Incorporating students' probabilistic expectations into a peer-driven tutoring game*

Ari Bader-Natal

ari@cs.brandeis.edu
DEMO Lab, Brandeis University, Waltham, MA 02134

Abstract. Games provide a promising mechanism for intelligent tutoring systems in that they offer means to influence motivation and structure interactions. We have designed and released several game-based tutoring systems in which students learn to identify the best game strategies to adopt, and, in doing so, create for each other increasingly productive learning environments. Here, we first detail the core game underlying our deployed systems, designed to leverage human intelligence in tutoring systems through the tutor's identification of "appropriate" challenges for their tutee. While this game works well for task domains in which problem difficulty is known, it cannot be applied to domains if nothing is known about a problem beyond its correct solution. We introduce a second, more robust, game here capable of addressing this larger set of task domains. By incorporating player-generated probability estimates (in place of a difficulty metric), we show that a game can be designed to simultaneously elicit best-effort responses from tutees, honest statements of probability estimates from tutees, and appropriate challenges from tutors. We derive a set of constraints on the parameterized version of this game necessary for rational players to converge on this "Teacher's Dilemma" learning environment. Beyond providing a foundation for future tutoring systems, this work offers a new mechanism with which to simultaneously leverage and enhance the knowledge of peer learners.

1 Introduction

Games provide a promising mechanism for learning, in that they offer means to engage students, enable realistic simulations, and provide motivational structure [1–5]. Over the past several years, we have designed and released several game-based learning systems. The focus of these has been on games in which student grapple to identify the most effective strategies for play, and in doing so, they – perhaps unknowingly – create for one another increasingly productive learning environments.

In this paper, we first detail the core game underlying our deployed systems, designed to leverage human intelligence in tutoring systems through the tutor's identification of "appropriate" challenges for their tutee. We identify one shortcoming of the model: it cannot be applied to domains if nothing is known about a

* Brandeis CS Tech Report CS-08-269

problem beyond its correct solution, and introduce a second, more robust, game that compensates for the lack of a difficulty metric by incorporating player-generated probability estimates. We then show how a game can be designed to simultaneously elicit best-effort responses from tutees, honest statements of probability estimates from tutees, and appropriate challenges from tutors. The constraints on a parameterized model of this game are derived such that these criteria will always be met, yielding a robust second model for peer-driven learning in games.

2 Context: A Game-theoretic approach to learning games

We approach this question by focusing on the potential of games to motivate and direct the actions of learners. We use the framework of game theory [6] (and draw also from the literature in mechanism design [7, 8]) to describe, construct, and analyze multi-player games that, at equilibrium, provide environments highly conducive to learning. Each of these games forces peer-tutors to try to understand the abilities of their tutee and challenge them appropriately, by offering them what we call a "Teacher's Dilemma." The resulting sustained engagement with appropriate challenges provides the tutees opportunity to learn. One key assumption that we make in approaching game-driven learning is that players will act rationally, adopting strategies that they believe will maximize their utility (as measured by their score in the game). While this assumption may be entirely reasonable for game interactions among autonomous agents [9], it is not necessarily appropriate for human players. But while many of the results from game theory are lost when players are no longer assumed to be fully rational, a body of work on Evolutionary Game Theory has recovered many of these concepts by replacing rationality with repeated play over time [10]. While we assume rationality and only analyze one-shot games in this paper, the tutoring system in which we implement these games involve repeated play among pairs of players, and we believe that the dominant strategies identified below will remain as such for sub-rational players repeatedly playing with the same partner over time.

While game theory may model external decisions as due to Nature, and provides a construct – the "trembling hand" – to model situations in which a player does not have full control over the actions that they select [10], neither of these is entirely sufficient to represent the situation of the Student. While a Student may not have the ability to submit a guaranteed-correct response, they always *can* select a guaranteed-*incorrect* response. More generally, the Student's available strategy options are limited by their ability in one direction but not in the other. Depending on how the payoff functions are formulated, the Student may have a strategic reason to do just this, and purposefully act "below" their abilities, as a form of "gaming the system" [11, 12]. The games analyzed below have been designed both with these notions of ability-limited play and strategic under-performance in mind.

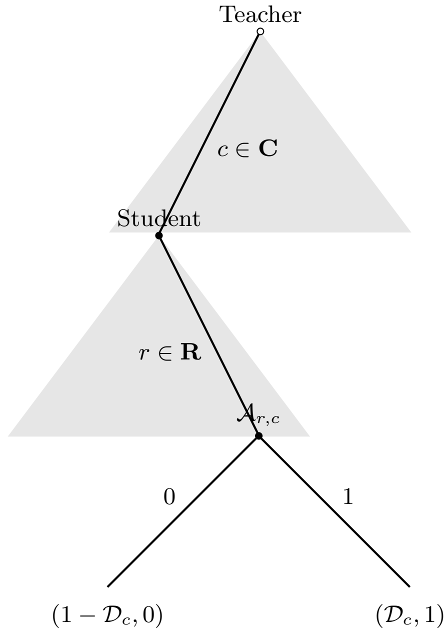

Fig. 1. TD-$\mathcal{D}$: A *difficulty*-based Teacher's Dilemma game. The Teacher first selects a challenge from the space of legal challenges. The Student select a response from the space of legal responses. The accuracy of the response is assessed (objectively), and the Teacher and Student are rewarded as a function of problem difficulty and response accuracy.

2.1 TD-$\mathcal{D}$: The *difficulty*-based Teacher's Dilemma game

Our current tutoring systems incorporate a game played between two players, in which one plays the role of Teacher and the other plays the role of Student.[1] The Teacher selects a challenge, the Student attempts to solve it, and both players receive payoffs based on challenge difficulty and response accuracy. The extensive form representation of this game, which we will refer to as TD-$\mathcal{D}$, is shown in Figure 1. We present this game primarily as a point of reference for the following sections. In this game, there is an objective (and assumed-correct) notion of problem difficulty: "Problem difficulty" $\mathcal{D}_c$ is an externally-defined estimate of the probability that the Student will provide an incorrect response to the challenge. Furthermore, we assume a "well-defined" task domain: for every legal challenge c and legal response r, accuracy $\mathcal{A}$ can be assessed (i.e. $\mathcal{A}_{r,c}$ is known.) The task domain, like with any Intelligent Tutoring System, is likely restricted to one topic that the students are trying to learn.

[1] These types do not correspond to classroom roles. In many of our existing tutoring systems, the two peers are constantly reversing roles. The names are only meant to indicate which player poses challenges and which player responds.

Definition 1. *An "appropriate challenge" for a Student is defined in terms of the probability of the Student providing a correct response to it. Here, we let* $\mathrm{P}_{appropriate}[\mathcal{A}_{r,c} = 1] = 0.5$ *(although other values can be adopted, given minor modifications to the players' payoff functions.)*

Definition 2. *A two-player game between a Teacher and a Student is considered a "Teacher's Dilemma" if the strategy that maximizes the Teacher's expected utility is to always provide the Student with "appropriate challenges."*

While we have previously argued that the structure of the difficulty-based game is sufficient to converge on appropriate challenges being posed, this argument assumes the existence of an accurate difficulty metric. While the existence of such a metric may be reasonable to assume or approximate for some task domains, it remains a limitation on the game's value for tutoring system applications. In response, we now introduce a second game, designed to provide a Teacher's Dilemma without the need for any external difficulty metric.

3 TD-$\mathcal{E}$: An *expectation*-based Teacher's Dilemma game

The intuition behind the expectation-based game that we will discuss is that the missing problem "difficulty" information (i.e. probability-based estimation of anticipated response accuracy) can come from a different source: the players themselves. Where TD-$\mathcal{D}$ assumed that the game can accurately assess difficulty and accuracy, TD-$\mathcal{E}$ instead assumes that the game can only assess accuracy, but that the players themselves can express a probability of expected accuracy. For the Teacher, this amounts to answering a student modeling question: "With what probability do you expect the tutee to accurately respond to the challenge question?" For the Student, it involves metacognitive reasoning: "With what probability do you believe your response to the challenge question is accurate?" However, the inherent complication in asking players to answer such questions is that players are free to misrepresent their true beliefs, and can be expected to do so whenever it provides strategic advantage. We therefore define two different notions of expectation:

1. A player's "true expectation" $\dot{\mathcal{E}}$ is an estimate of the probability that the Student's will be able to provide a correct response to the challenge, as privately believed by that player.
2. A player's "stated expectation" $\mathcal{E}$ is an estimate of the probability that the Student's will be able to provide a correct response to the challenge, as publicly stated by that player during the course of game-play.

In order to show that an expectation-based game is a Teacher's Dilemma (i.e. that the Teacher is motivated to choose appropriate challenges for the Student), we must first show that players have no incentive to misrepresent their true expectations when stating them. We must also be sure that Students are not purposefully providing known-incorrect responses (as discussed in Section **??**) in order to increase the accuracy of their stated probability-expectations.

Only if both of these conditions are met can we look at the Teacher's incentives in order to determine whether the game is a Teacher's Dilemma. Thus, an expectation-based Teacher's Dilemma game must simultaneously motivate three sub-strategies:

1. *Best-effort responses*: The Student is motivated to respond to the best of their abilities (i.e. they will never benefit from providing an incorrect response if the correct response is known.) A Student's payoff function π_s is said to be "effort-dominant" if, when asked to provide a response to a posed challenge, the strategy that produces the highest expected utility is for the Student to always provide their "best-effort" response (i.e. The Student has no alternative response for which they expect a higher likelihood of accuracy.)
2. *Honest statement of expectations*: Both the Teacher and the Student are motivated to state their true expectations (i.e. they never benefit from intentionally mis-representing their true beliefs.) A Student's payoff function π_s is said to be "truth-dominant" if, when asked to state their estimated probability of response accuracy, the strategy that always produces the highest expected utility is always honestly sharing one's true beliefs.
3. *Appropriate challenges*: The Teacher should be motivated to select challenges of "appropriate difficulty" for their tutee. A Teacher's payoff function π_t is said to be "appropriateness-dominant" if, when asked to select a challenge to pose to their tutee, the strategy that always produces the highest expected utility is select the one for which the Student expects to have a 50% likelihood of answering correctly.

3.1 Examples: Two expectation-based games

Figures 2 and 3 show two expectation-based games, differing only in the player payoffs. The first is not a Teacher's Dilemma, the second is. Figure 2 was designed to mimic the payoffs in the difficulty-based Teacher's Dilemma game shown in Figure 1. In this version, we can see that while the Student receives a higher payoff for correct responses (and is thus motivated to answer correctly whenever possible), the Student is provided with no motivation to state their expectations of problem difficulty *honestly*. The Student's payoff is entirely independent of the Student's stated expectation. As a result, there is no reason to believe that the statement is anything other than randomly generated, and so the Teacher's payoff structure (which is dependent on the Student's stated expectation) no longer motivates appropriate challenges. This game, therefore, does not meet the criteria of a Teacher's Dilemma. In Figure 3, on the other hand, the Student's payoff still favors correct responses (and thereby favors best-effort responses.) It is dependent on the Student's stated expectation, but in such a way that the Student's highest expected utility results from honestly stating their true expectations (i.e. $\mathcal{E}_s = \dot{\mathcal{E}}_s$.) This can be seen in Figure 5. For the Teacher, the expected utility of selecting an appropriate challenge (for which $\mathcal{E}_s = 0.5$) is higher than any other selection, as seen in Figure 4. Thus, all three conditions are satisfied and the game meets the criteria of being a Teacher's Dilemma.

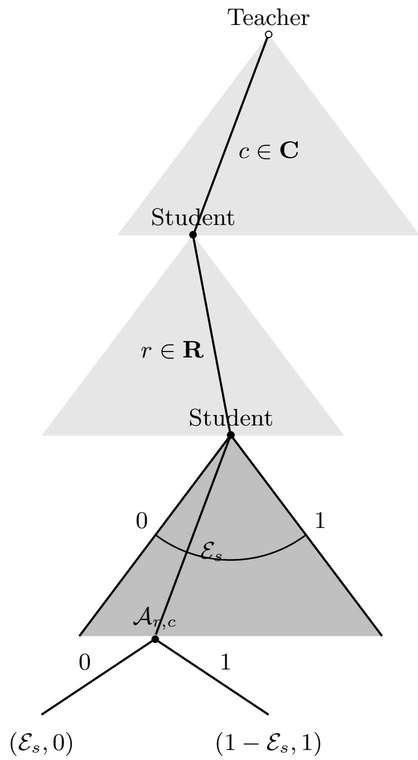

Fig. 2. Example Game $\mathcal{E}$-1: An expectation-based game that, despite it's similarity to the TD-$\mathcal{D}$ game shown in Figure 1, does not create a Teacher's Dilemma.

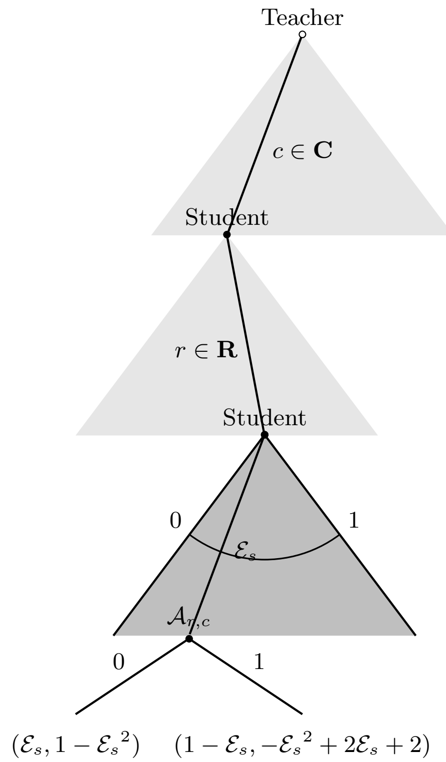

Fig. 3. Example Game $\mathcal{E}$-2: An expectation-based game that *does* create a Teacher's Dilemma.

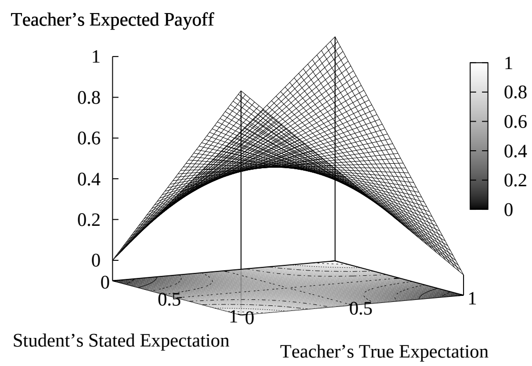

Fig. 4. The Teacher's expected utility for Example Game $\mathcal{E}$-2 (in Figure 3) is a function of their true expectation and the Student's stated expectation of response accuracy.

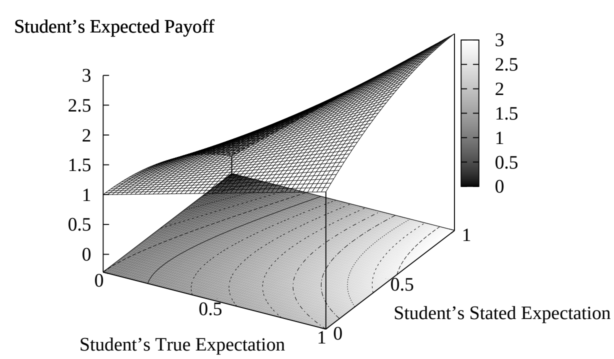

Fig. 5. The Student's expected utility for Example Game $\mathcal{E}$-2 (in Figure 3) is a function of their true and stated expectations of response accuracy.

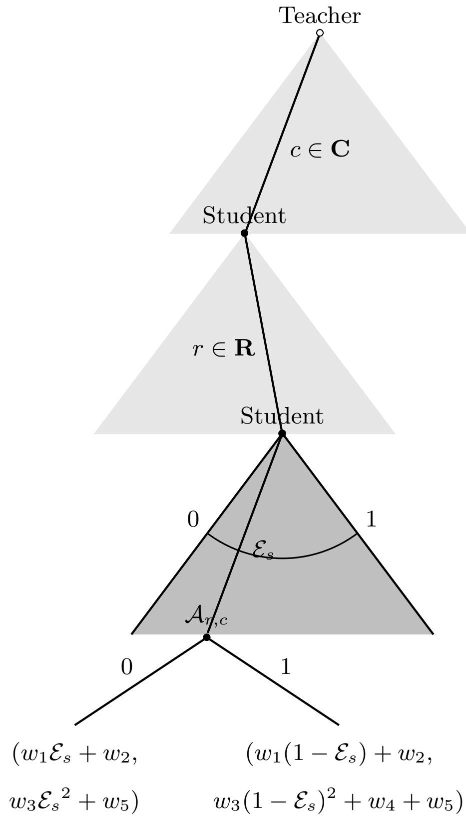

Fig. 6. TD-$\mathcal{E}$: A parameterized *expectation*-based Teacher's Dilemma.

3.2 Necessary conditions for a parameterized TD-$\mathcal{E}$ game to become a Teacher's Dilemma

Figure 6 shows a generalization of Example Game $\mathcal{E}$-2 (in Figure 3.) The payoffs in this game are parameterized, to allow for some flexibility in the actual payoff values.

$$\pi_t = \begin{cases} w_1\mathcal{E}_s + w_2 & \text{if } \mathcal{A}_{r,c} = 0 \\ w_1(1-\mathcal{E}_s) + w_2 & \text{if } \mathcal{A}_{r,c} = 1 \end{cases} \tag{1}$$

$$\pi_s = \begin{cases} w_3{\mathcal{E}_s}^2 + w_5 & \text{if } \mathcal{A}_{r,c} = 0 \\ w_3(1-\mathcal{E}_s)^2 + w_4 + w_5 & \text{if } \mathcal{A}_{r,c} = 1 \end{cases} \tag{2}$$

By definition of true expectation, the Teacher believes that $\mathrm{P}[\mathcal{A}_{r,c} = 1] = \dot{\mathcal{E}}_t$ and $\mathrm{P}[\mathcal{A}_{r,c} = 0] = 1 - \dot{\mathcal{E}}_t$, and the Student believes that $\mathrm{P}[\mathcal{A}_{r,c} = 1] = \dot{\mathcal{E}}_s$ and

$P[\mathcal{A}_{r,c} = 0] = 1 - \dot{\mathcal{E}}_s$. The expected utility of each player can be stated as a function of these probabilities and this payoff, as shown below and in Figures 4 and 5.

For the Teacher:

$$\begin{aligned} \mathrm{E}_{\pi_t} &= \sum_{i=0}^{1} \mathrm{P}\left[\mathcal{A}_{r,c} = i\right] (\pi_t | \mathcal{A}_{r,c} = i) \\ &= \left(1 - \dot{\mathcal{E}}_t\right) (w_1 \mathcal{E}_s + w_2) + \left(\dot{\mathcal{E}}_t\right) (w_1 (1 - \mathcal{E}_s) + w_2) \\ &= w_1 \left(\dot{\mathcal{E}}_t + \mathcal{E}_s - 2\dot{\mathcal{E}}_t \mathcal{E}_s\right) + w_2 \end{aligned} \tag{3}$$

For the Student:

$$\begin{aligned} \mathrm{E}_{\pi_s} &= \sum_{i=0}^{1} \mathrm{P}\left[\mathcal{A}_{r,c} = i\right] (\pi_s | \mathcal{A}_{r,c} = i) \\ &= \left(1 - \dot{\mathcal{E}}_s\right) \left(w_3 \mathcal{E}_s^2 + w_5\right) + \left(\dot{\mathcal{E}}_s\right) \left(w_3 (1 - \mathcal{E}_s)^2 + w_4 + w_5\right) \\ &= w_3 \mathcal{E}_s^2 - 2w_3 \dot{\mathcal{E}}_s \mathcal{E}_s + w_3 \dot{\mathcal{E}}_s + w_4 \dot{\mathcal{E}}_s + w_5 \end{aligned} \tag{4}$$

We can show that for TD-$\mathcal{E}$, the tutee's payoff function π_s can be truth-dominant and can be effort-dominant, and the tutor's payoff function π_t can be appropriateness-dominant. When the conditions necessary for each of these are simultaneously satisfied, TD-$\mathcal{E}$ becomes a Teacher's Dilemma. We now show the conditions necessary for this to happen.

Theorem 31 *π_s in TD-$\mathcal{E}$ is truth-dominant if $w_3 < 0$, regardless of the values of all other w_i.*

Proof. Depending on the value of w_3, the conditions required to satisfy this fall into one of three cases:

(i) When $w_3 < 0$, the expected utility value surface (as a function of stated and true expectations) opens downward. The maximum expected utility for each true expectation occurs when the partial derivative of E_{π_s} (from Equation 4) with respect to $\mathcal{E}_s$ is zero:

$$\begin{aligned} 0 &= \frac{\partial \mathrm{E}_{\pi_s}}{\partial \mathcal{E}_s} \\ 0 &= 2w_3 \mathcal{E}_s - 2w_3 \dot{\mathcal{E}}_s \\ \mathcal{E}_s &= \dot{\mathcal{E}}_s \end{aligned} \tag{5}$$

The Student attempting to maximize expected utility will therefore always state their expectation ($\mathcal{E}_s$) exactly as they truly believe it ($\dot{\mathcal{E}}_s$.) Thus, π_s is truth-dominant.

(ii) When $w_3 = 0$, $\mathrm{E}_{\pi_s} = w_4 \dot{\mathcal{E}}_s + w_5$. Since this is independent of $\mathcal{E}_s$, there is no strategic value in truthfully stating $\mathcal{E}_s$. In this case, π_s is not truth-dominant.

(iii) When $w_3 > 0$, the expected utility value surface (as a function of stated and true expectations) opens upward, rendering the $\mathcal{E}_s = \dot{\mathcal{E}}_s$ from case (i) into a *minimum* rather than a *maximum*. As such, π_s is not truth-dominant in this case.

Theorem 32 *π_s in TD-$\mathcal{E}$ is effort-dominant if $w_4 > w_3$, $w_4 > 0$, and $w_3 + w_4 > 0$, regardless of the values of all other w_i.*

Proof. Given a posed challenge c, call the tutee's potential response that they believe most likely to be correct their "best-effort" response, r_{BE}. Call another response that they believe less likely to be correct their "less-effort" response, r_{LE}. In the expectation notation used above, we denote this as $\dot{\mathcal{E}}_s(r_{LE}, c) < \dot{\mathcal{E}}_s(r_{BE}, c)$. We abbreviate this by using a tilde to denote "less-effort,"so we can restated this as $\tilde{\mathcal{E}}_s < \dot{\mathcal{E}}_s$. A payoff function is *effort-dominant* if the expected utility associated with a "best-effort" response is greater than that associated with any "less-effort" response (i.e. $\mathrm{E}_{\pi_s}(r_{BE}, c) \geq \mathrm{E}_{\pi_s}(r_{LE}, c)$.) We can compare these expected utility values for the payoff function from Equation 2 in order to identify the conditions that must hold for the Student's payoff to be effort-dominant:

$$\mathrm{E}_{\pi_s}(r_{BE}, c) > \mathrm{E}_{\pi_s}(r_{LE}, c)$$
$$\left[w_3\mathcal{E}_s^2 - 2w_3\dot{\mathcal{E}}_s\mathcal{E}_s + w_3\dot{\mathcal{E}}_s + w_4\dot{\mathcal{E}}_s + w_5\right] > \left[w_3\mathcal{E}_s^2 - 2w_3\tilde{\mathcal{E}}_s\mathcal{E}_s + w_3\tilde{\mathcal{E}}_s + w_4\tilde{\mathcal{E}}_s + w_5\right]$$
$$(-2w_3\mathcal{E}_s + w_3 + w_4)\left(\dot{\mathcal{E}}_s - \tilde{\mathcal{E}}_s\right) > 0 \tag{6}$$

Given that $\dot{\mathcal{E}}_s > \tilde{\mathcal{E}}_s$, we can say that there exists some $\Delta > 0$ such that $\dot{\mathcal{E}}_s = \tilde{\mathcal{E}}_s + \Delta$. Substituting this into 6, we get:

$$(-2w_3\mathcal{E}_s + w_3 + w_4)\left((\tilde{\mathcal{E}}_s + \Delta) - \tilde{\mathcal{E}}_s\right) > 0$$
$$(-2w_3\mathcal{E}_s + w_3 + w_4)\,\Delta > 0$$
$$w_3\mathcal{E}_s < \frac{w_3 + w_4}{2} \tag{7}$$

Depending on the sign of w_3, the conditions required to satisfy this inequality fall into one of three cases:

(i) *When* $w_3 < 0$. For this to hold for all possible values of $\mathcal{E}_s$, we can solve Equation 7 for the value of $\mathcal{E}_s$ that most restricts the inequality, which, in this case, is $\mathcal{E}_s = 0$. Rearranging the terms reveals that $w_3 + w_4 > 0$ must hold.
(ii) *When* $w_3 = 0$, $w_4 > 0$ must hold.
(iii) *When* $w_3 > 0$. For this to hold for all possible values of $\mathcal{E}_s$, we can solve Equation 7 for the value of $\mathcal{E}_s$ that most restricts the inequality. Here, that is $\mathcal{E}_s = 1$. Rearranging the terms reveals that $w_3 < w_4$ must hold.

Theorem 33 *π_s in TD-$\mathcal{E}$ is appropriateness-dominant if players' expectations agree and $w_1 > 0$, regardless of the values of all other w_i. (If players' expectations differ and $w_1 > 0$, TD-$\mathcal{E}$ converges on becoming appropriateness-dominant through repeated play between a pair of players, if both players update their expectations over time to reflect observed accuracy outcomes.)*

Proof. The optimal Teacher strategy depends on the relationship between the accuracy expectations of the Teacher and Student:

- If the Teacher believes that they will agree with the Student's stated expectation of response accuracy (i.e. $\dot{\mathcal{E}}_t = \mathcal{E}_s$), the conditions required vary according to the value of w_1:
  (i) *When* $w_1 > 0$. Since the Teacher and Student agree on the expectation of response accuracy, we can rewrite the Teacher's expected utility purely in terms of their own expectations, and then solve for the level that corresponds with the highest expected utility:

$$\begin{aligned} \mathrm{E}_{\pi_t} &= w_1 \left(\dot{\mathcal{E}}_t + \mathcal{E}_s - 2\dot{\mathcal{E}}_t \mathcal{E}_s \right) + w_2 = w_1 \left(2\dot{\mathcal{E}}_t - 2\dot{\mathcal{E}}_t^{\,2} \right) + w_2 \\ \frac{d\mathrm{E}_{\pi_t}}{d\dot{\mathcal{E}}_t} &= 0 = w_1 \left(2 - 4\dot{\mathcal{E}}_t \right) \\ \dot{\mathcal{E}}_t &= 0.5 \end{aligned} \tag{8}$$

  In this case, the optimal strategy is to select challenges for which $\dot{\mathcal{E}}_t = 0.5$, making π_t appropriateness-dominant.
  (ii) *When* $w_1 = 0$, expected utility varies only with weight w_2, and is thus not sensitive to challenge appropriateness.
  (iii) *When* $w_1 > 0$, the case is similar to when $w_1 < 0$, accept that $\dot{\mathcal{E}}_t = 0.5$ marks the *minimum* rather than *maximum* expected utility. Thus, in this case π_t is not appropriateness-dominant.
- If the Student states expectations truthfully (i.e. $\mathcal{E}_s = \dot{\mathcal{E}}_s$, but the Teacher believes that their own expectations will differ with those of the Student (i.e. $\dot{\mathcal{E}}_t \neq \dot{\mathcal{E}}_s$), In this case, the Teacher sees the opportunity to out-perform the appropriateness strategy by selecting the challenge that maximizes the anticipated difference between $\dot{\mathcal{E}}_t$ and $\dot{\mathcal{E}}_s$. In this case, either the Teacher or the Student has provided a poor estimation of expectation, and the results from the following accuracy assessment will lead that player's expectations back in line. In this sense, the one-shot game may not be appropriateness-dominant, but the repeated game converges to an appropriateness-dominant payoff, as the two players independently converge on increasingly accurate expectation models. Feedback from off-diagonal challenges serve to sharpen either the Teacher's student model, the Student's metacognitive skills, or both.

Theorem 34 *TD-$\mathcal{E}$ is a Teacher's Dilemma when $w_3 < 0$, $w_3 + w_4 > 0$, and $w_1 > 0$.*

Proof. For TD-$\mathcal{E}$ to become a Teacher's Dilemma, three conditions must simultaneously hold: π_s must be truth-dominant, π_s must be effort-dominant, and π_t must be appropriateness-dominant. Since the proof of Theorem 31 requires that $w_3 < 0$ for truth-dominace, we only consider case (i) from Theorem 3.2 for effort-dominance. Here we require that $w_3 + w_4 > 0$. Finally, for appropriateness-dominance, we require only that $w_1 > 0$. When all three of these conditions are met, the resulting TD-$\mathcal{E}$ game is necessarily a Teacher's Dilemma. The Teacher's Dilemma example presented in Figure 3 is a form of this game, where $w_1 = 1$, $w_2 = 0$, $w_3 = -1$, $w_4 = 2$, and $w_5 = 1$.

4 Conclusion

In showing that the payoff functions for the TD-$\mathcal{E}$ could be restricted in such a way as to arrive at a Teacher's Dilemma, we surpassed one of the significant limitations of the prior TD-$\mathcal{D}$ game: the need for an existing accurate difficulty metric. We were able to organize the game in such a way as to obtain approximations of this information from the players themselves. Errors in their estimations serve as opportunities for learning, and dissipate as that learning occurs. The next frontier for generalizing the model would be to remove the only remaining assumption: that the accuracy of any response to any challenge can be assessed. In future work, we plan to build a new Teacher's Dilemma game in which accuracy is also effectively assessed by the participating players. In doing so, we will have achieved a very flexible and general framework for peer-driven learning, equally applicable for tutoring systems in well-defined and ill-defined domains.

References

1. Manske, M., Conati, C.: Modelling learning in educational games. In: Proceedings of the 12th International Conference on AI in Education. (2005)
2. Squire, K., Barab, S.: Replaying history: engaging urban underserved students in learning world history through computer simulation games. In: Proceedings of the 6th international conference on Learning sciences, Santa Monica, California (2004) 505–512
3. Malone, T.W., Lepper, M.R.: Making learning fun: A taxonomy of intrinsic motivations for learning. In Farr, R.E.S..M.J., ed.: Aptitute, Learning and Instruction: III. Conative and affective process analyses. Erlbaum., Hillsdale, New Jersey (1987) 223–253
4. Tan, J., Biswas, G.: Simulation-based game learning environments: Building and sustaining a fish tank. In: Proceedings of the First IEEE International Workshop on Digital Game and Intelligent Toy Enhanced Learning, Jhongli, Taiwan (2007) 73–70
5. Gómez-Martín, M.A., Gómez-Martín, P.P., González-Calero, P.A.: Game-driven intelligent tutoring systems. In: Proceedings of the Third International Conference in Entertainment Computing – ICEC 2004. (2004) 108–113
6. von Neumann, J., Morgenstern, O.: Theory of Games and Economic Behavior. Princeton University Press (1947)

7. Hurwicz, L.: Optimality and informational efficiency in resource allocation processes. In Arrow, Karlin, Suppes, eds.: Mathematical Methods in the Social Sciences. Stanford University Press (1960)
8. Nisan, N.: Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, E.T., Vazirani, V., eds.: Algorithmic Game Theory. Cambridge University Press, Cambridge (2007)
9. Weiss, G.: Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence. MIT Press (1997)
10. Gintis, H.: Game Theory Evolving. Princeton University Press, Princeton, New Jersey (2000)
11. Baker, R., Corbett, A., Koedinger, K.: Detecting student misuse of intelligent tutoring systems. In: Proceedings of the 7th International Conference on Intelligent Tutoring Systems. LNCS 3220, Berlin Heidelberg, Springer-Verlag (2004) 531–540
12. Magnussen, R., Misfeldt, M.: Player transformation of educational multiplayer games. In Smith, J.H., Sicart, M., eds.: Proceedings of the Other Players Conference, Copenhagen, Denmark, IT University of Copenhagen (2004)