FULL PAPER | LUDUS 2026

Difficulty Curves as Rhetorical Devices: Pacing in Roguelike Design

Ada Lovelace[1], **Grace Hopper**[2]

[1]Department of Game Design, University of Somewhere
[2]School of Interactive Media, Institute of Computing
Corresponding: a.lovelace@example.ac.uk

Abstract

Roguelikes are frequently described as "punishing," yet their difficulty curves are rarely analysed as a deliberate communicative device rather than an incidental design cost. We argue that permadeath-driven pacing functions rhetorically: escalating failure states teach players to read a run's risk profile before committing to it, turning death into feedback rather than punishment. Drawing on case studies of three commercially released roguelikes, we identify three pacing patterns —front-loaded attrition, plateaued escalation, and spike-and-recover —and show that each communicates a different relationship between player skill and player luck. We conclude that difficulty curve design is best understood as an argument the game makes about what kind of failure is meaningful.

Keywords: roguelike; difficulty curve; procedural rhetoric; game pacing; permadeath; player feedback

1. Introduction

Roguelikes make failure routine: a run ends, progress resets, and the player begins again with only the knowledge gained, not the items or levels lost. This design is often framed as punishing by contrast with genres that preserve player progress unconditionally. We take a different view: a roguelike's difficulty curve is not simply "how hard the game is" but a structured argument about which kinds of failure the game treats as meaningful information versus arbitrary bad luck.

1.1. Why Pacing, Not Just Difficulty

Difficulty alone is a scalar; pacing is a curve, and a curve has shape. Two roguelikes can share an identical average failure rate while communicating entirely different things about player agency, depending on *when* in a run failure becomes likely and how sharply the odds shift once it does.

2. Three Pacing Patterns

We analysed run-length and cause-of-death data from three commercially released roguelikes spanning distinct design lineages, and identify three recurring pacing shapes.

2.1. Front-Loaded Attrition

Difficulty is highest in the opening minutes, when the player has the fewest resources to respond to it, then eases as build-defining choices compound. This pattern rewards players who treat early runs as information-gathering rather than progress, echoing the "legible failure" argument in procedural rhetoric [1].

2.2. Plateaued Escalation

Difficulty rises in discrete steps rather than continuously, holding steady across several encounters before jumping at fixed narrative or structural checkpoints. Players can therefore read their current plateau as a signal of what the game expects them to already know, similar to the environmental legibility cues described by [2].

2.3. Spike-and-Recover

Isolated high-variance encounters interrupt an otherwise gentle curve, after which difficulty drops back below the pre-spike baseline. This pattern communicates that a single bad outcome is a local event, not evidence the run itself is doomed —the closest of the three to Murray's notion of a game world that remains legible and navigable even when hostile to the player [3].

3. Discussion

Table 1 summarises how each pattern differs in where difficulty concentrates within a run and what it implies about the source of failure.

Table 1. Three difficulty-curve pacing patterns and what each communicates about failure.

Pattern	Peak timing	Implied cause
Front-loaded	Early	Player unpreparedness
Plateaued	Stepped	Structural checkpoint
Spike-recover	Isolated	Local bad luck

None of these patterns is inherently more punishing in aggregate terms —they can share an identical overall failure rate —but they differ

Figure 1. Illustrative reference image used as a figure placeholder in this starter; replace with your own.

sharply in what a player is entitled to conclude from a death. A front-loaded curve tells the player to treat early deaths as tuition; a spike-and-recover curve tells them not to over-read a single bad encounter. Conflating these as "just difficulty" erases a design decision that is doing real communicative work.

4. Conclusion

Treating a roguelike's difficulty curve as a rhetorical structure rather than a tuning parameter reframes what "fair" failure means in the genre: a curve is fair not because it is gentle, but because it accurately signals which failures are lessons and which are noise. Future work could extend this framework beyond roguelikes to other permadeath-adjacent genres, and test whether players' own accounts of "unfair" deaths track the pacing pattern we identify here or diverge from it.

Acknowledgments

We thank playtesters who shared run logs used to identify the pacing patterns discussed here.

References

[1] Ian Bogost. *Persuasive Games: The Expressive Power of Videogames*. MIT Press, Cambridge, MA, 2007.

[2] Henry Jenkins. Game design as narrative architecture. In Noah Wardrip-Fruin and Pat Harrigan, editors, *First Person: New Media as Story, Performance, and Game*, pages 118–130. MIT Press, Cambridge, MA, 2004.

[3] Janet H. Murray. *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*. MIT Press, Cambridge, MA, 1997.