

Learning Through Dissonance: Modeling Cognitive Conflict in Narrative-Driven Educational Games

Jingyu Wang and Hazra Imran

Khoury College of Computer Sciences, Northeastern University, Vancouver, BC, Canada

wang.jingyu6@northeastern.edu

h.imran@northeastern.edu

Abstract: In many serious games, psychological concepts function as narrative backdrop rather than as structures players can actively engage with. Cognitive dissonance is a clear example: players may sense the tension between belief and action, yet the dynamics through which that tension forms and resolves remain implicit in the mechanics. This work operationalises cognitive dissonance as playable dynamics within an interactive narrative system. Belief Strength tracks how firmly a player holds a declared stance, Cognitive Strain accumulates from the divergence between that belief and in-game actions, and three labelled reduction pathways (behavior change, belief revision, and justification) are instantiated as choosable options rather than abstract explanations. A Persona Dominance Ratio captures which coping pattern has stabilised across the session. The prototype examines whether experiencing dissonance through interactive decision-making supports conceptual understanding of the theory, and whether engaging with these mechanics prompts reflection on personal patterns of rationalization and belief change.

Keywords: Cognitive dissonance, Game-based learning, Serious games, Procedural rhetoric, Computational modelling, Reflective learning

1. Introduction

Serious games and game-based learning (GBL) increasingly draw on psychological concepts such as stress, empathy, and moral conflict as the substance of their educational ambition. Yet across a growing body of design work, these concepts surface as narrative texture rather than as playable structures. Bogost (2007) argues that games can mount arguments procedurally, through the rules and processes they embody, but in practice many serious games carry their psychological content in dialogue, framing, and aesthetic mood while leaving the underlying mechanics generic. Gee (2003) similarly argues that the educational power of games lies in the learning principles built into their design, not in the content layered on top, yet most serious games invert this priority. Wright and Denisova (2024) confirm the pattern empirically, documenting a paucity of research on the specific mechanics that scaffold moral decision-making and finding that even prominent moral-choice titles often leave the underlying psychological process untouched by the rule system. The gap is not that psychology is absent from these games; it is that the mechanics do not model the psychology.

Cognitive dissonance (Festinger, 1957) is a particularly clear example of this missed opportunity. It is precisely the kind of concept whose explanatory power lies in its dynamics: a belief and an action conflict, the resulting strain pushes the agent toward one of several reduction strategies, and the route taken depends on the strength of the belief and the constraints of the situation. Players of moral-choice games can certainly feel the tension between belief and action, but the underlying dynamics are never made visible or interactive. How Belief Strength modulates strain, how repeated reduction stabilises into a recognisable coping persona: these remain opaque. Dissonance is treated as something the player should feel, not something the player can probe. That gap is also an opportunity.

This paper presents a prototype that operationalises Festinger's theory as a computational model running beneath the play. Three core quantities, Belief Strength, Cognitive Strain, and a Persona Dominance Ratio, work together with three labelled reduction pathways (behaviour change, belief revision, justification) instantiated as in-game choices. Players make decisions, the model updates, and the consequences feed back into subsequent affordances. The organising question is whether experiencing dissonance through this kind of interactive decision-making supports conceptual understanding of the theory itself.

Section 2 reviews related work. Section 3 specifies the prototype and computational model. Section 4 develops the theoretical rationale. Sections 5 and 6 address the planned evaluation and limitations, and Section 7 concludes.

2. Related Work

2.1 Cognitive Dissonance Theory

Festinger (1957) introduced cognitive dissonance as the psychological discomfort that arises when a person holds two cognitions, or a cognition and a behaviour, that are inconsistent. The theory predicts that this discomfort motivates change: either the behaviour shifts, the underlying belief is revised, or the inconsistency is reframed so that it no longer matters. Cooper and Fazio (1984) refined this account by specifying that dissonance requires free choice over the inconsistent act, personal responsibility for an aversive consequence, and foreseeability of that consequence. Harmon-Jones and Mills (1999) consolidate subsequent work and identify three principal reduction strategies: changing the behaviour to align with the belief, changing the belief to align with the behaviour, and trivialising the conflict by reducing the importance of one of the cognitions. These three strategies map directly onto the dissonance-reduction pathways implemented in the prototype. Despite its centrality to social psychology, dissonance theory has rarely been translated into the rule systems of interactive learning environments.

2.2 Constructivism and Reflective Learning

Festinger's idea that internal inconsistency drives change has a close parallel in Piaget's (1952) account of cognitive development, where new experiences disrupt existing schemas and force either assimilation or accommodation. In both frameworks, learning is provoked by disequilibrium rather than by the smooth accumulation of facts. Kolb (1984) builds on this constructivist tradition with an experiential learning cycle in which concrete experience leads to reflective observation, which yields abstract conceptualisation, which is then tested through active experimentation. The prototype's core loop, in which the player encounters a scenario, makes a choice, and is then invited to reflect on the resulting strain and persona drift, is designed to traverse precisely this sequence within a single play session. Mezirow (1991) offers a complementary lens at the level of the adult learner, arguing that situations whose demands cannot be met by existing assumptions are the trigger for genuine perspective change. Cognitive dissonance, on this reading, is a microcosm of such a dilemma; making it playable invites the kind of structured reflection on which transformative learning depends.

2.3 Computational Modelling in Educational Games

Player and learner modelling has a substantial history in game-based learning, but its dominant concerns lie elsewhere than the modelling of belief or value. Knowledge tracing, originally developed for intelligent tutoring systems (Corbett and Anderson, 1994) and more recently extended to educational games (Hooshyar et al., 2022), estimates player mastery over a defined set of skills. These systems excel at tracking what a player knows or can do, but they treat the learner as a knowledge state to be filled, not as a self with revisable commitments. A second strand, affective computing in educational games (Calvo et al., 2015), models the emotional state of the player and uses it to adapt difficulty, feedback, or pacing (D'Mello and Graesser, 2012). This work has produced sophisticated detection methods drawing on facial expression, dialogue cues, and physiological signals, but the affective state is typically inferred from the player rather than represented within the game's diegetic logic.

Continuous belief or stance variables remain uncommon in narrative GBL. Role-playing and simulation games sometimes expose alignment meters or reputation scores, but these tend to be summary outputs rather than dynamical variables that condition future affordances. Work on player modelling in interactive storytelling (Thue et al., 2008) adapts narrative branches to player preferences without modelling the psychological mechanism by which those preferences shift. The contribution of the present work sits in this gap: rather than detecting affect or tracing knowledge, it represents cognitive dissonance itself as computable dynamics, with Belief Strength and Cognitive Strain feeding back into the option space the player sees. To our knowledge, no prior serious game has operationalised dissonance reduction as a transparent, interactive system in this way.

3. System Design

3.1 Overview and Onboarding

Learn From Dissonance (v1.0) is a mobile-first interactive narrative delivered as a progressive web application. After registration (name, email, consent), the player completes a two-step onboarding: entering a display name, then selecting exactly three beliefs from a curated list of six, each phrased in first-person behavioural terms (Responsibility: *I follow through*; Kindness: *I help when I can*; Honesty: *I tell the truth even when it costs*; Self-control: *I can stick to my plans*; Fairness: *I try to be open-minded and fair*; Safety in love: *love should not feel*

frightening). The three-belief constraint forces prioritisation, bounds the active belief set, and makes the player's value hierarchy explicit and personally staked before any scenario begins (Figure 1).




Figure 1: Belief elicitation screen

3.2 Core Variables

Belief Strength (0–100) expresses how strongly the player holds each declared belief. Initialised at onboarding and updated by every decision, it weights how much a conflicting action contributes to strain.

Cognitive Strain (0–100, in multiples of 5) is the primary affect variable, displayed as a persistent indicator throughout the session. Strain bands are Calm (less than 25), Uneasy (less than 55), Tense (less than 80), and Overloaded (80 or above). Each decision applies a base scenario delta plus a Dissonance Engine delta computed for every belief whose tags overlap the action and whose strength is 60 or above: $\text{contribution} = \text{strength} \times |\text{beliefImpact}| \times 0.8 \times (0.7 + \text{currentStrain}/100 \times 0.8)$, capped at 25 per belief. The compounding term means existing strain amplifies subsequent conflicts, modelling escalating psychological cost.

Persona Dominance Ratio tracks inner voice weights across four dimensions (Honesty, Security, Social, Attachment), whose dynamic weights (0–1) are reshaped by the player's accumulated coping pattern across scenarios.

3.3 Scenario Structure

The game comprises 10 scenarios (Table 1). Each presents an episode title, the focal declared belief, a second-person situation paragraph, and four inner voices offering competing first-person rationalisations, tagged by psychological dimension (Figure 2).




Figure 2: Scenario Screen

Table 1: Scenario Inventory

ID	Title	Belief	Theme
1	Missed Deadline	Responsibility	Procrastination
2	Helping Hand	Kindness	Helping others
3	White Lie	Honesty	Small lies
4	Diet Promise	Self-control	Self-control
5	Identity Clash	Fairness	Fairness and bias
6–10	Intimate Relationship arc	Safety in love	Coercive control

Scenarios 1–5 each address a distinct belief domain. Scenarios 6–10 form a thematically linked arc on coercive control, covering physical intimidation, surveillance, financial control, social isolation, and apology cycling. Players who prioritised Safety in love encounter five consecutive scenes testing the same belief, creating sustained pressure on one dimension of the system.

3.4 Decision Point

The Decision screen presents exactly two options, each carrying an action label and a one-line description making the belief relationship explicit. The choice is irreversible within the scene, replicating how real-world commitments close off alternatives (Figure 3). The conflict path applies a negative belief impact and positive strain delta; the aligned path applies a smaller positive belief impact and negative strain delta. The two options represent the player's first-order action; the three reduction pathways are instantiated separately on the Coping screen.




Figure 3 Decision Screen

3.5 Strain and Coping Screens

Strain screen (Figure 4). After a decision, the system surfaces the dissonance through a qualitative title scaled to magnitude (Very light to Heavy strain), a one-line interpretive body, the net strain delta, the focal belief hit, and a gradient bar labelled Calm / Uneasy / Tense / Overloaded. The screen is informational; the player taps Reduce Strain to continue.

Coping screen, gating (Figure 5). The Coping screen appears only when (a) the chosen action lies on the conflict path and (b) the Dissonance Engine's contribution reaches a 15-point threshold (only beliefs at strength 60 or above contribute, see Section 3.2). Otherwise, it is bypassed. The bypass mirrors Cooper and Fazio's (1984) finding that dissonance arousal depends on perceived stakes: low-stakes or already-eroded beliefs generate no tension that demands resolution.

Coping screen, content and effects. Three labelled options instantiate the classical reduction strategies: Change Behaviour, Adjust Belief, and Justify Behaviour, each carrying a first-person rationalisation. Justification's strain relief is computed at runtime to exactly cancel the scene's dissonance, so it always feels costless. Counter-intuitively, Justification erodes Belief Strength most (minus 10, versus minus 5 for Adjust Belief): the player keeps the verbal commitment while letting the lived standard slip. The selected pathway personalises the Reflection screen and reweights the inner voices on the next scene, operationalising the Persona Dominance Ratio (Section 3.2).




Figure 4: Strain screen after "Missed Deadline"




Figure 5: Coping Screen

3.6 Reflection and Feedback Loop

After a Result screen reporting the episode's quantitative deltas, the player reaches Reflection (Figure 6), a card with three fields.

What happened. A single sentence naming the focal belief, the concrete action taken, and the coping mechanism selected (e.g. "You hold 'I am responsible,' but you chose 'Keep scrolling.' Justification: you added a reason that made the action feel okay without changing either.").

Insight. A one-sentence theoretical takeaway personalised to the coping type (e.g. "Each time justification wins, the shortcut gets faster while the gap between belief and action stays open.").

Ask yourself. A metacognitive question naming the specific choice, coping strategy, and focal belief, with a future-orientation cue.

This structure traces Kolb's cycle: What happened delivers concrete experience and reflective observation; Insight delivers abstract conceptualisation; Ask yourself primes active experimentation. On hosted deployments, reflection text is LLM-generated, conditioned on the scenario, choice, coping type, strain trajectory, and running justification count, producing output uniquely responsive to each player's decisions.




Figure 6: Reflection screen after "Missed Deadline"

4. Theoretical Framework

4.1 Dissonance as Computable Structure

Festinger's (1957) original formulation treats dissonance as quantitative: its magnitude is a function of the ratio of dissonant to consonant cognitions, weighted by the importance of those cognitions to the person. This formula is more often paraphrased than implemented. The prototype takes it seriously. Belief Strength operationalises importance, attaching a continuous weight to each elicited belief that determines how much a conflicting action contributes to strain. Cognitive Strain operationalises the magnitude side of Festinger's formula: at each decision point, the system identifies which declared beliefs are activated by the player's choice, filters for those the action conflicts with, weights each by Belief Strength, and accumulates their contributions into the strain value. Consonant cognitions do not reduce strain directly; they instead appear in the base scenario delta, which is lower for aligned choices. This is a model, not a metaphor. Many serious games invoke dissonance as mood or atmosphere, but nothing in the rule system tracks that discomfort. Here the discomfort has a number, the number has consequences, and the consequences are visible to the player.

4.2 Agency Modulation as a Learning Mechanism

The Persona Dominance Ratio is the bridge between the underlying state and the playable surface. Dissonance theory predicts that under sustained strain, certain self-justifying frames begin to dominate and the felt range of acceptable responses narrows. The prototype enacts this directly. As a particular reduction pathway is repeated, its associated inner voice is reweighted on the next scene, so the same situation is interpreted through a progressively more partial set of rationalisations. Cognitive Strain itself compounds: the strain formula in Section 3.2 amplifies later conflicts in proportion to existing strain, making the same belief violation increasingly costly the more often it has been ignored. The Coping screen's gating reinforces the asymmetry: only conflicts above the threshold open the three reduction pathways at all. Players do not simply read about how dissonance constrains agency; they experience the constraint as it tightens around their own decision space. This is the

central pedagogical claim of the paper, and, in the absence of empirical data to date, it remains a design hypothesis rather than a demonstrated effect until tested by the evaluation described in Section 5.

4.3 The Three Pathways as Conceptual Objects

Harmon-Jones and Mills's (1999) reduction strategies typically appear in textbooks as a taxonomy, a list of three things people might do under dissonance. The prototype takes that taxonomy off the page and turns each strategy into a labelled, choosable action. When the player faces a scenario in which their behaviour contradicts a strongly held belief, the response is not free-form; the system presents three explicitly framed options corresponding to behaviour change, belief revision, and justification. A player who selects justification is not merely picking a narrative beat. They are enacting a psychological coping pattern that ordinarily operates below the threshold of awareness, and, because the option is labelled with its theoretical name, they have a chance of recognising it as such. This connects directly to Mezirow's (1991) account of transformative learning, on which genuine perspective change depends on bringing tacit assumptions into explicit view. Naming the pathway is what allows the play to become reflective.

4.4 The Learning Loop

The ten-scenario play session is structured as an explicit loop. It begins with Belief Elicitation, in which the player commits to and weights an initial set of beliefs. Each scenario then introduces a Dissonance Trigger: a situation that pits action against belief, with strain computed from the activated cognitions. The player selects a Coping Choice from the three pathways, after which a brief Guided Reflection asks them to articulate what just happened and what it cost. The system then updates Belief Strength and carries the new state forward. The loop maps onto Kolb's (1984) experiential cycle: the trigger is concrete experience, reflection is reflective observation, the labelled pathways scaffold abstract conceptualisation, and subsequent scenarios are the active experimentation phase. The loop iterates with memory: beliefs and strain accumulate across the ten scenarios, so a player whose earlier choices have driven their Persona Dominance Ratio toward justification will meet later scenarios with a narrower option set than a player who has alternated. The same scenario is, in effect, not the same scenario for two different players; the variables have done structural work.

4.5 Prototype Evaluation (Planned)

This section sets out the evaluation strategy to be carried out following the colloquium. The study has not yet been run; what follows describes the design, instruments, and analytic plan that have been prepared.

Research Questions

RQ1. Does interacting with the dissonance mechanics, namely Belief Strength, Cognitive Strain, and the three reduction pathways, support players' conceptual understanding of cognitive dissonance theory?

RQ2. Do players interpret the decision structures as representations of dissonance-reduction strategies, or do they experience them purely as narrative choices?

RQ3. Does engaging with the mechanics prompt reflection on personal patterns of rationalisation and belief change?

Planned Method

Participants. Adult learners recruited through the institution's research participant pool, with a target sample of approximately 20 to 25. Participants will be screened to exclude those with formal training in social psychology, since the study is concerned with naive conceptual learning rather than expert calibration.

Design. The study uses a pre/post design with an embedded qualitative component. Each session lasts approximately 60 minutes and combines three instruments: (a) a brief conceptual knowledge quiz on cognitive dissonance theory administered before and after play, capturing changes in understanding of the core mechanism, conditions of arousal (Cooper and Fazio, 1984), and reduction strategies (Harmon-Jones and Mills, 1999); (b) a concurrent think-aloud protocol during play; and (c) a semi-structured post-game interview of approximately 15 minutes focused on how participants interpreted the on-screen variables and whether play surfaced reflection on personal beliefs.

A specific concern for the interview and think-aloud protocol is distinguishing genuine identification with the modelled coping mechanisms from choices made for strategic or narrative reasons. To address this, the interview will include direct comparison prompts for individual decisions the participant made, asking whether the choice reflected how they believed they would act in a comparable real-life situation or was made primarily

out of curiosity about the resulting narrative branch. Think-aloud utterances will additionally be coded into outcome-exploration language (e.g. wanting to see what a choice leads to) versus self-recognition language (e.g. describing a choice as consistent with how the participant actually behaves), so that the balance between these two readings can be reported rather than assumed.

Analysis. Pre/post quiz scores will be compared using paired-samples statistics with effect sizes reported. Think-aloud transcripts and interview data will be analysed using reflexive thematic analysis, with an initial coding frame organised around the three research questions and open coding to allow unanticipated themes to emerge. Triangulation across the three instruments will assess whether quiz gains, in-play interpretation, and post-game reflection converge.

5. Discussion and Limitations

A design-stage prototype paper benefits from being explicit about what it cannot yet claim, and the present work has four limitations that bear directly on how its contribution should be read.

- **No empirical data yet** - The design is theoretically grounded and internally coherent, but every claim about learning outcomes in this paper remains a hypothesis. Until the planned study described in Section 5 has been conducted and analysed, statements about whether players actually develop conceptual understanding of dissonance theory through play, or whether the mechanics support reflection on personal rationalisation patterns, must be read as design predictions rather than findings.
- **Binary decision structure** - Each of the ten scenarios presents exactly two options. This was a deliberate simplification that made the mapping between choice and reduction pathway legible, but it is a poor reflection of how dissonance episodes unfold in life, where coping repertoires are richer and more layered. Future versions will explore richer decision architectures in which reduction strategies can be partially combined.
- **Belief elicitation validity** - The opening phase asks players to select beliefs from a pre-defined list. It is an open question whether this procedure produces genuine psychological ownership, or whether it elicits socially desirable responses with limited motivational force. If the latter, the strain computed in subsequent scenarios would be technically correct but psychologically thin. Comparing pre-set selection against open-ended elicitation is a priority for future iterations.
- **Individual differences.** Tolerance for dissonance and characteristic reduction styles vary between individuals (Harmon-Jones and Mills, 1999). The current system applies a single set of strain parameters to every player, which is appropriate for an initial prototype but means that the same on-screen state may correspond to quite different subjective experiences. A more developed system would adapt these parameters to baseline measures of dissonance tolerance

6. Conclusion

Serious games and game-based learning have built a rich repertoire for evoking psychological states, but the underlying psychology has tended to remain in the narrative layer rather than in the rule system. Concepts such as cognitive dissonance are felt by players, not modelled by the game. This paper has argued that this is a missed opportunity, and that closing it requires treating psychological theory as a source of computable mechanics.

The prototype presented here operationalises Festinger's theory of cognitive dissonance through four interlocking elements: Belief Strength, Cognitive Strain, three labelled reduction pathways, and a Persona Dominance Ratio. These variables are not decorative. They structurally narrow what the player can do as strain accumulates, and the design is intended to make the theoretical claim that dissonance constrains agency directly experienceable in the player's own decision space. Whether players actually experience the constraint in this way is, at this stage, a design prediction rather than a demonstrated outcome, and is the target of the evaluation proposed in Section 5.

The contribution to GBL is a design pattern that other researchers and designers can adapt to other psychological constructs whose explanatory power lies in their dynamics. The pattern has three moves: identify a continuous quantity that the theory says matters, expose it to the player as a tracked variable, and let it feed back into the option set. Empathy gradients, attentional load, moral disengagement, and social conformity are all candidates.

Moral disengagement (Bandura, 1999) illustrates how the mapping would work in practice. A Disengagement Score could rise each time a player selects a reasoning option that diffuses responsibility or reframes a harmful action in euphemistic terms. As the score climbs, later scenarios could present increasingly severe actions in

progressively softened language, making visible to the player, in real time, how repeated disengagement dulls sensitivity to escalating harm. The three moves used for dissonance in this paper, a tracked quantity, a visible indicator, and feedback into the option set, would carry over largely unchanged.

Even without empirical data, the framework offers a replicable design pattern, namely the operationalisation of a psychological theory as interactive computable dynamics rather than narrative texture, together with a set of testable hypotheses that the GBL community can take up, contest, or extend. The planned study described in Section 5 will provide the first empirical test of whether this approach supports conceptual understanding.

Ethics Declaration: This paper describes the design of a prototype game-based learning system and does not report empirical research involving human participants. Ethical clearance was therefore not required for the work presented here. The planned pilot study described in Section 5 will be conducted under full institutional ethical review prior to data collection.

AI Declaration: AI writing assistance tools were used in the preparation of this manuscript to support drafting and editing of prose. All theoretical framing, design decisions, and intellectual content are the authors' own. The authors take full responsibility for the accuracy and integrity of the work presented.

References

- Bandura, A. (1999) 'Moral disengagement in the perpetration of inhumanities', *Personality and Social Psychology Review*, 3(3), pp. 193–209.
- Bogost, I. (2007) *Persuasive Games: The Expressive Power of Videogames*. Cambridge, MA: MIT Press.
- Calvo, R.A., D'Mello, S.K., Gratch, J. and Kappas, A. (eds.) (2015) *The Oxford Handbook of Affective Computing*. New York: Oxford University Press.
- Cooper, J. and Fazio, R.H. (1984) 'A new look at dissonance theory', in Berkowitz, L. (ed.) *Advances in Experimental Social Psychology*, Vol. 17. New York: Academic Press, pp. 229–266.
- Corbett, A.T. and Anderson, J.R. (1994) 'Knowledge tracing: modeling the acquisition of procedural knowledge', *User Modeling and User-Adapted Interaction*, 4(4), pp. 253–278.
- D'Mello, S. and Graesser, A. (2012) 'AutoTutor and affective AutoTutor: learning by talking with cognitively and emotionally intelligent computers that talk back', *ACM Transactions on Interactive Intelligent Systems*, 2(4), Article 23.
- Festinger, L. (1957) *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford University Press.
- Gee, J.P. (2003) *What Video Games Have to Teach Us About Learning and Literacy*. New York: Palgrave Macmillan.
- Harmon-Jones, E. and Mills, J. (1999) 'An introduction to cognitive dissonance theory and an overview of current perspectives on the theory', in Harmon-Jones, E. and Mills, J. (eds.) *Cognitive Dissonance: Progress on a Pivotal Theory in Social Psychology*. Washington, DC: American Psychological Association, pp. 3–21.
- Hooshyar, D., Yousefi, M. and Lim, H. (2022) 'GameDKT: deep knowledge tracing in educational games', *Expert Systems with Applications*, 196, 116670.
- Kolb, D.A. (1984) *Experiential Learning: Experience as the Source of Learning and Development*. Englewood Cliffs, NJ: Prentice-Hall.
- Mezirow, J. (1991) *Transformative Dimensions of Adult Learning*. San Francisco, CA: Jossey-Bass.
- Piaget, J. (1952) *The Origins of Intelligence in Children*. New York: International Universities Press.
- Thue, D., Bulitko, V., Spetch, M. and Wasylshen, E. (2008) 'Interactive storytelling: a player modelling approach', *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 4(1), pp. 43–48.
- Wright, S. and Denisova, A. (2024) "'It's a terrible choice to make but also a necessary one": exploring player experiences with moral decision making mechanics in video games', *Computers in Human Behavior*, 159, 108342.