Replay Penalties in Cognitive Games

Matthew W. Easterday and I. Yelee Jo

School of Education and Social Policy, Northwestern University, Evanston IL USA

Abstract. Replay penalties that punish players by making them repeat progress are ubiquitous in video games yet noticeably absent from tutors, creating a dilemma for designers seeking to combine games and tutors to maximize interest and learning. On the one hand, replay penalties can be frustrating and waste instructional time, on the other, they may increase excitement and prevent gaming the system. This study tested the effects of replay penalties on learning and interest. In a randomized, controlled experiment with a two-group, between subjects design, 100 University students played two versions of Policy World, an educational game for teaching policy argument, with and without penalties that forced students to replay parts of the game. Results showed that replay penalties decreased learning and interest. These findings suggest a minimize penalties principle for designing cognitive games.

Keywords: intelligent tutoring, educational games, serious games, penalties.

1 Introduction

Can *cognitive games*—educational games with embedded intelligent tutoring, promote learning as effectively as tutors [1] and be as fun to play as games? Cognitive games may not be able to maximize both learning and fun—by attempting both, they might achieve neither. In this study, we examine the effect of *penalties* on learning and interest to develop empirically supported principles for designing cognitive games.

How do we design cognitive games? Unfortunately, we cannot simply add tutors to stand-alone games—tutors and games are designed differently and for different goals. As a result, designers are forced to choose which game-like and tutor-like features to use, some of which are compatible, some of which are not.

Some of these differences *are* compatible. For example, tutors often lack fantasy environments. In most tutors, a learner is more likely to find himself solving a textbook problem than battling aliens. But we can easily design a cognitive game with both a fantasy environment and intelligent tutoring. Recent studies on game-like elements in tutors have focused on compatible features that do not directly affect tutoring, like 3D graphics [2] or narrative, visual presentation, and rewards [3].

Other differences between tutors and games are *incompatible*. Tutors provide more assistance than games, and they make it easy for the learner to figure out what to do by giving scaffolding and feedback on each *step*. Imagine the first-person shooter *Halo* giving the same level of assistance: not only would it tell you whether you've hit or been hit by an enemy, it would tell you what kind of weapon to choose, which

S. Trausan-Matu et al. (Eds.): ITS 2014, LNCS 8474, pp. 388–397, 2014.

enemy to target, how to point the weapon, when to shoot, the enemy's weakness, and so on. Whereas players make their own game guides and walkthroughs for entertainment games, tutors provide these answers for free via hints. Tutors also minimize penalties—after incorrect steps, tutors often allow learners to try again. Imagine *Halo* with minimal penalties: being hit wouldn't reduce your health; after missing an enemy, the alien would patiently wait for you try again. These conflicting approaches to assistance and penalties means that it is unclear whether cognitive games can simply add tutors to normal games to maximize learning and fun—adding tutors may increase learning at the expense of fun.

Here we are interested in penalties that directly affect tutoring, specifically replay penalties, where the game punishes players by making them restart at an earlier point. Replay penalties are ubiquitous across a wide variety of single-player video games such as Angry Birds, Halo, and Tetris. Replay penalties are ubiquitous because they make single-player games fun—losing lives or progress after a mistake creates pressure to make the right choice—which increases the excitement of making the choice and the satisfaction of choosing correctly.

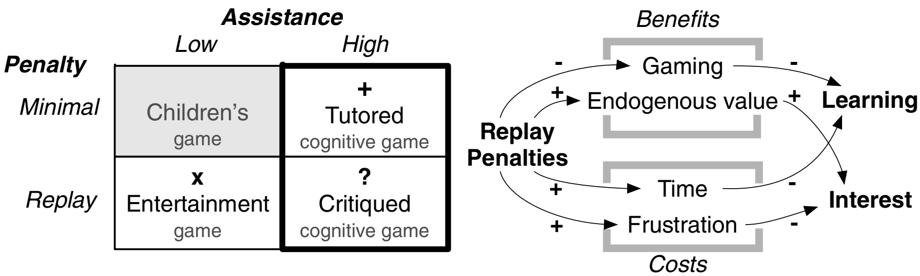

Fig. 1. Cognitive game design types (left) and possible causal effects penalties (right)

To explore the design space at the intersection of tutors and games, Easterday, Aleven, Scheines & Carver [4] compared two games: a *tutored* cognitive game with high-assistance and minimal penalties and an *entertainment* game with low-assistance and replay-penalties (Figure 1). Intuitively, we might predict a tradeoff with the tutored game better for learning and the entertainment game better for interest. In fact, the tutored game led to greater learning and competence, which in turn increased interest. So if entertainment game conventions are not effective, feedback promotes learning after all, how might a *critiqued game* with replay penalties and high feedback fare? In this study, we examine the role of replay penalties in cognitive games.

The case for replay penalties. Penalties are "rewards in reverse," such as points, resources and time that are taken away for making a mistake [5, p. 192, 6, p. 94]. Game designers consider penalties essential because they create the challenge and meaning needed to generate excitement. First, penalties *create challenge* by removing a resource needed to achieve a game goal, such as removing one of the player's limited number attempts or lives, forcing the player to replay part of the game, or reducing the player's points (needed to achieve a high score). Designers use penalties to make an easy game more challenging to prevent boredom. Second, penalties create *endogenous value* [5, pp. 31-33] or *meaningful play* [7, pp. 353-355] by establishing the relationship between the players' actions and game outcomes—penalties and rewards communicate to the player whether her actions move her closer to, or further from the goals of the game. Third, the combination of *challenge* and *endogenous value* are necessary for generating the interest/excitement/pleasure the player experiences when overcoming a challenge to reach a meaningful goal [5, p. 192], [7, p. 346].

Penalties might also increase learning by decreasing gaming. Intelligent tutors are susceptible to the *gaming the system* phenomenon, when learners "attempt[] to succeed in an interactive learning environment by exploiting properties of the system rather than by learning the material" [8]. For example, when hints give the learner the correct answer after a given number of requests, learners often rapidly click the hint request button until they receive the answer, rather than think about the problem. Penalties that impose a cost to random guessing or hint abuse might prompt students to think about the problem.

The case against penalties. On the other hand, penalties might decrease learning by wasting instructional time. Easterday et al. [4] found that an intelligent tutor embedded in a game-like environment increased learning and interest compared to a version that provided less feedback and stronger penalties, as is more typical of games, although this game-like tutor also provided less assistance, so the causal effect of penalties was unclear.

Second, replay penalties might not be necessary for creating interest. Entertainment games designed for children such as *Lego Star Wars* are immensely popular and impose extremely minimal penalties: when a player *dies* in Lego Star Wars, he drops all his *pieces* (points and money) but immediately reappears on the screen and is given several seconds to pick up the dropped pieces. While children's games have penalties that do not affect tutoring such as losing points, they suggest that *replay* penalties may not be necessary for generating challenge and interest.

Hypotheses. In this study, we compared how two cognitive games with either *replay* or *minimal penalties* affected learning and interest. The *replay penalty* version required students to replay parts of the game after an error, while the *minimal penalty* version allowed immediate error correction. The outcome measures were *learning*, which measured the policy analysis skills taught by the game, and *interest*, as measured by the Intrinsic Motivation Inventory [9]. Assuming that penalties make games more challenging, there are several plausible hypotheses:

1. *Null*: Replay penalties have minor, floor, or ceiling effects on learning and interest.
2. *Reduced gaming*: Replay penalties increase learning by reducing gaming (caused by low levels of interest), but have little effect on low levels of interest.
3. *Tutored game*: Replay penalties decrease learning and interest, because they waste instructional time and are unnecessary for generating interest.
4. *Critiqued game*: Replay penalties increase interest by making the game more challenging and, at best, equal learning by providing identical assistance.
5. *Painful game*: Penalties decrease interest by making the game too challenging.

We predicted support for either the *null* or *painful game hypothesis* based on the motivational importance game designers place on penalties and our previous finding that a *minimal penalties* version of Policy World increased learning and aspects of interest more than "game-like" version with minimal feedback and penalties [4]–possibly suggesting that lack of feedback in the game-like version decreased learning and masked the motivational effects of penalties.

2 Policy World

Policy World [4, 10] is a cognitive game designed to teach policy argumentation [11, 12]. In Policy World (Figure 2), the learner plays a policy analyst who must defend the public against the handsome but unscrupulous corporate lobbyist *Mr. Harding* by persuading the *Senator* to adopt policies based on evidence on topics such as carbon emissions, national health care and childhood obesity. The story employs an empowerment theme in which the young policy analyst, after typically failing an initial job interview (a disguised pre-test), is recognized as having great potential by *Ms. Cynthia Stark*, the head of a policy think-tank. The learner is guided through a grueling training by two mentor characters: *Molly*, another young but more senior analyst, and a sharp-tongued virtual *Tutor* that teaches the learner to analyze policies. At the end of the game, the player is tested in "real" senate hearings (posttests) in which the player must debate two policies with Mr. Harding to save the think tank's reputation and defend the public against Mr. Harding's corrupt agenda.

Fig. 2. Policy World screenshots of player and tutor characters

Policy World's fantasy environment follows *anime adventure/visual-novel* genre conventions that use dialogue boxes and hand drawn images of characters representing the speaker against backgrounds that display the character's location. The fantasy environment is heavily based on the game *Phoenix Wright* where the player stars as a defense attorney who "...must collect evidence, weed through inconsistent testimonies, and overcome corrupt agendas to ensure that justice prevails" [13], and which is one of Capcom's top-10 best-selling series [14]. Learners routinely comment positively on the similarities between the games.

Most Policy World levels include three broad activities: searching for policy information, analyzing that information, and debating policy recommendations against a computer opponent. During search, learners use a fake Google interface to find 3-7 newspaper-like reports, typically 3-5 paragraphs in length, containing causal claims from various sources like the New York Times, scientific journals, and bloggers that have varying levels of credibility and evidential support. At any time during search, learners can select a report to analyze, which requires them to comprehend, evaluate, diagram, and synthesize the evidence about the causal claims in the report using causal diagramming tools. Once learners have completed searching for evidence and constructing their causal diagrammatic analysis, they proceed to the final debate phase. During debate, learners make a policy recommendation, explain how the policy will affect a desired outcome, and provide evidence for their position by citing reports. The computer opponent (either Molly or Mr. Harding depending on the level) will argue against the player, attacking his recommendations, mechanism and evidence by providing alternate recommendations mechanism and evidence.

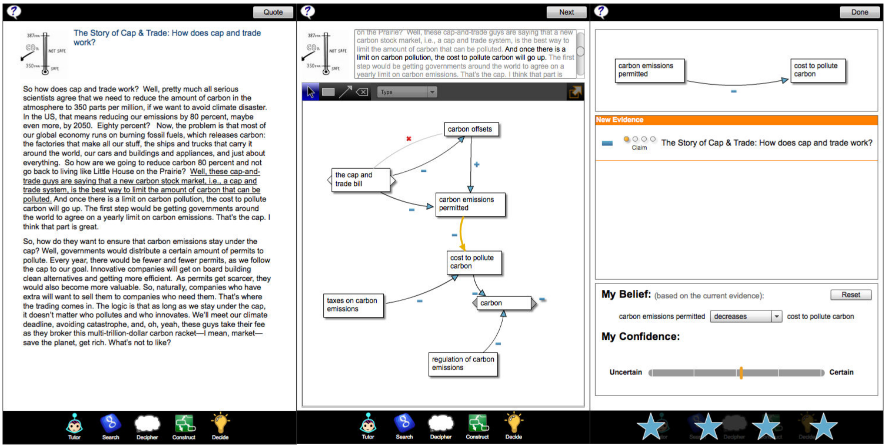

Fig. 3. Policy World comprehension, diagramming and synthesis screens

In this study, we focus on the analysis skills: comprehension, evaluation, diagramming and synthesis (Figure 3) described in [4] and repeated here for coherence:

- *Comprehend*. After selecting a report to analyze, the learner attempts to highlight a causal claim in the text such as: "the Monitoring the Future survey shows that 21 minimum drinking age laws decrease underage consumption of alcohol."
- *Evaluate*. The learner then uses combo boxes to identify the evidence type (experiment, observational study, case, or claim) and strength of the causal claim. Strength is rated on a 10-point scale labeled: none, weakest, weak, decent, strong, and strongest. The evaluation was considered correct if: (a) the evidence type is correctly specified, and (b) the strength rating roughly observes the following order taught during training: experiments > observational studies > cases > claims.
- *Diagram*. The learner next constructs a diagrammatic representation of the causal claim using boxes to represent variables and arrows to represent an increasing, decreasing, or negligible causal relationship between the two variables. The learner also "links" the causal claim in the report to the new diagram arrow which allows him to reference that report during the debate by clicking on that arrow.
- *Synthesize*. The learner then synthesizes his overall belief about the causal relationship between the two variables based on all the evidence linked to the arrows between those variables up to that point. The synthesis step requires the learner to

specify which causal relationship between the two variables is best supported by the evidence, and his confidence in that relationship on a 100 point slider from uncertain to certain. During training, a synthesis attempt is considered valid if: (a) the learner moves his belief in the direction of the evidence, assuming the learner's description of the evidence was correct, and (b) the learner's belief mirrors the overall evidence, assuming the learner's description of the evidence was correct.

Assistance. During training, errors in analysis are flagged by animated red stars and an explanation for the error. Errors in debate are also flagged and followed by Socratic questions that walk the learner through the steps involved in reading the diagram produced by analysis and citing evidence linked to the diagram.

3 Method

Design. The study used a two-group, between subjects, randomized, controlled, experimental design that compares a *replay penalties* version with a *minimal penalties* version of the game. During training, the *replay* penalties version of Policy World erased learners' progress upon making a mistake. When the learners made errors on an analysis step for a particular causal claim, they were sent back to the first analysis step. When learners received 5 debate strikes, they had to replay the level. The minimal penalties version allowed learners to correct errors with no loss of progress.

Participants. 100 university students were recruited through campus flyers and email. Students were compensated $16 for completing the study and an additional $4 for beating posttest 1 and an additional $4 for beating posttest 2.

Procedure. Students first took a pretest on either the drinking age (5 causal claims) or obesity (7 causal claims). During the pretest, students were allowed to search and analyze as many or as few reports as they liked before continuing to the debate. Students were then randomly assigned to the *replay* or *minimal penalties* training. Each group completed 3 training problems on video game violence (4 causal claims), organic foods (5 causal claims), and vaccines (4 causal claims). During training, replay penalties students received penalties for errors while minimal penalties students did not. Since it was possible that replay penalties students might take much longer on training, they were allowed to advance to the test levels after 1 hour on the training levels. After training, students completed the intrinsic motivation inventory survey [9] with sub-scales measuring perceived competence, effort, pressure, choice, value and interest. Finally students played two test levels without replay penalties or tutoring. The debate test (on cap-and-trade, with 8 causal claims) was a debate-only level that provided a completed diagram (to test hypotheses about debate skills outside the scope of this paper). Students then took a posttest identical to, and counterbalanced with, the pretest.

4 Results

Analysis 1: Do replay penalties affect learning? To examine how penalties affect learning we examined students' pre/post test analysis skill across the minimal/replay penalties groups using a two-way, repeated measures (mixed) ANOVA. Both groups improved on all four skills. The minimal penalty group showed significantly greater improvement than the replay penalty group on comprehension, evaluation and diagramming and a (not significantly) greater improvement on synthesis, (Table 1-2).

Table 1. Both groups learned analysis but the minimal penalties group learned more

		Pretest		Posttest	
Analysis skill	**Penalties**	**M**	**SD**	**M**	**SD**
Comprehend	*Replay*	2.68	1.92	3.50	1.79
	Minimal	2.24	1.82	4.26	1.64
Evaluate	*Replay*	1.72	1.58	2.38	1.59
	Minimal	1.68	1.63	3.10	1.72
Diagram	*Replay*	2.26	1.77	3.36	1.79
	Minimal	1.94	1.68	4.08	1.68
Synthesize	*Replay*	2.76	2.19	4.00	2.06
	Minimal	2.66	2.50	4.60	2.34

Table 2. The ANOVA showed a significant increase on all analysis skills for both groups and a greater increase on 3 out of 4 skills for the minimal penalties group

	Test (pre/post)				Penalty				Test-penalty interaction			
	df	**F**	**p**	**GES**	**df**	**F**	**p**	**GES**	**df**	**F**	**p**	**GES**
Comprehend	1 98	53.4	7.5E-11 *	0.138	1 98	0.28	0.60	0.002	1 98	9.53	2.6E-03 *	0.028
Evaluation	1 98	36.4	2.9E-08 *	0.094	1 98	1.51	0.22	0.011	1 98	4.86	3.0E-02 *	0.014
Diagram	1 98	70.9	3.2E-13 *	0.183	1 98	0.48	0.49	0.003	1 98	7.31	8.1E-03 *	0.022
Synthesize	1 98	39.2	9.8E-09 *	0.110								

Analysis 2: Do penalties affect intrinsic motivation? To examine how penalties affect interest we asked students to complete the well-validated intrinsic motivation inventory [9], immediately after the three training levels and analyzed the results with pair-wise t tests. The minimal penalties group felt significantly more competent, found the game more interesting and more valuable for learning policy (Table 3).

Table 3. Replay penalties decreased perceived interest, competence and value

	Replay		Minimal						
	M	**SD**	**M**	**SD**	**t**	**df**	**p**	**ll**	**ul**
Interest	3.44	1.32	3.93	1.24	1.89	97.62	0.061 .	-0.02	0.99
Effort	4.83	1.06	4.83	1.09	-0.02	97.88	0.985	-0.43	0.42
Choice	3.41	0.82	3.50	0.87	0.57	97.59	0.567	-0.24	0.43
Competence	3.45	1.43	4.17	1.20	2.71	94.91	0.008 **	0.19	1.24
Pressure	3.74	1.64	3.74	1.06	0.01	84.13	0.988	-0.54	0.55
Value	3.88	1.56	4.41	1.34	1.80	95.91	0.075 .	-0.05	1.10

Analysis 3: How are penalties, learning and interest related? To better understand the causal relationships between penalties, training, interest and analysis, we constructed a path model using the GES algorithm implemented in Tetrad 4 [15, 16] which searched for equivalence classes of unconfounded causal models consistent

with the correlations in Table 4 and our prior knowledge that: (a) penalties were determined before any other factor, (b) training was completed next, (c) intrinsic motivation was measured next, and (d) the posttest was completed last. Figure 4 shows the model discovered by Tetrad that we consider highly plausible and shows an excellent fit to the data. A chi-squared test of the deviance of the path model from the observed values showed we cannot reject this model a significance level of .05, $\chi 2$ (53, n=100)=48.41, p>.65, (here larger p-values indicate better fit and values *above* 0.05 indicate that we can't reject the model at a significance level of .05).

Table 4. Correlations between penalties, training success, analysis skills, and IMI subscales

			Analysis skills				Intrinsic Motivation Inventory							
	Penalty	**Train**	**Compr**	**Evalua**	**Diagra**	**Synth**	**Effort**	**Interest**	**Choice**	**Compt**	**Pres**	**Val**	**M**	**SD**
Penal	1												0.5	0.5
Train	-.50 ***	1											0.7	0.2
Com	-.22 *	.41 ***	1										0.6	0.2
Evalu	-.22 *	.48 ***	.74 ***	1									0.4	0.2
Diagr	-.20 *	.37 ***	.97 ***	.71 ***	1								0.5	0.2
Synth	-.15 ·	.33 ***	.82 ***	.63 ***	.83 ***	1							0.6	0.3
Effor	.00	.09	.04	.07	.02	.02	1						4.8	1.1
Inter	-.19 ·	.43 ***	.43 ***	.40 ***	.38 ***	.35 ***	.41 ***	1					3.7	1.3
Choic	-.06	.15 ·	.06	.17 ·	.05	.12	-.07	.17 ·	1				3.5	0.8
Com	-.26 **	.54 ***	.52 ***	.55 ***	.46 ***	.45 ***	.16 ·	.55 ***	.26 **	1			3.8	1.4
Press	-.00	-.20 *	-.23 *	-.26 **	-.21 *	-.20 *	.07	-.21 *	-.25 **	-.45 ***	1		3.7	1.4
Value	-.18 ·	.37 ***	.45 ***	.40 ***	.40 ***	.43 ***	.35 ***	.78 ***	.16 ·	.54 ***	-.21 *	1	4.1	1.5

*p<.05 **p<.01 ***p<.001

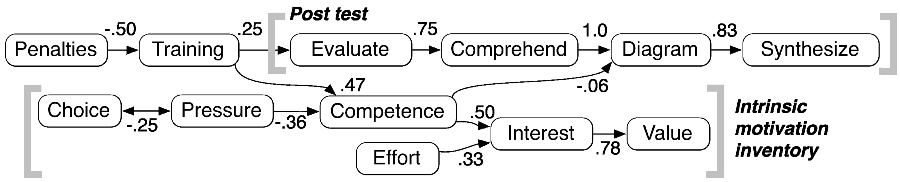

Fig. 4. Penalties decrease learning directly and by reducing perceived competence, which also decreases interest and perceived value of the game

In this model, replay penalties decreased training performance. Training performance affected posttest analysis (by increasing evaluation skills, which increased comprehension, which increased diagramming, which increased synthesis) and also influenced motivation (by increasing perceived competence, which increased interest, which increased value). Motivation in turn affected analysis by increasing diagramming. Choice was correlated with pressure, but it is not clear which caused the other.

Analysis 4: How do penalties affect training time? Students in the replay penalties version took longer to complete the training (M=20.8 min, SD=9.5) than students in the minimal penalties version (M=16.4 min, SD=4.5), t(70)=-2.96, p<.004.

5 Discussion

The results show that replay penalties decrease learning and interest in cognitive games. They do so by decreasing training performance, which directly impacts learning, and by decreasing motivational factors (specifically perceived competence which affects learning and interest and in turn value), which indirectly impact learning.

While these results may contradict our intuitions about the motivational effects of penalties, they are consistent with the effects on learning of previous work on combining tutors and games, which found that greater assistance also increased learning and motivation through similar mechanisms [4]. What is surprising is that game designers seem to so consistently and ubiquitously use a feature that seems to decrease interest across a wide variety of entertaining single-player video games.

The (apparent) contradiction is resolved by appealing to *balance*. Entertainment game designers often use (entertainment) tasks that are cognitively simple and add replay penalties to make them more challenging. Replay penalties don't create excessive frustration because players are likely to succeed if they keep trying. Educational game designers often begin with learning tasks that are cognitively complex and add assistance to make them easier. Replay penalties here make a complex task *too* frustrating. Of course, education game designers could use less assistance and easier, more gradated problems, but this would lengthen learning time.

Our intuitions about the motivational effects of games may be misleading because they are biased by our experience of players who have *voluntarily selected* to play a given game. Furthermore, entertainment games are not designed to promote learning that transfers out of the game, so there is no reason to think that cognitive games will succeed by mimicking their conventions. Entertainment games are designed to create the *illusion* of competence in a fake world, not actual competence in the real world [17].

Contribution: the *minimize penalties principle*. Thus the contribution of this work is support for a *minimize penalties principle*—that cognitive games should reduce replay penalties to increase learning and interest. Like the children's game *Lego Star Wars*, it is possible to maintain interest in cognitive games when the only penalty is a halt in progress (the most minimal possible). This leads to a design implication for educational games quite different from entertainment games: if tutoring is provided, it is better to balance a game by providing minimal penalties on a complex problem than replay penalties on a simple problem. This is the best possible result: embedding tutors in game environments increases learning and interest with *no tradeoff.*

If we are to make educational games that are effective for fun and learning, we must take advantage of what we have already learned about intelligent tutoring. While there are many proposed principles for games [5] and educational games [e.g., 18] and even some with empirical support [19], there are none that help designers resolve the conflicts that arise when applying intelligent tutoring techniques to games. Our previous work provided support for adding tutors to games (*tutoring principle*), to which we now add the *minimize penalties* principle. Future work must generalize and expand upon these principles if we are to apply intelligent tutoring research to realize the full potential of educational games.

References

1. VanLehn, K.: The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems. Educational Psychologist 46(4), 197–221 (2011)
2. Lane, H.C., Hays, M.J., Auerbach, D., Core, M.G.: Investigating the Relationship Between Presence and Learning in a Serious Game. In: Aleven, V., Kay, J., Mostow, J. (eds.) ITS 2010, Part I. LNCS, vol. 6094, pp. 274–284. Springer, Heidelberg (2010)
3. Rai, D., Beck, J.E.: Math Learning Environment with Game-like Elements: An Experimental Framework. International Journal of Game-Based Learning (IJGBL) 2(2), 90–110 (2012)
4. Easterday, M.W., Aleven, V., Scheines, R., Carver, S.M.: Using Tutors to Improve Educational Games. In: Biswas, G., Bull, S., Kay, J., Mitrovic, A. (eds.) AIED 2011. LNCS (LNAI), vol. 6738, pp. 63–71. Springer, Heidelberg (2011)
5. Schell, J.: The Art of Game Design: A Book of Lenses. Morgan Kaufmann, Burlington (2008)
6. Konzack, L.: Computer Game Criticism: A Method for Computer Game Analysis. In: Mayra, F. (ed.) Proceedings of Computer Games and Digital Cultures Conference, pp. 89–100. Tampere University Press, Tampere (2002)
7. Salen, K., Zimmerman, E.: Rules of Play: Game Design Fundamentals. MIT Press, Cambridge (2004)
8. Baker, R.S.J.D., de Carvalho, A., Raspat, J., Aleven, V., Corbett, A.T., Koedinger, K.R.: Educational Software Features that Encourage and Discourage "Gaming the System". In: Dimitrova, V., Mizoguchi, R., du Boulay, B., Graesser, A. (eds.) Proceedings of the 14th International Conference on Artificial Intelligence in Education: AIED 2009, pp. 475–482. IOS Press, Amsterdam (2009)
9. University of Rochester: Intrinsic Motivation Inventory, IMI (Web page) (1994), retrieved from `http://www.psych.rochester.edu/SDT/measures/IMI_description.php`
10. Easterday, M.W.: Policy World: A Cognitive Game for Teaching Deliberation. In: Pinkwart, N., McLaren, B. (eds.) Educational Technologies for Teaching Argumentation Skills, pp. 225–276. Bentham Science Publishers, Oak Park (2012)
11. Easterday, M.W., Aleven, V., Scheines, R., Carver, S.M.: Constructing Causal Diagrams to Learn Deliberation. International Journal of Artificial Intelligence in Education 19(4), 425–445 (2009)
12. Easterday, M.W., Aleven, V., Scheines, R., Carver, S.M.: Will Google Destroy Western Democracy? Bias in Policy Problem Solving. In: Dimitrova, V., Mizoguchi, R., du Boulay, B., Graesser, A. (eds.) Proceedings of the 14th International Conference on Artificial Intelligence in Education: AIED 2009, pp. 249–256. IOS Press, Amsterdam (2009)
13. Amazon: Phoenix Wright: Ace Attorney: Video Games. Amazon (Web page) (2013), retrieved from `http://www.amazon.com/Phoenix-Wright-Ace-Attorney-Nintendo-DS/dp/B000B69E96`
14. Capcom: Total Sales Units (Web page) (2012), retrieved from `http://www.capcom.co.jp/ir/english/business/salesdata.html`
15. Spirtes, P., Glymour, C., Scheines, R.: Causation, Prediction, and Search, 2nd edn. MIT Press, Cambridge (2000)
16. Tetrad: Tetrad (Computer Software) (2008), retrieved from `http://www.phil.cmu.edu/projects/tetrad/`
17. Thompson, C.: Halo 3: How Microsoft Labs Invented a New Science of Play. Wired Magazine 15(09) (2007)
18. Gee, J.P.: Learning by Design: Games as Learning Machines. E-Learning 2(1), 5–16 (2005)
19. Mayer, R.E.: Multimedia Learning and Games. In: Tobias, S., Fletcher, J.D. (eds.) Computer Games and Instruction, pp. 281–305. Information Age Publishing, Charlotte (2011)