illustrated by the lack of transfer of learning to new games. Despite many experiments on learning in games – in which subjects converge towards equilibrium when playing a game repeatedly with prompt outcome feedback – there is very little evidence that what is learned transfers to new strategically similar games (i.e., games in which a meaningful principle, such as dominance or backward-induction, applies non-trivially to both games). For instance, Ho et al. (1998) explicitly test for transfer in two closely-related dominance-solvable games and find no transfer from the first game to the second. Similarly, in a series of papers, Cooper and Kagel (2003, 2005, 2008) find that transfer of learning does not occur when subjects play two signaling games sequentially under typical feedback conditions, unless there are experimental treatments (such as meaningful context or team play) that prime subjects to think more deeply about the game they are playing.

Given the correspondence between strategy and implicit learning, manipulations commonly employed in psychology experiments to inhibit implicit learning and stimulate

explicit learning could serve an analogous function in game theory experiments and could facilitate meaningful learning. One such manipulation involves the amount of feedback participants receive about task performance. Counterintuitively, psychologists have often found that explicit learning is more likely to occur when people are given minimal or delayed feedback than when people are routinely given prompt feedback (e.g., Salmoni et al., 1984; Winstein and Schmidt, 1990; Goodman, 1998; Lurie and Swaminathan, 2009).

Indeed, there is also some evidence in the economics literature that people can engage in a type of learning inconsistent with strategy learning when playing games repeatedly without *any* feedback. Weber (2003) conducted an experiment in which subjects played a dominance-solvable game 10 times without any outcome information between plays of the game. Across several treatments, significant learning occurred – behavior converged towards equilibrium.¹ Such learning cannot be considered strategy learning, which requires regular feedback. Although it is impossible to conclude that the feedback-free learning observed in Weber (2003) was meaningful without examining whether such improvements in performance transfer to a new game, these findings combined with the relevant psychological research suggest that withholding feedback in games may stimulate meaningful learning.

The main hypothesis for our research is that *withholding feedback in a game played repeatedly will produce meaningful learning that will transfer to the first period of a new but strategically similar game.*² We focus on iterated dominance as the

¹ A handful of other papers also provide support for the notion that people can learn in environments where they make repeated choices without feedback (Grether, 1980; Cason and Mui, 1998; Rapoport et al., 2002). However, none of these studies directly explores this phenomenon or focuses on learning.

² We focus on “immediate transfer” (to the first period of the new game) because such one-shot behavior is the focus of much of the research on iterated rationality (e.g., Costa-Gomes, Crawford and Broseta 2001) and also because it is a good way to measure prior learning in games (Merlo & Schotter 1999).

principle that is learned and transferred. We conduct three experiments, all of which demonstrate that the feedback-free learning that occurs in earlier dominance-solvable games transfers to later dominance-solvable games. Moreover, our second and third experiments demonstrate that the kind of (strategy) learning that occurs with regular outcome and payoff feedback fails to yield similar transfer. We interpret this difference in transferability as evidence that the kind of learning produced by feedback-free repetition is more meaningful than the learning produced with immediate feedback.

We also explore, in our third experiment, *why* withholding feedback might stimulate meaningful learning (or, equivalently, why providing feedback impairs the development of meaningful learning). Goodman (1998) proposed that “external feedback,” or learning the correct answer to a problem one just attempted to solve, impairs meaningful learning by diverting attention from what is learned by thinking about and performing a particular task (see also Anzai and Simon, 1979). Similarly, in a meta-analysis of studies that examined the effectiveness of external feedback, Kluger and DeNisi (1996) found that the ability of external feedback to induce meaningful learning decreases as it moves attention away from important features of the task being performed. Beyond merely distracting, feedback can also reduce the motivation to think carefully about a task (Einhorn, 1980).

If feedback impairs meaningful learning by reducing the ability and motivation to think carefully about the task (or game), then interventions that force people to think more deeply about their decisions should reduce the negative influence of feedback on the development meaningful learning. Requiring people to generate “self-explanations,” or engage in verbalization during problem-solving, is one commonly employed method

for inducing deeper thought and developing deeper learning (Gagne and Smith, 1962; Chi et al., 1994).

Therefore, in our third experiment we vary in a first game both whether feedback is withheld and whether players are required to engage in self-explanation as they play the game (in the form of written explanations for behavior). We again find that, in the absence of self-explanation, feedback-free repetition produces meaningful learning that transfers while learning with feedback does not. Adding self-explanation to feedback-free learning does not produce any increase in meaningful learning (measured by transfer). However, when subjects play the first game with feedback and provide written explanations, they demonstrate transfer to the first period of a new game that is almost identical in magnitude to the transfer produced by withholding feedback. Thus, withholding feedback and requiring self-explanation have very similar effects on the development of meaningful learning, and the presence of both factors does not yield any increase in meaningful learning. This interaction suggests that feedback-free repetition generates meaningful learning through similar mechanisms as those produced by self-explanation (presumably, deeper thought).

Experiment 1

Experiment 1 examined whether it is possible to produce meaningful learning, measured by transfer across games, when playing games repeatedly without feedback. Subjects played four normal-form games repeatedly without any feedback until the end of the experiment. They played each game 20 times before proceeding to the next game,

and the order of games varied across sessions. Table 1 presents the four games used in the experiment. The payoffs represent points, with 200 points equal to \$1.

Game A is a symmetric stag-hunt game with three Nash equilibria: the Pareto-dominant equilibrium (*Top/Left*), the risk-dominant equilibrium (*Bottom/Right*) (Harsanyi and Selten, 1988), and a mixed-strategy equilibrium in which subjects play *Top/Left* with probability 0.56. Games B, C, and D, on which we will primarily focus, have unique equilibria resulting from iterated deletion of dominated strategies.³ In Game B, *Bottom* is dominated (step 1) and then *Left* is dominated once *Bottom* is removed (step 2), producing the equilibrium prediction of *Top/Right*. For Game C, which is symmetric, *Bottom/Right* is dominated (step 1) and *Top/Left* is then subsequently dominated (step 2), resulting in the equilibrium of *Middle/Middle*. Finally, in Game D *Left* is dominated (step 1), and then subsequently so are *Bottom* (step 2), *Right* (step 3), and *Top* (step 4), leaving the equilibrium where both players choose *Middle*.

A. Experimental Design

There were four sessions. In each session 18 to 20 subjects played the above four games in 20-period blocks. The sequence of games is presented in Table 2.

Each subject sat at a computer monitor. Subjects' roles (Row or Column) were fixed and each subject was anonymously paired with someone of the opposite role for the duration of the experiment. Matching and role assignment were done by random

³ We included Game A to test one possible interpretation of feedback-free learning – that a player's change in behavior results from best responding to her own prior choices. Since we find little evidence of this phenomenon (players in fact regularly move away from a best-response to their own prior strategies), we focus our attention primarily on the other three (dominance-solvable) games.

assignment of participant numbers. At the beginning of a session, subjects received extensive instruction in how to interpret game matrices with generic payoffs.⁴

At the beginning of each 20-period block, the computer displayed the game matrix for those periods. The experimenter read all of the payoffs in the matrix aloud. Subjects then proceeded through each of the 20 periods by clicking on a choice. After each choice, the computer screen froze, displaying the subject's choice and the matrix for 20 seconds. Subjects received no payoff information until the end of the experiment.

Subjects were recruited from an e-mail list of graduate and undergraduate students at the University of Pittsburgh. The experiment lasted approximately 1.5 hours. Subjects were informed at the beginning of the experiment that their earnings would be determined exclusively by the points they accumulated during the 80 rounds of play. At the end of the experiment, subjects were privately paid one at a time.

B. Results

We first focus on learning within games, pooling across sessions, to examine whether we also find evidence of the feedback-free learning observed in Weber (2003). We then turn our attention to whether there is any transfer of learning across the dominance-solvable games.

1. Learning within games

The aggregate choice frequencies – by 5-period blocks – are presented in Table 3. Table 4 presents logistic regressions that systematically examine changes in behavior

⁴ Instructions and the complete datasets for all experiments are available from the authors.

across periods. The first four regressions explore the extent to which subjects play Nash equilibrium in the three dominance solvable games and the risk-dominant equilibrium in Game A. The change in behavior across periods is significant in all four games.

Subjects clearly learn in our experiment – their behavior shows systematic change in the direction of the predictions of rationality. However, it is also important to understand whether what they are learning corresponds to meaningful principles. The dominance-solvable games allow us to explore this issue.

To explore whether the above learning involves the acquisition of iterated dominance, we consider behavior consistent with the first two steps of this principle in the three dominance-solvable games. If subjects learn to avoid dominated strategies (D1) the frequency of *Bottom* choices in Game B, *Bottom/Right* choices in Game C, and *Left* choices in Game D should decrease. If they learn to avoid strategies that are dominated once one eliminates an opponent's dominated strategies (D2), subjects should be less likely to play *Left* in Game B, *Bottom/Right* or *Top/Left* in Game C, and *Bottom* in Game D.⁵ The second and third sets of regressions in Table 4 examine D1 and D2 violations across periods. The frequency of D1 violations decreases significantly for all three games, while for D2 the decrease is significant in Games C and D.

We find that subjects learn when they play games repeatedly without feedback, and that this learning appears to involve the acquisition of iterated dominance. But to

⁵ In Games B and D, the second step of iterated dominance only applies to players in one of the two roles (column in B, row in D), so we can look exclusively at behavior consistent with the second step. In Game C, which is symmetric, the principle implies that subjects should neither play dominated strategies (*Bottom/Right*) nor strategies dominated in the second step (*Top/Left*). An alternative approach for Game C would be to label a subject as “not violating D2” if she chose either *Bottom* or *Right*, but this would allow a violation of D1 to count as not violating D2.

more definitively determine whether subjects are acquiring a meaningful understanding of iterated dominance, we next examine whether this acquisition transfers across games.

2. Transfer across games

Games B, C, and D all have in common the applicability of two steps of iterated dominance. To test whether learning transfers across games, we compare the frequency with which subjects violate these principles by games' positions within sessions. Table 5 presents, by game position, the frequency with which subjects violated the first two steps of iterated dominance. The table compares behavior in exactly the same three games.⁶

The frequencies of both kinds of violations decrease with game position. For both D1 and D2, roughly 10 percent fewer choices violated the principle in the last game than in the first.⁷ This can also be seen in Table 6, which presents logistic regressions of the frequency with which subjects violate each of the two principles by game position within a session (i.e., the "Game Position" variable takes on values from 1 to 4). The coefficients on Game Position are negative and significant; thus, subjects are less likely to violate the two principles when playing games later in a session.⁸

Another consequence of meaningful learning is that, once a subject learns to identify dominated or iteratively dominated strategies, (s)he should cease to play such strategies for the remainder of the experiment. Therefore, we also consider the number of subjects in the first and last games who never violated each of the first two steps of

⁶ The number of subjects in each cell varies because session 4 had a different number of participants.

⁷ The magnitude of the transfer between the first and last game in a session is roughly 1.5 times greater than the within-game learning we found between the first and last five blocks in a game.

⁸ A reviewer correctly noted that the decrease could correspond to subjects simply making fewer errors later in the experiment, after gaining experience. This interpretation, while plausible here, is inconsistent with the results of Experiments 2 and 3, where we find greater transfer without feedback than with feedback, and with comparable levels of experience between the two conditions.

iterated dominance. For D1, 12 of 37 subjects (who had an opportunity to do so) never played a dominated strategy in their first game. When we look at the fourth game, however, this proportion goes up to 21 of 36 ($\chi^2(1) = 4.94$, $p = 0.026$). Similarly, every subject violated D2 at least once in their first game (0 subjects with no violations). However, in the last game, 9 of 36 subjects never played such strategies ($\chi^2(1) = 10.55$, $p = 0.001$). It appears that many subjects – at some point in the experiment – learned to apply iterated dominance, and continued to apply this principle throughout the remainder of the experiment.

We also find evidence of “complementary acquisition” of the two steps of iterated dominance (i.e., subjects need to acquire D1 before they can acquire D2). We classify subjects according to whether they “never violated,” “stopped violating,” or “never acquired” a principle throughout the experiment (with the last classification for subjects who violated the principle in at least one of the last 5 choices in which they could do so). The number of subjects who never violated D1 but violated D2 at least once (19) is considerably higher than the number who never violated D2 but violated D1 at least once (1). Similarly, of the 19 subjects who never acquired D1, only 1 never violated or stopped violating D2, but of the 44 subjects who never acquired D2, 26 never violated or stopped violating D1. Both of these comparisons are consistent with the fact that acquisition of D2 requires prior understanding of D1.

C. Discussion

When playing different games repeatedly without feedback, subjects learn to stop violating the first two steps of iterated dominance and transfer such learning to later

games. This transfer of learning stands in contrast to the limited evidence of cross-game transfer in previous experiments that provide regular outcome and payoff feedback. The results thus suggest that withholding feedback might stimulate the kind of meaningful learning that transfers to new games.

Of course, a limitation of Experiment 1 is that it did not include a comparison of the kinds of learning that occur with and without feedback. Such a comparison allows us to determine whether withholding feedback is in fact more likely to stimulate meaningful learning than providing feedback. We turn to this comparison next.

Experiment 2

Experiment 2 examines whether meaningful learning is more likely to occur when an earlier game is played with no feedback than when an earlier game is played with payoff and outcome feedback at the end of every period. We measure meaningful learning by the extent to which learning transfers to the first period of a new game and by the extent to which it is consistent with the acquisition of iterated dominance.

The experiment uses a procedure similar to that of Ho et al. (1998). We use two versions of Nagel's (1995) competitive guessing ("p-beauty contest") game. In the game, N players each choose a number in a given range ($s_i \in [\underline{s}, \bar{s}]$). The average of the N numbers is then multiplied by a constant (p) to obtain a target number. The player whose choice is the smallest absolute distance from the target number wins a fixed prize.

We use one version of the game with $p < 1$ ($p = 0.7$, $s_i \in [0, 100]$) and another with $p > 1$ ($p = 1.3$, $s_i \in [100, 200]$), and follow Ho et al. in referring to the former as the "Infinite Threshold" (IT) game and the latter as the "Finite Threshold" (FT) game.

Iterated deletion of dominated strategies selects unique symmetric equilibria in these two games. In the IT game, infinite iterations of multiplying 0.7 times the upper bound of 100 yields the Nash equilibrium of $s_i^* = 0$. In the FT game, three iterations of multiplying 1.3 times the lower bound of 100 yields the Nash equilibrium of $s_i^* = 200$. Thus, both games are solvable by iterated deletion of dominated strategies. In spite of this strategic similarity, however, Ho et al. found no evidence of transfer in the first period of the second game when the two games were played sequentially with feedback.

As in Ho et al.'s study, subjects in our experiment played both games, in sequence, for 10 periods. Half of the subjects received the games in one order (IT→FT) while the other half received the other order (FT→IT) (see Table 7).

We varied the feedback provided between plays of the first game. In the Feedback treatment, subjects received outcome and payoff information at the end of each period, as in Ho et al. In the No Feedback treatment, subjects did not receive any feedback between periods of the first game – they found out first-game outcomes only after completing the second game.

Following 10 periods of the first game, either with or without feedback, subjects played the second game with feedback. Since iterated dominance applies in both games, we are interested in whether what is learned in the first game produces immediate transfer to the second game. More precisely, we focus on the extent to which initial (Period 11) choices in the second game deviate from Nash equilibrium. Our main prediction is that initial choices in the second game will deviate less from Nash equilibrium in the No Feedback treatment than in the Feedback treatment.

A. Experimental Design

Participants were Carnegie Mellon and University of Pittsburgh students. There were 16 sessions. In each session 8 to 10 subjects played two versions of the competitive guessing game – IT and FT – for 10 periods each. The winner in each period received \$3, and in case of a tie this amount was equally divided among the winners. Subjects played one game first and then received instructions for the other, and the order of games was counterbalanced. Aside from the game order, sessions also varied by whether the first game was played with or without feedback. Table 7 presents the different kinds of sessions. The second game was always played with feedback.

At the beginning of the experiment, subjects received instructions describing general procedures and the first version of the game. After having the opportunity to ask questions, subjects proceeded to the first period. In each period, subjects recorded their choices on a table at the bottom of their instruction sheet. The experimenter recorded their choices and entered them in a laptop, which computed the outcome and payoffs for that period.

In the Feedback treatment, subjects received feedback after every period – the average, target number, and participant number(s) of the winner(s) were written on the board and read aloud, and subjects recorded this information. In the No Feedback treatment, the experimenter determined the outcome, but this information was not revealed to subjects. Instead, the experiment proceeded to the next period.

In both the Feedback and No Feedback treatments, the first game was played for 10 periods. After the 10th period, subjects were handed a new instruction sheet that

described the other version of the game. They then played 10 periods of the second game, receiving feedback at the end of every period.

Following the second game, subjects in the No Feedback treatment received all of the outcome information for the first game. Subjects were then paid privately.

B. Results

In presenting the results, we first briefly examine whether learning occurred in the first game. We then explore whether there is transfer of learning to the first period of the second game and whether such learning differs by feedback treatment.

1. Learning in first competitive guessing game

Figures 1a and 1b present the mean choices across periods for the two versions of the competitive guessing game, by order and feedback treatment. The left side of each figure presents the first 10 periods, when the first game was played either with feedback (F) or without any feedback (NF) between plays of the game. As both figures reveal, there is convergence towards equilibrium in the first 10 periods, when subjects played the first version of the competitive guessing game, both with and without feedback.

Table 8 presents average choices, by treatment, in the first and last periods of the first game (1 and 10) and the first period of the second game (11). In the Feedback treatments (first and third columns of data), the average choices move significantly in the direction of equilibrium between the first and tenth periods (IT: 34.5, $t_{37} = 6.96$, $p < 0.001$; FT: 36.6, $t_{38} = 9.29$, $p < 0.001$). Thus, not surprisingly, learning occurs with feedback in the first game played.

In the No Feedback treatment, however, we also observe learning. Average choices move significantly towards equilibrium between the first and tenth periods (IT: 19.9, $t_{37} = 5.43$, $p < 0.001$; FT: 21.4, $t_{36} = 5.70$, $p < 0.001$). Thus, the results of the first game are very similar to those of Weber (2003) – subjects learn both with and without feedback, though the change in behavior is greater with feedback.

Table 9 presents the frequency of violations of the first two steps of iterated dominance at different stages of the experiment. For each combination of periods, the table presents how many subjects played a strategy that violated the principle at least once in those periods.

In Periods 1-3 of both treatments, roughly 18 percent of subjects play a dominated strategy at least once. In both treatments, however, this percentage decreases to 4 percent by Periods 8-10 (of the first game), and this change is significant for both treatments (Feedback: $\chi^2(1) = 6.97$, $p = 0.008$; No Feedback: $\chi^2(1) = 8.03$, $p = 0.005$). Similarly, in Periods 1-3 roughly half of subjects violate the second step of iterated dominance at least once, but this proportion decreases significantly in both treatments for Periods 8-10 – to 4 percent in the Feedback treatment ($\chi^2(1) = 40.72$, $p < 0.001$) and 19 percent in the No Feedback treatment ($\chi^2(1) = 22.34$, $p < 0.001$). Thus, while we observe greater learning with feedback, there is clear evidence that subjects learn even without it.

2. Transfer to second competitive guessing game

To explore transfer across the strategically similar competitive guessing games, we compare the behavior of “inexperienced” subjects, who have never encountered either game before, to that of “experienced” subjects who have previously played the other

game. To measure the choices of inexperienced subjects, we use choices in the first period (at the beginning of the experiment) for both games.⁹

Transfer in the Feedback treatment. Similarly to Ho et al. (1998), we find no immediate transfer to the second game when the first game is played with feedback. The average first-period (inexperienced) choices are 43.3 (IT) and 167.4 (FT). As Table 8 reveals, the respective average Period 11 choices in the Feedback treatment (at the beginning of the second game) are 43.3 and 163.5.

Moreover, Table 9 reveals no decrease in the Feedback treatment in violations of either the first or second steps of iterated dominance across the two games. For the first step (dominance), 17 percent of subjects violated the principle in Periods 1-3, but this percentage increases to 22 percent in the first three periods of the second game (11-13). For the second step, 49 percent of subjects violate the principle in Periods 1-3, and this percentage decreases slightly to 44 percent in Periods 11-13. In both cases, the changes are statistically insignificant (first step: $\chi^2(1) = 0.66$; second step: $\chi^2(1) = 0.24$). Thus, we find no evidence of immediate transfer in the Feedback treatment.¹⁰

Transfer in the No Feedback treatment. In the No Feedback treatment, we find clear evidence of immediate transfer. Recall that the average first-period (inexperienced) choices in the two versions of the competitive guessing game are 43.3 (IT) and 167.4

⁹ Thus, we compare first-period choices in the second game from one sequence (when subjects have experience with the other game), to first-period choices in the first game from the other sequence (in which subjects play the same game, but without any prior experience). We also pool first-game choices from the two feedback treatments (i.e., pooling the first row of Table 8 by game), which proceeded identically up to the end of the first period (when one treatment received feedback and the other did not). If we compare first-period choices by feedback treatment, we find no significant differences between the Feedback and No Feedback treatments (IT: $t_{74} = 0.79$; FT: $t_{74} = 1.53$).

¹⁰ In later periods of the FT→IT treatment, choices move away from Nash equilibrium. In a few sessions, subjects began experimenting with choices of 100 (perhaps out of boredom or to try to gain a strategic advantage – cf. Weber, 2003). In some cases this occurred in consecutive periods, raising the average substantially and persisting into subsequent periods. In the No Feedback treatment, these choices of 100 started to occur somewhat earlier in two sessions (contributing to the slower convergence towards zero).

(FT). The average choices in the No Feedback treatment at the start of the second game (Period 11) are 30.9 and 176.1 (see Table 8), which are both significantly closer to Nash equilibrium than the corresponding choices of inexperienced subjects (IT: $t_{111} = 2.63$, $p = 0.005$, one-tailed; FT: $t_{112} = 1.86$, $p = 0.03$, one-tailed).

We also observe evidence of transfer in the No Feedback treatment in Table 9 (right column). In Periods 1-3 of the first game, 19 percent of subjects violated the first step of iterated dominance at least once and 56 percent did so for the second step. In Periods 11-13, however, these proportions decrease to 12 percent and 31 percent, respectively. While the decrease for the first step is not significant ($\chi^2(1) = 1.28$), the decrease in violations for the second step is ($\chi^2(1) = 9.80$, $p = 0.002$).¹¹

Thus, we find immediate transfer to the second game when the first game is played without feedback, even though we observe no such transfer when the first game is played with feedback. One caveat worth highlighting is that, as Figure 1 demonstrates, there is evidence of transfer in the Feedback treatment if transfer is defined more broadly as the speed with which subjects converge towards equilibrium. Mean choices in *both* treatments converge to equilibrium more quickly in the second game than in the first game, which was also observed by Ho et al. But a comparison of No Feedback/Feedback means alone does not shed much light on whether Feedback subjects are just taking longer to transfer a meaningful understanding of iterated dominance to the second game, or whether they are perhaps simply correcting a misapplied heuristic (initially believing that choices will converge to the same boundary as in Part 1, but then realizing that choices will converge to the opposite boundary and adjusting quickly in that direction).

¹¹ We also observe greater equilibrium play in Period 11 of the No Feedback treatment (IT: 2/37; FT: 10/38) than in the Feedback treatment (IT: 1/39; FT: 4/38). Combining the two game order treatments, this difference is marginally significant ($\chi^2(1) = 3.46$, $p < 0.07$).

To shed light on this question, we can examine whether learning in the first game is more strongly related to evidence of learning in the second game in the No Feedback treatment than in the Feedback treatment. That is, to what extent is learning at the beginning of the second game correlated with learning at the end of the first game? A positive relationship would suggest that subjects are indeed transferring something meaningful from the first to the second game, while the absence of a relationship would indicate that subjects are not really transferring any knowledge from the first game to the second, and that anything that looks like transfer of learning in early periods of the second game must be something else.

Table 10 presents OLS regressions exploring the determinants of Period 11 choices (the first period in the second game). The dependent variable is the distance of a subject's Period 11 choice from the Nash equilibrium in the corresponding game (i.e., x_i in IT and $200-x_i$ in FT). The first regression confirms the above effects of withholding feedback on immediate transfer: Period 11 choices are significantly closer to equilibrium, on average, in the No Feedback treatment.

The remaining regressions explore the relationship between Period 11 choices (second game) and behavior in the first game, separately for each treatment. There is no significant relationship between Period 10 and Period 11 choices in the Feedback treatment, but there is a strong and statistically significant relationship in the No Feedback treatment. That is, only in the No Feedback treatment do subjects who played strategies close to equilibrium at the end of the first game also do so at the beginning of the second game. This is robust to controlling for subjects' Period 1 choices (how much knowledge they had at the beginning of the experiment). Thus, subjects who appear to

“learn” a principle by the end of the first game are likely to behave consistently with that principle in the second game in the No Feedback treatment, but not in the Feedback treatment. We observe a similar pattern if we examine the relationship between violations of iterated dominance in the last three periods of the first game and the first three periods of the second game.¹²

Thus, when we move beyond a simple comparison of means in the second game, we find little evidence that subjects in the Feedback treatment are applying something they learned in the first game to the beginning of the second game, while for No Feedback subjects we do find such a relationship. This indicates that the learning that occurred in the No Feedback treatment was more meaningful than the learning that occurred in Feedback treatment.

C. Discussion

Consistent with much previous research on strategy learning, the kind of learning obtained with feedback does not immediately transfer across strategically similar games. But the kind of learning produced by repeated play without any feedback transfers to the first period of the second game. Thus, we find that meaningful learning is more likely to occur when feedback is withheld than when feedback is given.

The results obtained thus far suggest that feedback-free repetition results in meaningful learning that transfers across games (Experiments 1 and 2), even when

¹² In the Feedback treatment, 74 subjects never violated dominance in the final three periods of the first game (see Table 9). Of these, 17 (23 percent) violated this principle at least once in the first three periods of the second game. However, in the No Feedback treatment this proportion is lower (9 of 72, or 13 percent; $\chi^2(1) = 2.73$, $p = 0.10$). For the second step of iterated dominance, the proportion of Period 8-10 non-violators who violated this principle in Periods 11-13 is also higher in the Feedback treatment (32 of 74; 58 percent) than in the No Feedback treatment (14 of 61; 23 percent; $\chi^2(1) = 6.13$, $p = 0.01$).

feedback-based learning does not (Experiment 2). Our final experiment sheds light on why feedback-free repetition yields meaningful learning.

Experiment 3

Experiment 3 examines why meaningful learning is more likely to occur when feedback is withheld than when feedback is given. Prior psychological research has proposed that feedback can reduce people's ability and motivation to think carefully about a task (Einhorn, 1980; Goodman, 1998). If so, then interventions that force people to think carefully about the task at hand should moderate the influence of feedback on meaningful learning. "Self-explanation" is one such intervention: people who explain to themselves why they are doing what they are doing tend to think more deeply about the task at hand than people who do not generate such explanations (e.g., Chi et al. 1994). Such explanations also tend to facilitate the development of knowledge that transfers across related tasks (Gagne and Smith, 1962).

Experiment 3 examines whether self-explanation (in the form of written explanations for one's own decisions and the decisions of others) moderates the influence of feedback on meaningful learning. If self-explanation has a greater effect on meaningful learning when feedback is provided than when feedback is withheld, then we would have evidence indicating that withholding feedback stimulates meaningful learning by encouraging players to spontaneously engage in a process similar to self-explanation. Put differently, if withholding feedback primes deeper thinking in the same way self-explanation does, then their effects should be similar and substitutable.

In the experiment, subjects initially played 10 periods of the Finite Threshold (FT) competitive guessing game from Experiment 2. We varied whether or not subjects received feedback at the conclusion of each round, as well as whether or not subjects were asked to explain their decisions and the (revealed or predicted) decisions of others during each period. At the conclusion of these 10 periods, subjects played a single period of the Infinite Threshold (IT) game. As in Experiment 2, we measure meaningful learning by the extent to which choices in the second game deviate from Nash equilibrium. If feedback-free repetition stimulates people's ability and motivation to think carefully about a game, then its effects should be similar to those produced through self-explanation. Therefore, we predict that the presence of either feedback-free repetition or self-explanation should facilitate meaningful learning, but that their combination should produce no additional benefits.

A. Experimental Design

Participants were University of Pennsylvania students. In each session 7 to 10 subjects played 10 periods of the FT competitive guessing game, followed by one period of the IT competitive guessing game. The winner in each period received \$2, and in case of a tie this amount was equally divided among the winners.

At the beginning of the experiment, subjects received instructions describing general procedures and the FT game. After having the opportunity to ask questions about the game, subjects proceeded to the first period. The experiment was computerized, using the z-Tree software (Fischbacher, 2007), and in each period subjects typed their choices and recorded them on a table at the bottom of their instruction sheet.

In the first game, we varied whether or not subjects received feedback, as well as whether or not subjects were asked to explain their own decisions and the (revealed or predicted) decisions of others. In the Feedback (F) treatments, the computer revealed the average, target number, and the amount of money earned at the end of each round. Subjects recorded this information on a record sheet before proceeding to the next round. In the No Feedback (NF) treatments, the computer asked subjects what they thought the average choice was at the end of each round. Subjects typed this number and recorded it on a record sheet before proceeding to the next round.¹³

In the Explanation (E) treatments, subjects were asked to enter a choice and then, on the same screen, to “type a sentence or two indicating why you made the particular choice you made. You can write what led you to think of this choice or why you think this choice may be a good one that is likely to win the prize.” After either receiving feedback or guessing the average choice in the round, subjects were asked to “type a sentence or two indicating why you believe this [was / is likely to be] the average in this round. You can write what [made this number likely to be the average / led you to think of this as your guess of the average] or why you think others may have chosen numbers that [made this / make this likely to be] the average.” In the No Explanation (NE) treatments, subjects made their decisions and either received feedback or guessed the average, but were not prompted to explain their own decisions or the decisions of others.

Following the 10th period, subjects were handed a new instruction sheet that described the IT game. After having the opportunity to ask questions about the game,

¹³ We asked No Feedback subjects to guess the average because, as we explain below, we asked subjects in the Explanation treatments to explain why they thought other subjects chose what they chose. Weber (2003), in a similar competitive guessing game, varied whether or not No Feedback subjects guessed the average choice in each round and found that making these guesses did not influence learning.

subjects played one period of the game. Subjects did not explain their decisions or the decisions of others, and all subjects received feedback at the end of the period.

After the second game, we administered the short form of the Need for Cognition scale (Cacioppo, Petty, and Kao, 1984), which measures individual differences in the tendency to engage in and enjoy thinking.¹⁴ Next, No Feedback subjects received all of the outcome information for the FT game. Subjects were then paid privately.

B. Results

Table 11 presents the number of sessions and subjects per treatment.¹⁵ We first briefly examine whether learning occurred in the first game. We then explore whether there is transfer of learning to the second game and, more importantly, whether self-explanation moderates the influence of feedback on transfer.

1. Learning in the first competitive guessing game

Figure 2 presents the mean choices across periods, by treatment. As the figure reveals, there is convergence towards equilibrium in all treatments. Average choices move significantly in the direction of equilibrium between the first and tenth periods in all treatments (NF+NE: 11.0, $t_{49} = 2.17$, $p < 0.04$; F+NE: 33.9, $t_{56} = 9.60$, $p < 0.0001$;

¹⁴ Since individuals high in Need for Cognition may be more likely to spontaneously engage in self-explanation, we included the scale to control for this potentially important individual difference.

¹⁵ We excluded one session (N = 10) from the analysis. Partway through this session, the experimenter had to leave abruptly, and the remainder of the session was administered by a lab assistant. One subject in this session, either due to boredom, misunderstanding, maliciousness, or some combination of these factors, consistently chose numbers whose individual digits summed to or included seven (e.g., round 6: 131.11, reason given: “seven. I was trying to reach seven”; round 8: 124, “1+2=3 3+4=7”; round 11: 77.77). Because this was a Feedback session, this subject’s erratic behavior coupled with the experimenter leaving raised suspicion among other subjects about whether the game was real (e.g., subject 5 in round 9: “Perhaps there is a computer player choosing random numbers”; subject 1 in round 10: “This is real, right?”). An analysis of data from this session clearly reveals it to be an outlier relative to other sessions.

NF+E: 16.1, $t_{41} = 4.31$, $p = 0.0001$; F+E: 33.1, $t_{35} = 6.23$, $p < 0.0001$). Consistent with Experiment 2, we observe a significant main effect of feedback when we pool across self-explanation treatments (mean difference between Period 1 and Period 10 choices in NF treatments: 13.3, mean difference in F treatments: 33.6; $t_{183} = 4.61$, $p < 0.0001$). Also, Explanation generally had little effect on choices in the Feedback treatment (no F+NE / F+E difference reaches $p < .10$) or the No Feedback treatment (only the period 5, 6, and 9 NF+NE / NF+E differences reach $p < .10$).

2. Transfer to second competitive guessing game

Figure 3 presents the mean choices in Round 11, by treatment. Lower choices are closer to equilibrium. Consistent with Experiment 2, mean choices under No Explanation are closer to equilibrium under No Feedback than under Feedback (31.9 vs. 41.9, $t_{105} = 1.63$, $p = 0.05$, one-tailed).

Our hypothesis for Experiment 3 was that Explanation and No Feedback would have substitutable effects on transfer. That is, we expected that the presence of either treatment would enhance meaningful learning, measured by transfer in Period 11, but that the presence of both treatments would produce no additional benefit. The pattern in Figure 3 is consistent with this hypothesis. Average Period 11 choices are very similar when either No Feedback (NF) or Self-Explanation (E) is present, ranging from 27.2 to 31.9, and are statistically indistinguishable. However, the average choices in all three of these conditions are well below the average choice in F+NE (41.9) and these differences are at least marginally statistically significant (vs. NF+NE: $t_{105} = 1.63$, $p = 0.05$; vs. NF+E: $t_{97} = 1.53$, $p = 0.06$; vs. F+E: $t_{91} = 2.16$, $p = 0.02$, all one-tailed).

Table 12 presents OLS regressions that systematically explore these differences. No Feedback yields lower Period 11 choices, as does Explanation. However, the interaction of No Feedback and Explanation produces a net effect that is similar to the effect of No Feedback alone. That is, as we hypothesized and as can also be seen in Figure 3, adding Explanation to Feedback moves choices significantly closer to the equilibrium, but adding Explanation to No Feedback has no effect. Model 2 reveals that the results are unchanged when controlling for individual differences in Need for Cognition. Need for Cognition has a negative influence on Period 11 choices ($p < .04$, one-tailed), suggesting that the more people are naturally inclined to engage in deep thinking, the closer their Period 11 choices are to the Nash equilibrium.

C. Discussion

Experiment 3 sheds light on the process by which withholding feedback stimulates meaningful learning. We find that self-explanation significantly enhances meaningful learning and transfer, but only in the Feedback treatment. Under No Feedback, self-explanation has no effect. Given that self-explanation has previously been demonstrated to deepen thinking, this interaction suggests that the enhancement of meaningful learning produced in the No Feedback treatments was the result of a process similar to that which was produced by self-explanation. Put differently, people in our No Feedback treatments appear to engage in the same kind of deeper thinking that is produced through self-explanation.

Conclusion

Although economists have long focused on understanding how players figure out what strategy works best in a particular game, interest in how players acquire generalizable knowledge that transfers to strategically similar settings has only recently received significant attention (e.g., Stahl, 2000a,b; Rankin et al., 2000; Cooper and Kagel, 2003, 2005, 2008). We contribute to this growing literature by exploring whether and why feedback-free repetition enhances such meaningful learning. We find that withholding feedback in repeated play of games stimulates meaningful learning and cross-game learning transfer in dominance-solvable games (Experiments 1 and 2), that withholding feedback is more likely to stimulate meaningful learning than providing regular outcome and payoff feedback (Experiment 2), and that the effects of withholding feedback are very similar to those produced by self-explanation, another intervention known to facilitate deeper thinking (Experiment 3).

Our work is important for several reasons. First, we demonstrate that withholding feedback, a manipulation previously employed in game theory experiments to eliminate learning (e.g., Costa-Gomes and Crawford, 2006), can have counterintuitive effects, namely the facilitation of meaningful learning. While games played only once or twice each (as in Costa-Gomes and Crawford) might not yield meaningful learning when played without feedback, we demonstrate that such learning occurs with 10 or 20 repetitions of a game.

Second, since our work is based on a large body of research in psychology, we present an opportunity for economics and game theory to further integrate knowledge from other disciplines. For instance, one goal of subsequent research should be to further develop theoretical models that account for the two kinds of learning. Turning to the

existing empirical and theoretical literature in psychology can be helpful in this regard. For instance, John Anderson's well-known (in psychology) ACT-R model (Anderson and Lebiere, 1998) allows for the two kinds of learning, based on distinctions between types of knowledge and how they are stored. While a direct translation of this model to economic theory is not straightforward, it presents a potential starting point.

Finally, our work also produces implications for the measurement and study of learning. While learning may appear "better" in environments with feedback – when measured by the degree of adjustment towards equilibrium – we demonstrate that other ways of measuring learning, such as immediate transfer, produce quite different conclusions. Measuring transfer, as we do here, facilitates greater insight into the type and depth of learning taking place.

Of course, the procedure we use – feedback-free repetition – is only one way to develop meaningful learning, and there may be other, potentially better, methods. For example, would providing time for introspection function as well as feedback-free repetition? Some models in economics propose that players can develop improved reasoning ability by introspecting prior to playing a game (Goeree and Holt, 2002; MacLeod, 2002; Capra, 2003). However, it is also possible that time alone might not produce meaningful learning. For example, empirical research in psychology suggests that people tend to postpone thinking concretely about specific aspects of a situation until it is imminent (Trope and Liberman, 2003). This work suggests that, no matter how long people have to think about a game they are about to play, they may not think about it carefully until they actually begin to play. Some economists have similarly observed that "learning by doing" (Arrow, 1962) is more likely to occur than "learning by thinking."

For example, Smith (1991, p. 12) noted that “many years of experimental research have made it plain that real people do not solve decision problems by thinking about them in the way we do as economic theorists. Only academics learn primarily by reading and thinking.” Thus, it is not entirely clear whether introspection without decision-making would stimulate meaningful learning, but we leave this important open question for future research.

Moreover, there may be conditions under which meaningful learning develops even when feedback is provided. Experiment 3 demonstrated one such condition (self-explanation), but there may be others. In Experiments 2 and 3, subjects played one game and were then “surprised” by a second, strategically similar game (cf. Merlo and Schotter, 1999). It would be interesting to examine whether warning subjects in advance that they will play a different but related game in the future allows meaningful learning even with feedback. Moreover, previous research in psychology demonstrates that delayed feedback is useful for producing the kind of learning that transfers, and a considerable literature explores the optimal timing of feedback (e.g., Schmidt et al., 1989; Erev et al., 2006). Therefore, subjects in our experiments might demonstrate even greater meaningful learning if presented with feedback after some such “optimal” delay.

Several other interventions found in psychological research to enhance deeper thinking could also stimulate meaningful learning in economic settings. Positive affect, for example, has been found to stimulate creativity and improved problem-solving (Isen, Daubman, and Nowicki, 1987), and it could similarly facilitate acquisition of deeper principles in games. Separately, recent research suggests that reading disfluent text (i.e., text presented in a difficult-to-read font) stimulates deeper thinking than reading fluent

text, holding actual text constant (Alter et al., 2007). Thus, manipulations as simple as presenting game payoffs or feedback disfluently may stimulate meaningful learning.

Future work should also examine the extent to which the benefits of feedback-free repetition generalize across contexts. For example, in contexts where opponents' characteristics are important – such as public goods or prisoner's dilemma games (Fischbacher et al., 2001; Parks and Rumble, 2001), feedback-free repetition may facilitate a greater understanding of deeper principles *without* facilitating better performance. That is, when the optimal strategy is highly contingent on the type of opponent encountered, meaningful learning may not improve performance. Future work should also examine whether feedback-free repetition facilitates meaningful insight into game-theoretic principles other than dominance (e.g., backward-induction).

It is also worth considering whether games can be classified according to whether meaningful learning is required to play them well. For example, games like the competitive guessing game are contexts in which players do not need meaningful learning to perform well when feedback is provided. These are precisely the kinds of contexts in which providing feedback should inhibit transfer, and therefore our feedback-free environments are useful for facilitating meaningful learning. By contrast, we should observe less of an advantage for feedback-free learning in games and environments where meaningful learning is required to perform well even when feedback is provided.

Our experiments provide a useful starting point – along with other recent work – for further exploring distinctions among different kinds of learning in games, and the source and consequences of such a distinction. But much remains to be done before this distinction can be fully incorporated into economics and game theory.

References

- Alter, A. L., Oppenheimer, D. M., Epley, N., and Eyre, R. N. (2007). Overcoming Intuition: Metacognitive Difficulty Activates Analytic Reasoning,” *J. Exp. Psych.: General* **136** (4), 569-576.
- Anderson, J. R., and Lebiere, C. (1998). *The Atomic Components of Thought*. Mahwah, NJ: Erlbaum.
- Anzai, Y., and Simon, H. A. (1979). “The Theory of Learning by Doing,” *Psychological Rev.* **86** (2), 124-140.
- Arrow, K. J. (1962). “The Economic Implications of Learning by Doing,” *Rev. Econ. Studies* **29** (3), 155-173.
- Berry, D. C., and Broadbent, D. E. (1984). “On the Relationship between Task Performance and Associated Verbalizable Knowledge,” *Quart. J. Exper. Psych.* **36A**, 209-231.
- Cacioppo, J. T., Petty, R. E., and Kao, C. F. (1984). “The Efficient Assessment of Need for Cognition,” *J. Personality Assessment* **48** (3), 306-307.
- Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Camerer, C., and Ho, T.-H. (1999). “Experience-Weighted Attraction Learning in Normal Form Games,” *Econometrica* **67** (4), 837-874.
- Capra, M. (2003). “Reasoning Chains in One-Shot Games,” Working Paper.
- Cason, T. N., and Mui, V.-L. (1998). “Social Influence in the Sequential Dictator Game,” *J. Mathematical Psych.* **42** (2-3), 248-265.
- Cheung, Y.-W., and Friedman, D. (1998). “A Comparison of Learning and Replicator Dynamics Using Experimental Data,” *J. Econ. Behav. Organ.* **35** (3), 263-280.
- Chi, M. T., De Leeuw, N., Chiu, M., and LaVancher, C. (1994). “Eliciting Self-Explanations Improves Understanding,” *Cognitive Sci.* **18**, 439-477.
- Cooper, D. J., and Kagel, J. H. (2008). “Learning and Transfer in Signaling Games,” *Econ. Theory* **34** (3), 415-439.
- Cooper, D. J., and Kagel, J. H. (2003). “Lessons Learned: Generalizing Learning Across Games,” *Amer. Econ. Rev. Papers Proceed.* **93** (2), 202-207.

- Cooper, D. J., and Kagel, J. H. (2005). "Are Two Heads Better Than One? Team versus Individual Play in Signaling Games," *Amer. Econ. Rev.* **95** (3), 477-509.
- Costa-Gomes, M. A., and Crawford, V. P. (2006). "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *Amer. Econ. Rev.* **96** (5), 1737-1768.
- Costa-Gomes, M. A., Crawford, V. P., and Broseta, B. (2001). "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica* **69** (5), 1193-1235.
- Einhorn, H. (1980). "Learning from Experience and Suboptimal Rules in Decision Making," in *Judgment under Uncertainty: Heuristics and Biases*, (D. Kahneman, P. Slovic, and A. Tversky, Eds.), pp. 268-283. Cambridge, UK: Cambridge University Press.
- Erev, I., Luria, A., and Erev, A. (2006). "On the Effect of Immediate Feedback," Working Paper.
- Erev, I., and Roth, A. E. (1998). "Predicting how People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *Amer. Econ. Rev.* **88** (4), 848-881.
- Fischbacher, U. (2007). "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics* **10** (2), 171-178.
- Fischbacher, U., Gächter, S., and Fehr, E. (2001). "Are People Conditionally Cooperative? Evidence from a Public Goods Experiment," *Economics Letters* **71**, 397-404.
- Fudenberg, D., and Levine, D. (1998). *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Gagne, R. M., and Smith, E. C. (1962). "A Study of the Effects of Verbalization on Problem Solving," *J. Exp. Psych.* **63** (1), 12-18.
- Goeree, J. K., and Holt, C. A. (2004). "A Model of Noisy Introspection," *Games Econ. Behav.* **46** (2), 365-382.
- Goodman, J. S. (1998). "The Interactive Effects of Task and External Feedback on Practice Performance and Learning," *Organizational Behav. Human Dec. Processes* **76** (3), 223-252.
- Grether, D. M. (1980). "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quart. J. Econ.* **96** (3), 537-557.
- Harsanyi, J., and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.

- Haruvy, E. and Stahl, D. (2008). "Learning Transference Between Dissimilar Symmetric Normal-Form Games," Unpublished manuscript.
- Hayes, N. A., and Broadbent, D. E. (1988). "Two Modes of Learning for Interactive Tasks," *Cognition* **28**, 249-276.
- Ho, T.-H., Camerer, C., and Weigelt, K. (1998). "Iterated Dominance and Iterated Best-Response in p -Beauty Contests," *Amer. Econ. Rev.* **88** (4), 947-969.
- Holyoak, K. J., and Spellman, B. A. (1993). "Thinking," *Annual Rev. Psych.* **44**, 265-315.
- Isen, A. M., Daubman, K. A., and Nowicki, G. P. (1987), "Positive Affect Facilitates Creative Problem Solving," *J. Personality Social Psych.* **52** (6), 1122-1131.
- Kluger, A. N., and DeNisi, A. (1996). "The Effects of Feedback Interventions on Performance: A Historical Review, a Meta-Analysis, and a Preliminary Feedback Intervention Theory," *Psychological Bull.* **119** (2), 254-284.
- Lurie, N. H. and Swaminathan, J. M. (2009). "Is Timely Information Always Better? The Effect of Feedback Frequency on Decision Making," *Organizational Behav. Human Dec. Processes* **108** (2), 315-329.
- MacLeod, W. B. (2002). "Complexity, Bounded Rationality and Heuristic Search," *Contributions Econ. Anal. Pol.* **1** (1), Article 8.
- Mandler, J. M. (2004). *The Foundations of Mind: Origins of Conceptual Thought*. New York, NY: Oxford University Press.
- Merlo, A., and Schotter, A. (1999). "A Surprise-Quiz View of Learning in Economic Experiments," *Games Econ. Behav.* **28** (1), 25-54.
- Nagel, R. (1995). "Experimental Results on Interactive Competitive Guessing," *Amer. Econ. Rev.* **85** (5), 1313-1326.
- Nissen, M. J., and Bullemer, P. (1987). "Attentional Requirements of Learning: Evidence from Performance Measures," *Cognitive Psych.* **19**, 1-32.
- Parks, C. D., and Rumble, A. C. (2001). "Elements of Reciprocity and Social Value Orientation," *Personality Social Psych. Bull.* **27** (10), 1301-1309.
- Rankin, F. W., Van Huyck, J. B., and Battalio, R. C. (2000). "Strategic Similarity and Emergent Conventions: Evidence from Similar Stag Hunt Games," *Games Econ. Behav.* **32** (2), 315-337.
- Rapoport, A., Seale, D. A., and Parco, J. E. (2002) "Coordination in the Aggregate Without Common Knowledge or Outcome Information," in *Experimental Business*

Research, (R. Zwick and A. Rapoport, Eds.), pp. 69-99. Norwell, MA: Kluwer Academic Press.

Reber, A. S. (1967). "Implicit Learning of Artificial Grammars," *J. Verbal Learning Verbal Behav.* **6**, 855-863.

Reber, A. S. (1989). "Implicit Learning and Tacit Knowledge," *J. Exper. Psych.: General*, **118** (3), 219-235.

Salmoni, A. W., Schmidt, R. A., and Walter, C. B. (1984). "Knowledge of Results and Motor Learning: A Review and Critical Reappraisal," *Psychological Bull.* **95** (3), 355-386.

Schmidt, R. A., Young, D. E., Swinnen, S., and Shapiro, D. C. (1989). "Summary Knowledge of Results for Skill Acquisition: Support for the Guidance Hypothesis," *J. Exper. Psych.: Learning, Memory, Cognition* **15** (2), 352-359.

Smith, V. L. (1991). "Experimental Economics: Behavioral Lessons for Microeconomic Theory and Policy," 1990 Nancy Schwartz Lecture, KGSM, Northwestern University.

Stahl, D. O. (2000). "Action Reinforcement Learning versus Rule Learning," Working Paper.

Stahl, D. O. (2000). "Rule Learning in Symmetric Normal-Form Games: Theory and Evidence," *Games Econ. Behav.* **32** (1), 105-138.

Trope, Y., and Liberman, N. (2003). "Temporal Construal," *Psychological Rev.* **110** (3), 403-421.

Weber, R. A. (2003). "Learning' with no feedback in a competitive guessing game," *Games Econ. Behav.* **44** (1), 134-144.

Winstein, C. J., and Schmidt, R. A. (1990). "Reduced Frequency of Knowledge of Results Enhances Motor Skill Learning," *J. Exper. Psych.: Learning, Memory, Cognition* **16** (4), 677-691.

Game A

	Left	Right
Top	90, 90	0, 50
Bottom	50, 0	50, 50

Game B

	Left	Right
Top	75, 35	40, 40
Bottom	65, 100	35, 10

Game C

	Left	Middle	Right
Top	30, 30	35, 40	100, 35
Middle	40, 35	45, 45	65, 40
Bottom	35, 100	40, 65	60, 60

Game D

	Left	Middle	Right
Top	50, 50	40, 75	75, 55
Middle	20, 25	50, 65	65, 45
Bottom	90, 55	25, 30	60, 80

Table 1. Games used in Experiment 1

	Game 1	Game 2	Game 3	Game 4	N
Session 1	D	A	C	B	18
Session 2	C	B	A	D	18
Session 3	A	D	B	C	18
Session 4	B	C	D	A	20

Table 2. Sequence of games by session (Experiment 1)

	Periods			
	1-5	6-10	11-15	16-20
Game A				
Top/Left	0.765	0.668	0.665	0.659
Bottom/Right	0.235	0.332	0.335	0.341
Game B				
Top (<i>Equil.</i>)	0.741	0.768	0.827	0.811
Bottom (<i>D1 viol.</i>)	0.259	0.232	0.173	0.189
Left (<i>D2 viol.</i>)	0.654	0.654	0.611	0.627
Right (<i>Equil.</i>)	0.346	0.346	0.389	0.373
<i>Equilibrium</i>	0.543	0.557	0.608	0.592
Game C				
Top/Left (<i>D2 viol.</i>)	0.200	0.186	0.224	0.189
Middle (<i>Equil.</i>)	0.603	0.624	0.630	0.678
Bottom/Right (<i>D1/D2 viol.</i>)	0.197	0.189	0.146	0.132
Game D				
Top	0.530	0.508	0.584	0.589
Middle (<i>Equil.</i>)	0.092	0.151	0.103	0.119
Bottom (<i>D2 viol.</i>)	0.378	0.341	0.314	0.292
Left (<i>D1 viol.</i>)	0.076	0.049	0.049	0.049
Middle (<i>Equil.</i>)	0.243	0.308	0.270	0.314
Right	0.681	0.643	0.681	0.638
<i>Equilibrium</i>	0.168	0.230	0.186	0.216

Table 3. Choice frequencies by 5-period blocks (Experiment 1)

Logistic regression with clustering of standard errors by subject

	Equilibrium <i>Game</i>				D1 violations <i>Game</i>			D2 violations <i>Game</i>		
	<i>A (B/R)</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>B</i>	<i>C</i>	<i>D</i>
Period	0.031 ^{***} (0.009)	0.015 ^{**a} (0.009)	0.020 ^{**a} (0.010)	0.016 ^{**a} (0.011)	-0.032 ^{***a} (0.013)	-0.035 ^{***a} (0.015)	-0.044 ^{*a} (0.031)	-0.008 (0.016)	-0.020 ^{**a} (0.010)	-0.026 ^{**a} (0.011)
Constant	-1.130 ^{***} (0.226)	0.146 (0.186)	0.345 ^{**} (0.177)	-1.556 ^{***} (0.199)	-0.974 ^{***} (0.292)	-1.262 ^{***} (0.190)	-2.399 (0.352)	0.646 ^{***} (0.248)	-0.345 [*] (0.173)	-0.433 [*] (0.262)
Obs.	1480	1480	1480	1480	740	1480	740	740	1480	740
L.L.	-912.1	-1007.8	-970.0	-430.7	-381.6	-661.7	-157.2	-484.8	-970.0	-468.0

Standard errors in parentheses

D1 violations: Choices inconsistent with first step of iterated dominance (B in Game B, B/R in Game C, L in Game D)

D2 violations: Choices inconsistent with second step of iterated dominance (L in Game B, B/R or T/L in Game C, B in Game D)

* – $p < 0.1$; ** – $p < 0.05$; *** – $p < 0.01$; ^a – one-tailed

Table 4. Change in behavior across periods within games (Experiment 1)

	Position of game within session			
	1st game	2nd game	3rd game	4th game
D1 violations	0.205 (37)	0.150 (38)	0.123 (37)	0.122 (36)
D2 violations	0.459 (37)	0.454 (38)	0.441 (37)	0.343 (36)

D1 violations: Choices inconsistent with first step of iterated dominance

D2 violations: Choices inconsistent with second step of iterated dominance

Number of subjects making relevant decision in parentheses

Table 5. Frequency of violations of iterated dominance by game position (Experiment 1)

	Logistic regression with clustering of standard errors by subject	
	<i>D1 Violations</i>	<i>D2 Violations</i>
Game Position	-0.220 ^{*a} (0.152)	-0.148 ^{**a} (0.090)
Constant	-1.206 ^{***} (0.396)	0.064 (0.223)
Obs.	2960	2960
L.L.	-1241.9	-2008.5

Standard errors in parentheses

D1: First step of iterated dominance (B in Game B, B/R in Game C, L in Game D)

D2: Second step of iterated dominance (L in Game B, T/L in Game C, B in Game D)

* – $p < 0.1$; ** – $p < 0.05$; *** – $p < 0.01$; ^a – one-tailed

Table 6. Change in behavior across games (Experiment 1)

Game order	Feedback treatment (first game)	Number of sessions	Number of subjects
IT → FT	Feedback	4	38
	No Feedback	4	38
FT → IT	Feedback	4	39
	No Feedback	4	37

Table 7. Description of treatments and sessions (Experiment 2)

Period	IT→FT		FT→IT	
	Feedback	No Feedback	Feedback	No Feedback
1 (Game 1)	41.2 (24.8)	45.3 (19.8)	163.4 (24.6)	171.8 (23.2)
10 (Game 1)	6.7 (15.7)	25.4 (17.0)	200.0 (0.2)	193.2 (14.6)
11 (Game 2)	163.5 (27.9)	176.1 (22.2)	43.3 (35.4)	30.9 (25.4)

Standard deviations in parentheses

Table 8. Average choices by treatment and game order (Experiment 2)

		Feedback (n=77)	No Feedback (n=75)
First step of iterated dominance (IT: $x_i > 70$; FT: $x_i < 130$)	Periods 1-3 (Game 1)	13 (17%)	14 (19%)
	Periods 8-10 (Game 1)	3 (4%)	3 (4%)
	<i>Periods 11-13 (Game 2)</i>	<i>17 (22%)</i>	<i>9 (12%)</i>
Second step of iterated dominance (IT: $x_i > 49$; FT: $x_i < 169$)	Periods 1-3 (Game 1)	38 (49%)	42 (56%)
	Periods 8-10 (Game 1)	3 (4%)	14 (19%)
	<i>Periods 11-13 (Game 2)</i>	<i>34 (44%)</i>	<i>23 (31%)</i>

Table 9. Number of subjects violating iterated dominance across periods (Experiment 2)

	Both treatments	Feedback treatment		No Feedback treatment	
No Feedback	-12.59 ^{***a} (4.58)				
Distance from Equil. in Period 10		0.19 (0.32)	0.19 (0.32)	0.43 ^{***a} (0.14)	0.41 ^{***a} (0.16)
Distance from Equil. in Period 1			0.05 (0.15)		0.04 (0.13)
Constant	39.92 ^{***} (3.22)	39.29 ^{***} (3.80)	37.35 ^{***} (7.01)	20.41 ^{***} (3.53)	19.31 ^{***} (5.06)
R²	0.05	0.01	0.01	0.09	0.11
Obs.	152	77	77	75	75

Standard errors in parentheses

* – $p < 0.1$; ** – $p < 0.05$; *** – $p < 0.01$; ^a – one-tailed

Table 10. Factors influencing distance from equilibrium of Period 11 choice (OLS) (Experiment 2)

Self-Explanation treatment	Feedback treatment	Number of sessions	Number of subjects
Explanation (E)	Feedback (F)	4	36
	No Feedback (NF)	5	42
No Explanation (NE)	Feedback (F)	6	57
	No Feedback (NF)	5	50

Table 11. Description of treatments and sessions (Experiment 3)

	Model 1	Model 2
No Feedback	-9.98** (5.64)	-11.06** (5.17)
Explanation	-14.68*** (4.49)	-14.88*** (4.04)
No Feedback x Explanation	14.22** (7.62)	14.41** (7.26)
Need for Cognition		-0.34** (0.18)
Constant	41.91*** (4.09)	68.60*** (14.66)
R²	0.04	0.06
Obs.	185	185

Standard errors in parentheses

* – $p < 0.1$; ** – $p < 0.05$; *** – $p < 0.01$; all one-tailed

Table 12. Factors influencing distance from equilibrium of Period 11 choice (OLS with clustering of standard errors by session) (Experiment 3)

Figure 1a. Mean choices across rounds (IT - FT), Exp. 2




Figure 1b. Mean choices across rounds (FT - IT), Exp. 2




Figure 2. Mean choices across rounds 1-10 (FT), Exp. 3




Figure 3. Mean Round 11 Choice (IT), Exp. 3

