

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

Roberto A. Weber

Department of Social & Decision
Sciences
Carnegie Mellon University

Scott Rick

Department of Operations and
Information Management
The Wharton School,
University of Pennsylvania

February 2008

Working Paper # 2008-02-20

Risk Management and Decision Processes Center
The Wharton School, University of Pennsylvania
3730 Walnut Street, Jon Huntsman Hall, Suite 500
Philadelphia, PA, 19104
USA
Phone: 215-898-4589
Fax: 215-573-2130
<http://opim.wharton.upenn.edu/risk/>

THE WHARTON RISK MANAGEMENT AND DECISION PROCESSES CENTER

Established in 1984, the Wharton Risk Management and Decision Processes Center develops and promotes effective corporate and public policies for low-probability events with potentially catastrophic consequences through the integration of risk assessment, and risk perception with risk management strategies. Natural disasters, technological hazards, and national and international security issues (e.g., terrorism risk insurance markets, protection of critical infrastructure, global security) are among the extreme events that are the focus of the Center's research.

The Risk Center's neutrality allows it to undertake large-scale projects in conjunction with other researchers and organizations in the public and private sectors. Building on the disciplines of economics, decision sciences, finance, insurance, marketing and psychology, the Center supports and undertakes field and experimental studies of risk and uncertainty to better understand how individuals and organizations make choices under conditions of risk and uncertainty. Risk Center research also investigates the effectiveness of strategies such as risk communication, information sharing, incentive systems, insurance, regulation and public-private collaborations at a national and international scale. From these findings, the Wharton Risk Center's research team – over 50 faculty, fellows and doctoral students – is able to design new approaches to enable individuals and organizations to make better decisions regarding risk under various regulatory and market conditions.

The Center is also concerned with training leading decision makers. It actively engages multiple viewpoints, including top-level representatives from industry, government, international organizations, interest groups and academics through its research and policy publications, and through sponsored seminars, roundtables and forums.

More information is available at <http://opim.wharton.upenn.edu/risk>.

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

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

Scott Rick
The Wharton School
University of Pennsylvania

February 20, 2008*

Abstract:

We explore a distinction among different ways in which people learn, in the context of games. Psychologists have long recognized two kinds of learning: one that is relatively shallow, operates through repeated association of stimuli and outcomes, and is domain-specific; and another that is deeper, involves the acquisition of concepts or meaning, and transfers 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.

In two experiments, we demonstrate that meaningful learning – the acquisition of iterated dominance – occurs in the absence of any feedback. We also demonstrate that, consistent with the psychological literature, meaningful learning transfers to new domains (games), and that such transfer does not occur with the kind of “strategy” learning routinely observed in repeated games played with consistent feedback.

* 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 Colin Camerer, George Loewenstein, Ted O’Donoghue, seminar participants at Texas A&M, Case Western, York University, and participants at the 2004 American Economic Association meetings, 2004 New and Alternative Directions on Learning Conference, 2005 Economic Science Association meetings, 2005 Judgment and Decision Making meetings, and 2006 Academy of Management meetings for helpful comments and suggestions.

I. Introduction

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 (Stahl 2000,a,b; Rankin et al., 2000), 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 performance improvement exceeds their ability to articulate or generalize their knowledge (Berry and Broadbent, 1984; Hayes and Broadbent, 1988; Nissen and Bullemer, 1987). 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. 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 abstract signaling games sequentially.

Given the correspondence between strategy and implicit learning, manipulations commonly employed in psychology experiments to prevent 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. The typical finding in psychology is that explicit learning is most likely to occur when participants are given minimal or delayed feedback (e.g., Salmoni et al., 1984; Winstein and Schmidt, 1990). Goodman (1998) attempted to account for such results by proposing that “external feedback” (e.g., learning the correct answer to a problem one just attempted to solve or, in a game-theoretic context, learning the outcome of a game just played) can divert attention from “response-produced feedback” (what is learned by thinking about and performing a particular task; cf. Anzai and Simon, 1979). Specifically, Goodman (1998, p. 231) proposes that external feedback inhibits learning by “distracting [subjects] from mentally picturing the task execution process.” Similarly, in a meta-analysis of studies that examine the effectiveness of external feedback, Kluger and DeNisi (1996) found that the ability of external feedback to induce explicit learning decreases as it moves attention away from features of the task at hand. Beyond merely distracting, feedback can also reduce the motivation to think carefully about the task at hand (Einhorn, 1980).

Recent theoretical work in economics on “fact-free” learning (Aragones et al., 2005) similarly suggests that people can learn meaningfully in the absence of new information. In their model, learning occurs by re-organizing one’s existing knowledge to find relationships and patterns in what one already knows. Although the psychological process by which such re-organization occurs is unclear (Aragones et al. (2005) do not cite, or ostensibly draw upon, any work from the psychology literature), the authors share the psychological intuition that meaningful learning can occur in the absence of feedback.

Thus, empirical research in psychology and theoretical research in economics suggests that meaningful learning occurs under limited or no feedback. The notion that

withholding feedback can induce meaningful learning is likely counterintuitive from the standard economic perspective, as withholding feedback in repeated games is normally intended to eliminate learning (e.g., Costa-Gomes and Crawford, 2006). However, there is already 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 (that varied the priming received by subjects to think about strategic elements of the game), significant learning occurred – behavior converged towards equilibrium.¹ Such learning cannot be considered strategy learning, which requires regular feedback. However, without examining whether such improvements in performance transfer to a new game, it is impossible to conclude that the feedback-free learning observed in Weber (2003) was meaningful.

Our experiments

We extend previous work by Weber (2003) by examining whether learning that occurs in repeated games played without feedback transfers to new domains, and whether such learning transfers more readily than does learning that occurs in the presence of consistent feedback.²

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 (Rapoport et al., 2002; Grether, 1980; Cason and Mui, 1998). However, none of these studies directly explores this phenomenon or focuses on learning.

² It is worth considering whether providing time for introspection would serve the same purpose as feedback-free repetition. Indeed, some models in economics allow players to develop improved reasoning ability by introspecting prior to playing a game (Goeree and Holt, 2002; MacLeod, 2002; Capra, 2003). Therefore, it is possible that subjects forced to think about a game, with some cues likely to prime deep thinking, would learn similarly to how our subjects learn with feedback-free repetition. However, it is also possible that time alone might not produce meaningful learning. For example, temporal construal theory in

In the first experiment, subjects play four normal-form games, 20 times each and in sequence, without any feedback until the end of the experiment. We find that subjects learn both within games (they converge towards equilibrium and decrease violations of iterated dominance) and that this learning transfers across games (they are significantly less likely to violate iterated dominance in later games than in earlier ones). Thus, subjects engage in learning that is consistent with meaningful learning.

In the second experiment, we directly compare the extent to which the two kinds of learning transfer to a new game. In the first stage, subjects play one of two versions of a competitive guessing (“p-beauty contest”) game (Nagel 1995), either with or without feedback. In the second stage, they play a variant of the same game that is superficially different but strategically very similar. We replicate the results of Ho et al. (1998), who find that learning in the first game does not transfer to the first round of the second game, when the first game is played with feedback (i.e., under strategy learning). However, we find significant immediate transfer to the second game when the first game is played without any feedback. Thus, we provide compelling evidence that the learning that occurs in the absence of feedback is meaningful – and that in at least one important way it is better than the kind of (strategy) learning that occurs with immediate feedback.

II. Theory

We use a simple model to illustrate how the two learning mechanisms operate in games, and how our experiments allow us to test for meaningful learning.

psychology (Trope and Liberman 2003) proposes that people think about distant situations abstractly, in terms of general goals, and postpone thinking about more concretely about specific aspects of the situation until it is imminent.

Let G denote a game with a set $N = \{1, \dots, n\}$ of players. In every period, each player $i \in N$ selects a strategy, $s_t^i \in S^i = \{s_1^i, \dots, s_{m_i}^i\}$, according to the mixed-strategy profile $\sigma_t^i = (p_{1it}, \dots, p_{m_it})$. Let $s_t = (s_t^1, \dots, s_t^n)$ represent the combined strategy choices and $\pi_t = \{\pi_t^1(s_t), \dots, \pi_t^n(s_t)\}$ the resulting payoffs.

Assume that a player's mixed strategy in a period is influenced by two components. First, H_{it} represents the history of outcomes observed by the player through period $t-1$. This includes all payoff information and any information regarding the actions of other players. For example, players may observe all strategy choices ($H_{it} = \{s_1, \dots, s_{t-1}\}$) or only their own payoffs ($H_{it} = \{\pi_1^i(s_1), \dots, \pi_{t-1}^i(s_{t-1})\}$).

The other component corresponds to a player's meaningful knowledge, K_{it} , which represents player i 's general knowledge about playing games. Such knowledge is given by a vector, $K_{it} = \{k_{i1}, \dots, k_{iJ}\}$, where each of the J elements, k_{ij} , is a binary variable indicating player i 's knowledge of a corresponding principle. For example, suppose k_{i1} refers to dominance. Then, $K_{it} = \{1, \dots\}$ might represent a player who in period t knows to avoid dominated strategies, while $K_{it} = \{0, \dots\}$ represents one who does not.

We allow for the possibility of learning through both kinds of knowledge. Therefore, a player's mixed strategy profile in any period t is the result of both that player's experienced history and meaningful knowledge in period t , or $\sigma_t^i = f(H_{it}, K_{it})$. A player's first-period strategy profile is a function only of the player's meaningful knowledge prior to the start of the game, or $\sigma_1^i = f(\emptyset, K_{i1})$.

Current models of strategy learning in games assume that all learning operates through H_{it} . That is, these models can be generally represented as $\sigma_t^i = f(H_{it}, \theta)$ or

$\sigma_t^i = f(H_{it}, K_i, \theta)$, with the latter case corresponding to models with heterogeneity in player-specific initial conditions, and where θ corresponds to a set of parameters characterizing the strategy learning process across periods. For example, in Camerer and Ho's (1999) Experience Weighted Attraction (EWA) model, which includes several other strategy learning models as special cases, player i 's mixed strategy profile in period 1 (σ_1^i) is determined by i 's idiosyncratic initial attraction to different strategies, which we would represent as K_i in our model,³ and by a parameter measuring noise (λ). Subsequent strategies are determined by the idiosyncratic initial conditions (K_i), experienced outcomes (H_{it}), the noise term (λ), and three additional parameters indicating sensitivity to different kinds of outcome information (θ), or $\sigma_t^i = f_{EWA}(H_{it}, K_i, \theta, \lambda)$.

In contrast to these models, we assume that K_{it} can change over the course of playing a game. That is, we allow the possibility that $k_{it-1j} \neq k_{ij}$, with the value changing from zero to one when a subject acquires the corresponding meaningful principle of game theory. Such variation can be important for determining behavior in a game.

For example, suppose that $K_{it} = \{k_{it1}, \dots\}$, where the first element corresponds to dominance, as in the example above. Then, suppose player i learns to identify and avoid dominated strategies from period $t-1$ to period t – meaning that $k_{it-11} = 0$ and $k_{it1} = 1$. Beginning with period t , that player will cease to play any strategies ruled out by K_{it} , i.e., any dominated strategies.

Let $\tilde{S}^i(j) \subseteq S^i$ represent the subset of i 's strategies permissible by the j -th element of K_{it} . If a player knows principle j , then some strategies may be ruled out. (In

³ In EWA, initial conditions (K_i) are characterized by players' initial attractions to each strategy and by the weight placed on these initial attractions. Camerer and Ho describe these as representing "pregame experience, either due to learning transferred from different games or due to introspection" (p. 830).

the above example, $\tilde{S}^i(1)$ is the set of undominated strategies in S^i .) Then, *meaningful learning* of principle j in period t^* implies that $k_{ij} = 0$ for all $t < t^*$, and $k_{ij} = 1$ and $s_t^i \in \tilde{S}^i(j)$ for all $t \geq t^*$.⁴

Our experiments explore the possibility of meaningful learning by eliminating the possibility of strategy learning. We do so by eliminating all outcome information (i.e., $H_{it} = \emptyset$) in repeated play. We also explore how such learning transfers by having subjects play different games.

To generate hypotheses for our experiments, we extend the above model to a situation where learning may transfer across games. Let $\Gamma = \{G, G'\}$ represent different games. Next, we describe two distinct experimental environments:

- E_1 corresponds to an environment in which first G is played for T periods and then G' is played for another T periods, and
- E_2 corresponds to an environment in which the order of games is reversed (first G' is played for T periods and then G is played for another T periods).

Let σ_{t,E_1}^i and s_{t,E_1}^i , respectively, represent player i 's mixed and pure strategy choices in period t of E_1 (which is game G if $t < T$ and game G' if $t \geq T$) and let H_{it,E_1} represent the history observed by player i through period t in E_1 .

Our model allows us to compare two sets of principal hypotheses. In each case, the null hypothesis predicts meaningful learning will not occur in our feedback-free

⁴ Standard strategy learning models could be easily modified to incorporate meaningful learning. For example, EWA (Camerer and Ho, 1999) operates through attractions to different strategies ($A_i^m(t)$). The attraction updating process could proceed as in current EWA, but with an attraction set to zero in any period where that strategy is ruled out by a meaningful principle in K_{it} . (i.e., $A_i^m(t) = 0$ if there exists a j such that $k_{ij} = 1$ and $s_m \notin \tilde{S}^i(j)$).

environments – either because only strategy learning is relevant in games or because our no-feedback treatment fails to produce meaningful learning. Conversely, the alternate hypothesis predicts what should occur if, as we expect, feedback-free repetition produces meaningful learning.

First, we address learning within a game. Under the null hypothesis, behavior should not change across periods if players receive no outcome or payoff information. Under the alternate hypothesis, behavior will change from early periods to late ones and such change will involve players behaving consistently with meaningful principles of game theory (for example, by ceasing to play dominated strategies).

H1(null): If $H_{it,E_1} = \emptyset$ for all $t \leq T$, then $\sigma_{t,E_1}^i = \sigma_{t',E_1}^i$ for all t and $t' \leq T$.

H1(alternate): If $H_{it,E_1} = \emptyset$ for all $t \leq T$, then there exists some t^* , such that for any $t' < t^*$ and $t'' \geq t^*$, $\sigma_{t',E_1}^i \neq \sigma_{t'',E_1}^i$. Moreover, there exists some j such that $s_{t',E_1}^i \notin \tilde{S}^i(j)$ for some $t' < t^*$ and $s_{t'',E_1}^i \in \tilde{S}^i(j)$ for all $t'' \geq t^*$.

The next two hypotheses address transfer across games. To do so, we define strategically similar games as games in which at least one meaningful principle of K_{it} applies non-trivially to both games. That is, $G, G' \in \Gamma$ are *strategically similar* if there exists some j such that $\tilde{S}^{i,G}(j) \subset S^{i,G}$ and $\tilde{S}^{i,G'}(j) \subset S^{i,G'}$. Under the null hypothesis, there should be no transfer from a strategically similar early game to a later game if the earlier game is played without feedback. However, the alternate hypothesis predicts immediate (first-period) transfer when the first game is played without feedback.

H2(null): Assume that G and G' are strategically similar. If $H_{it,E_1} = \emptyset$ for all $t \leq T$, then

$$\sigma_{T+1,E_1}^i = \sigma_{1,E_2}^i.$$

H2(alternate): Assume that G and G' are strategically similar. If $H_{it,E_1} = \emptyset$ for all $t \leq T$, then $\sigma_{T+1,E_1}^i \neq \sigma_{1,E_2}^i$. Moreover, there exists some j such that $s_{T+1,E_1}^i \in \tilde{S}^i(j)$ and, for some $t, t' \leq T$, $s_{t,E_1}^i \notin \tilde{S}^i(j)$ and $s_{t',E_2}^i \notin \tilde{S}^i(j)$.

That is, if transfer occurs due to meaningful learning, then we expect behavior in period $T+1$ of E_1 and period 1 of E_2 to differ, even though in both cases it is the first period of game G' . Additionally, we expect the player not to violate a principle in period $T+1$ of E_1 that is violated at least once in the first T periods of both E_1 and E_2 .

Experiments 1 and 2 test both sets of hypotheses above. In both experiments, subjects play games repeatedly without feedback. We explore the extent to which they learn within those games (H1) and the extent to which such learning transfers to a subsequent strategically similar game (H2). In both cases, we identify learning through systematic changes in behavior consistent with the acquisition of meaningful game-theoretic principles.

Additionally, Experiment 2 directly compares the amount of immediate transfer when the earlier game is played either with or without feedback. Studies in psychology show that immediate feedback can limit meaningful learning by interfering with people's ability or motivation to imagine a task's underlying structure (Einhorn, 1980; Goodman, 1998; cf. Salmoni et al., 1984; Winstein and Schmidt, 1990; Kluger and DeNisi, 1996). In our framework, this implies that if $H_{it} \neq \emptyset$, then $K_{it} = K_{it-1}$, meaning that no meaningful learning occurs with feedback. This produces the following hypothesis, which we explore in Experiment 2.

H3: Assume that G and G' are strategically similar. If, for all t , $H_{it,E_1} \neq \emptyset$, then

$$\sigma_{T+1,E_1}^i = \sigma_{1,E_2}^i.$$

That is, if subjects receive feedback in every period of an experiment, then we predict no transfer to the first period of a second game. This contrasts with H2(alternate) above, which predicts positive transfer when the first game is played without any feedback.

III. Experiment 1

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. The main purpose of this experiment was to demonstrate that subjects engage in learning when playing games repeatedly without feedback (H1) and that such learning transfers across games (H2), indicating that it is meaningful learning

The four games used in the experiment are shown in Table 1. The payoffs represent points, with 200 points equal to \$1. Game A is a symmetric stag-hunt game with three Nash equilibria: two pure strategy equilibria (*Top/Left* and *Bottom/Right*) 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. 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. 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 assignment of participant numbers.

Each subject sat at a computer monitor. At the beginning of a session, subjects received extensive instruction in how to interpret game matrices with generic payoffs.⁵ They were presented with several examples of how to calculate payoffs and allowed to ask questions. The experiment then proceeded to the first of the four 20-period blocks.

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.

⁵ Instructions and the complete datasets for both experiments are available from the authors.

B. Results

We first explore results at the game level, pooling across sessions, to test whether learning occurs during the 20 periods that a game is played (H1). We then turn our attention to whether there is any transfer of learning across games (H2).

1. Learning within games

The aggregate choice frequencies – by 5-period blocks – are presented in Table 3. While behavior appears to change across periods, we are interested in systematic changes consistent with learning.

In Game A, Nash equilibrium does not make a directional prediction since both strategy choices correspond to equilibrium behavior. However, the frequency of *Top/Left* choices decreases from the first block to the last, by roughly 11 percent, indicating that subjects are moving away from the Pareto-optimal equilibrium and toward the mixed strategy and risk-dominant equilibria (Harsanyi and Selten, 1988).⁶

Games B through D are all dominance-solvable and yield unique equilibrium predictions. In Game B, we observe a 5 percent increase in equilibrium play from the first to last five-period blocks; in Game C we see an 8 percent increase; and in Game D we see a 5 percent increase. Thus, in all games we find changes in behavior, consistent with movement towards Nash equilibrium.⁷

⁶ Comparing the first and last five-period blocks, using the individual average as the unit of observation, this change is statistically significant ($t_{73} = 3.22$, $p = 0.002$, two-tailed).

⁷ While the changes in behavior are not very large in magnitude (5 to 11 percent), all are at least marginally significant (A: $t_{73} = 3.22$, $p = 0.002$, B: $t_{73} = 1.56$, $p = 0.06$; C: $t_{73} = 2.11$, $p = 0.02$; D: $t_{73} = 1.79$, $p = 0.04$; respectively, all one-tailed except for A). Finally, since there is likely heterogeneity in whether subjects learn, the change in behavior among those actually engaging in learning is likely higher (see Panel A of Table 4).

Table 4 presents regressions exploring changes in behavior across periods. Both panels report logistic regressions, with the top panel including subject fixed effects (omitting subjects whose behavior was invariant on the dependent variable). The first four regressions in each panel 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 prediction is clear: the frequency of *Bottom* choices in Game B, *Bottom/Right* choices in Game C, and *Left* choices in Game D should decrease. A comparison of the first five periods with the last five in Table 3 indicates that this is true for all three games.⁸ Overall, the frequency of dominance violations decreased from 18 percent in the first five periods of a game, to 13 percent in the final five periods.

We can also perform a similar comparison for the second step of iterated dominance (D2). Acquiring this principle implies subjects should be less likely to play

⁸ These changes are significant for Games B ($t_{36} = 2.49$, $p = 0.009$, one-tailed) and C ($t_{73} = 2.34$, $p = 0.01$, one-tailed), but not for D ($t_{36} = 1.15$, $p = 0.13$, one-tailed) in which the frequency of dominance violations starts off close to zero (0.076).

Left in Game B, *Bottom/Right* or *Top/Left* in Game C, and *Bottom* in Game D.⁹ Looking at Table 3, we see that all three of these predictions are supported in a comparison of the first five periods versus the last five.¹⁰ Overall, the frequency of violations of the second step of iterated dominance decreases from 46 percent in the first five periods of a game to 39 percent in the last five.¹¹

The second and third sets of regressions in Table 4 explore violations of iterated dominance across periods. The frequency of dominated strategies decreases significantly for all three games. Violations of the second step of iterated dominance decrease across periods in all games, and significantly so in Games C and D.

2. Transfer across games

The above results reveal that subjects learn when they play games repeatedly without feedback, and that this learning involves the acquisition of iterated dominance. However, if subjects are really learning meaningful principles of iterated dominance, then we expect their acquisition to 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.

⁹ 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. Note that since Game C is dominance-solvable in two steps, complying with D2 is analogous to playing Nash equilibrium.

¹⁰ The difference is not significant for Game B ($t_{36} = 0.48$), but is significant for Game C ($t_{73} = 2.11$, $p = 0.02$, one-tailed) and for Game D ($t_{36} = 2.21$, $p = 0.02$, one-tailed).

¹¹ The third and fourth steps of iterated dominance only apply to Game D. The prediction for the third step is that column players should be less likely to play *Right*, which is supported in Table 3. The prediction for the fourth step – that Row players should play *Top* less often – does not hold. This is not surprising, since subjects do not frequently acquire more than two steps in experimental games (Camerer, 2003, Chapter 5).

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.¹² Therefore, a decrease in the frequency of violations means that subjects are more likely to obey the corresponding principle over the course of an experimental session – thus demonstrating transfer.

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.¹³ 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, as predicted by H2.

Another consequence of meaningful learning is that, once a subject learns to identify dominated or iteratively dominated strategies, she 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 iterated dominance. For the first step, 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 the second step of iterated dominance at least once in their first game (0 subjects with no violations). However, in the last game, 9 of 36 subjects never played

¹² The number of subjects in each cell varies because session 4 had a different number of participants.

¹³ For the third and fourth steps of iterated dominance, there is no decrease in the frequency of violations. This is expected, both because there is no evidence that subjects acquire these principles within games, and also because the principles only apply to Game D, which is only encountered once by each subject.

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.¹⁴

C. Discussion

Experiment 1 demonstrates a kind of learning inconsistent with strategy learning – a kind of learning that can operate without feedback. Moreover, this kind of learning appears to be meaningful, in that it transfers to new contexts and corresponds to important principles of game theory.

The transfer of learning is particularly important, as it indicates that the kind of learning occurring in Experiment 1 is more meaningful than that typically obtained in experiments with feedback. As we note in the Introduction, previous studies rarely find transfer across games (and when they do it is because there is some manipulation that makes learning “meaningful”). Here, we do so by limiting subjects’ ability to engage in strategy learning, and forcing them to substitute towards a more cognitive kind of learning in which they acquire important strategic principles. This suggests an important comparison: does learning transfer between games when earlier games are played without feedback (H2), even when it does not transfer when earlier games are played with feedback (H3)? Experiment 2 conducts this comparison.

¹⁴ There is also evidence of “complementary acquisition” of the two steps of iterated dominance (i.e., subjects need to acquire the first step (D1) before they can acquire the second (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 indicate that it is considerably easier to acquire D1 without acquiring D2 than vice versa, which is consistent with the fact that D2 requires understanding of D1.

IV. Experiment 2

This 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, the games are strategically similar, with respect to 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 \rightarrow FT) while the other half received the other order (FT \rightarrow 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. ($H_{it,E_1} \neq \emptyset$ for all t). 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 ($H_{it,E_2} = \emptyset$ for all $t \leq T$).

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 choices in the second game deviate from Nash equilibrium. In the Feedback treatment, we predict that subjects will engage in strategy learning and will therefore be unable to transfer what they learn to the first period of the second game (as Ho et al. found and as in H3). In the No Feedback treatment, however, we predict that subjects will learn in the first game (as in Weber, 2003, Experiment 1, and H1), and that this learning will be meaningful and will transfer to the first period of the second game (H2).

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 competitive guessing game. After having the opportunity to ask questions about the game, 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 came by and 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 consider whether learning occurred in the first game of the No Feedback treatment (H1). We then explore whether there is transfer of learning to the first period of the second game (H2), and whether such learning differs by feedback treatment (H3).

1. Learning in first competitive guessing game

Figures 1a and 1b present the median 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 No Feedback treatment, however, we also observe learning. In both games, average choices move significantly towards equilibrium between the first and tenth periods, though by not quite as much as with feedback (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 related 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.¹⁵

Transfer in the Feedback treatment. Recall that Ho et al. (1998) found no immediate transfer to the second game when the first game was played with feedback. We replicate this finding. 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

¹⁵ 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 see no significant differences between the Feedback and No Feedback treatments (IT: $t_{74} = 0.79$; FT: $t_{74} = 1.53$).

Feedback treatment (at the beginning of the second game) are 43.3 and 163.5. Thus, there is no immediate transfer for the FT→IT transition and there is slight *negative* transfer for the IT→FT transition.¹⁶

Moreover, Table 9 reveals no decrease in the Feedback treatment of 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. 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, as predicted by H3.^{17 18}

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 (FT). The average choices in the No Feedback treatment at the start of the second game

¹⁶ Ho et al. (1998) also found negative immediate transfer. This is potentially because some subjects apply a flawed heuristic (such as “choose the lowest number”) that might have worked well in the first game but does poorly in the second. For instance, only 2 of 38 subjects playing the IT game at the beginning of the experiment (Period 1 in the Feedback treatment) chose 100 (which is the Nash equilibrium for the FT game), but 7 of 39 did so when playing the IT game after playing the FT game.

¹⁷ 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 toward equilibrium. Median 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. In the Feedback treatment, this is consistent with the transfer of a heuristic that choices converge towards the closest boundary rather than meaningful learning. This is supported, as we note later, by evidence that 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 (see Table 10).

¹⁸ In later periods of the FT→IT treatment, choices increase, moving away from Nash equilibrium. This is because, in a few sessions, subjects began experimenting with choices of 100 (perhaps out of boredom or to try to gain a strategic advantage – this was also observed in Weber (2003)). In some cases this occurred in consecutive periods, raising the average substantially, and leading several subjects to increase their choices in subsequent periods. Moreover, in the No Feedback treatment, these choices of 100 started to occur somewhat earlier in two sessions (contributing to the slower convergence towards zero).

(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 ($\chi^2(1) = 9.80$, $p = 0.002$).¹⁹

Thus, as we anticipated in H2 and H3, 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. However, aside from demonstrating that greater transfer occurs in the No Feedback treatment, we can also explore whether such transfer is consistent with subjects learning in the first game.

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 feedback on immediate transfer: Period 11 choices are significantly farther from equilibrium, on average, in the Feedback treatment than 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

¹⁹ 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 significant ($\chi^2(1) = 3.46$, $p < 0.07$).

relationship between Period 10 and Period 11 choices in the Feedback treatment, but there is a strong and statistically significant relationship for 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.²⁰

C. Discussion

Consistent with 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. As in Experiment 1, this transfer is consistent with subjects acquiring the first two steps of iterated dominance by the end of the first game, and then applying those principles at the beginning of the second game. In the No Feedback treatment, it is primarily those subjects who “learn” in the first game that exhibit immediate transfer to the second game. Therefore, our results are consistent with meaningful learning in the No Feedback treatment and the occurrence of strategy learning in the Feedback treatment, as predicted by psychological research and our hypotheses.

²⁰ We observe a similar pattern if we consider violations of iterated dominance. 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 a principle in Periods 11-13 is also higher in the Feedback treatment (32 of 74; 58 percent) than in No Feedback (14 of 61; 23 percent; $\chi^2(1) = 6.13$, $p = 0.01$).

V. Conclusion

Game theorists interested in how people learn in strategic environments have long focused on understanding how players figure out what strategy works best in a specific game, a process we refer to as strategy learning. However, categorizing this process as a “meaningful” kind of learning is questionable, as movement toward equilibrium in repeated games played without feedback rarely transfers to new strategically similar games (e.g., Ho et al., 1998). The ability to transfer and generalize knowledge to new domains is often considered a hallmark of such deeper learning (Schmidt, 1991; Schmidt and Bjork, 1992). Although some economic research has focused on such meaningful learning (Stahl 2000,a,b; Rankin et al., 2000), there is still a very limited understanding of the conditions that are most likely to stimulate meaningful learning in games.

In this paper, we attempt to advance economists’ understanding of meaningful learning by importing relevant insights from psychology, a field that has long recognized a distinction between two kinds of learning. One factor commonly found to influence the development of meaningful learning is feedback about performance. Specifically, immediate feedback limits meaningful learning by diverting attention from information gleaned from thinking about and performing a task and interfering with people’s ability or motivation to think about a task’s underlying structure (Einhorn, 1980; Goodman, 1998; Salmoni et al., 1984; Winstein and Schmidt, 1990; Kluger and DeNisi, 1996). In our experiments, we use feedback-free environments in which subjects play games repeatedly without any outcome or payoff information to foster meaningful learning and to hinder strategy learning. We demonstrate the depth of such meaningful learning by

showing that it produces immediate transfer to new games and that it operates through subjects acquiring iterated dominance.

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, our demonstration that such learning occurs with 10 or 20 repetitions of a game shows that there are conditions under which attempting to eliminate learning by eliminating feedback may fail.

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. The simple model presented here does not explain precisely how meaningful learning occurs. 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 kinds 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. Focusing on learning that transfers, as we do here, is an important way to understand the kind and depth of learning taking place, and the extent to which learning is likely to prove useful across domains.

It is also worth mentioning that the procedure we use – feedback-free repetition – is only one way to develop meaningful learning, and there are other, potentially better, methods. 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 (Schmidt et al., 1989; Schooler and Anderson, 1990; Erev et al., 2006). Therefore, subjects in our experiments might demonstrate even greater meaningful learning if presented with feedback after some such “optimal” delay. Moreover, interventions unrelated to the timing of feedback also influence the kind of learning that occurs and the degree to which such learning transfers. For instance, both making decisions in groups and verbalizing mental processes facilitate meaningful learning (Cooper and Kagel, 2003 and 2005; Budescu and Maciejovsky, 2005; Alevan and Koedinger, 2002).

Thus, future work should examine the limits and robustness of our effects to other contexts and learning treatments. There are likely highly complex situations – such as learning a programming language – in which feedback may be essential for developing meaningful learning. It is also worth examining what game-theoretic principles other than dominance transfer across domains in an experiment like ours. Also, in addition to situational factors, some individual differences may also moderate the relationship between feedback and the development of meaningful learning. For example, feedback

may be less likely to prevent meaningful learning among individuals that have a tendency to engage in cognitive reflection (Frederick, 2005; cf. Cacioppo and Petty, 1982).

Our experiments demonstrate a distinction in how people learn, and the consequences of such a distinction, in an economic context. However, much remains to be done in economics to fully incorporate this distinction. Our work is a useful first step.

References

- Aleven, V., and Koedinger, K. R. (2002). "An Effective Meta-Cognitive Strategy: Learning by Doing and Explaining with a Computer-Based Cognitive Tutor," *Cognitive Sci.* **26** (2), 147-179.
- 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.
- Aragones, E., Gilboa, I., Postlewaite, A., and Schmeidler, D. (2005). "Fact-Free Learning," *Amer. Econ. Rev.* **95** (5), 1355-1368.
- 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.
- Budescu, D. V., and Maciejovsky, B. (2005). "The Effect of Payoff Feedback and Information Pooling on Reasoning Errors: Evidence from Experimental Markets," *Manage. Sci.* **51** (12), 1829-1843.
- Cacioppo, J. T., and Petty, R. E. (1982). "The Need for Cognition," *J. Personality Social Psych.* **42** (1), 116-131.
- 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.
- 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.
- 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.
- Frederick, S. (2005). "Cognitive Reflection and Decision Making," *J. Econ. Perspect.* **19** (4), 24-42.
- Fudenberg, D., and Levine, D. (1998). *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- 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.
- 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.

- 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.
- 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.
- 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.
- 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. (1991). "Frequency Augmented Feedback can Degrade Learning: Evidence and Interpretations," in *Tutorials in Motor Neuroscience* (J. Requin and G.E. Steimach, Eds.), pp. 59-75. London: Kluwer.
- Schmidt, R. A., and Bjork, R. A. (1992). "New Conceptualizations of Practice: Common Principles in Three Paradigms Suggest New Concepts for Training," *Psychological Sci.* **3** (4), 207-217.
- 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.

Schooler, L. J., and Anderson, J. R. (1990). "The Disruptive Potential of Immediate Feedback," in *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, Cambridge, MA, pp. 702-708.

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)

Panel A. Logistic regression with subject fixed effects

	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	0.084 ^{***} (0.017)	0.034 ^{***a} (0.014)	0.034 ^{***a} (0.012)	0.021 ^{**a} (0.013)	-0.097 ^{***a} (0.028)	-0.051 ^{***a} (0.015)	-0.050 ^{**a} (0.030)	-0.012 (0.016)	-0.034 ^{***a} (0.012)	-0.050 ^{***a} (0.019)
Obs. (N)	640 (32)	880 (44)	1120 (56)	940 (47)	280 (14)	860 (43)	260 (13)	600 (30)	1120 (56)	460 (23)
L.L.	-279.4	-382.7	-484.1	-430.7	-87.3	-344.1	-87.3	-277.7	-484.1	-203.1

Panel B. 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. (N)	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 subject fixed effects		Logistic regression with clustering of standard errors by subject	
	<i>D1 Violations</i>	<i>D2 Violations</i>	<i>D1 Violations</i>	<i>D2 Violations</i>
Game Position	-0.142 ^{**a} (0.075)	-0.127 ^{**a} (0.056)	-0.220 ^{*a} (0.152)	-0.148 ^{**a} (0.090)
Constant			-1.206 ^{***} (0.396)	0.064 (0.223)
Obs. (N)	2040 (51)	2680 (67)	2960	2960
L.L.	-775.9	-1393.0	-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	
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	27.34 ^{***} (3.26)	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)

Figure 1a. Median choices across rounds (IT - FT)




Figure 1b. Median choices across rounds (FT - IT)

