EMERGENT STORYTELLING IN ROGUELIKE GAMES:

FROM DIGITAL SYSTEMS TO PHYSICAL PLAY

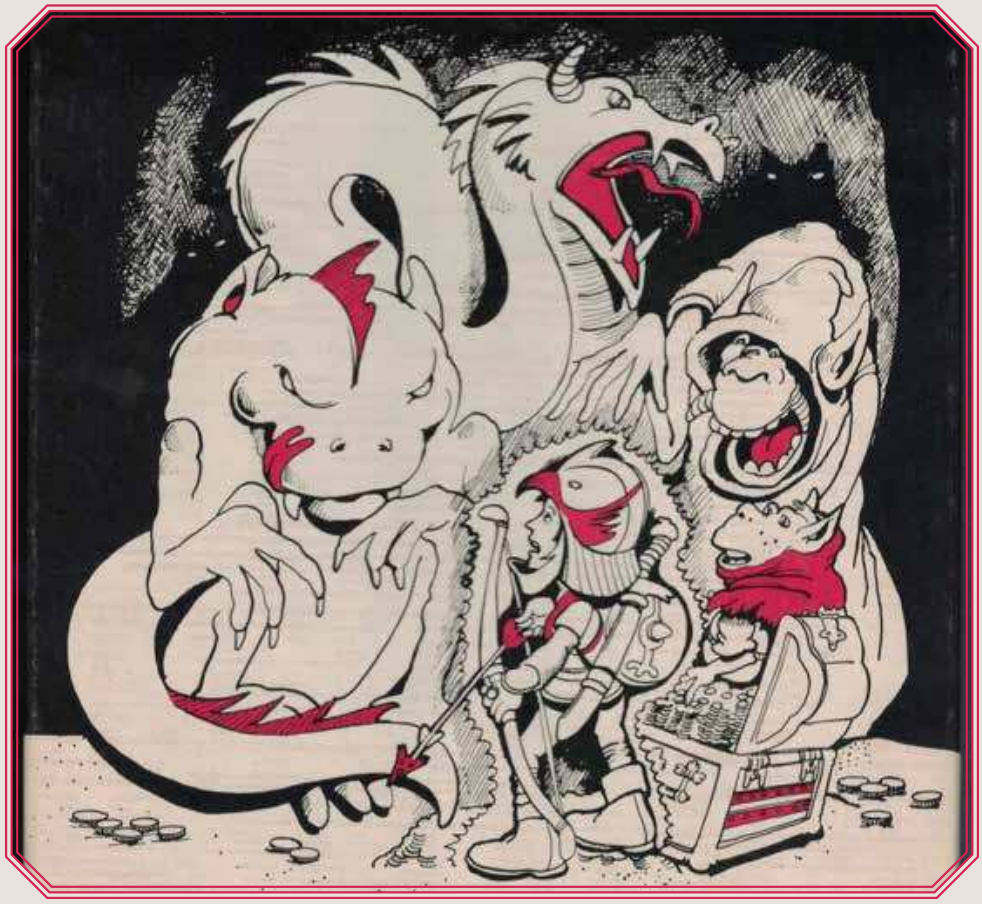

CAN UYGURÇETIN

I would like to thank my dear friends who helped bring The Final Exam in life. Eda contributed the illustrations and visual identity of the game. Ali supported me in shaping its structure. Buğra, Denizhan and Umut helped through playtesting and by sharing valuable ideas that improved the design.

Politecnico
di Torino

Bachelor's Degree Programme in
DESIGN AND COMMUNICATION

a.y. 2025/2026

Emergent Storytelling in Roguelike Games: From Digital Systems to Physical Play

Supervisor:
Andrea Di Salvo

Candidate:
Can Uygurçetin

CONTENTS

9 Kings, 2025

INTRO-DUCTION

Roguelike and roguelite games have shaped a distinct way of thinking about play. They also reshaped my love of games and the way I see games. Their use of procedural generation, unforgiving resets and constantly shifting situations has turned repetition into a new way of play and a new way of storytelling. These games do not guide players through a fixed storyline. They rely on systems that can surprise both the designer and the player, which is why they have remained relevant for more than forty years.

Most of this experimentation happens in digital environments where code handles every calculation and keeps the world in motion. A computer can reshuffle a map, adjust difficulty and track dozens of interactions without breaking flow. Tabletop games cannot lean on that invisible machinery. They depend on the people at the table. They are slower, more physical and more heavily shaped by social negotiation. This creates an interesting design challenge: if roguelikes owe much of their identity to computational processes, what happens when those processes must be recreated by hand?

This thesis explores that question. It examines the history of the roguelike tradition, the rise of roguelites, and the academic work on systemic and emergent storytelling. It then asks how these ideas function once they leave the digital domain and adopt cardboard. The research is paired with the development of a tabletop game titled The Final Exam, which serves as a testing ground for the concepts discussed in the earlier chapters.

The goal is not to imitate digital complexity with cardboard. The goal is to understand how procedural variation, failure loops and emergent narrative behave when they are carried by physical components and human interaction. Some elements translate cleanly. Others require adaptation. Some transform completely. The project aims to make these shifts visible and to show how the spirit of the roguelike can persist even when the medium changes.

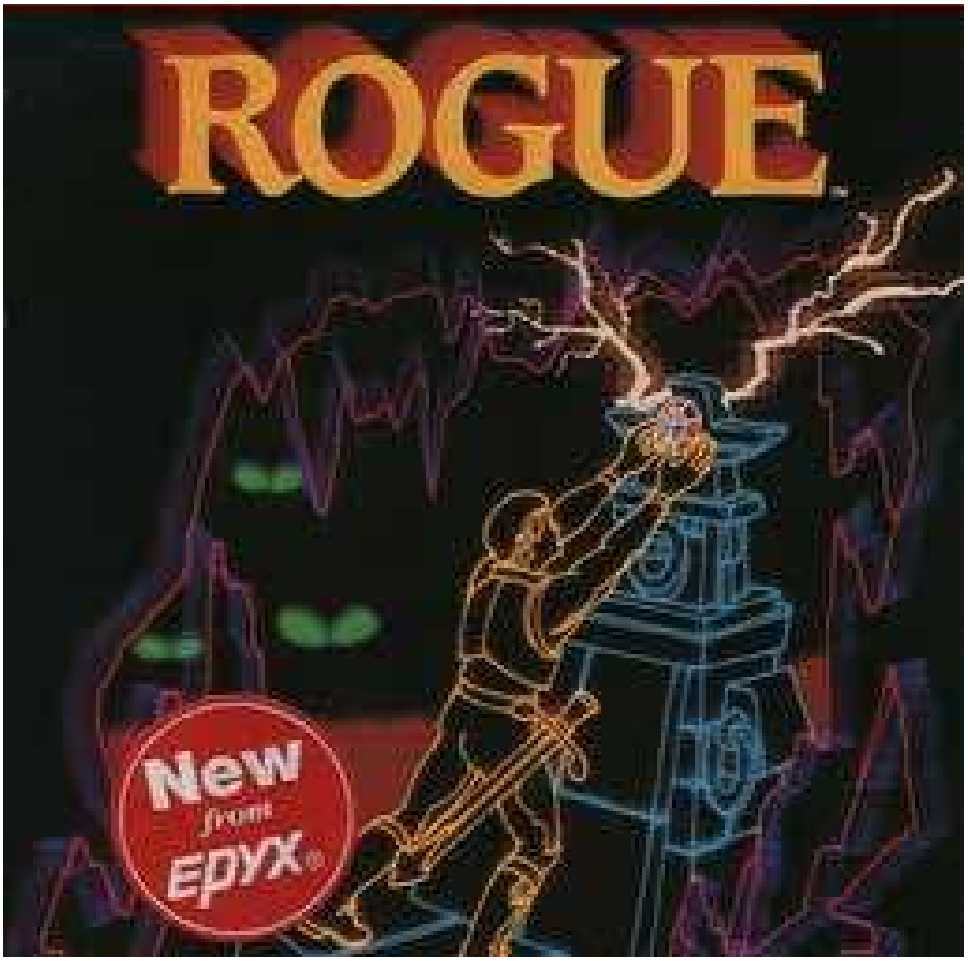

Rogue, 1980

CHAPTER 1
ROGUELIKE/ ROGUELITE: DEFINITIONS, HISTORY AND EVOLUTION

1.1 The Origin and Spread of the Term “Roguelike”

The term roguelike comes from Rogue (1980), a turn-based dungeon crawler created by Michael Toy, Glenn Wichman, and Ken Arnold at the University of California, Berkeley would quietly start a design tradition that still shapes games more than forty years later. It consists of three ideas :

1. Generated dungeons
2. Permanent death
3. Turn based gameplay

Mark R. Johnson (2015) reminds us that Rogue’s famous ASCII graphics were not only a technical constraint but also an aesthetic statement. Players navigated a world made entirely of symbols like “@” for the hero, letters for monsters, punctuation for the walls and doors and somehow it worked. They read the game as much as they played it. Johnson calls this a “semiotic” form of play: meaning arises from interpretation, not depiction. Looking at the screen, one did not see a dragon; one understood a “D.” Roguelikes in order to find their distinctive texture they used the gap between symbol and imagination.

Still, Rogue did not appear from nowhere. Earlier experiments such as Beneath Apple Manor (1978) had already flirted with algorithmic dungeons, but Rogue achieved a more elegant balance between chaos and coherence. Its most significant innovation, though, might have been social rather than technical. After being distributed freely through the Berkeley Software Distribution (BSD) of Unix, it spread across university networks where anyone could study or modify the code. Ho and Carter (2019) describe this ecosystem as a “genealogical web of influence,” a living network where games inherited and recombined design traits much like genetic material. As a result procedural generation, resource scarcity and permadeath became heritable features that mutated over time. In the following years, Rogue would become something more than a video game, a pioneer of an arising genre.

NetHack, 2006

Angband, 2005

From those ideas Rogue introduced, came an entire family. Hack (1984) and NetHack (1987) turned minimalism into enthusiasm, filling their worlds with hundreds of objects and unexpected interactions. Angband (1990) and Ancient Domains of Mystery (1994) expanded both scope and story, adding Tolkien-like mythology and overworld maps. Ho and Carter's network mapping shows how these descendants formed recognizable clusters, each a different interpretation of the same genetic code. The important thing is how open this evolution was. Without commercial pressure, hobbyist programmers freely exchanged ideas, creating a genre that was effectively co-authored by its community.

The habit of naming genres after landmark games is not unique to roguelikes. Game history is full of such nicknames: Metroidvania for Metroid + Castlevania, or Soulslike for the many descendants of Dark Souls. These labels are shorthand, but they also bring light to l something deeper about how game culture remembers itself. A single title often becomes an emblem for an entire design philosophy. In the same way that Doom defined the early first person shooter or Diablo defined the action RPG, Rogue came to signify a specific attitude toward uncertainty and consequence.

Johnson (2015) situates Rogue firmly within the intellectual culture of early computing. At Berkeley, students used mainframes not only for programming but also for play and experimentation. Limitations were severe but they encouraged creativity. ASCII, he notes, was part of a design mindset that valued clarity and imagination over spectacle. The result was a game that expected its players to do creative work in their own minds. Even in today's graphically lavish games, many designers still go after the same balance of transparency and depth that Rogue achieved with a few keyboard characters.

Procedural generation –one of Rogue's defining mechanics– would eventually escape the dungeon and colonize the rest of game design. Gellel and Sweetser (2020) trace this practice from its early use in Rogue to its sophisticated role in contemporary games. They argue that the original algorithm's appeal was conceptual rather than technical: it replaced static design with possibility space.

They talk about “authored randomness,” the deliberate shaping of procedural outcomes so that the unexpected still feels meaningful. Yet the underlying fascination remains the same creating a world that surprises even its creators.

Through the 1990s, the roguelike scene developed in small corners of the internet on bulletin boards, mailing lists and shareware archives. Because these games were rarely commercial, developers experimented freely. Ho and Carter (2019) note that this open source model made roguelikes resemble collaborative software projects more than entertainment products. Each new version, each “variant,” was a public conversation about design. This culture of repetitive improvement foreshadowed the ethos of the later indie movement, where small teams and personal expression would again redefine what games could be.

When digital distribution exploded in the 2000s, those old ideas resurfaced with new energy. Spelunky (2008) turned Rogue’s turn-based deliberation into a fast, physical dance of real-time platforming. In order to explore grotesque themes of guilt and childhood, The Binding of Isaac (2011) used repetition and randomness. FTL: Faster Than Light (2012) reimagined the dungeon as a spaceship, every encounter a new procedural crisis. For Gellel and Sweetser (2020), these games demonstrate how procedural generation evolved from a background process into an expressive way of telling stories through systems. These interwoven systems were replacing scripts for storytelling.

Looking across these decades, one sees that the roguelike never truly settled into a single form. It merged with other genres and resurfaced in new shapes, yet its logic risk, randomness and repetition remained intact. Ho and Carter (2019) call this combinatorial resilience, the ability of certain design principles to survive endless recombination. That resilience explains why the name “roguelike” still circulates. It no longer describes just ASCII dungeons but an attitude toward design: trusting algorithms to collaborate with players in producing surprise.

Today, the roguelike feels both ancient and current. Its code may have been born in an age of terminals, but its philosophy, balancing control and chaos still defines much of contemporary game design. Johnson (2015) suggests that its endurance lies in how it divides authorship between human and machine: designers write the rules, algorithms write the moment. Each playthrough becomes an act of co-creation. The same unpredictability that once thrilled students on monochrome monitors now animates sprawling modern worlds. In that sense, the roguelike promises room for surprise, failure and imagination.

Spelunky, 2012

```
==High value factors==

====Random environment generation====

The game world is randomly generated in a way that increases
replayability. Appearance and placement of items is random.
Appearance of monsters is fixed, their placement is random.
Fixed content (plots or puzzles or vaults) removes randomness.

====Permadeath====

You are not expected to win the game with your first character. You
start over from the first level when you die. (It is possible to save
games but the savefile is deleted upon loading.) The random
environment makes this enjoyable rather than punishing.

====Turn-based====

Each command corresponds to a single action/movement. The game is not
sensitive to time, you can take your time to choose your action.

====Grid-based====

The world is represented by a uniform grid of tiles. Monsters (and
the player) take up one tile, regardless of size.
```

Part of the original text from
The Berlin Interpretation, 2008

1.2 The Berlin Interpretation: A Community Definition and Its Limits

The roguelike's open, evolving nature and its tendency to invite endless remixing eventually produced a paradox. The same qualities that made it so generative also made it difficult to define. By the early 2000s, communities that had grown around NetHack, Angband, and ADOM began to debate what the term roguelike really meant. After all, when a design language spreads as widely as Rogue's had, some begin to worry that its meaning might dissolve altogether. Out of this anxiety for clarity came an unusual attempt at codification: the Berlin Interpretation of 2008.

The Interpretation emerged from a group of developers and enthusiasts who gathered at the International Roguelike Development Conference. Their aim was simple but ambitious to describe what, exactly, makes a game "roguelike."

The resulting document listed nine "high value" and six "low value" factors. Random environment generation and permadeath ranked at the top. It canonized a few titles like Rogue, NetHack, Angband, Crawl, and ADOM as the measure of authenticity. At the time, this consensus felt necessary: the indie scene was growing, new hybrids were appearing and long-time players feared the loss of a clear identity. The Interpretation gave the community a shared vocabulary and for a brief period.

What was meant as a clarification soon started to feel like a border wall. Johnson (2015) points out that trying to fix a genre born from procedural unpredictability is almost contradictory. The roguelike's creative energy depends on variation; each run, each new project, proves that rules are meant to be bent. Turning those rules into commandments risks freezing the very process that sustains the genre. Even the use of ASCII, once a technical necessity and later an aesthetic tradition, was treated by some purists as a sign of authenticity. Yet, as Johnson reminds us, ASCII was never sacred. It was a solution for its time.

Ho and Carter (2019) provide a broader view of this tension through their "ancestry network" of 639 roguelikes. Their visualization discloses that mechanical features evolve in clusters rather than straight lines: a mechanic disappears for years, then resurfaces in a new context; ideas recombine across distant branches. In such a landscape, any attempt to impose a stable classification feels random. Genres, they argue, act like ecosystems rather than family trees with predictable branches. The Berlin Interpretation, by contrast, treated the roguelike as a static organism.

Critics within the community soon voiced their concern. Among them was developer Darren Grey, whose provocatively titled essay "Screw the Berlin Interpretation!" captured the mood of frustration (Grey, 2013). Grey argued that

codifying a creative movement born in basements and terminals risked turning it into a museum piece. What mattered, he said, was not adherence to form but commitment to spirit: curiosity, experimentation and the willingness to embrace failure. Many players agreed. To call something a roguelike was supposed to mean that it behaves like Rogue in spirit offering uncertainty and discovery. Players in the community weren't wanting a checklist of items.

As independent game development expanded in the late 2000s, the limits of the Berlin framework became increasingly clear. Platforms such as Steam, Itch.io and digital consoles allowed small studios to reach global audiences. Designers were experimenting with procedural generation in real time genres, mixing it with platforming, storytelling or deck building. Ho and Carter (2019) note that their network data captures precisely this moment of divergence: around 2010, the number of hybridized roguelikes grew exponentially. Rather than branching neatly, the graph explodes outward, an image of creative entropy. The old boundaries simply could not contain what the genre had become.

This philosophical change opened a new door into endless possibilities for indie developers. Gellel and Sweetser (2020) observe that procedural generation itself has matured from simple randomness to a sophisticated craft. Designers explored that they could guide algorithms with human intention, a practice they call hybrid generation. In this sense, the roguelike's evolution mirrors the shift towards a conscious design control. Developers were less interested in recreating Rogue's exact mechanics than in exploring its logic: how rules could interact to produce surprise. The Berlin Interpretation, with its binary of "high value" and "low value" traits, had little room for such nuance.

A historical need appears : to articulate identity during a time of transition. It reminded designers where the genre came from, even as it failed to predict where it was going. Genres, Johnson (2015) suggests, are partly about community memory. The Berlin document may have overreached, but it also signaled a collective awareness that roguelikeness was worth preserving, even if no one could agree on what it was.

After 2008, debates over authenticity gave way to a more flexible vocabulary. Developers began using the term roguelite to describe games that borrowed the structural loop of Rogue short runs, randomization, failure followed by retry but abandoned its strict conventions. Either way, it acknowledged the genre's fragmentation. Spelunky (2008), The Binding of Isaac (2011) and later Hades (2020) offered experiences that felt roguelike in spirit yet played nothing like NetHack.

In the end, the Berlin Interpretation stands as both a milestone and warning. It demonstrates how communities, in seeking to protect a legacy, can inadvertently stifle it. Ho and Carter (2019) show that the roguelike's history is defined by mutation. Each new title rewrites its own origin and as Gellel and Sweetser (2020) remind us, procedural design is inherently forward looking: its beauty lies in unpredictability.

Games of Grey

Musings of a rogue roguelike developer

Games of Darren Grey

Home > Design, Musings > Screw the Berlin Interpretation!

Screw the Berlin Interpretation!

May 14th, 2013

In the year 2008 several men and women came together in Berlin to create the last, best definition of a roguelike. It failed...

Or at least in my view it did. The Berlin Interpretation as it became known was a set of high and low value factors for what constitutes a roguelike. These were based largely on the major roguelikes of the day. They ranged from the obvious like random content to the downright nonsense, like being set in a dungeon or using ASCII.

Let me make this really clear – adding ASCII to your game does not in any way make it more roguelike. Taking ASCII away does not make it any less roguelike. It's absurd to place value on this beyond an aesthetic choice. It's like saying platformers have to have pixel graphics because all the old platformers had pixel graphics. This is just one of several utterly nonsense features that the Berlin Interpretation terms roguelike.

The Interpretation comes with a disclaimer stating "The purpose of the definition is for the roguelike community to better understand what the community is studying. It is not to place constraints on developers or games." However prefacing a definition with this line is as futile as saying "I'm not racist, but..." Of course it puts constraints on developers and games! Devs want their games to fit in, and so tweak their works to score highly on this system. Gamers want to be exclusive about their community, and so rail against anything that doesn't fit the letter of their newfound law. This has been happening for all the years since the Interpretation was made public, and has reared its ugly head again this week in a poorly written PA Report on "What the hell is a roguelike?" The comments on the post and the related reddit thread are at least polite by internet standards, but there are still many in them that stick dunderheadedly to the idea that the Berlin

"Screw the Berlin Interpretation!", 2013

Dead Cells, 2018

1.3 From ASCII Dungeons to Modern Roguelites: Continuities and Shifts

By the time the Berlin Interpretation was debated, the roguelike had already slipped its textual skin. Its underlying logic, procedural generation, repetition and consequence was migrating into spaces far removed from the green on black terminals of the 1980s. Each new form, from minimalist dungeons to lavishly animated worlds, carried forward the same heartbeat of uncertainty.

In the early roguelikes, visual austerity was a statement. Johnson (2015) explains that ASCII graphics created a unique kind of literacy: players learned to read the screen rather than simply look at it. The dungeon was composed of symbols like letters, punctuation, fragments of code that the player mentally assembled into

meaning. This act of interpretation became part of the pleasure. The eye supplied only hints; imagination filled in the gaps. As Johnson observes, this symbolic minimalism encouraged a deeper engagement with the system itself. Without visual distraction, attention shifted to probabilities, timing and the fragile logic of survival. Every "@" or "g" carried the weight of risk. Maybe the world was not outstanding, but it was quite precise.

What began as necessity soon hardened into identity. The use of ASCII, Johnson (2015) notes, persisted long after hardware had expanded the need for it. As of today, the original roguelikes –Rogue, Hack, ADOM and many more– still use ASCII graphics despite the technology advanced the need for it. It became a sign of authenticity for the roguelike's roots in code and computation. Still, even among purists, the limitations of that form were evident. The absence of graphics created a cognitive intimacy but also a cultural boundary. Players needed to share a vocabulary of symbols to read the world correctly, and that vocabulary excluded many. As the medium matured, the roguelike faced an implicit question: could it evolve visually without losing its soul?

The answer came slowly, through experimentation. When developers began to replace ASCII with tiles, sprites or full animation, the change signaled a rethinking of how information, emotion, and rhythm could coexist. Gellel and Sweetser (2020) argue that these technological shifts paralleled transformations in procedural design itself. Early algorithms in Rogue relied on simple random placement; modern techniques use hybrid generation, combining chance with intention. Instead of pure randomness, designers now craft rule sets that guide algorithms toward meaningful variety.

This subtle balance defines much of the roguelike's modern evolution. Gellel and Sweetser (2020, pp. 6–8) describe how current practices integrate human judgment into automated processes, ensuring that procedural worlds follow a rhythm of challenge and rest. The algorithm serves as a collaborator. Designers tune its parameters to bring the mind curiosity and tension, sometimes even to suggest narrative pacing. A well-built generator can mimic an author's sense of

structure, while still producing surprises that no writer could have planned. Randomness has become one of storytelling's most productive allies.

The emergence of the roguelite in the late 2000s illustrates this transformation vividly. Spelunky (2008) transposed Rogue's deliberate turns into fast, physics driven motion. Its caves were still algorithmic, but the language had changed from numbers to momentum. Each fall or misjudged jump replaced the quiet calculation of a terminal command with physical intuition. The same grammar of risk persisted one life, one mistake but expressed through kinetic immediacy. In Dead Cells (2018), procedural rooms link into a fluid, interconnected map; Slay the Spire (2017) reshapes randomness into a deck of probabilities; Hades (2020) folds it into narrative dialogue. In every case, the genre's mechanical DNA remains undamaged even as its aesthetic surface changes.

Johnson's semiotic reading helps explain why this continuity feels natural rather than forced. The roguelike, he writes, has always been a language of systems more than images. The developers translate that language into color and motion, expanding its vocabulary. The "@" becomes a hero drawn in light and shadow, yet the underlying conversation with the algorithm, the constant negotiation between prediction and surprise remains unchanged. The move from text to texture therefore marks not a loss of authenticity but a widening of expression.

Early procedural worlds often felt indifferent, even cruel, because their randomness lacked curation. Modern roguelites aim for a different balance. Designers use limitations as creativity, shaping unpredictability into something that feels purposeful. This makes the games readable. Each death, each restart, becomes a legible sentence in a larger grammar of learning. It gives the player more opportunity to learn and explore while still giving them something new every time.

The refinement of procedural generation has also changed how we think about authorship. In traditional storytelling, meaning flows from writer to reader; in roguelites, meaning emerges between system and player. Johnson (2015) suggests that the act of interpretation once applied to ASCII symbols now extends

to the procedural system itself. Players learn to read the generator, to conclude its logic through experience. Over time, they develop what might be called algorithmic literacy: a sensitivity to patterns, probabilities and the subtle fingerprints of design. Every run is a lesson in systems thinking.

In this way, the journey from ASCII to modern roguelites reveals more than an aesthetic makeover or just consequences of technological advancements. It shows how a design philosophy born from limitation adapted to abundance. The same impulse such as economy, iteration, experimentation now operate on many larger scales. Gellel and Sweetser (2020) describe this as the natural outcome of a medium learning to speak through its tools. Procedural generation has matured from a trick of efficiency into a language of expression, capable of producing tension, humor, and even emotion. Its roots in Rogue's terminal code remain visible, but the branches have reached into nearly every genre imaginable.

Looking back, it becomes clear that the roguelike reinterpreted its origins by becoming more accessible to the everyday player. Johnson's analysis of ASCII as an enduring aesthetic, combined with Gellel and Sweetser's account of hybrid procedural design, reveals a continuity of mindset across decades. Both the old @ on a black screen and the vividly animated hero of Hades exist on the same spectrum: systems designed to surprise. The roguelike has learned new ways to speak, but it still says the same thing that unpredictability, when carefully shaped, can be a form of meaning.

Screenshot of Rogue (1980) on the left and Zagreus from Hades

Hades, 2020

1.4 Roguelike vs. Roguelite: Terminological Distinctions

By the early 2010s, the term roguelike had started to stretch beyond recognition. The old definition of turn based, grid based, permadeath, and procedural levels no longer described the growing variety of games borrowing its logic. Developers and players needed a new word to express this divergence. As a result, the word roguelite emerged. It is a label that at first sounded faintly dismissive, almost like a diet version of something more authentic, but soon became an accepted shorthand. Nowadays, both terms are used almost interchangeably.

The distinction between roguelike and roguelite is philosophical. Chen (2023) observes that classic roguelikes such as NetHack or ADOM treated death as absolute. Every mistake erased minutes, (sometimes hours) of progress, return-

ing the player to the beginning with nothing but experience and frustration to show for it. The punishment was harsh but honest. Roguelites, by contrast, began to soften that rigidity. They introduced meta progression: systems that allowed players to carry something across death, currency, upgrades or fragments of story. In Hades (2020), the protagonist dies repeatedly but grows stronger each time, turning failure into rhythm rather than rupture. This eases new players' adaptation to the game. Roguelites use death –so repetition– as a learning tool. As Chen puts it, the roguelite "disciplines the player through repetition," transforming loss into a learning loop rather than an ending.

The shift reflects a change in mindset in modern game design. Johnson (2015) reminds us that early roguelikes cultivated a severe engagement. Their ASCII minimalism removed spectacle, directing the player's focus toward systems and decisions. There was no cushioning, no narrative justification for failure, only the silent logic of the algorithm. Roguelites, in contrast, invite emotional investment. They build worlds that respond, characters that remember and narratives that progress one death at a time. The player no longer simply reads the system but becomes part of its unfolding drama. This movement from abstraction to empathy marks one of the genre's most significant shifts.

At the mechanical level, the difference is easy to see. Traditional roguelikes reset the game entirely upon death. Progress is internal, existing only in the player's memory and skill. The more you play the game, the better you get. No other way around. Roguelites allow fragments to persist such as new weapons, abilities or knowledge of the world. This continuity changes how time feels. Where roguelikes emphasized the tension of total loss, roguelites build momentum through gradual gains. Chen (2023) frames this transformation through Foucault's concept of discipline: modern roguelites turn repetition into a tool of self-regulation. Players are encouraged to persist, to learn, to adapt, finding satisfaction in improvement rather than endurance. Each death becomes both consequence and encouragement.

Representation plays a role here too. Johnson (2015) characterizes ASCII's se-

semiotic minimalism as a form of detachment. Players interpret symbols. In Rogue, the hero is an "@", a cipher without biography. Modern roguelites replace that distance with immediacy. Dead Cells animates each movement with speed and weight; Hades fills its underworld with color, dialogue and personality. The emotional economy changes: where roguelikes engage through abstraction, roguelites engage through empathy. Yet the underlying logic of uncertainty persists. Even as visuals grow rich, the player still confronts randomness, risk, and repetition.

The term roguelite thus captures both continuity and rupture. It signals a genre that remembers its ancestry. However it refuses its austerity. Chen (2023) suggests that this evolution mirrors broader cultural shifts in how games imagine challenge. In early roguelikes, difficulty was punitive. This is still evident in modern roguelikes that don't possess a meta progression system (progress between each run), even if they don't have ASCII graphics. In The Binding of Isaac – a roguelike game published in 2012– there is no meta progression other than unlocking new characters. In modern roguelites, difficulty becomes productive. Players are not punished for failing but rewarded for returning. This inversion aligns with Foucault's idea that power need not repress it can motivate. Hades trains the player through repetition much like a disciplinary routine: progress is earned by participation. The system teaches.

Accessibility has also become part of this distinction. Classic roguelikes were opaque, their controls arcane, their interfaces uninviting. For many, that opacity was a shared language among insiders. Roguelites, on the other hand, streamline interfaces, visualize information and introduce tutorials without erasing depth. This democratization expands the audience while preserving complexity. Johnson's (2015) notion of the roguelike as a "self referential genre" which means one that speaks primarily to those who already know its language. Roguelites speak outward, translating procedural play into a grammar anyone can learn.

Ho and Carter's (2019) network visualization suggests that hybridization has always been the roguelike's natural state. From Rogue's university terminals to

Hades' vibrant art direction, the line from past to present is continuous. Roguelites are translations of modern dialects of an old procedural language.

Seen from this perspective, both share the same heartbeat: procedural unpredictability as a means of creating meaning. If the roguelike asks, "Can you bear the unknown?" the roguelite asks instead, "What will you learn from it?" Johnson (2015) and Chen (2023) show that the genre's evolution is an expansion of its purpose : from ascetic simulation to expressive experience, from terminal command to emotional loop.

The Binding of Isaac, 2012

Slay the Spire, 2017

1.5 How Hades, Slay the Spire and The Binding of Isaac Contributed to the Genre's Evolution

If classic roguelikes taught players to absorb loss, modern roguelites taught them to grow from it. The shift reshaped how designers thought about failure, repetition and the stories that emerge across multiple attempts. Games like Hades, Slay the Spire and The Binding of Isaac show how repetition can become something other than punishment. In these works, returning to the start works as a continuation. It becomes coherent with the story.

Chen (2023) treats Hades as the clearest example of this new approach. Supergiant Games designed it around a loop of death and return, yet nothing about that loop feels empty. Each escape attempt resets the level structure, but the

world remembers what happened. Characters comment on past failures, dialogue branches open and relationships slowly accumulate. The pace of the narrative follows the pace of defeat. Chen uses Foucault's idea of discipline to describe how Hades trains the player. Failure becomes instructive rather than punitive. Each death invites another round of adjustment, a chance to rethink weapons, boons and strategies. Even the quiet walk through the House of Hades after a failed run becomes a moment to absorb what changed. The game guides without coercing; repetition encourages patience and gradual mastery.

The emotional weight of this structure matters as much as the mechanics. Hades frames loss through narrative context and music. When Zagreus fails, the game responds with sympathy and curiosity. It creates a form of procedural storytelling in which authored writing reacts to algorithmic variation. Chen (2023) sees this as a turning point in the roguelike tradition. Instead of systems that merely test players, these newer designs seem to acknowledge them. They give the impression of a world that listens.

Slay the Spire (2017) offers a different angle on the same evolution. Its design blends the uncertainty of roguelikes with the structure of deck-building. Every run constructs a new set of probabilities. The player chooses a path, collects relics, builds a deck and hopes the next fight aligns with the tools at hand. No two attempts feel alike. Storytelling arises from decision-making rather than dialogue. The drama comes from balancing long-term plans with short-term improvisation. Ho and Carter (2019) position games like this within the roguelike family tree because they preserve Rogue's modular logic and permadeath while applying them to new genres. What Slay the Spire adds is the idea that randomness can be expressive. A deck that barely holds together can create tension as vivid as any narrative script. Choices become stories, even when no character speaks them.

Where Hades and Slay the Spire treat repetition as discipline or strategy, The Binding of Isaac (2011) turns it into a metaphor for memory. Ruiz and Mota (2022) read the game as an allegory of trauma, with each descent into the basement re-

presenting another confrontation with fear. The monsters echo guilt, religious anxiety and distorted childhood memories. In their interpretation, the roguelike loop mirrors the psychological process of returning to the same unresolved thoughts. Failure here has narrative meaning. Every run is a new attempt to process something old. The child's journey repeats not because the system demands it, but because the story does.

Ruiz and Mota describe this structure as an "iterative narrative," where understanding emerges through cycles rather than linear progression. The basement changes slightly each time, much like memory reshapes itself. What was a procedural challenge in Rogue becomes procedural confession in Isaac. Randomness and death –once the markers of difficulty– become narrative devices that support introspection.

Taken together, these three games show that the evolution of the roguelike is emotional as well as structural. Earlier titles focused on survival by calculation. Their modern descendants explore survival through meaning. Chen (2023) views Hades as a blend of discipline and care. Ruiz and Mota (2022) see Isaac as a study of fear, recovery and the repetition of old wounds. Even Slay the Spire, which has no characters in the traditional sense, encourages reflection through its cycle of planning, risk and revision. These games keep the old building blocks – procedural generation, permadeath, short runs– but their purpose has changed. They now generate stories about persistence, empathy and understanding.

This shift also reflects a broader change in how difficulty is understood. In early roguelikes, difficulty acted like a wall. You either climbed it or gave up. Modern roguelites treat difficulty as a conversation. According to Chen's (2023) idea of "cyber discipline," the player becomes both the subject of challenge and an active participant in shaping their response. A failure is more of a suggestion than a verdict. "Try again, but differently." Each return becomes part of the learning curve. Ruiz and Mota (2022) show that repetition can carry emotional weight as well. Dying again and again can reflect a process of growth rather than a measure of inadequacy. Because of this, roguelites have become laboratories for emoti-

onal design. Systems that once felt cold now express warmth and reflection. Hades uses repetition to build empathy. The Binding of Isaac uses it to explore self-analysis. Slay the Spire uses it to reveal the drama hidden inside numbers and choices. What connects these experiences is the idea that systems can tell stories by shaping possibilities rather than dictating outcomes. The roguelite takes Rogue's experiment and expands it. The symbols have become characters; the randomness is more deliberate; the repetition feels more humane. Yet the core remains unchanged: meaning emerges through interaction.

Through games like Hades, Slay the Spire, and The Binding of Isaac, the genre links code to feeling. Chen (2023) and Ruiz and Mota (2022) show that repetition can be a method of storytelling, a pattern that lets players and systems meet halfway. In that ongoing exchange, the spirit of Rogue persists. The loop remains, always beginning and beginning again, always searching for what might be discovered next.

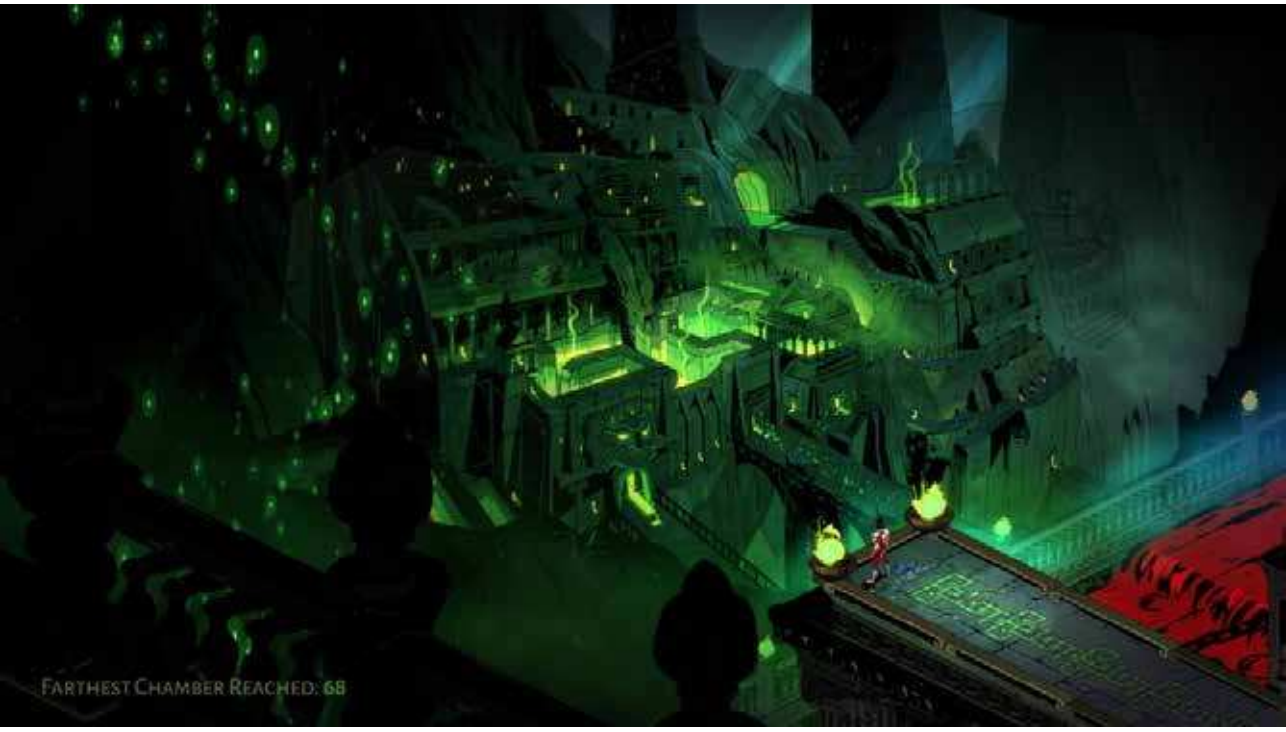

Hades, 2020

Red Dead Redemption 2, 2018

CHAPTER 2

NARRATIVE STRUCTURES IN VIDEO GAMES

2.1 Main Narrative Structures in Video Games

Narrative structure in games is about how stories are built, delivered, and experienced by players. Since video games are not recorded pieces of media, they balance authored storytelling (what designers and scenarists plan) with interactivity through systems (how players interact with the story). So, narrative structures in video games are natural consequences of this tension between storytelling and interactivity. Theorists have long debated whether narrative and interactivity can truly coexist. Some argue that story is fixed and passive, while play is dynamic and unpredictable, meaning games constantly negotiate between control and freedom (Ip, 2011a, pp. 104–105).

In order to grasp the concept of narrative structure, it helps to separate three key terms: story, plot, and narrative.

"Story" is the full sequence of events in the fictional world, what happened, and in what order (Ip, 2011a, p. 105). "Plot" is how those events are causally connected (Ip, 2011a, pp. 105–106). "Narrative" refers to how those events are actually presented to the player through pacing, structure, and emotional framing (Ip, 2011a, pp. 105–107). In opposition to classic media such as books, movies, series and such, video games give both designers and players different opportunities through systematic game design. For example, while we play Red Dead Redemption 2, we observe the conflict between characters through a cutscene. This is a passive act just like watching a movie or reading a book. We, watching this scene as merely observers, witness everything without drawing any characters' attention like a piece of dust. As the argument escalates and tensions rise, the camera slowly descends toward the back of the main character, Arthur. Just as the guns are drawn, the player suddenly finds themselves in the middle of the conflict. Their heartbeat quickens; their palms begin to sweat as if they had reached for the weapon themselves. This constant tension between being a passive observer and staying on alert is what keeps us from ever getting bored of a video game that could last for weeks.

Linear Structure

Linear narratives are the most traditional form, where the game unfolds in a set sequence of levels or chapters (Ip, 2011a, p. 109). Developers decide when to introduce tension, tragedy, or resolution. Linear narrative structure offers clear pacing and emotional control since there is minimum room for systemic variance or controlled randomness. Thus, linearity restricts agency. Players often feel they are merely following instructions rather than shaping outcomes (Ip, 2011a, pp. 108–110). Despite this, linearity is one of the most common narrative structures in video games. It guarantees coherent drama and storytelling. Last of Us is a masterpiece in textbook linear storytelling. Every scene, encounter, and emotional beat happens in a fixed order. You, as the player, can help Joel and Ellie's journey unfold, not choose or decide. The game decides the pacing, tone, and resolution.

Detroit: Become Human, 2018

Last of Us, 2013

Branching and Nonlinear Structures

Branching structures attempt to restore agency by giving players decisions that lead to diverging outcomes (Ip, 2011a, pp. 109–110). The appeal lies in control: choices feel meaningful, and players believe they are co-authoring the story. At first it sounds like there can be endless outcomes but this model is limited. Every branch requires new dialogue, art, and testing, quickly becoming unsustainable. Developers often "funnel" branches back into the main path, making freedom largely an illusion (Meadows, as summarized in Ip, 2011a, p. 66; Ryan, as summarized in Ip, 2011a, pp. 109–110). Detroit: Become Human is a good case study for this narrative structure. Every dialogue choice can ripple into massive consequences or collapse back into the same funnel. The illusion of choice sells the fantasy of agency in the first playthrough. However, when you replay it, you notice how many branches loop back. For this particular reason, full branching video games can be unsustainable from a design point of view.

Cutscene-Driven Structure

A more controlled hybrid alternates gameplay and cutscenes (Ip, 2011a, p. 108). Cutscenes explain events, move players to new settings, show consequences of actions or sometimes, they hide defects, such as buying time when a new part of the map is downloading. They solve the tension between interactivity and storytelling by briefly pausing play to deliver exposition. However, this can disrupt immersion and frustrate players by breaking flow (Ip, 2011a, pp. 108–109). In any case, cutscenes are used very commonly because of how efficient its cinematic control and predictability are. Some games mitigate this by offering interactive cutscenes dialogue choices or quick actions to make the transition smoother (Ip, 2011a, pp. 109–110). Kojima's Metal Gear Solid saga embodies the cutscene-heavy hybrid structure. Entire plotlines are delivered through cinematic sequences that interrupt gameplay for 20+ minutes at a time. It's theatrical and controlled, perfect for exposition, but often criticized for killing flow and agency mid-mission.

Dark Souls, 2016

Metal Gear Solid Master Collection Vol. 1

Environmental and Emotional Structures

More recent design trends shift focus toward environmental storytelling: conveying narrative through world design, ambient sound, and visual cues rather than explicit dialogue (as discussed in Ip, 2011a, pp. 109–110, 204–208). These cues let players feel the story. Emotional design plays a key role like games use camera framing, animation, and sound to express emotion without words (Ip, 2011b, pp. 203–210). According to Ip, major games repeat emotional peaks such as fear, loyalty, grief about every 20–40 minutes, pacing emotion like beats in film (Ip, 2011b, pp. 204–210). Still, most rely on external emotions like anger or revenge, rarely exploring introspection or moral ambiguity (Ip, 2011b, pp. 207–210). Dark Souls series deliver this environmental and emotional depth so perfectly with its dark and grim world. Ruins, corpses, item descriptions, and sound design whisper the story. Dark Souls builds emotion through space, silence, and decay. While you feel trapped in this grim world of death, you often find yourself admiring how beautifully this world is built, especially for its time. The world itself is the narrative, forcing the player to feel the story instead of being told it.

Mythic and Archetypal Structures

Many games borrow from Joseph Campbell's Hero's Journey, with recognizable stages: Call to Adventure, Trials, Ordeal, Reward, and Return (as summarized in Ip, 2011a, pp. 111–112). Ip found that most games emphasize the middle "Tests and Allies" stage, often consuming 98% of total runtime, while setup and resolution receive minimal attention (Ip, 2011b, pp. 213–214). Overemphasizing challenge and progression fit gameplay loops. It helps the player feel in charge and relate to the character. However, as Ip said, in cases where the resolution is merely 2% of the total gameplay, it can reduce emotional closure. Similarly, games use archetypes as hero, mentor, shadow, ally in order to keep stories legible (as discussed in Ip, 2011a, pp. 111–112). The Legend of Zelda: Ocarina of Time (1998) is the Hero's Journey turned interactive. Link's "Call to Adventure," trials, defeat, transformation, and final return all follow Campbell's arc. The archetypes (hero, mentor, villain, ally) are pure myth distilled into play. Its charm is in its simplicity.

Kernel and Satellite Structure

This is the typical structure of most RPGs (roleplaying games). Cohan and Shires distinguished between kernels (essential story events) and satellites (optional moments that add depth), a distinction presented through Ip's analysis (Ip, 2011a, pp. 112–113). This translates into games' "main" and "side" quests. Kernels (main quests or main storyline) tell the story chapter by chapter, just like in a novel. They ensure the plot progression. Players reveal their character's background, game's setting, the conflict and many more by completing the main quests. Satellites (side quests) allow emotional pacing and player driven exploration. These may vary depending on the game's setting like, helping other NPCs (non playable characters), or revealing hidden parts of the map. They usually do not contribute to the main storyline but give more agency, freedom and resources to players. Games rich in satellites encourage reflection and attachment; those with only kernels can feel urgent but emotionally flat (Ip, 2011b, pp. 203–204). In The Witcher 3: Wild Hunt, one of every gamer's beloved fantasy RPG, the main storyline (finding Ciri, confronting the Wild Hunt) is the kernel, essential to progress. But its emotional depth comes from the satellites: optional side quests like "The Bloody Baron" or "A Towerful of Mice," which expand characters, themes, and world texture.

Witcher 3, 2015

In summary, game narrative structures: linear, branching, cutscene-driven, environmental, archetypal, and kernel/satellite, can give us useful templates to work with but they do not restrict recipes that game developers should follow or fit in one specific type. They are merely ongoing negotiations between control and freedom; story and play. Story demands structure; play demands agency (Ip, 2011a, pp. 104–105). Each structure balances those forces in their own unique way by shaping how players live the story rather than simply watch it.

Minecraft, 2012

2.2 What is emergent storytelling and how does it work in interactive media?

Emergent storytelling usually refers to how various systems interact with each other and from each other to produce some form of narrative or gameplay (as described in Kybartas, Verbrugge, & Lessard, 2020). This term is also used in tabletop games, historically the most famous example being Dungeons and Dragons (as discussed in Spierling, Grasbon, Braun, & Iurgel, 2002). The moderator, called DM (Dungeon Master for short), guides the players and provides a different storyline each game run. Emergent storytelling creates narrative not from fixed scripts but from the interaction of players, systems, and rules. Instead of predetermined scenes, stories arise through simulation, chance, and player improvisation (Kybartas, Verbrugge, & Lessard, 2020). This design choice is also

crucial for some digital video game genres. Roguelike/roguelite games like Binding of Isaac, Spelunky, Dead Cells, or open-world "sandbox" games such as Don't Starve or Minecraft.

In Minecraft, the player spawns in the middle of a land where there are no instructions, no rules, no tutorial, no explanation. A player who has never heard or watched Minecraft before must discover how to play the game in order to survive. The clueless player will eventually discover that they can interact with things around them such as trees and dirt, collect those resources, make tools and weapons. When the sun goes down and monsters arise, they will get caught and permanently die in their first run. They will start from the beginning again with nothing in hand. The world of Minecraft, where you are the only person trying to survive, might feel empty or meaningless even. But the more you play it becomes evident that there is a way out from this cruel and lonely world. One of the appealing aspects of these types of games is that in a world of endless possibilities, the systemic design of the game gives players the illusion of total agency.

Stories Emerging from Simulation

Traditional narrative defines what happens and when. Emergent systems define why things can happen. Developers establish agents, goals, and relationships, then let interactions unfold autonomously (Kybartas et al., 2020). A strong emergent narrative system, like Minecraft, produces outcomes that feel authored. Players recount these outcomes as unique but they put up with the same challenges each game that emergent simulation provides. Thus, they improve their gameplay and eventually "unlock" new challenges and experiences. Success, then, is measured not by how well players follow a plot, but by whether they create memorable events worth retelling (Kybartas et al., 2020). Characters must not only act but interact meaningfully with alliances, betrayals, and conflicts that generate drama (Kybartas et al., 2020). The richer these systems are, the greater the narrative potential is.

Player Agency and Shared Authorship

In emergent systems, authorship is distributed between designers, algorithms, and players. Designers shift from dictating "what happens" to defining "why things happen" (Spierling, Grasbon, Braun, & Iurgel, 2002). Players become co-authors through their decisions, consciously or not. Spierling et al. describe interaction across layers: high-level story choices, mid-level scene outcomes, and low-level behaviors like tone or gestures. Even identical setups can lead to distinct personal narratives depending on player input.

Cavazza, Charles, and Mead (2002) expand this model: their AI-driven characters use hierarchical planning to pursue goals, dynamically reacting to player interventions. If the player hides an object or delivers new information, the characters' plans shift, altering the unfolding plot. Unlike branching paths, this system does not choose between "ending A" or "ending B" but rewrites causality on the fly: A true collaboration between player and system (Cavazza et al., 2002).

Tension and Conflict

Maintaining drama without pre-authored control is one of the biggest challenges in emergent storytelling setups. Too much player freedom risks dull outcomes; too much control breaks interactivity. Ip (2011, as cited in Miller, 2004/2011) describes this as the balance between control and conflict. Kybartas et al. (2020) propose modeling this through "tension space": a map of character goals and values that inevitably collide, producing conflict and drama. This concept allows game designers to anticipate which interactions will naturally generate engaging story beats without direct scripting.

Aphrodite in Hades 2, 2025

2.3 Why are roguelike games well-suited to emergent narrative design?

The roguelike genre is, by nature, ideal for emergent storytelling. The core elements of procedural generation, permadeath, and randomness ensure that the story is not dictated by a fixed scenario or a linear sequence, but instead arises organically from the game's internal systems. Of course, as in every genre, there are exceptions within the roguelike family. One of the most striking examples — and the one that inspired me to write this thesis — is Hades, the acclaimed title by Supergiant Games.

In Hades, Zagreus, the son of the god of the Underworld, attempts to escape his father's realm with the help of the Olympian gods. His goal is to find his mother, Persephone, and uncover the truth about himself and his family. Time and again, he is defeated by the monsters and creatures his father has set to guard the Underworld, but being immortal like all gods and demigods, Zagreus is reborn after each failure.

Hades meticulously blends Greek mythology with contemporary storytelling to deliver an enthralling experience. Unlike most roguelike games, the narrative in Hades takes a central role so much so that the game's script exceeds 300,000 words, more than the combined total of The Iliad and The Odyssey (Mathys, 2020). Supergiant fully voiced this monumental script, pairing it with stunning art direction to create an unforgettable experience for players. Hades successfully merges the systemic game design where each weapon, boon, aspect, encounter and more can quite affect the feel of each run; with meticulously written, rich script. This well-done combination of both worlds, as a gamer, mesmerized me like never before.

Of course, it's worth emphasizing once again: as mentioned at the beginning of this section, Hades is one of the rare exceptions that successfully merges both aspects. Looking at the roguelike/roguelite genre as a whole, it would be more accurate to address the question posed at the start of the chapter "Why are roguelike games well-suited to emergent narrative design?" through the genre's core components: procedural generation, permadeath, systems, repetition, and persistence.

Procedural Generation Creates Personal Histories

Different from fixed narratives, roguelikes generate new worlds each run. In The Binding of Isaac, every floor assembles rooms, items, and enemies from probabilistic pools, making "the same game with different results every time" (McMillen, 2011, as cited in Goulin & Mota). Similar to Hades, instead of choosing weapons in the beginning of each run, in The Binding of Isaac the player gets to choose one of different characters. From that point on, the player has no idea what they will encounter each run. The procedural generation provides players with unique experiences each run. Instead of saying, "Here's the plot," roguelikes make players say, "Here's what happened to me." Gustafsson, Holme, and Mackay (2020) call such singular moments "first-time and unique events" that players treasure and retell. Similar phenomena appear in long-term strategy games, where players narrate emergent dynasties or betrayals as personal chronicles (Burgess & Jones, 2021). Roguelikes achieve the same effect in miniature: combinations of enemies and items become memorable sequences that feel authored but aren't. Procedural remixing therefore sustains narrative freshness rather than chaos.

Permadeath Adds Weight and Meaning

Permadeath turns failure into story closure. It is evident that many other games from various genres also adopted the same or a similar approach to character death. For example, in Diablo series, a multiplayer action RPG, the player can toggle on the "Hardcore" mode option: When your character dies, all progress is lost and you can not play that character anymore. In a game like Diablo, where the player puts in literally hours and hours of gameplay to increase a character's level, unlock new equipment and beat bosses, it is much heavier, even enraging to start all over again. Thus, dying feels like a serious consequence of unthoughtful actions. It increases realism, strategic depth, and emotional impact. But most importantly, presents an actual challenge.

In The Binding of Isaac, our main character Isaac's mother begins to experience religious delusions. God, she believes, whispers to her that Isaac is being corrupted by the devil. As her visions intensify, she takes away all of Isaac's belongings like first his toys, and eventually even his clothes. At one point, she completely loses her sanity and, convinced she is obeying God's command, decides to

sacrifice her son. Knife in hand, she corners Isaac. Just as she is about to strike, Isaac lifts a trapdoor in the floor and throws himself into the unknown darkness below to escape her.

From that moment on, we as the players try to help Isaac survive the horrors lurking underground. Each attempt becomes an episode with its own beginning, escalation, and loss. The permadeath mechanic aligns with the game's themes of trauma and memory: "the roguelike nature of the game [is] a narrative tool to explore ... Isaac processing his trauma" (Goulin & Mota). Drawing on Piaget's theory of preoperational childhood, the authors explain that memories at Isaac's age are volatile and ego-centered (Piaget, 1929, 1951, as discussed in Goulin & Mota). Each failed attempt shows this instability: subjective recollection reframed as new reality. Endings "connected to the same bigger narrative ... trespass individual playthroughs," building meaning through repetition rather than continuity (Goulin & Mota). Therefore, permadeath helps the player grow emotionally with the game's story, rather than being just a simple reset.

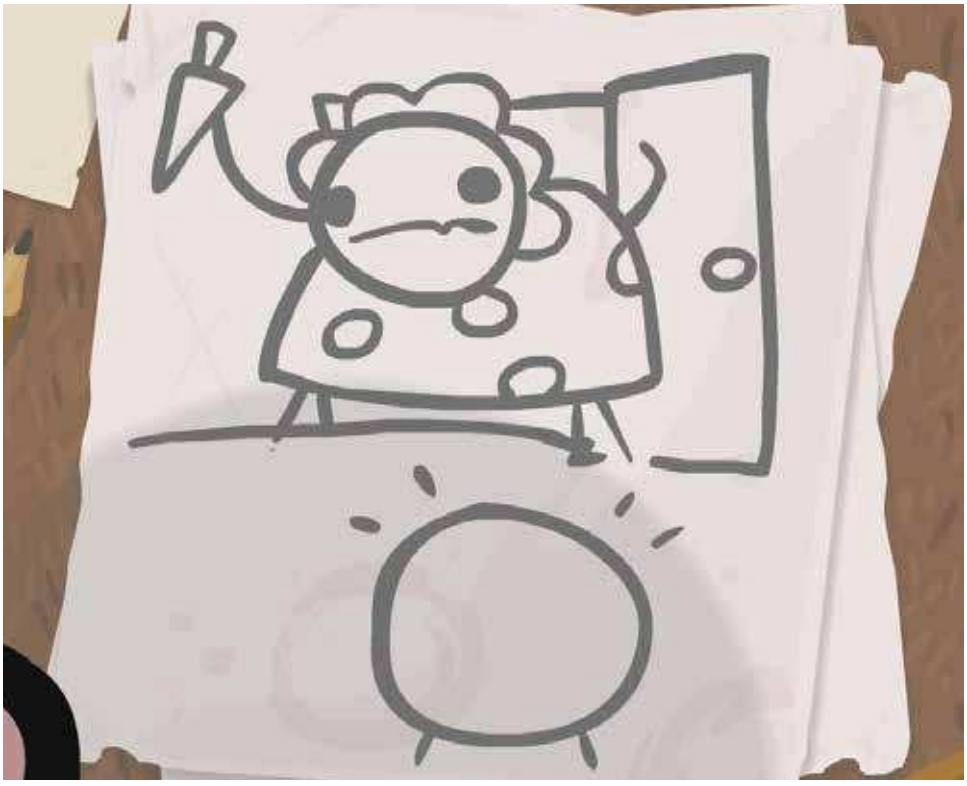

The Binding of Isaac's opening scene, 2012

Systems Enable Self-Expression

Emergent storytelling might also depend on interpretation. Roguelikes externalize the player state so that mechanical changes express identity. Isaac, in the beginning of each run, starts only with himself and his nude body. Just like how he escaped from his mother. He attacks monsters with the only thing he has left: his tears. Every item he finds in dungeons visibly mutates his body—tears of blood, horns, masks—each carrying metaphorical weight. "Dad's Key," for instance, isn't just a utility item but a symbol of authority and control within an unstable family (Goulin & Mota). Other items, like "Divorce Papers," carry similar emotional charges.

The player reads builds (a character's specific combination of abilities and gear in a game) as psychological portraits. Every character variation Judas, Eve, Samson is "still Isaac," reflecting a child's blurred self-concept (Piaget, 1929, as cited in Goulin & Mota). This mechanic allows co-authorship. The system provides expressive symbols and players interpret their meanings. Comparable dynamics appear in Total War campaigns where fans write "After Action Reports" turning emergent gameplay into dynastic lore (Burgess & Jones, 2021). Roguelikes formalize that impulse of self-expression through systemic game design.

Persistence Turns Runs into Lore

Emergent stories carry a risk for the average player that cannot be ignored: the risk of the story vanishing every time the character dies and the game "resets". Gustafsson et al. (2020) propose Narrative Substrates frameworks that "represent, manage, and persist traces of player activity as unique, interactive content". Players "highly value first time and unique events" and want them remembered. The Binding of Isaac echoes this at the meta level: endings across runs "are connected to ... the same bigger narrative" (Goulin & Mota). By tracking progress and unlocked memories, the game effectively reifies past play as lore for what Gustafsson et al. (2020) call Story Events.

Gaming communities often extend this persistence themselves. They build wikis, logs, YouTube documentaries and many more to "persist traces of their actions". Roguelikes could integrate such systems directly so players "feel like legends ... remembered and influential in the game world".

The Binding of Isaac: Rebirth, character select screen, 2014

Repetition as Reflection

Repetition in roguelikes gives room of improvement and introspection to the player. From the storytelling perspective, each attempt deepens the connection between the player and the story. Meanwhile, the player mechanically gets better in-game, allowing them to outplay their opponents and overcome harder challenges.

Goulin & Mota stresses that replays are "not just trial and error ... but connected ... to the same bigger narrative". As new layers unlock, difficulty and symbolism increase (McMillen, 2011, as cited in Goulin & Mota). The process mirrors descent into Isaac's increasingly-worsened, deeper and darker subconscious.

In summary, roguelikes are well-suited for emergent storytelling for a reason: systemize uniqueness, consequence, interpretation, memory, and reflection. They shift narrative authorship from developer to player, turning each run into a lived myth.

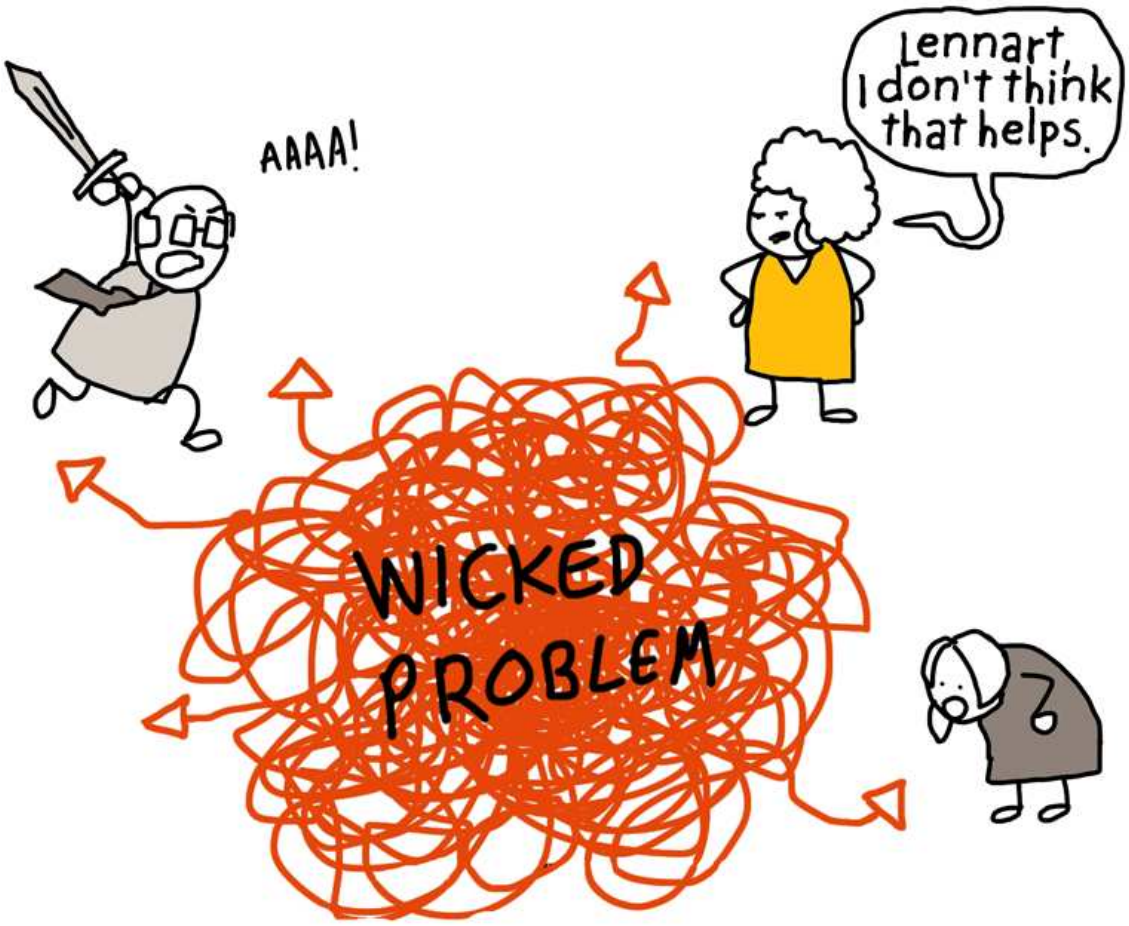

Wicked Problems illustration by businessillustrator.com

CHAPTER 3

SYSTEMIC DESIGN AND STORY EMERGENCE

3.1 What is "systemic game design," and how is it defined in academic research?

When game studies talk about design today, it's not just "making a game." Like in every other branch of design, a lot of recent work frames game design as building and tuning interconnected, complex systems. Games are a group of dynamic structures with a bunch of interacting parts. Game design aims to accomplish this complexity by deciding how those interacting parts connect, behave and produce meaning over time for the player, often in ways you can't fully script in advance. This is what I was referring to as "systemic game design" in previous chapters.

First, we need to understand what a "system" is. A system is "a group of interacting, interrelated, or interdependent elements forming a complex whole" (Salen & Zimmerman, 2004, p. 50, as cited in Akcaoglu & Green, 2018). In their breakdown, a system includes: (a) objects or variables, (b) attributes, (c) internal relationships, and (d) context. In a game like Pac-Man, the objects are things like Pac-Man, the ghosts, pellets, fruit, etc. Those objects have attributes (movement speed, vulnerability windows, point values), and those attributes interact. For example, eating a power pellet changes the ghost's behavior and the player's level of risk. There are also internal relationships, like how player skill and the game's difficulty curve shape each other. The important point here is that none of these parts mean much alone. What matters is how they affect and reshape each other over time.

Systems thinking on the other hand, is an aspect that a game designer should have. It is the cognitive skill of seeing and reasoning about those relationships. It's not just naming the components but mapping how they interact, change, loop back, and create new states (Assaraf & Orion, 2005; Sweeney & Sterman, 2000, as cited in Akcaoglu & Green, 2018). In other words, systems thinking is being able to say "if I change this variable here, this other thing over there will behave differently five turns later, and that will force the player to adapt." In game terms, in-game balancing, pacing, learning curve, randomness, resource management etc. are challenges that game designers should overcome in order to produce a well working game. Research argues that this kind of thinking is high-level and

difficult, because it involves working with feedback loops and sometimes invisible causal links, not just linear cause effects (Akcaoglu & Green, 2018).

These two: game design and systems thinking are also linked in the literature. Games are described as complex systems "composed of various elements that interact with one another in many different and complex ways" (Crawford, 1984, 2003; Fullerton, 2008; Salen & Zimmerman, 2004, as cited in Akcaoglu & Green, 2018). In this context designing a game means planning, sketching, and creating those systems. Crawford even says that the core design challenge is "figuring out how to distill the fantasy of the goal and topic into a workable system" (Crawford, 1984, p. 55, as cited in Akcaoglu & Green, 2018). It reframes design away from just writing a story or drawing assets. It asks a more intangible but important question: what rules and interactions will generate the experience I want?

Thus, we can draw parallels between game design and Rittel & Webber's "wicked problems". A game design problem just like the field of Social Design doesn't have one correct answer. The systems within the game are interconnected, where every problem is a symptom of a bigger problem. The designer does not just pick from a menu of right solutions; you actually construct the space in which "solutions" even make sense (Jonassen, 2000; Nelson, 2003, as cited in Akcaoglu & Green, 2018). Designers frame the problem by asking various questions: What are the objects in this world? What are their behaviors? How do they react to the player? What are the reward loops? Where does difficulty come from? That process identifying variables, mapping relationships, testing and revising is literally system definition (Akcaoglu & Green, 2018).

Another useful lens comes from how game design research talks about agency. Salen and Zimmerman famously describe game design as "second-order design," meaning the designer doesn't directly design the experience but they design the rules that will then shape the player's experience (Salen & Zimmerman, 2004, as cited in Kultima, 2015). In other words, designers can not force the player to feel the tension but they can define enemy speed, ammo scarcity, checkpoint distance, damage output. These different parameters together can shape the player's perception of the game and make them feel in a certain, intended way. The designer works on the underlying rule structure (first order), and the felt experience emerges (second order). Eric Zimmerman later even suggested

that maybe “design is always about second order problems,”. It’s about designing conditions so that certain behaviors or meanings emerge, instead of hard coding those meanings directly (Kultima, 2015). For roguelike/roguelites this feature maps perfectly. Rather than writing and rendering a cutscene about the main character’s desperation, roguelike designers let the resource scarcity or permadeath loop generate that emotion on its own.

Due to this indirectness, design research positions game design as a kind of systems level intervention in player behavior and meaning making. The challenge here is not merely creating content but creating dynamics that constantly evolve once the game is in the player’s hands. Thus, iterative testing, feedback loops, and tuning are very important. Designers continually adjust variables, watch how the system responds, and then rebalance. Design literature describes this as: design is “transformation of existing conditions into preferred ones” (Simon, as cited in Kultima, 2015). Game designers analyze how the current system behaves, decide what “preferred” looks like (more tension, less boredom, more fairness, etc.), and change the system in order to obtain desired outcomes.

Just like there is no “single and simple solution” in Wicked Problems case, there is no “correct” answer to how systems should be constructed and how they should function. Designers carry different priorities and values, sometimes even contradictory ones, and they shift those values depending on the project (Kultima, 2015). This variety of different approaches and points of view are the main drivers of creating different and unique games. So “systemic design” at least in game design as far as I’m concerned is not a neutral engineering activity but it is: political, personal, aesthetic, cultural and social. Choosing what systems to build (for example, punishing failure harshly vs. giving generous recovery mechanics) already encodes what kind of play style, and honestly what kind of player you’re supporting.

Contradictory to that, systemic game design is also framed as a way of thinking that can be taught and measured. When students like myself build games, they have to identify parts of a system, model how those parts influence each other, and then revise when the system doesn’t behave as expected. I will come on to that later in my Design Proposal. Studies that treat game creation as a learning activity describe how learners practice “analyzing and understanding systems,”

then “creating these from scratch,” while also reflecting on consequences, adjusting parameters, and making causal maps of how changes ripple through the design (Akcaoglu & Green, 2018). This tells me something very important for my thesis and my project: systemic game design is not just a “fancy and academic” way to talk about games; it is also a cognitive practice of mapping causes and effects within the limits of certain rules, predicting emergent outcomes and adjusting them accordingly to the right direction.

In summary, academic research defines systemic game design as designing games by constructing and tuning interconnected systems of rules, objects, relationships, and feedback loops that generate player experience indirectly. It treats games as complex systems and game design as the act of shaping those systems so that specific behaviors and meanings emerge during play (Salen & Zimmerman, 2004, as cited in Akcaoglu & Green, 2018; Kultima, 2015).

This view also implies especially in roguelikes and roguelites that storytelling in games can emerge from the systems behaviour. Even in Hades’ example where it has a really successful and rich script that makes the game stand out in its genre, having a script is not enough for good storytelling in games. Because the player is not just an observer, the player is the hero.

The Legend of Zelda: Breath of the Wild, 2017

Don't Starve, 2013

3.2 How do complex systems create emergent narratives in games?

As we discussed in Chapter 2, emergent narrative is not stories written ahead of time but they're cumulative outcomes of how the game's system behaves during play. Instead of an author saying "this happens, then this happens," the story is produced by rules, AI behaviors, world simulation, and player action colliding in interesting ways (Kybartas, Verbrugge, & Lessard, 2020, p. 1). This matters a lot in roguelike and roguelite settings where procedural generation, unpredictability and repetition are central but also in larger open and sandbox settings like Minecraft or Don't Starve Together where people regularly share their story of what happened "in their world", even though nobody scripted that story (Chauvin, Levieux, Donnart, & Natkin, 2015, pp. 1–2).

Let's take Hades' gameplay as an example: you are low on health, you are desperately running away from the monsters while trying to hit them at the same time, you step on a trap and die to the spikes in a miserable and ridiculous way. This whole sequence is "just the system working" but you retell it to your friends

like a story. This is what Kybartas et al. (2020) describe as emergent narrative working almost like “nonfiction”. Players talk about these events as things that actually happened. That feeling is really valuable. It’s part of why people are so attached to their runs and why my heart is racing when I finally make it to the final boss.

So, complex systemic design contributes to story on two levels at once:

1. It enables many possible futures (possibility space).
2. It lets the player feel like they authored the path they actually got (agency).

It lets the player feel like they authored the path they actually got (agency).

Those two together are already narrative. Chauvin et al. (2015) call this “codirection,” meaning the final story experience is shaped both by the designer’s systems and the player’s actions (pp. 1–2, 9).

Last of Us, Sarah death scene, 2013

Bottom-up storytelling vs. top-down control

Traditional narrative in video games has a problem: story wants control, but interactivity wants freedom. Let's take Last of Us for example. One of the story's main pillars is the relationship between Ellie and Joel. We see the character development from both sides and how they form a father daughter bond in a heart warming way. The game does this by introducing Joel's past as one of pure tragedy. Joel loses his beloved daughter Sarah in the beginning of the zombie outbreak. It is obviously very impactful and emotionally heavy from a storytelling perspective. If you really want the player to feel "the pain of losing a child," the surest way is to script that loss and force it to happen. But then the player doesn't really have a say, which breaks agency (Chauvin et al., 2015, p. 5). Emergent narrative design tries not to force moments like that. Instead it tries to set up systems that can produce intense drama without literally hardcoding it.

In interactive storytelling research, there are two classic approaches to dealing with this tension: bottom-up and top-down (Chauvin et al., 2015, p. 6).

- In a bottom-up approach, autonomous characters (NPCs, AI agents) act according to their own goals and rules, and stories emerge from how they and the player interact. There's no "director" forcing specific plot beats. This gives huge agency, but it can also create stories that are technically unique but kind of dull or incoherent (Chauvin et al., 2015, p. 6).

In The Sims series, or in other simulation games, there is no linear story at all. Sims follow AI routines based on traits and relationships. The "narrative" comes from emergent events like jealousy, death, or career success. All consequences come from the player-system interactions.

- In a top-down approach, a "drama manager" watches what the player is doing and reshapes the world so that certain dramatic events happen. This keeps the story focused and emotionally meaningful, but it reduces freedom and makes all players' stories feel more alike (Chauvin et al., 2015, p. 6).

The famous zombie shooter Left 4 Dead's system adjusts pacing, enemy placement, and music based on player performance. It keeps tension and flow across

multiple sessions. The plot remains emotionally structured even though gameplay changes.

There are emergent games like Minecraft or Dwarf Fortress where neither approach works on their own. The worlds are too open, too generative and too underdetermined for a top-down drama manager. It needs well defined casts, props, locations and plot actions to work properly. But in Minecraft, what the system "knows" is basically: you can eat, mine blocks, craft, and try not to die. NPCs might not even have complex goals (Chauvin et al., 2015, p. 6).

So designers are looking for hybrid approaches to overcome the challenge of maintaining the openness while shaping the story in meaningful ways.

The rest of this section looks at three of those strategies:

1. Interpretation engines and narrative processes.
2. Persistence and narrative substrates.
3. Systemic character conflict / tension modeling.

Drawf Fortress, 2006

In systemic games, the raw data is extremely low level. Most of the time the developers and designers can look through player data and see logs like "player picked up item at (x,y,z)". This is merely telemetry, not story. Chauvin et al. (2015) propose an "Interpretation Engine" that continuously watches the live game world and tries to translate these tiny objective events into higher-level story concepts like places, relationships, and threats (pp. 7–8).

They separate what happens in the game into two layers:

- The Objective Story, which is just a chronological log of events ("character A pushed character B at time T, in position P"), and
- The Subjective Story, which is the interpreted version, where the system tries to understand meaning, like "this character is hostile to that character," or "this hut is the player's home base" (Chauvin et al., 2015, p. 7).

The Interpretation Engine basically manufactures story structure out of chaos so that other systems can act on it in a narratively meaningful way.

The enormous 300,000-word script of Hades is a clear testament to how meticulously its storytelling was crafted. The narrative unfolds as the player interacts with the game's systems. When Zagreus offers nectar to a god, titan, or demon as a gift (that is, an Objective Story event), their relationship deepens; new lines of dialogue are "unlocked," and the story progresses (the Subjective Story). Similarly, every event whether it's dying, discovering a new weapon, or defeating a boss serves as an Objective Story that fur-

Chaos featured in the 'Origination' Arcana Card

As a result, both the pace and the direction of the narrative vary from player to player. For instance, one player might try to maintain balanced relationships with all characters, while another may choose to favor a single god or demon. Even though both are playing the same game, the system adapts to their choices and offers each of them a uniquely personal experience.

The next question is: What do you do with that interpreted story? Chauvin et al. (2015) introduce "Narrative Processes," which are scripts players can trigger that alter the world in dramatic ways, in real time (p. 6). A narrative process might spawn a rival, send an NPC neighbor to move in, or threaten the player's home base. The key idea is that instead of a hidden drama manager secretly forcing events, the player can choose to fire off these higher level narrative tools. This further enhances the previous Hades example: the player is not just reacting to authored content; they're actively steering which kind of story tension enters their world (Chauvin et al., 2015, pp. 6–7, 9). This is the innovative part about roguelike/lite design and their approach to top-down drama management. It is like handing players a bunch of keys and doors that match those keys, so players can choose which door to unlock first. The system supports emotional arcs (like rivalry, danger, betrayal) without fully taking away the sandbox feel.

Chaos in Hades 2, 2025

Making stories persist: narrative substrates and player memory

Another issue with emergent stories is that they disappear once the player logs out. Since there are no save files, the world resets every time. You have this "legendary run" or a huge PvP battle, but after you turn off the game, it is like they never happened. From the game's point of view, it's like you were never there. In many MMOs, only your character data persists, not all the stuff you did to the world (Gustafsson, Holme, & Mackay, 2020, p. 1).

From my experiences and what studies suggest that players don't like that. Some online PvP games such as League of Legends or Counter Strike have the option to rewatch your old games but most games do not adopt this feature. Studies show that players highly value one of a kind, first time experiences, and they want those moments to leave marks on the world (Gustafsson et al., 2020, pp. 1–2). When the world can't acknowledge what happened, players move outside the game and build wikis, spreadsheets, videos, clips, edits and war stories on forums to preserve their own "legends" (Gustafsson et al., 2020, pp. 4–5). In other words, players generate emergent narrative anyway, but the game doesn't know how to hold onto it.

Gustafsson et al. (2020) propose "Narrative Substrates," which is basically an architectural layer that records player activity as structured story elements, then turns those elements into in-game artifacts that other players can actually encounter later (pp. 2–3, 5).

The pipeline looks like this:

1. The system logs Story Events: who did what, where, and when.
2. It groups those events around meaningful things, creating Story Artifacts. Example: a sword that "remembers" every duel it was used in, including famous kills, so the sword itself carries history (Gustafsson et al., 2020, pp. 5–6).
3. Those artifacts persist in the shared world, so future players can literally inspect them, see the log, and understand that "this item was part of that drama." The past becomes explorable content, not just memory (Gustafsson et al., 2020, pp. 5–6).

They tested this idea in a live MMORPG prototype called We Ride, where rare items evolve over time and publicly "tell" about the fights they've been in. The game even has "town criers," NPCs who announce noteworthy emergent events to everyone, basically acting like in-world storytellers (Gustafsson et al., 2020, pp. 5–6).

From a design perspective, it solves two problems. First, it acknowledges player action as canon, further encouraging players to create more stories and connecting them with the game's community. Second, it gives designers a way to curate and surface interesting emergent stories without having to pre-write them. The system can detect relationships like "this player repeatedly killed high-ranking enemies" or "this location has become meaningful because people keep returning there" and then expose those threads to others in a way that feels intentional, not random (Gustafsson et al., 2020, pp. 3–5).

This is essential for roguelite-style storytelling because they rely on repeat play and repeat play has the risk of feeling disposable. Narrative Substrates points toward a design pattern where your past runs can literally become part of the shared fiction and material culture of the world.

Hades 2 end-run screen, 2025

Children of Morta, 2019

Modeling social tension as narrative fuel

A third way complex systems generate story is by simulating characters with goals, values, and conflicts, then letting those conflicts play out. Kybartas et al. (2020) describe emergent narrative as "stories... created through a simulation of characters in a virtual storyworld," where the interesting parts come from chains of behavior over time (p. 1). The designers define characters, their priorities, their relationships and possible sources of conflict. The resulting simulated outcomes such as clashes, alliances, betrayals are the plot.

This is an extremely powerful tool but equally hard to author since the designer has to build an ecosystem of interacting motivations, not one hero's arc (Kybartas et al., 2020, p. 1). The number of possible outcomes can be exponential: every new added system multiplies what can happen. The upside, though, is that the narrative space becomes "pregnant" with potential drama; the world is set up so that meaningful conflict is always one step away (Kybartas et al., 2020, p. 1).

I think this direction brings us back to the central point of this section: emergent narrative is narrative distributed across systems. Complex systems generate story when:

- the world reacts coherently to what the player does,
- the traces of those reactions persist and matter, and
- the player is allowed to treat those traces as meaningful.

A procedurally generated dungeon map in the video game NetHack

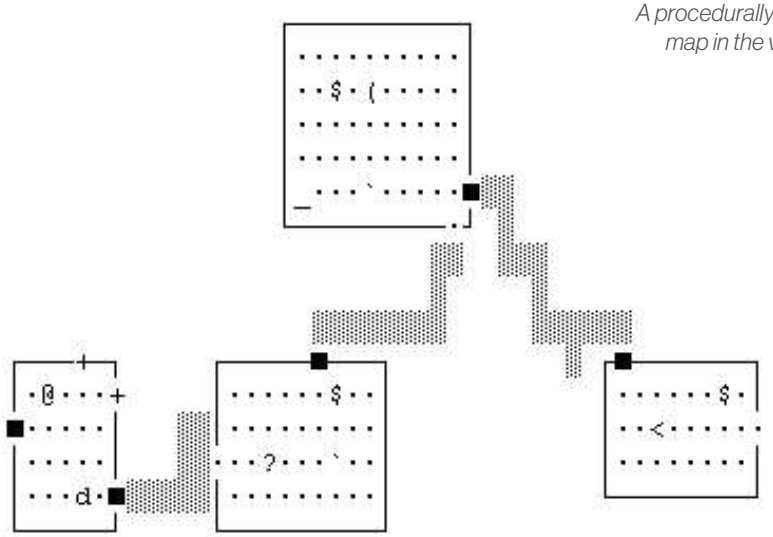

```
Joe the Tenderfoot      St:14 Dx:9 Co:13 In:14 Wi:17 Ch:8  Chaotic
Dlvl:1  $:0  HP:14(14) Pw:3(3) AC:7  Exp:1 T:82
```

3.3 Procedural Content Generation and Storytelling

Procedural content generation (PCG) is usually described as the algorithmic creation of game content like levels, quests, maps, characters, or even rules, either while the game is being played or during development (Risi & Togelius, 2020, p. 1). There are few main reasons why PCG is a practical choice in modern game development: it keeps games replayable, it cuts production costs, and it lets designers build worlds that would be impossible to author by hand, like the endless planets of No Man's Sky or the constantly changing dungeons of Rogue (Risi & Togelius, 2020, p. 1–2; Pereira de Araujo & Souto, 2017, p. 444). As discussed before, in emergent storytelling and systemic game design how "things" create meaning for the player and for the game. PCG is not just an easy shortcut to create more things, it is a tool to create consistent assets and situations that matter for the story. PCG can feed the constant interaction between the system and players in three main ways: by generating worlds with implied history, by generating quests and goals that react to the simulated world and by generating characters or situations that feel personal to the player.

Procedurally generated world of No Man's Sky, 2020

First, PCG can make the world feel like it already has a past. A world that looks like it has been lived in gives players material to interpret. This is important because interpretation is also storytelling. A good example is what some researchers call "generative archaeology" or "archaeological storytelling," where a game builds a ruined space and asks the player to reconstruct what happened there by exploring it (Smith Nicholls & Cook, 2023, p. 1–2). In their study of Nothing Beside Remains, the game uses procedural rules to generate a deserted village. All buildings, objects, environmental traces, even climate signs like dried water sources are procedurally generated. Players are then asked to write their version of the village's downfall based only on what they can observe (Smith Nicholls & Cook, 2023, p. 2–3). The interesting part is that the game is not giving an explicit "main quest" to solve the mystery. Instead, it generates material evidence such as architecture, object placement, environmental damage and the player concludes cause and effect from that evidence. The researchers frame this as "generative archaeology games," where the player basically behaves like an archaeologist, building narrative from environmental clues that were themselves generated by code (Smith Nicholls & Cook, 2023, p. 2).

In traditionally authored stories, story objects are placed to communicate one

canonical truth. In Red Dead Redemption 2, after Arthur defeats his enemies and searches through the hut, the objects that he needs to investigate are highlighted. Here, some of the content is authored (like a recurring statue that anchors every village seed), but other parts are produced by simulation, like the state of the ecosystem and which buildings still stand (Smith Nicholls & Cook, 2023, p. 2–3). That mix can cause tension in interpretation. In fact, Smith Nicholls and Cook point out that even a glitch influenced how players told the story of the village, because players tried to "explain" the glitch diegetically, as if it were part of the fiction (2023, p. 3–4).

We can connect that with how PCG is used for large-scale simulated histories. PCG is also capable of creating more "intangible" assets (needless to say there is no such thing as intangible for computers but more intangible for us humans), such as political events, lineages, wars, and myths. This produces what players often treat as "emergent narrative": memorable events that feel like personal stories even if they were not written by a human writer in advance (Prins et al., 2023, p. 1; Pereira de Araujo & Souto, 2017, p. 444).

The second way PCG supports storytelling is by generating quests. Whether it's a traditional cutscene driven storytelling game or an emergent one, quests have always been one of the most important aspects of video games. They're used in various ways. In online PvP (person vs. person) games they function as requirements to obtain rewards: "Slay 150 monsters, get this". Nowadays almost every online game has a "battle pass" system, which undermines the fun of completing quests in my opinion. In story driven role playing games, quests function as chapters in novels. They tell the players what matters right now. Classic RPGs often hand author quest lines in branches. But branching by hand is expensive and it still tends to collapse back into a few endings. The question is: can PCG generate quests that stay coherent with the current world state and still feel narratively meaningful?

Recent work says yes, or at least partly yes. StoryWorld is one example of a procedural quest generation system designed to create quests that are both

believable and varied (Prins et al., 2023, p. 1). The system simulates a world with characters, locations, items, relationships, and even memory and desire. Characters have things they want, things they remember, and opinions about other characters. The system then uses those factors to generate tasks for the player, and those tasks are delivered as quests (Prins et al., 2023, p. 1–2). In good stories where quests are treated as chapters, quests can't be random errands like "fetch 3 items". They should be motivated so it motivates the player as well. An NPC asks you to retrieve an item because that character actually remembers losing it or desires it, and those internal states exist in the simulation (Prins et al., 2023, p. 1–2). When you complete or fail the quest, the world state permanently changes, including relationships and knowledge, so later quests are generated in a world that "remembers" what you did (Prins et al., 2023, p. 1).

That memory link matters for narrative and character development. It holds us accountable as a player and forces us to pay attention to our actions. Prins et al. argue that this can make the world feel "more believable and thus more immersive," because NPCs act like social agents who react in sensible ways instead of just vending missions (2023, p. 1–2). So PCG quests act as dynamic plot hooks that emerge from simulation and then feed back into it. This is very close to emergent narrative design, but with the added structure of goals and consequences, which helps avoid pure chaos.

"Random NPC" is a commonly used term these days. It became a meme on social media. People claimed that they are the main character and everyone else they see on a daily basis are non-playable characters or NPCs for short. It is used almost like an insult to tell people that they don't have any motivations, emotions or feelings. Obviously, actual NPCs in video games have none of that but they must trick our brains to make them feel that they are real for an immersive experience. Thus, in order to have more successful and immersive storytelling, characters must feel less NPC-like.

In this regard, the third contribution of PCG to storytelling is character level personalization. Some systems use procedural generation not just to make "a ran-

dom NPC," but to generate a specific someone with traits, needs, and context that feel narratively grounded. StoryWorld, for instance, attaches memories and "Desires" to each character, and quests are partly driven by those internal states (Prins et al., 2023, p. 1–2). Other work uses PCG to generate patients with different physical traits and personal situations in a medical training game, and then builds short narrative scripts around them so each encounter feels like a convincing case, not just a repeated prompt (Pereira et al., 2019, p. 192–194). This personalization creates stakes. If two players meet "the same" quest type, but the NPC involved is different, with different history and different emotional framing, then the event can feel like "my story," not just "the quest everyone gets."

However, there are limits and risks in PCG models. Repetition is one of key drivers of these systems and roguelikes in general but when overdone, it can feel really boring really fast. PCG systems can generate thousands of assets like buildings, people, enemies, politics, rules etc. but players will still notice patterns and similarities. This is what Compton calls the "10,000 bowls of oatmeal" problem: content that is technically varied but emotionally flat, because it follows the same structure and tone over and over (Prins et al., 2023, p. 1). Another issue is that procedural systems can easily break believability. Pereira de Araujo and Souto point out that pure PCG can create huge worlds, or infinite quests, but that doesn't mean those worlds or quests stay interesting, or even coherent, over long play (2017, p. 444–447). Players will start to see through the trick if the system can't maintain meaning.

There's also the problem of control. "Giving narrative to an algorithm" can mean losing authorship, pacing, and tone for designers and writers. Prins et al. explicitly acknowledge that designers are often afraid of procedural narrative because it feels like losing control of the story (2023, p. 1). The research response is not to remove control entirely, but to structure it differently. For example, de Lima et al. (2022) use "story arcs" and planning constraints to guide what the branches should look like. StoryWorld uses desires, memories, and social relationships as constraints on what kinds of quests can even be proposed (Prins et al., 2023, p. 1–2). In other words, the human author doesn't write every quest, but they do de-

sign the social and dramatic rules that make quests feel motivated.

Finally, there's the issue of player interpretation. The human mind tends to fill the incompleteness. When the system gives you incomplete information like a ruined village and a few suggestive props, the player fills in the blanks. Sometimes they even explain bugs as lore. Smith Nicholls and Cook read this not as a flaw but as a tool: players are already doing interpretive labor similar to actual archaeology, building a story out of fragmented evidence (2023, p. 2–4).

Across all these cases, the shared idea is that procedural generation can act as a narrative engine. It can simulate causes and consequences so that when something happens in the game, it feels like it happened for a reason in that world. It can surface those reasons through quests, worldbuilding, and character motivations. And it can hand the player enough structure to make sense of what happened without forcing them down a single prewritten path. When it's used not as a replacement for good storytelling, but as a tool to enhance a good story like in Hades it can help designers to create content rich and not overly repetitive emergent games. For roguelikes and roguelites, which reset the world constantly and lean on replay, this is crucial. If every run can suggest a slightly different "why," and not just a different "where," then replay stops being mechanical grinding and starts to feel like alternate histories.

A procedurally generated dungeon in Enter the Gungeon, 2016

Bag of Dungeon, 2018

CHAPTER 4
BRIDGING DIGITAL AND PHYSICAL: DESIGNING A ROGUELIKE BOARD GAME

4.1 Systemic Design in Digital vs. Tabletop Games

Systemic game design, as discussed earlier in this chapter, is about building sets of rules, agents, resources, and interactions that produce experiences during play. However a system's look varies depending on the medium. Whether it is a digital roguelite game or a narrative-heavy tabletop RPG, they all rely on interacting systems. But they distribute that systemic work very differently between the designer, the software (or lack of software), and the players. In this section, I will discuss what are the main differences and similarities of systems in digital vs. physical games and why those differences matter for storytelling and experience.

One big difference is how rules are enforced. In digital games, the code is the authority: it tracks states, applies constraints, and delivers feedback in real time. The digitalization of rules via code allows games to coordinate a lot of moving parts, and it supports complex learning and decision making environments. For example, large-scale reviews of digital games for learning have shown that digital games can create "complex gaming environments for learning," where players work toward goals, get feedback from the system, and develop conceptual understanding, process skills, and even epistemological framing (Clark et al., 2016, pp. 79–80). Because the rules are embedded in software, designers can build layered systems and trust that the game will apply them consistently for every player.

On the other hand, tabletop games do not have automatic layered systems. The players themselves (like the Banker in Monopoly) or a facilitator like the dungeon master in Dungeons and Dragons, have to instantiate and maintain the system. Martinho and Sousa (2023) point out that, in an analog tabletop game, "the player's agency tends to be higher. Without direct player activation of the game mechanisms, the game would not function" (p. 2). This is an important statement because unlike digital games, in tabletop play the system isn't just what's printed in the rule book. The system also becomes the social contract, the shared attention, the willingness to enforce or bend rules and sometimes the mood at the

table. Let's take a famous card game example: UNO. After being forced to draw +2 or +4 cards, being able to play or not is strictly written in the rule book. Nonetheless players willingly bend and shape the rules within smaller social groups. They may choose what makes more sense to them. Martinho and Sousa (2023) argue that a tabletop game "depends as much on the design as on the players and the context of play," and that two sessions of the same game with the same people can still feel different because the social and physical context shifts (Martinho & Sousa, 2023, p. 2). Shortly, the system in tabletop is partly and directly authored by the players every time.

The direct authorship by the players causes a direct impact on story emergence. In digital games, story tends to emerge from internal simulation. The player's actions collide with the game's systems, and the system responds. Because all those responses are computed and tracked, designers can tie them to goals, rewards, pacing, and difficulty in a controlled way. Clark et al. (2016) describe work where "augmented game designs" (basically, versions of a game with extra structures) helped players learn more than standard versions, with a measurable positive effect on learning outcomes (p. 79). That means the system can be tuned to pull players through a sequence: teach concept, test concept, escalate challenge, face consequence, repeat. This tightly managed adaptation is much easier to guarantee in a digital system than around a table.

Tabletop systems adapt this sequence with social interaction and materiality, instead of coding. Kosof (2021) looks at the role of the dungeon master (DM) as something more than a mere storyteller. The design of tabletop play itself gives the DM a social designer's role: a facilitator, a systems designer in real time. The DM has to create dungeons, puzzles, and boss fights that escalate challenges "in a manner similar to the relationship between lectures, coursework, and exams in classrooms," even if they aren't professionally trained designers (Kosof, 2021, pp. 3–4). Kosof argues that tabletop campaigns demand "a great deal of preparation" and that DMs basically function like live game designers who have to balance pacing, mechanics, and dramatic payoff (Kosof, 2021, p. 3). Keeping the system flexible here is a choice. The DM can react to player creativity, bore-

Devon Chulick: "How I Became a Professional Dungeon Master", 2020

dom, or chaos in the moment, and adjust difficulty or narrative weight just like a roguelike's PCG engine reacts to player behaviour and creates different assets on the spot. Digital games offload this mission to code, while tabletop games offload it to players.

Where motivation comes from is also another difference between digital and tabletop systemic games. In digital games, motivation is often structured through mechanics and progression systems. The game tracks and rewards the player and this keeps them moving. Meta analyses of game based learning show that these structured systems don't just keep players engaged, they can help them build both knowledge and intrapersonal outcomes (like persistence or confidence), compared to non-game instruction (Clark et al., 2016, pp. 94–96).

In tabletop play, motivation is more entangled with social presence and imagination. Martinho and Sousa (2023) build a motivation model for tabletop players and find five main dimensions:

- Competitive interaction,
- Social challenge,
- Sensory experience,
- Intellectual challenge
- Imaginative experience (p. 1).

What's interesting is that not all of these come "in the box". These five dimensions show us that tabletop players often care about other activities around the game. They like discussing it, customizing components, even photographing parts as much as (or more than) just "winning" (Martinho & Sousa, 2023, pp. 1–2). It's more about social dynamics and collaboratively building fiction. That means the system of a tabletop game bleeds outside the match itself. The prep, the table talk, the house rules, and even the after-action debrief all become part of the ongoing narrative system. Digital games or publishing platforms sometimes get close to that (for example, Steam Community where players can share mods, sell and buy items etc.) but the baseline design of a digital game doesn't assume you'll rewrite its rules mid session. Tabletop design basically assumes you will.

Finally, there's a cultural difference in how the two spaces think about authorship. In digital systemic games the rules are bound to what our character or the world is capable of. The game world's physics matter: if you fall you lose HP (hitpoints) or click on "W" to go forward. There are also rules related to timing, positioning, turns and more but they're not necessarily considered as rules. Because the design tends to hide its machinery behind interfaces and worldbuilding. Players mostly experience the consequences of their actions, instead of considering staying within the rules. In tabletop, the rules are often visible, argued about, and changed. Martinho and Sousa (2023) note that hobby tabletop communities constantly talk about "new game systems and game mechanisms," hack rules, and even add "house ruling and modifications" as a normal part of play (pp. 2–3). It is safe to say that systemic literacy is part of the hobby. Digital systemic design often treats the player as a subject inside the system. Tabletop systemic design often treats the player as a collaborator on the system.

So, if we're thinking about roguelike / roguelite storytelling and ultimately connecting this knowledge to create a tabletop game, this distinction actually matters. Roguelikes generate narrative moments by leaning on procedural systems and creating combat loops, resource scarcity, and risk/reward. That lineage is closer to digital systemic design's "closed but reactive world." Tabletop roleplaying leans on social systems such as trust, improvisation, negotiated canon to

generate emergent narrative. That lineage is closer to analog systemic design's "open but co-authored world." They're two different ways of distributing systemic control across humans and rules. Understanding that distribution is going to be important for thinking about how emergent narrative actually forms in systemic games.

Goblin Cafe, a tabletop social club and store in Istanbul, Türkiye

4.2 How is procedural generation implemented in tabletop games?

The roots of procedural generation reach much further back than algorithms. Pen and paper systems relied on randomness and rules to produce variation. In tabletop contexts, procedural generation simply becomes a human-executed algorithm. Players shuffling cards, rolling dice, drawing from a deck of cards, assembling components according to predefined conditions to create new content each session. The principle of PCG, complexity emerging from simple rules, remains the same as in digital games but its execution is mechanical rather than computational (Hafis, Supianto, & Tolle, 2019, pp. 309–312).

At its core, procedural generation provides variability and replayability. When we directly compare digital and tabletop games, the correlation between them becomes more apparent. For example, in Monopoly, the physical game board remains the same but the "map" of the game changes depending on which properties are still free or which ones are taken. Let's say, in the previous run player A had the full set of green properties so it was safer for player A to move on the last quarter of the board. However in this run, player B got all the green properties. Even though there is minimal decision making and more luck factor in this case, player A needs to plan accordingly to get through the last quarter of the board. Freiknecht and Effelsberg (2017) describe procedural systems as "algorithms that generate assets with minimal human input while following structural constraints" (p. 3). Translated to board and card games, this algorithm is replaced by physical constraints and chance operations (like in Monopoly). It allows players to construct a new play space from the same set of materials. The designer specifies a limited set of components such as properties, tiles, cards, dice, and defines how they can combine. Players' decisions and actions generate the outcome in real time. In this sense, tabletop procedural generation maintains the spirit of the digital approach but grounds it in tangible interaction.

The Settlers of Catan

Modular boards

The most direct form of procedural generation in tabletop games appears through modular boards. Games like The Settlers of Catan came to my mind immediately. It shares similar features of roguelike games. At the beginning of each Catan game, the resource tiles and their associated dice roll numbers are randomized on your board. If you play with the Seafarers expansion, even your board itself can be randomized as well with land tiles and sea tiles being mixed up around your board. Once you start the game, you're competing until somebody gets a certain number of victory points, in which case the game ends and your run is over. The next time you play, you're completely resetting the game with new randomized layouts and resources, and resetting your points to zero. Finally, it is turn based. Everybody gets to make their decision on their turn and then pass their turn to the next player. Catan ensures coherence with each layout being playable, while constantly altering the distribution of resources. According to Freiknecht and Effelsberg (2017, pp. 4–5). This approach corresponds to constructive generation in digital terms, it is fast, predictable and easy to administer, which makes it ideal for physical formats.

Monopoly Deal

Dice-driven tables and flowcharts

Another widespread (but kind of antique) method relies on dice tables and flowcharts. In early tabletop role-playing games, long before computers handled randomization, dungeon masters (or moderators in other games) rolled dice to determine dungeon layouts, encounters, or treasure. Each roll functioned like a digitalized branching algorithm. Each result led to a different narrative path or spatial configuration. Risi and Togelius (2020, pp. 6–8) note that these early pen-and-paper systems already implemented procedural generation in a primitive but effective way, using probability and conditional logic to build new situations. There are still many modern board games that still use these types of structures. Each roll effectively becomes a "function call" that generates new content without human authorship. This demonstrates that procedural principles can exist entirely offline.

Card stacks as generators

Card-based tabletop games extend this logic further by turning decks into physical generators. Cards can determine events, rooms, or narrative outcomes when shuffled and drawn under certain conditions. In Betrayal at House on the Hill, for example, players explore a mansion by drawing room tiles and event cards, revealing new areas and story fragments that combine into unpredictable scenarios. Similarly, the game 504 uses modular rule books and components that recombine into entirely different rule sets each session. Risi and Togelius (2020, pp. 9–11) describe these practices as "rule-based content assembly," where modular pieces behave like variables in a procedural grammar. Each draw or combination acts as a deterministic process bounded by probability, producing variety without compromising coherence.

Monopoly Deal, successfully uses a card-based generation approach. It dematerializes the original form of the well-known board game by removing the board, houses, hotels, money, dice and all other pieces, just leaving a thick deck of cards and the players. The deck contains physical components from the original board game like properties, houses and hotels, but also action cards instead of pieces and dice. Every player starts with five cards and draws two cards each turn. Every player can do a maximum of three actions per turn: putting down properties, stealing other players' sets, putting cards on the side to use as money, drawing cards etc. First player to complete three sets, wins the game.

In my opinion, this version works perfectly for a few reasons:

1. It cuts the playtime significantly and encourages them to play more runs.
2. It opens up more possibilities. Instead of one dice-roll, and one action, players can take three actions.
3. It gives the chance to counterplay. In Monopoly Deal, there are rare "Reject" cards that allow players to mess up someone's plans.
4. It assures more interaction. Unlike the classic version, the players can play action cards on every player without impediments or requirements.
5. It is portable. The game reduces its physical components and fits everything in a small carton packaging. This feature alone yields much more playability by allowing players to bring the game with them.

The Keeyp, 2021

Balatro, 2024

Tile grammars and adjacency rules

Procedural generation also operates through adjacency and grammar-based rules. In digital design, spatial generators use grammar systems or cellular automata to create coherent worlds. Physical games achieve the same through placement constraints. Individual components are not just bound by the rules but they are also visually connected to each other like puzzle pieces. Tiles must connect through open doors, corridors must align, and hazards cannot block progression. McCooey (2021, pp. 45–48) connects this principle to constructive design models, showing that local constraints (rules about how one piece can sit next to another) can yield globally consistent structures. The result is a playable map produced step by step, much like an algorithm assembling a maze.

Entropy control: seeds, ranges, and difficulty bands

According to the definition in Cambridge Dictionary website, entropy is described as: "the amount of order or lack of order in a system" (Cambridge University Press, n.d.). In this regard, whether it is a digital or physical game, entropy control is key for procedural systems, because unbounded randomness may lead to incoherent or unfair results. Hafis et al. (2019, pp. 313–315) stress that the success of procedural systems depends on balanced randomness. Randomness, widely known as RNG (random number generator), in games adjusts variation that surprises without disrupting structure. In tabletop games, this is achieved by setting fixed ratios (for instance, only two high-value tiles, three medium, and several low ones) or by phasing elements in over time. This controlled variability ensures that outcomes remain believable and that players experience both novelty and stability across sessions.

Campaign state as "dynamic seed"

Especially in deck-building roguelike games like Balatro, Slay the Spire or Monster Train, dynamic seeds are largely used in the genre. In Balatro , a poker themed deck building roguelite players can dial in a "seed code", guaranteeing to encounter certain Jokers, Planet and Tarot cards, and other shenanigans that lead the game towards a certain direction. These seeds can be found on the internet. Obviously this shouldn't be used as a cheating method to get easy wins

because it would suck the joy out of the game. Players generally use seed features for limit testing certain builds, speedruning, record breaking or making content. It is possible to see titles on YouTube such as: “Balatro Set Seed Glitched - 1m25s Speedrun” (RandomisedTV, 2025).

Long-term tabletop campaigns often adopt this layer of procedural logic through persistent states. Stickers, altered decks, or narrative consequences from previous sessions act as dynamic seeds that influence future generations. Freiknecht and Effelsberg (2017, pp. 6–7) note that digital procedural systems frequently use seed values to reproduce or modify worlds deterministically. Legacy-style board games borrow this concept. Each play leaves traces that bias future configurations. The record of past choices becomes a physical variable guiding the next iteration, instead of numerical seeds seen in digital games. Over time, this produces evolving play spaces that feel algorithmic but remain entirely material.

Procedural generation in tabletop games, therefore, is not a lesser imitation of its digital counterpart. It represents the same design philosophy executed through tangible rules and probabilistic structures. Both rely on constraint, randomness, and repetition to produce novelty. The difference lies in who performs the computation. In digital systems, algorithms process data invisibly; in physical games, players execute the algorithm through their actions. As Risi and Togelius (2020) suggest, procedural design ultimately aims to “delegate creative responsibility to systems capable of surprise” (p. 12). On the tabletop, those systems are made of cards, dice, and human hands. Procedural generation in physical games transforms algorithmic logic into tactile ritual. Each shuffle or roll stands in for code execution.

We can subtract important lessons regarding game design. Introducing these systems into tabletop games by balancing randomness with constraint and modularity with clarity can evoke the same sense of unpredictability and discovery that defines digital roguelike and roguelite experiences.

Ben and Ed, 2015

4.3 How can permadeath and progression systems be applied to physical games?

The narrative strength of permadeath lies in consequence and emotional impact. Translating this mechanic into a physical format means recreating that sense of loss and finality. Tabletop games do not have automatic save systems or programmed resets, so designers must embed irreversibility into the material form of the game. This can be achieved via physical components of the game like the board, cards, or via rules themselves.

In Pandemic Legacy: Season 2, players are instructed to tear, mark, or permanently modify components when major events occur. Characters can die permanently, locations can be destroyed, and the game board slowly changes across sessions. As Imbierowicz (2019) explains, this inability to "relive a chosen passage" transforms each campaign into a unique, unrepeatable story (p. 84). Once a card is crossed out or a sticker is placed, the old state cannot return. This

simple act of physical change mirrors the logic of permadeath: the world moves forward, and choices have weight. According to Imbierowicz, the permanent loss of options "increases the stakes of play" and makes the narrative outcome "matter more" precisely because it cannot be undone (pp. 85–87).

This irreversible structure changes how players relate to their characters – the consequence and the anxiety of dying makes us empathize with the character. The absence of a digital save system forces players to live with their mistakes. Juul, as cited by Imbierowicz (2019, p. 86), argues that this inability to reload heightens emotional investment because "the risk of permanent loss creates meaning."

However, we can see that permadeath alone is no longer sufficient in the modern gaming landscape. Without progression systems, players might perceive permanent loss as punishment rather than consequence. As mentioned in Chapter 1, the roguelikes' descendants (the new roguelites) addressed this issue by integrating meta progression into their design. Progression counterbalances death by providing a sense of growth and continuity. The characters, NPCs and the world keep evolving after death. Legacy style board games handle this by letting players alter rules, unlock new regions, or enhance surviving characters after each session. Imbierowicz (2019) notes that in Pandemic Legacy: Season 2, players can "enhance their characters by adding stickers with new abilities" and open sealed compartments to reveal new cards or mechanics (pp. 83–84). This design blends loss and growth into a single rhythm. The repetition becomes a sense of continuity and traps us players (in a good sense) in a loop.

This approach mirrors what West (2020) describes in her study on permadeath and player psychology. She found that permanent death can generate feelings of grief and reflection rather than frustration when the game allows the player to attach meaning to the loss. According to West, the strength of attachment determines the emotional intensity of death: "the stronger the identification, the stronger the grief, and the greater the appreciation that follows" (pp. 97–99). Seeing the mark of that death on the world remembers what makes these experiences meaningful.

Permadeath in tabletop games also discloses something about human practice. It is not a mechanical act to write a character's name and then cross it out. It is a symbol that shows loss. West (2020) links this to grief story, arguing that physical rituals help players process loss by making it tangible (p 103) in a video game character's death is often followed by a restart. However, in a physical game, the destroyed component becomes a permanent memorial. The emotional hit comes from material permanence. The scar remains visible on the table.

The balance between loss and advancement is a way of keeping campaigns alive. Every irreversible event becomes a part of a larger narrative arc, forcing players to adapt rather than restart. Physical progression systems like unlocking new abilities, "eudaimonic appreciation arises when players perceive suffering as meaningful within the story's structure" (p. 112). Physical progression systems like unlocking new abilities, revealing hidden regions, or upgrading surviving heroes provide the context that makes sacrifice worthwhile.

This balance is clearly seen in games in combining character sheets, campaign logs, sealed decs. These components act like manual save files. When a player unlocks a new ability or suffers a permanent injury, they record it directly on

Pandemic Legacy: Season 2, 2018

the sheet. Over time, the document becomes a chronicle of victories and losses. Imbierowicz (2019) explains that this method “prevents repetition of the same events” while creating “continuity between sessions” (p. 86). The result is not endless replayability, but continuity of consequence, each decision builds upon what came before.

In order to achieve the balance, designers use some techniques :

1. To make alteration material : They ensure that consequences are visible and binding. For example, a sticker can be more effective changing a rule than a note in a manuel. Because it physically overwrites the past.
2. To provide meaningful upgrades after irreversible uses :They transform grief into motion. Losing a character might unlock new missions, reveal hidden cards, or permanently after the map turning defeat into narrative evolution.
3. To refuse allowing replays or do-overs : They preserve authenticity. Imbierowicz (2019) argues that the inability to “relive” a campaign is essential to its emotional gravity, since replay would erase the uniqueness of the experience (pp. 84–85).

Permadeath and progression give physical games a rare quality : That is memory. The board itself becomes a living archive of past actions. Scratches, stickers, and missing cards tell the story long after the session ends. West (2020) notes that this lingering evidence of mortality creates “lasting appreciation rather than fleeting excitement,” allowing players to reflect on their own agency within the game world (p. 118). Based on this, the mechanics of death and growth become tools for storytelling. They give shape to emotion, permanence to consequence, and authenticity to play.

In conclusion, adapting permadeath and progression systems to physical games is not about imitating software functions, it is about transferring meaning. Where a digital game deletes a file, a tabletop game writes over its own surface. Where a computer saves data, the players themselves preserve memory through ink, scars, and altered rules. Both forms concern about consequence, but the physical form turns results into ritual. It asks players not only to play but to remember.

D&D game at the Brooklyn Strategist, New York Times, 2022

4.4 How does social interaction influence emergent storytelling in tabletop games?

As discussed before, emergent storytelling in tabletop games is the direct outcome of people negotiating a shared fiction together. Around the table (or a virtual tabletop), players constantly shift between three "frames" of activity

1. Narrative (in-character talk and scene description)
2. Game (rules, dice, tactics)
3. Social (banter, logistics)

and these rapid shifts are exactly where stories gain texture and momentum (Webb & Cesar, 2019). When the group smoothly moves across frames, arguing rules for a second, then retreating into character to deliver a line, the fiction progresses in ways no single author could have planned. When those transitions snag, stories stall.

Frame-shifting as a story engine

Ethnographic work on distributed tabletop play shows that technology and media create "seams" where frames intersect voice channels, private messages, dice bots, music queues and players actively exploit those seams to steer narrative emphasis (Webb & Cesar, 2019). A whispered side-channel lets a rogue split from the party without breaking pacing for everyone else; a synchronized sound cue pulls attention to a wounded NPC; a rules macro pushes the table briefly into "game frame" to resolve stakes, then hands the spotlight back to the fiction. The point is social coordination tools are forces that grow tension, redirect focus, and seed new story beats.

Negotiation, authority and the "soft canon" of play

Tabletop stories need enough stability in order to continue from week to week, but too much rigidity may kill discovery. Recent work on Dungeons & Dragons communities describes how groups manage this with "soft canon": a shared understanding that privileges emotionally resonant details over a rigid, single chronology (Wee, 2024). In practice, multiple versions of events (what could have happened and what did) as coexisting resources are used by players. A gift mentioned in play becomes a forum post, then a physical prop; all of these artifacts circulate and reinforce the group's world without forcing one authoritative record (Wee, 2024). Social interaction here is the mechanism that confirms which bits "stick" and which remain speculative. The result is a world that feels lived-in, precisely because its meaning is co-authored and constantly renegotiated.

Conflict as fuel and as a risk

Strong emotions and disagreements can be narratively productive (they raise stakes) but if unmanaged, they also can be fractured communities. A cross-modal ethnography of role-playing communities identifies recurring sources of social conflict: creative agenda clashes, GM player power imbalances, online miscommunication, and "bleed" (spillover between player and character emotions) (Bowman, 2014). They're structural pressures in any longform campaign. The point for design is twofold. First, conflict can be a necessary ingredient for story

progression, but it must be framed with fair leadership and explicit safety norms to avoid community schisms (Bowman, 2014). Second, tools that make status and consent visible session zero expectations, lines/veils, private channels for de-escalation aren't "extras"; they are story infrastructure because they keep the group functional enough to sustain emergent arcs.

Persistence and the social memory of play

Digital games solve continuity by saving to a database; however tabletop groups keep the story through social memory and artifacts. Research on "narrative substrates" argues that making traces of player action persist, turning them into first-class objects enables unique events to accrue meaning and drive future play (Gustafsson et al., 2020). Tabletop communities already do this informally through session notes, handouts, recap posts, house loresheets, even nicknames for recurring opponents. These traces are social: players curate, remix, and circulate them, which is why they're powerful. The lesson from narrative-substrate theory is to deliberately externalize memorable actions and let them feed forward into new scenes. That practice increases players' sense that the world "remembers," which in turn encourages risk-taking and richer emergent narratives (Gustafsson et al., 2020).

Physical adaptation of Dead Cells, 2024

D&D representation in the first season of Stranger Things, 2016

The GM as improvisational orchestrator

Emergence needs restrictions. Classic work transferring GM practice into interactive storytelling frames the core tension as a balance between authorial guidance and player freedom (Peinado & Gervás, 2004). Good GMs use heuristics to choose the obvious, the challenging, the surprising, or the most pleasing outcome, then check consequences to keep momentum while honoring agency (Peinado & Gervás, 2004). Socially, that looks like active listening, spotlight rotation, and quick negotiation of rulings. The effect on the story is immediate: players feel licensed to propose bold actions because the table's improvisational contract is clear. When the GM over-specifies, players become passive; when the GM withdraws, disagreements expand. The sweet spot is a visible judgement style that makes tradeoffs fair to everyone.

Human conversation as the narrative substrate

A complementary line of work literally models tabletop as dialogue turn taking between a "storyteller" role and player roles and shows that varying conversational dynamics (who initiates, who replies, who interferes) produces different story shapes (Tapscott et al., 2018). This maps neatly onto live play: side comments can leave the tracks or elevate a scene; a well timed interjection reframes a conflict; a private message unlocks a twist. The practical insight is mild but powerful: monitor conversation patterns. If two voices dominate, your emergent story will narrow; if you raise structured turn-sharing (without killing spontaneity), your fiction gains breadth.

Social interaction shapes emergent storytelling by:

1. Enabling rapid frame-shifts that inject energy into scenes (Webb & Cesar, 2019)
2. Maintaining a flexible, emotionally grounded "soft canon" that legitimizes multiple perspectives (Wee, 2024)
3. Channeling conflict productively while protecting the group from bleed and schisms (Bowman, 2014)
4. Externalizing memorable actions so the world appears to persist (Gustafsson et al., 2020)
5. Applying GM heuristics that balance surprise with coherence (Peinado & Gervás, 2004)
6. Treating conversation flow itself as a designable system (Tapscott et al., 2018).

Example : A secret Discord whisper leads to a solo rescue that later becomes a table wide legend; a recap thread "canonizes" a player's dramatic confession because everyone quotes it; a GM's "most challenging" ruling escalates a heist into a moral dilemma the group still references months later. None of these moments are in the rulebook. They exist because people, together, decide what matters and then keep playing it forward.

Munchkin Deluxe, 2025

4.5 How can randomness be balanced in a physical game environment?

Randomness in tabletop design is about shaping uncertainty that makes players feel tension without losing control. Physical games rely on analog sources such as dice, cards, bags, spinners and on human decision making layered over those sources. The main problem is variety management: That is keeping outcomes amazing while ensuring skill, planning, and storytelling still matter.

The main point, keep in mind, is that ; humans are bad at producing truly random sequences. Under pressure, people switch too often and avoid long streaks; in competitive contexts with feedback they can look statistically close to a pseudorandom generator, but that effect doesn't reliably transfer outside the game context (Wong et al., 2021). This shows us for any design that leans on "players will behave randomly," like bluffing or hidden movement. If players naturally underproduce runs, a game can unintentionally become predictable. So we as designers don't expect players to provide randomness but structure it.

Distribution shaping, not just "rolling dice"

Balance starts by shaping probability distributions. Multiple small dice approximate normal like bell curves; single wide dice produce flat distributions. If you want swingy, cinematic highs and lows, use a flat distribution. If you want reliability with the occasional miss, stack narrow distributions (e.g., 2d6 for consistent mid-results). In competitive settings, you can even quantify the randomness quality of result streams by compressibility metrics (e.g., Lempel–Ziv complexity) during playtests to see whether players perceive swing vs. skill the way your math predicts (Wong et al., 2021).

Card decks are even stronger because of controlling without replacement. Cycling a deck reduces stacks and makes frequencies converge over time great for fairness and for signaling. Shuffling schemes or "discard-then-reshuffle" windows let you tune how fast variance normalizes. In other words, you impose "interval forecasts" rather than single-point odds: you design for a range of expected frequencies over a cycle, not a fixed, memoryless chance every trial (De Cooman & De Bock, 2021).

Mixtures and knobs: combining random sources

Many robust tabletop systems effectively run mixture distributions : a base chance plus situational modifiers that gate or reweight outcomes. In the game-theory, this is called "chance-constrained" decision spaces: choose strategies that stay within acceptable risk bounds under mixtures of distributions (Peng et al., 2021). In tabletop design, without eliminating uncertainty players manage the risks by exposing the knobs. Examples include "advantage/disadvantage" dice, reroll tokens, or converting a miss into a weaker effect. These are player facing constraints that cap catastrophic outliers while keeping suspense alive practically, you're implementing a joint chance constraint on failure rates at the table.

When local randomness becomes stable globally

Random outcomes can feel fair at session scale even if they're noisy moment to moment. In stochastic systems, local noise often "averages out" into predictable macro behavior, a theme formalized by homogenization results for random

walks in balanced random environments (Guo, Peterson & Tran, 2022). For tabletop loops, the analogy is: if your encounter table or loot deck is balanced and recurs often enough, the campaign's long-term resource curve will stabilize, even if individual draws spike. Designers can lean on this by (1) limiting deck sizes, (2) forcing periodic reshuffles, and (3) spreading critical effects across multiple independent pulls so that no single spike leaves the track.

Imprecision is a feature, not a bug

Classic probability assumes precise, stationary odds; real tables don't. Players learn, house rules creep, and table culture shifts. Treating randomness as inherently imprecise defined by intervals rather than exact points anticipates this drift and keeps systems resilient (De Cooman & De Bock, 2021). Practically: publish ranges (e.g., "boss crit rate ~8–12% across phases"), use soft caps instead of hard ones, and make safety valves (pity timers, guaranteed success after N attempts) scale within tolerances rather than at exact thresholds.

Load-balancing uncertainty in multiplayer

In competitive or semi-cooperative games, random shocks should not systematically benefit late movers or specific roles. Dynamic-games research with randomly arriving participants shows equilibrium quality depends on how uncertainty and entry timing interact (Bernhard & Deschamps, 2021). For board games, that means : avoid sequencing random boons immediately before a player's turn; prefer end-of round resolution or shared event queues that everyone can react to. If events must release mid-round, consider a common buffer ("forecast row") so all players can plan around the same stochastic horizon.

Tools that keep luck fair

- Deck-smoothing: Use finite decks with controlled composition; reshuffle only after depletion. Mulligan rules and "burn a card face down" steps can decorrelate streaks faster than naive shuffling (reduces the human tendency to underproduce long runs; Wong et al., 2021).
- Reroll currencies: Limited rerolls transform tail risk into strategic resource allocation, an applied chance constraint that bounds failure probability at the campaign level (Peng et al., 2021).

- Advantage/Disadvantage bands: Rolling two dice and taking the best/worst narrows or widens effective tails without changing the visible component set clean UX for risk control.
- Pity timers & thresholds: Guarantee an effect after N misses. This is homogenization by design: over enough trials, observed success rates converge to a promised floor (Guo et al., 2022).
- Public clocks: Global tracks that schedule hazards convert bursty randomness into paced uncertainty and reduce sequencing bias across players (Bernhard & Deschamps, 2021).

Short examples from existing games

- Gloomhaven replaces dice with modifier decks that cycle, include a “bless/curse” mixture, and reshuffle at defined triggers. Variance is high in the moment but converges across scenarios through controlled deck memory.
- The Binding of Isaac: Four Souls uses large loot/event decks and tradeable items (rerolls, cancels) that let players budget risk, effectively bounding catastrophic tails while preserving swingy moments.

Gloomhaven, 2022

The Binding of Isaac: Four Souls - Requiem Ultimate Collection, 2023

4.6 Translating Digital Roguelike Design into a Physical Game: The Binding of Isaac: Four Souls

The Binding of Isaac as a digital game builds its storytelling around roguelike structure: procedural generation, permadeath, randomized items and extremely hostile combat, all framed as the internal world of a traumatized five-year-old named Isaac (Ruiz S. Goulin & Ribeiro da Mota, 2024). The question for this section is: what happens when you try to move that design into a physical tabletop format? Specifically, how does The Binding of Isaac: Four Souls, the official card game adaptation, carry over not just mechanics but also the thematic work those mechanics were doing in the video game?

To answer that, we first have to be clear about what the digital game is doing. The

Binding of Isaac uses roguelike repetition to express Isaac's mental state. Each run is a new "version" of events, with altered rooms, altered bosses, altered items and sometimes altered endings. They are messy replays of the same trauma. The game's narrative "doesn't really end," it keeps circling around the same core wound: Isaac hiding from his abusive, religiously fanatic mother, and reframing everything he sees as monsters, dungeons, and bosses (Ruiz S. Goulin & Ribeiro da Mota, 2024; Bosman & van Wieringen, 2018).

This spiral structure matters because it connects mechanics and story. The roguelike format (procedurally generated rooms, random items, "permadeath," etc.) becomes a narrative device that represents a child processing trauma through fragmented memory and symbolic play (Ruiz S. Goulin & Ribeiro da Mota, 2024). The same enemies and pickups return in different combinations, like disturbing thoughts and recurring anxieties. Isaac is literally weaponizing his own tears. Even the character select screen reinforces this: you don't really pick totally different heroes, you pick different versions of Isaac, wearing different identities (like "Cain," "Judas," "Eve"), which are themselves loaded with biblical meaning. The game makes it pretty explicit that "they are all, canonically, Isaac," which is tied to the idea that at his age he still has an egocentric, unstable sense of self (Ruiz S. Goulin & Ribeiro da Mota, 2024). This is an important background for thinking about adaptation, because it shows that in Isaac, mechanics are already doing narrative work.

The Binding of Isaac: Four Souls, character cards

That choice already tells us something. A lot of physical games that draw from digital action games smooth out difficulty or stretch power progression across a campaign. But in Isaac's case, the structure is intentionally high risk, high variance, and fast. That matches the digital Isaac's "unforgiving difficulty," where even the smallest enemy can kill you and you're always one or two bad rooms away from losing the run (Guzsvinecz, 2022). In Soulslike design (which overlaps with Isaac in terms of punishment and repetition), high lethality is paired with a tight loop of death; retry; learn; patterns; push further (Guzsvinecz, 2022). Four Souls keeps that rhythm by letting player characters die, lose resources and still jump back in. Death in Four Souls –unlike its digital game– is used as a resource sink and a setback. That mirrors what the roguelike video game is saying narratively: Isaac can't "escape," he just keeps replaying versions of the same fear. In the tabletop version, you can't really bow out of the system either. You're stuck in it until someone wins. That persistence of the loop is part of how it carries story tone.

Another important thing Four Souls adapts is itemization. In the video game, picking up items is basically character rewriting. Your stats change, your appearance mutates and sometimes the logic of your attacks totally shifts (tears become lasers, bombs become blood bombs, etc.). The game encourages you to build weird, broken synergies. That's not just power fantasy, it also visually shows Isaac's body being reconfigured by guilt, shame, religion, illness, etc. (Bosman & van Wieringen, 2018). Those items are loaded with religious symbolism: rosaries, Bibles, relics, even relics of sin. The original game uses these objects to criticize abusive, weaponized religion. The mother figure is described as a "theomaniac single mom" who believes she's commanded by God to kill her child, and Isaac's world reframes that as dungeons, demons and cursed pickups (Bosman & van Wieringen, 2018).

Four Souls basically turns that system into card economy. You buy and steal passive items, you activate usable items, you hoard resources, you mess with each other. The symbolic charge is still there because many of the cards keep the exact same names and iconography from the video game such as biblical names, relics, references to sin, decay and body horror. Because tabletop games are social, this pushes Isaac's internal symbolic world into an openly negotiated space

between players. So instead of Isaac privately equipping “The Bible” to survive one more room, a group of players is literally arguing over who gets “The Bible,” who deserves it, and what it’s worth in trade. That move externalizes what in the digital game stayed mostly inside Isaac’s head. Theologically loaded objects become bargaining chips. The effect is that the critique of religion becomes a thing players actively manipulate to win. That lines up with readings of the original Isaac as a game that “incorporates criticism of religion” both narratively and mechanically (Bosman & van Wieringen, 2018).

There’s also the question of identity. In the video game, “all characters are Isaac,” just wearing different identities pulled from scripture and family memory (Ruiz S. Goulin & Ribeiro da Mota, 2024). This is connected to Piaget’s idea that at around five years old, a child is still in the preoperational stage, which is marked by egocentrism and unstable self/other boundaries. The child “confuses himself with the universe” because his sense of self is not fully separated yet (Piaget, as discussed in Ruiz S. Goulin & Ribeiro da Mota, 2024). The video game uses that to justify why Isaac can “be” Judas, Eve, Cain, etc. Four Souls keeps that cast. You don’t just play “generic hero”; you play as Isaac, or Magdalene, or Judas, etc. The important part is that the game doesn’t correct that framing. It doesn’t say “these are different people in a shared world,” it treats them as selectable play identities with different stats and abilities, the same way the digital game does. That means the tabletop game is also repeating the idea that identity is modular, performable and unstable which again fits the child-psychology framing of the original (Ruiz S. Goulin & Ribeiro da Mota, 2024).

“Mom’s Knife” in The Binding of Isaac and its tabletop adaptation

One interesting difference, though, is how they managed to control the tone of the game. In the digital Isaac, tone is mostly handled by the game itself: dark humor, gore, gross out visuals, religious horror. The player feels it but can't rewrite it. In a physical card game, tone is partly controlled by the table. Players can lean into comedy, or spite, or cruelty. They can "grief" each other, form alliances, betray each other, block each other's boss kills and so on. That kind of emergent social cruelty lines up weirdly well with the emotional core of Isaac, which is about being trapped in systems of punishment you didn't design. In games influenced by Soulslike combat design, players actually enjoy intense difficulty and even "pain and loss," and this becomes part of the appeal (Guzsvinecz, 2022). Four Souls takes that same pleasure in suffering energy and turns it into PvP spite. You're making other players suffer, just like Isaac's mother makes him suffer. So the cruelty that, in the video game, comes mostly from Isaac's mother and from the hostile dungeon, in the physical game can also come from your friends. That's a shift, but it's thematically consistent: Isaac is about abuse cycles. The tabletop version basically lets players reproduce cycles of harm on each other as a core play pattern.

The Binding of Isaac: Repentance,
Isaac vs The Beast, 2021

The Binding of Isaac: Four Souls, The Beast and Isaac!

Finally, we should talk about how win conditions translate. The Binding of Isaac as a video game technically has endings, but part of the analysis is that those endings are not stable. Even the so-called "victory" cutscenes often reveal darker implications, like ongoing abuse, religious trauma, self blame or even death (Bosman & van Wieringen, 2018). The message is: you don't just beat your trauma and walk away. Four Souls, in contrast, has a more concrete win condition (collect four souls before the others). That could look like a simplification, trauma turned into victory points. But I'd argue it's still in line with the world of Isaac, because "souls" in Isaac are literally trophies taken from bosses you kill. You're accumulating proof that you've destroyed something broken. The race to collect souls in the card game keeps that moral ambiguity: you're encouraged to be ruthless and the resource you win by isn't innocent.

So, does The Binding of Isaac: Four Souls successfully translate digital roguelike mechanics into a physical game? Mechanically, yes: it keeps core pillars like high lethality, fast power escalation through item drops, boss rush structure, and constant threat of loss, which are all key to Isaac's identity and are also central to modern "Souls-like" taste for punishing but rewarding loops (Guzsvinecz, 2022). But more importantly, it also keeps the narrative logic behind those systems. Death is still part of the cycle, not the endpoint. Items are still symbols of guilt, faith, sin and survival. Identity is still unstable and performative. Religion is still a dangerous force embedded directly into rules. The only big shift is that trauma, which was interior and solitary in the video game, becomes social and negotiable in the card game.

CHAPTER 5

DESIGN PROPOSAL:

THE FINAL ETAM

5.1 Concept Overview and Design Goals

The Final Exam is a board game about university life, especially Politecnico di Torino, presented in a caricatured way and built around the design philosophies of roguelike games. While designing it, I drew inspiration from many digital and physical games I love playing, and my goal was to create an experience where every university student can recognize a part of themselves, laugh at the struggles they normally complain about, and feel right at home in the chaos.
The game is the practical outcome of the theoretical work about roguelike/lite games, developed in the previous chapters. It takes some of the core roguelike ideas such as procedural variation, repeated runs, meaningful failure, emergent storytelling and grid based movement while leaving some ideas out like permadeath and meta progression between runs. The game intends to rebuild these core principals and ideas in a physical, card driven, multiplayer format. Instead of adventurers delving into caverns, players embody stressed students trying to survive a chaotic campus. In this tabletop dungeon crawler experience that masks itself as a university campus, the first player who reaches 10 Credits wins the game.

Playtesting the first prototype of The Final Exam, 2025

At its core, the game revolves around a shared Campus Grid Board with a 6×6 layout. At the beginning of the game, the grid is empty. During play, it is gradually populated with Environment Tiles drawn from a face-down pool of 32 tiles: 8 Classrooms, 8 Corridors, 8 Cafeterias, and 8 Bathrooms. Each turn, a player first draws and places a tile, then moves onto an adjacent space and interacts with that environment. It creates a constantly evolving campus. No two runs share the same layout, and the board itself becomes a record of the group's decisions and risks.

The second layer of the design is driven by cards. Each environment on the board has a matching Environment Deck. When a player steps onto a tile, they draw from the corresponding deck and get one of several card types: Encounters, Traps, Events, Actions, or Items. Encounters are our main antagonists: assignments, group projects, quizzes, exams and sometimes our beloved professors (they are the strongest bosses); typically everything that troubles a student on a daily basis. They are the main engine of progression; resolving them successfully always grants +1 Credit unless the card states otherwise. Traps and Events create immediate negative or positive effects, disturbing the rhythm of play. Actions and Items go to the player's hand and can be used strategically within a strict card play limit of three cards per turn. Together these elements structure the core loop: explore the campus, face risks, gain advantages, and race toward 10 Credits.

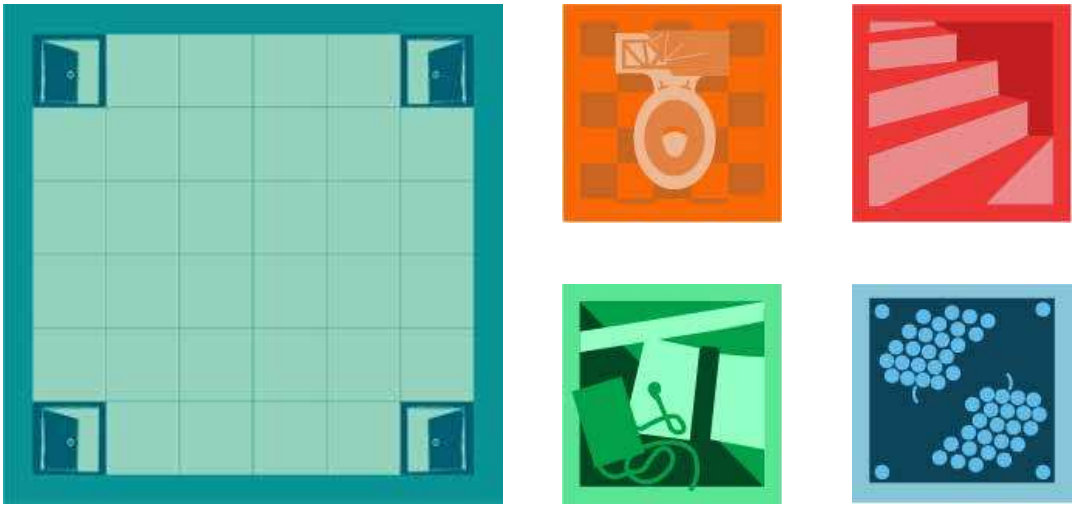

The Final Exam, campus grid and four Environment Tiles

The Final Exam,
Aura Student Card in work

Each player controls a unique Student Card, which provides a permanent passive ability and defines their role within the group. Student Cards cannot be traded, stolen, or discarded. They sit outside the usual constraints of card limits and Items. This light asymmetry is important because it adds further variation in terms of gameplay by pushing players into different strategic identities (more supportive, more aggressive, more opportunistic) without adding complex subsystems. Combined with the fact that the game is designed for 2–4 players, ages 12+, with an estimated duration of 60–90 minutes, this structure targets groups who are familiar with modern board games but do not necessarily want a long, rules-heavy experience.

The rulebook explicitly defines The Final Exam as a "chaotic, negotiation-heavy, semi-competitive crawl." It is not just thematic. It is pretty similar to our university life. The game is semi-cooperative in that players often need each other to overcome difficult encounters. At the same time, sabotage is possible through cards that interfere with another player's Encounter. This structure deliberately encourages bargaining, promises, bluffing, and betrayal. The system creates opportunities for social interaction (and potentially the end of some friendships), letting table politics become part of the emergent narrative.

From a mechanical perspective, the game adopts a clear and minimal stat model. In the first draft of The Final Exam, the game had four different stats and an Energy system on top of that! This alone made the game insufferable and I had negative feedback on our playtests in this regard. Thus, I found the solution by subtracting and simplifying these aspects. Instead of multiple character attributes, Encounters are resolved by comparing their Difficulty value to the sum of the player's Credits + Knowledge. Knowledge is provided by equipped Items and temporary bonuses from Classroom tiles. This keeps resolution fast and legible while still allowing meaningful upgrades and temporary boosts.

Tile passives follow the same philosophy: each tile type has exactly one short, one sentence effect. Classrooms grant +5 Knowledge for the turn, Cafeterias temporarily increase the card play limit, Corridors offer extra movement and Bathrooms allow Item draws. These effects are easy to remember but shift the tempo and risk profile of a turn in subtle ways.

The win condition further aligns the design with roguelike logic. Players track Credits on a shared track so everyone can see and plan to sabotage other players. The game can only be won by completing an Encounter. No Event, Item, or Action can ever move a player from 9 to 10 Credits. To graduate, a player must reveal or play an Encounter on their own turn, face it, and succeed. This ensures that the climax of the game always occurs inside the core risk–reward loop, not through a passive bonus or random event. Failure is punishing but not terminal: losing an Encounter costs 1 Credit (if possible) and sends the player back to their starting corner, but nobody is eliminated. The design imitates roguelike "death" as a setback rather than a full reset, keeping players engaged until the end of the session.

Overall, the main design goals of The Final Exam can be summarized as follows. First, to offer a systemic translation of roguelike principles into a physical game: procedural campus generation, repeated risk loops through Encounters, meaningful failure, tight resource limits, grid based movement. Second, to create a socially driven, negotiation-heavy experience where cooperation and sabotage naturally generate emergent stories. Third, to maintain accessibility and readability through a strict turn structure, simple tile passives, a single core resolution formula (Credits + Knowledge vs. Difficulty), and clear card timing rules. Finally to

ensure strong replayability without legacy mechanics, relying instead on changing campus layouts, shuffled decks and varying player combinations to produce new runs.

The following sections of this chapter unpack how these goals inform the design of the game's systems, the handling of randomness and difficulty, the process of prototyping and playtesting, and the visual and material choices that support clarity and player experience at the table.

The Final Exam,
four Environment Decks

5.2 Translating Digital Roguelike Principles into Tabletop Mechanics

In the previous section, I briefly talked about the concept and what I tried to achieve with my design choices. We can classify physical components as "outputs". In this section, I want to demonstrate how I drew parallels between the knowledge I gathered from literature research and implemented it into my game. By doing this we will connect our tangible outputs to intangible, short-term "outcomes". So, I am asking myself and to my project these questions: "What did I want to achieve in the end and what did I come up with to reach those goals?".

Designing The Final Exam required translating the logic of digital roguelikes into a physical system made of tiles, cards, and social interaction. Digital games automate these systems through code; tabletop games distribute that automation across components and players. The challenge for me was interpreting these principles in a way that it feels natural within the constraints and strengths of a physical format. Early versions of the game lacked the understanding of turning digital complexity into an interactive and social game environment. The more I proceeded with my research, the more I started drawing parallels.

The first major translation involves procedural generation and randomness control. In digital roguelikes, procedural generation creates structured variability through rule systems rather than handcrafted content (Risi & Togelius, 2020). In The Final Exam, Environment cards and tiles together mimic the digital roguelikes' PCG systems. The campus begins as an empty 6×6 grid. Each turn, players draw a tile from a shuffled pool and place it adjacent to their position. Furthermore, players draw cards from tiles' matching Environments. This produces a gradually revealed map that mirrors the "possibility space" created by digital generators. Because tile and card draws are random and placement depends on player movement, campus layouts differ from run to run. The design intentionally uses simple components. Four environments, four matching decks... The combinations of spatial placement, adjacent effects and timing create the same systemic variation that defines roguelike structure.

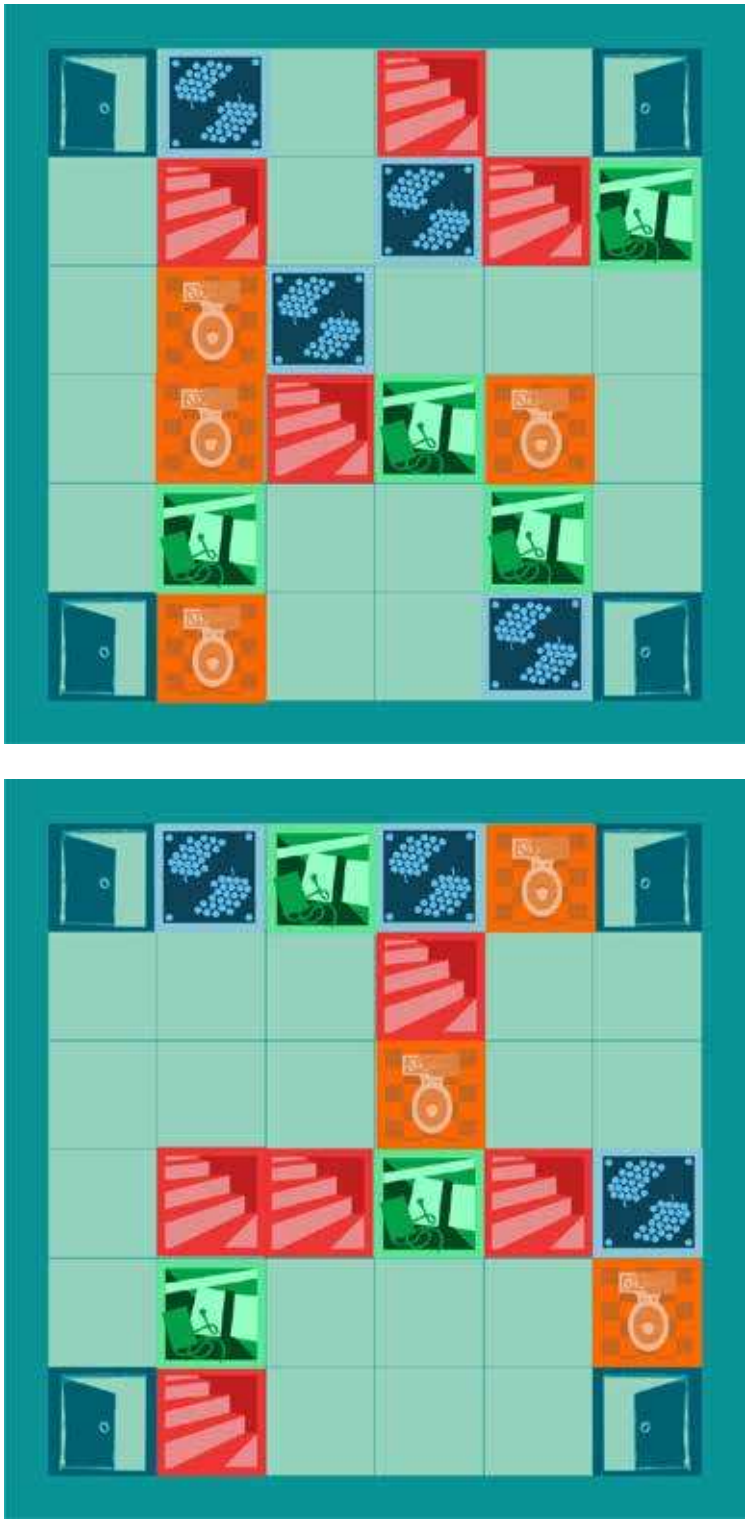

Visual representation of randomly generated Campus Map examples

The second foundational system is the Encounter loop, which functions as the tabletop equivalent of digital combat rooms or challenge nodes. In digital roguelikes, encounters provide pacing through repeated cycles of tension, resolution, and reward (Chauvin et al., 2015). In The Final Exam, this becomes the card-driven moment when a player steps onto a tile and draws from its deck. Encounters are the primary risk–reward mechanism: success grants +1 Credit, failure removes 1 Credit and sends the player back to their starting corner. In addition, overcoming an Encounter grants rewards to players, just like in digital roguelites such as Hades, The Binding of Isaac or Dead Cells. Losing one credit mirrors the logic of soft permadeath discussed in Chapter 1: setbacks matter but do not remove players from the run. The system interprets how modern roguelites transform failure into motivation.

Digital roguelikes rely on builds, synergies, and stat tuning to create emergent strategies. In Hades, players get to choose one of 6 unique weapons before each run. Each weapon presents 4 different aspects that completely change the gameplay. In addition, there are dozens of Keepsakes, Boons, Mirror upgrades, and more. Every good roguelite turns repeated gameplay cycles into a fun, memorable and unique experience by adding endless variety. The Final Exam delivers the same logic by adding six different Student Cards, each with unique passive abilities; and dozens of Items that have Knowledge bonuses and more passive abilities. As described in emergent narrative research, player interpretation of these systemic shifts becomes part of their story-making process (Kybartas et al., 2020). A strong Knowledge engine can turn previously risky encounters into opportunities, echoing the strategic rhythm of digital roguelite “build spikes” without heavy bookkeeping.

Both digital and physical roguelites depend on meaningful failure loops because players need that challenge to try again to win. Otherwise, repeating the same layout would become boring, really fast. Death or failure resets progress but reinforces learning. Tabletop games or multiplayer games in general cannot permanently put players out of the game because waiting for other players to finish kills the immersion and social interaction. For this reason, The Final Exam adopts

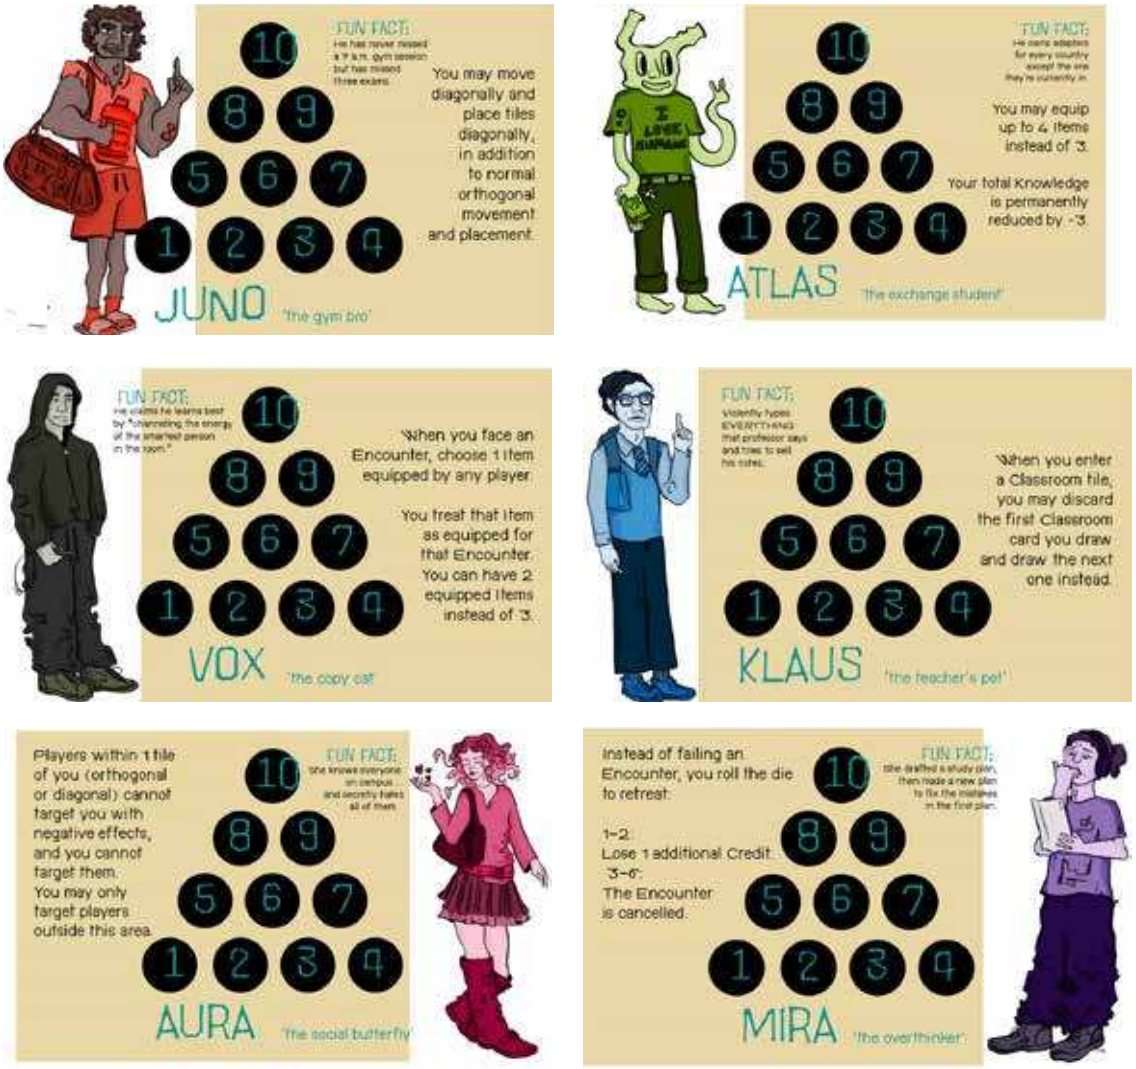

The Final Exam, Student Cards in work

a softer approach to failure. Failing an Encounter in The Final Exam results in a Credit loss and positional reset, echoing the emotional rhythm of frustration, re-assessment, and adaptation that university students often feel. As West (2020) notes, failure gains meaning when players attach emotion and reflection to its consequences. The design ensures that failure is significant enough to influence strategy but not severe enough to cause disengagement. This preserves the roguelite loop while respecting the social dynamics of analog play.

The most significant translation, however, lies in social emergent narrative. In dig-

ital roguelikes, unpredictability comes from procedural systems, AI behavior, and randomized itemization. In tabletop games, unpredictability often comes from players themselves. As discussed in Chapter 4, tabletop storytelling emerges from negotiation, timing, and shared authorship. The Final Exam embraces this concept by making cooperation and sabotage systemic. It also mimics the shenanigans of real student life. There is the sweet competitiveness of academic life but also real friendships made along the way. Players can help each other during Encounters, but only after agreeing on how rewards will be divided. Items can interrupt or manipulate another player's Encounter, creating tension and bluffing opportunities. Furthermore, every physical component in the game recalls past memories: different characters that we see everyday on campus, situations that we experience reflected on the cards and challenges like exams, assignments, projects that make us miserable sometimes. There is definitely no script per say but these components tell us much more than a pre-written story could.

In summary, translating digital roguelike principles into a tabletop format required identifying which genre elements rely on computational automation and replacing them with mechanical systems and social dynamics. Procedural maps become tile placement; RNG engines become deck structures; builds become Knowledge and Item synergies; soft permadeath becomes setbacks and tempo loss; AI-driven unpredictability becomes negotiation and timing. Through this translation, The Final Exam preserves the essence of digital roguelites while leveraging the strengths of physical play: tactility, social interaction, and shared storytelling.

Item deck card back

5.3 Emergent Storytelling in The Final Exam

One of the most striking things about The Final Exam is that its stories do not come from scripted plotlines or character arcs. They come from the players themselves. The game gives them a campus, some cards, and a few rules about cooperation and sabotage, then steps out of the way. What happens next is less about mechanics and more about emotional reflection. Players begin to see themselves, their friends, and their university experiences reflected in the situations the game produces. In that sense, The Final Exam behaves like a loose simulation of student life, filtered through a roguelite lens.

The characters already hint at this. They are still in development phase so their names can change but basically, they are archetypes: The Social Butterfly, The Overthinker, The Copycat, The Nerd, The Exchange Student, the "Athlete". Nonetheless, they are not just caricatures. They are versions of the people we meet on campus and versions of ourselves. Honestly, I am all of them at the same time depending on how I am feeling at the moment. Players often start by laughing at their assigned role, but as the game unfolds they slide into these identities almost without realizing it. Even without explicit roleplay, the game subtly pushes players to inhabit the rhythms of students they know intimately. This identification becomes the foundation for the stories that follow.

A run of The Final Exam tends to move through a recognizable emotional arc. It begins with curiosity: players piece together how the campus grows, how tile effects work, what Encounters look like. Once they understand the flow, strategy kicks in. They begin calculating risks, mapping routes, and watching what others are collecting. Mid-game is where the tone shifts as quiet alliances emerge. Losing players collaborate out of necessity. Leaders deliberately become less visible, hoping not to attract an improvised coalition aimed at slowing them down. Everything still feels playful, but the atmosphere starts to tighten. Someone asks for help on a tough Encounter. Another player offers a deal but adds a small threat behind their smile. Promises are made with no intention of being kept. A player at the bottom of the Credit track suddenly becomes the kingmaker. The unhindered competitiveness forces players to stab their best friend's back. These moments push the group into negotiations that often end in betrayal, desperation or

unexpected generosity. It is the kind of narrative texture digital roguelites achieve through systems. Here it emerges through human behavior.

The funniest stories tend to come from the game's randomness colliding with the personalities around the table. During early testing, a really dear friend of mine got trapped in a loop where a Cafeteria Trap card "Too Much Coffee" kept sending him back to the Bathroom and he kept drawing the "Doomscroll" card, which drained his stats and stunned him for turns. The rules alone weren't funny, but the timing was. The way he got stuck in the toilet for turns was the story. The group laughed for minutes, not because the card text was clever, but because the system had accidentally recreated a shared experience: being overwhelmed by small, stupid problems that pile up until the situation becomes absurd. Roguelikes have always been good at producing this kind of comedy.

Not every version of the game produced this energy. The earliest prototype was quietly dead. There were too many systems, too many stats, too little interaction. Everyone played in isolation, and nothing connected. The absence of story during that test revealed the core principle that now shapes the game:

Emergent narrative requires friction.

Not difficult. Not complexity. Friction. Moments where players need each other, rely on each other, or get in each other's way. After that realization, the game was stripped down and rebuilt around purposeful interaction. Post-game questionnaires and interviews showed that I needed to subtract before thinking about adding another layer. Many systems that overcrowded the game were removed: 20-faced die, Energy system, four different stats instead of one and more (they will be explained in more detail in the next chapter).

Helping Encounters, interfering with them, bargaining, timing card plays, trading advantages became the real engines of narrative. Story finally started appearing once the systems gave players a reason to look at each other instead of just looking at the board.

Without a doubt the final turns of a run produce the sharpest tension as players come close to victory. Players switch from collaboration to deepening competitiveness and distrust. Leaders pretend to be harmless. For example, the leader has 8 Credits while others are trying to catch up. he leader may bribe a weaker player by offering all Encounter rewards if they agree to help. This sneaky atte-

mpt masked as innocence will help the leader to reach 10 Credits and win the game faster. Any Encounter could become the winning moment, so every turn feels like a crossroads. Even when a player wins, the victory rarely feels solitary. It is the result of timing, bluffing, the table's shifting alliances and a little luck.

What matters most is what happens after the game ends. They talk about what happened between them:
"Dude, you always do the same."
"She had all the good items from the start."
"I won't help you in an Encounter ever again!"
All with good intentions and laughter of course. One of the desired outcomes is to ensure players will recall these memories, even in their next run, on another day.
These are micro stories: brief, chaotic and personal. They are exactly the kind of stories roguelites are built to produce. They give players the impulse to say, "Let's play one more time." And that, ultimately, is the form of meaning The Final Exam aims to generate.

The Final Exam, playtesting to first prototype

5.4 Randomness, Difficulty, and Fairness

The Final Exam is built as a controlled-luck game, getting inspiration from some of my favorite tabletop games like The Binding of Isaac: Four Souls or Munchkin. Chance shapes the texture of each run, but it does not govern the outcome on its own. Let's think of randomness as a spectrum. On one side there is a pure strategy game, like chess. Every square is created equal. Pieces serve different purposes but both of the players have the exact same setup, the exact same pieces and the exact time limitation. One turn for each player. Chess completely removes the luck factor. On the opposite side of the spectrum, there is Chutes and Ladders. The result of the game is solely determined by dice rolls: who is luckier, wins the game. The design of The Final Exam aims for a middle space where unexpected events give the session its personality, while player choices and interactions ultimately determine who reaches ten Credits first. In practice, this means that randomness creates the stories, but understanding how the systems interact creates the wins.

Every element of the game is balanced around stable ratios: four tile types with eight copies each, Environment Decks built around an even distribution of Traps, Events, Actions, Items, and Encounters. Different Student Cards, Items and passives offer asymmetry without dominance. Nothing is tuned to be overwhelmingly strong or disastrously weak. A player might draw a rough sequence of cards, but over a full session the distribution smooths out. Since the decks use physical memory (cards don't re-enter the deck until reshuffle) the randomness feels more predictable. As discussed in Chapter 4, this is "randomness without replacement," a form of structured chance that keeps outcomes surprising but statistically bounded (De Cooman & De Bock, 2021; Wong et al., 2021).

The Final Exam being a multiplayer tabletop game, unlike digital single player games, provides fairness through social correction. The rules deliberately give players opportunities to assist one another, and these helping interactions naturally prevent snowballing. A player who falls behind still has currency to offer: their future Item rewards, their political value, their willingness to be part of an al-

liance. A player who pulls ahead, however, becomes an immediate target. No one wants the leader to enter an Encounter with support. As a result, the group collectively elevates the difficulty for whoever appears close to winning. Other players' interventions punish the leading player instead of the game itself. This dynamic aligns with research showing that social coordination in tabletop groups generates natural sources of tension and pacing, shaping how stories emerge through interaction rather than rules alone (Webb & Cesar, 2019).

The game's difficulty curve mirrors the shape of many digital roguelikes like Hades and Isaac. Early turns are light: few tiles exist, Item pools are small and there are few variables to consider. As the map spreads and inventories grow, choices become more meaningful. The more players get experienced in the game, the more they start making educated guesses. They learn to go to the Bathroom to refresh Items, Corridors for faster movement, Classrooms to guarantee tougher Encounters and Cafeterias for making combos. The middle phase of a run tends to be the most strategic. Everyone now has enough tools to pursue their own agenda. Small decisions start to matter. Decisions like when to move, when to negotiate, what to attempt may begin to snowball into larger consequences.

Failure is part of this arc. Losing a Credit or being sent back to the starting corner is annoying, sometimes deeply frustrating but it never removes players from the experience. The actual emotional impact is not just about failing. It usually depends on the timing of failing. Losing a Credit at two or three points doesn't matter much. Losing one at nine feels like a punch. This scaling effect is intentional. A setback in late game carries narrative weight: other players were watching, waiting, perhaps even contemplating that very moment. This echoes broader analyses of emergent narrative systems, where dramatic tension often comes from conflict produced by agents with competing goals—human or artificial (Kybartas, Verbrugge, & Lessard, 2020). A player may be fully capable of beating an Encounter mathematically, but if the three opponents agree –explicitly or silently– that the leader should not win this turn, no amount of pure skill guarantees success. Most failures in The Final Exam come from timing and politics. In this regard, the game's approach to failure aims to never feel unfair. In digital roguelites,

difficulty often emerges from the environment and increases with levels. Here, difficulty emerges from the people at the table weaponizing the system.
Fairness in The Final Exam, on the other hand, comes from the presence of agency. A player can always influence their situation through strategy independently of whatever happens on the board. And because no one can win without facing an Encounter, it feels earned or stolen.

The game ultimately plays with randomness the same way roguelites do: as a storytelling catalyst. It never dictates outcomes. The unpredictability creates drama. The decisions create satisfaction. When the run ends, players rarely blame the system when they lose. They blame each other, laugh about the absurd chains of events and hopefully feel the pull to set the board again and see how the next run is going to be.

The Final Exam, player pawns

5.5 Prototyping, Playtesting and Iteration

5.5.1 Concept Genesis and Early Design Pillars

The origins of The Final Exam came from a moment in my life when everything was colliding at the same time. I had one last module left, Social Design, and I was starting an eight-hour-a-day internship that left me very little time to think about a thesis. In the background of all this, I was playing Hades obsessively with my roommate. After the tiredness of a work day, coming home to play Hades was like a reward for me. I had a simple realization that roguelikes had a specific rhythm, structure, and emotional shape that matched how I experience everyday life. Runs full of mistakes, victories, improvisation, repetition, and "one more try" matched not only the games I liked but the way I lived.

When I approached my advisor, prof. Andrea Di Salvo, he encouraged me to propose a project rather than write a purely theoretical dissertation. As I continued reading papers about roguelites, procedural generation, and emergent storytelling, something clicked: the more digital roguelikes I studied, the more I noticed their structural similarities with analog games like Dungeons & Dragons, Fighting Fantasy, and early dungeon crawlers. Authors like Salen and Zimmerman argue that systems produce meaning in games rather than scripts (2004). Fullerton describes prototyping as a process of exploring these systems physically, not mentally (2018). Hunicke, LeBlanc and Zubek's MDA framework highlights how aesthetic experiences emerge from the interaction of rules and dynamics (2004). All of this pushed me toward a hybrid project, which is a board game that behaves and feels like a roguelike.

Once I committed to the idea, I started sketching pillars of what such a game would require. I opened Figma and produced my first mind map. This map became the skeleton for everything that followed. It spelled out the roguelike DNA I wanted:

- Random environment generation
- Grid-based movement
- Turn-based structure
- Different characters

- Cards as PCG outputs
- Bosses as final tests
- Stats
- A clear win condition

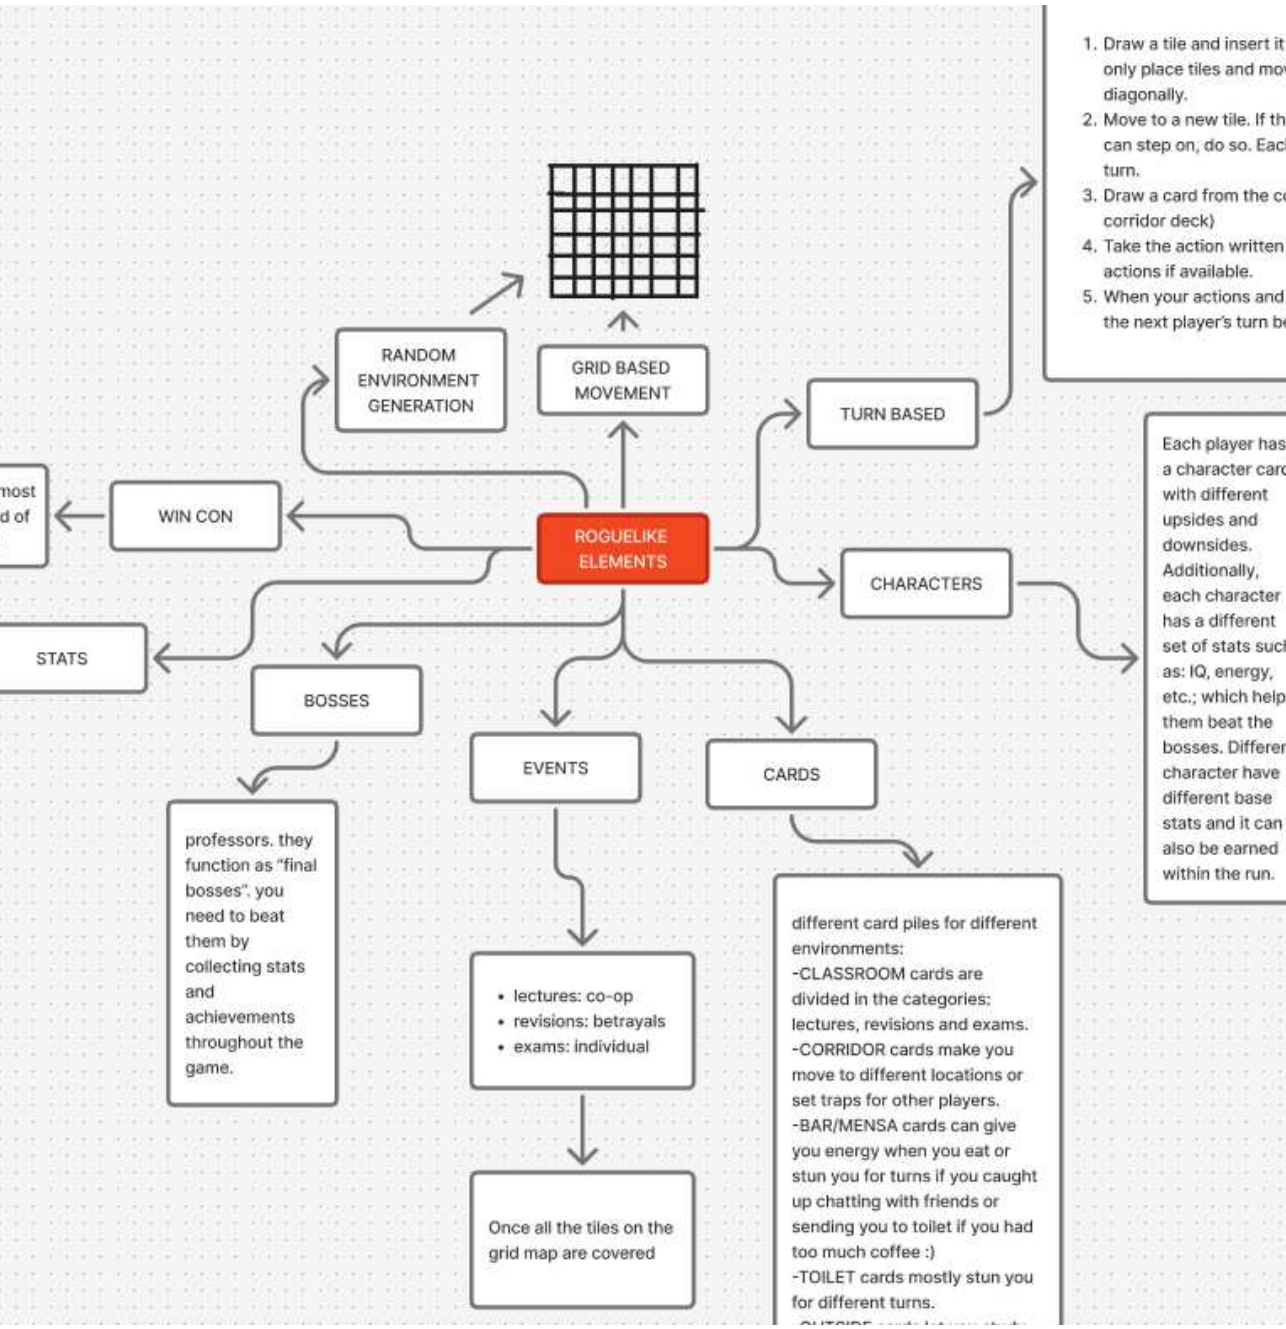

The initial mindmap of The Final Exam on Figma, 2025

At this point, I still didn't know the theme. I only knew the structure. It was still missing its "fantasy" and it needed a world. Something personal, something I knew. That's when the idea of a comically hostile university campus appeared. Dungeons became classrooms. Monsters became exams. Encounters became professors, deadlines, forgotten documents, lucky breaks and social disasters. It felt natural. Most importantly, it felt familiar.

This is where the autobiographical element entered the design. The characters I sketched were pieces of my own personality, my friends, the people I recognized on campus. Without even trying, I had turned the project into a quiet self-portrait. I wanted players to roleplay as versions of themselves. And as emergent narrative research argues, identification makes systems more meaningful because players interpret events through personal lenses (Kybartas, Verbrugge & Lessard, 2020).

I also realized I needed help. I contacted Eda Bolver, a friend from high school, because I knew her illustration style could give the game a humorous, warm identity that matched its tone. I reached out to Ali, someone I had played countless games with, because I trusted his intuition about games. These collaborations shaped the project as it was leaving to be merely an idea and becoming more of a product. Eda helped define the visual mood and Ali helped me see blind spots in the early system.

Some of Eda's sketches for Cafeteria deck

At this stage, everything existed in fragments: a grid, a handful of environment ideas, a rough progression curve and sketches for student archetypes. I still could only explain the game through ideas and my amateur diagrams on Figma. But conceptually, the foundation was already aligned with the roguelite tradition. It would rely on systems, risk, failure, and player-driven structure. It would be, in Salen and Zimmerman's terms, a space where "meaningful play" emerges from consistent feedback loops (2004). And, like most roguelikes, it needed to encourage repeat runs – an idea strongly connected to the "one more try" aesthetic described in design analysis of roguelite games (Hunicke et al., 2004).

What I didn't understand yet is how fragile this foundation was and the following subsections will cover that. The pillars made sense intellectually. They connected to the literature. They sounded elegant when written in a document. But as Fullerton warns, a design that works "on paper" is not a design that works in practice (2018). The next phase – building the first prototype – would expose exactly how much I misunderstood about translating digital systemic thinking into a physical multiplayer format.

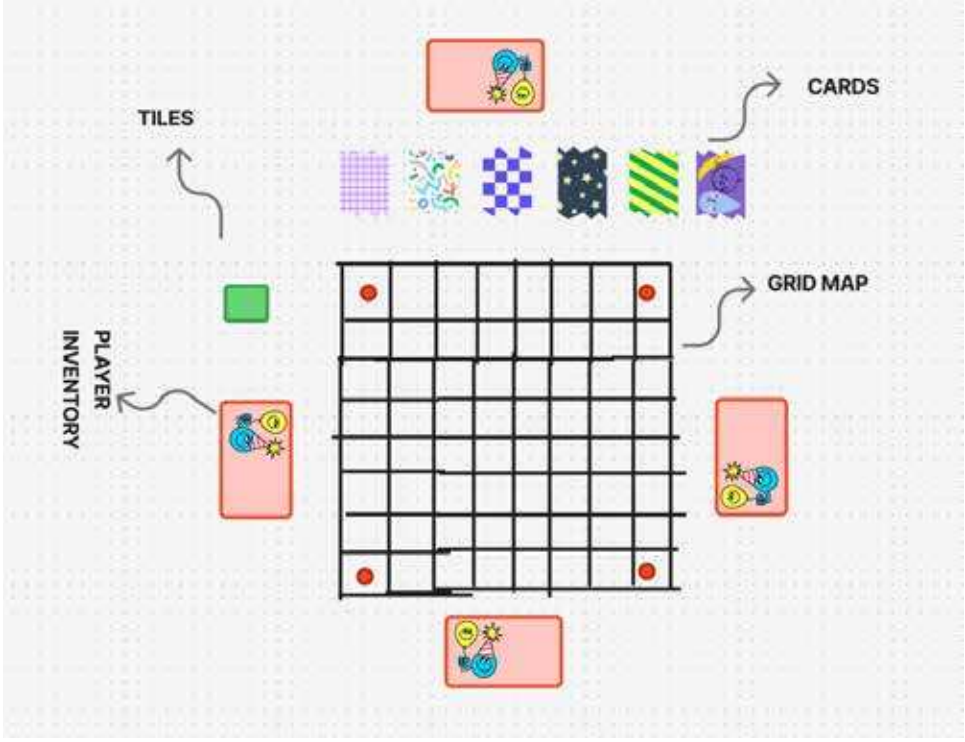

Early concepts of the 8x8 grid, student boards, decks and tiles

5.5.2 First Prototype and Playtesting: Building a Digital Roguelike on a Tabletop

The first playable version of The Final Exam marked the moment when all the theoretical work of Chapter 1–4 collided with reality. I had spent weeks –even months– shaping the system on paper, convinced it would translate smoothly into a tabletop experience. But as many authors warn (Fullerton, 2018; Salen & Zimmerman, 2004), a design that feels elegant in theory can collapse the moment players enter the system. My first prototype became the clearest proof of that warning.

When I finally assembled the first playable prototype of The Final Exam, I believed I had built something robust. I had all the components: a large 8×8 board, five different environment types, Student Boards with four stats, tiles, a Boss system, progression, movement rules and even an Energy mechanic that was supposed to regulate how much players could do each turn. From a distance, the game looked ambitious, dense and complete. I neatly designed every single card and it had to work.

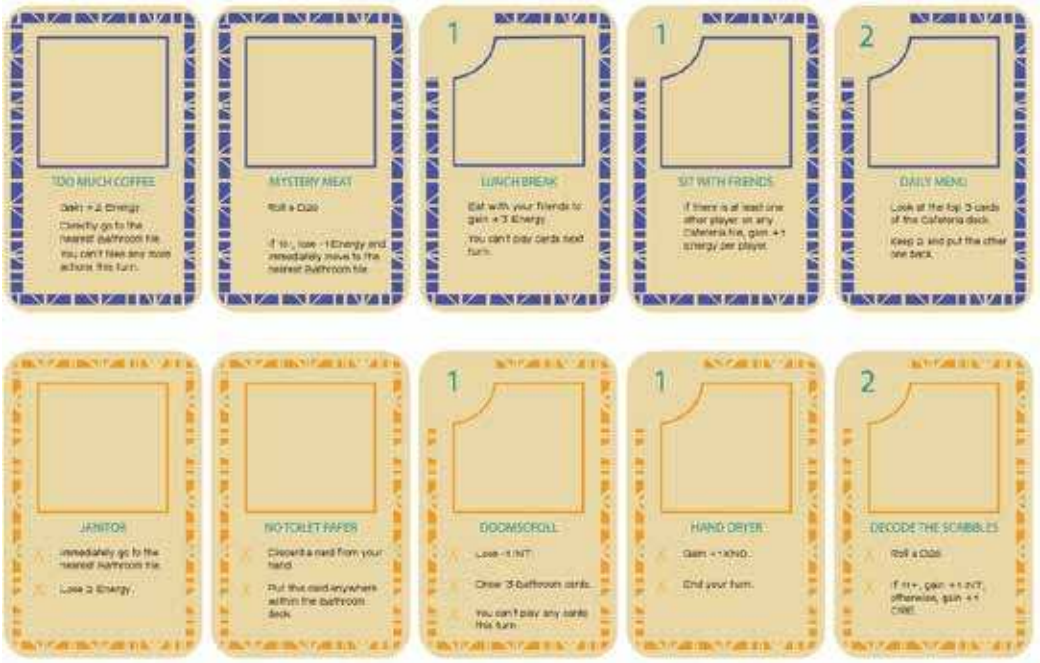

Some of the cards in the first prototype of The Final Exam

I genuinely believed this version would function because I had reverse-engineered it from digital roguelikes I admired. In Hades, The Binding of Isaac or Slay the Spire, multiple stats, layered systems and resource management feel natural. I assumed those structures were "roguelike essentials," so I transplanted all of them into cardboard without considering how physical play fundamentally differs from software automation. In hindsight, this was the root misunderstanding: I treated tabletop design like coding logic, assuming players could mentally process what a computer normally handles instantly.

But the moment I brought it to the table, I realized a fundamental truth that game design authors warn about repeatedly:
designing on paper and designing for play are two completely different activities. Salen and Zimmerman describe games as systems that only "become real" in interaction (2004).

I tried to carry the logic of a computer game directly into cardboard. I added too many systems without properly thinking and testing how they would function together. In digital games, resources, multiple stats, builds are calculated by computers instantly. Tracking these aspects automatically is drastically different than tracking them manually. In a physical multiplayer game, every number is a delay, every stat is a question and every variable is an invitation to confusion.

In this version, each Student had:

- Knowledge
- Charisma
- Intelligence
- Creativity
- Energy (tracked manually)
- A large grid to navigate
- And bosses that required calculating stat combinations

This meant that every single turn involved players asking:
"How much energy do I have left?"
"Did I update my Charisma correctly?"
"Wait, which stat applies to this encounter?"

And because the grid was 8×8, moving across the board took ages. The campus felt like an airport runway. Players wandered far away from each other in the beginning of the game, which unintentionally destroyed the social core the game would later rely on. At the end of the game, because of the old Boss fight system, players started to "camp" around certain tiles to win the game. The randomized layout of the map I imagined didn't serve its purpose.

The Boss system – designed to act as the climactic roguelike "Exams" – only triggered after all Classroom tiles were placed, which meant the real clash didn't begin until the late game. As a result, players spent the first 60–90 minutes of the run doing side tasks, gathering small advantages and waiting for the "real" game to start. By the time the Bosses appeared, the winner was already mathematically decided.

The Energy system was ,by far, the worst offender. My initial idea was to include a fun resource management mechanic that turned into the boredom of bookkeeping. Players were refreshing energy, spending energy, checking energy, re-checking energy, and forgetting energy. It was a mechanic that made sense in theory but collapsed immediately in physical space. It slowed every turn to a crawl. It also created a strange imbalance: players were neither threatened nor empowered by Energy; they were simply annoyed by it. Also, there wasn't enough energy recovery. We found ourselves sitting with a bunch of cards in our hands, not being able to play them because of lack of energy. The components were also rushed because I had just printed them before traveling. I designed the Student Boards in five minutes. The cards had mismatched backs. I hadn't prepared tokens, a die, or even proper pawns. We borrowed pieces from my friend's Monopoly. The moment the prototype hit the table, I realized the most obvious mistake: I had spent weeks designing a system in my head, but almost no time

preparing the physical experience. Thus, the emotional experience was lacking too. Hunicke, LeBlanc and Zubek's MDA framework explains that aesthetics — the emotional experience — emerge from how players interpret the system, not from the rulebook (2004). In this prototype, the aesthetics were misery: slowness, arithmetic, guesswork, low tension and almost no social interaction.

One small moment unintentionally highlighted how broken and yet promising the system was. A friend drew a series of cards –"Too Much Coffee" followed by "Doomscroll" – that repeatedly forced him into the bathroom and drained his intelligence. The loop was absurd, frustrating and hilarious at the same time. Even though the mechanic around it was badly designed, the feeling of academic chaos was exactly the tone I wanted the game to capture. It became one of the only genuinely fun moments in the entire session, and it helped me see that the fantasy itself was strong. But it just needed the right systems wrapped around it.

I had built a multiplayer game that behaved like a single-player simulation. Everyone played alone, silently optimizing their stats. When the game finally ended – 2.5 hours later– victory was anticlimactic, because the leading player had been far ahead for more than half the game.

My friend Buğra Onay, who is also studying Design at Politecnico di Torino, summarized it with only a sentence:
"Your game works in theory, but not in practice."

At the time, I felt the sting of that sentence instantly. But now, looking back, it was the most important moment of the project. The prototype wasn't bad because the idea was bad. It was bad because I had built the wrong kind of game. I had made a digital roguelike on paper instead of designing a social, emergent tabletop experience.

Yet even in this flawed form, a few small moments showed the potential buried underneath. The procedurally generated map worked. The tile placement phase was intuitive and enjoyable. The "draw tile place move draw card" rhythm

Some of the Items used in the first playtest of The Final Exam

was easy to grasp. Some card names produced laughter even without illustrations. These sparks were enough to show that the project wasn't doomed. The core fantasy worked. It was the systems around it that needed radical change.
The first prototype taught me a foundational truth that Fullerton writes about: you cannot refine a game you have not truly played (2018). And likewise, emergent narrative theorists remind us that stories appear only when systems create actionable, meaningful moments (Kybartas et al., 2020).

This prototype created none.

The failure of this prototype didn't discourage me. It clarified that the game wasn't a slow, stat-heavy simulation and pointed directly toward what it needed to become: a fast, interaction-driven roguelite where the story emerges from tension, not from numerical complexity. The next step was to tear the system apart and rebuild it from the ground up.

5.5.3 Survey Results: Quantifying Fun and Failure

After the first prototype collapsed under real play, I needed more than intuition to understand what had gone wrong. I distributed two forms, a Gameplay Feedback Survey and a Story & Experience Survey. The first one was quantitative, using a 1–5 scale (1 = strongly disagree, 5 = strongly agree). The other one was qualitative with open-ended questions and answers. It confirmed that the problems I felt at the table were systemic failures visible to every player.

Gameplay Outcomes: A System Too Slow to Support Tension

One of the clearest metrics in the Gameplay Feedback Survey was game length. The average score was extremely low (hovering close to 1.0), meaning players overwhelmingly felt the session lasted far too long. This aligned perfectly with the 2.5-hour duration of the playtest. The problem was that nothing in the system justified the length. There were no meaningful peaks, no small wins, no crises that escalated tension. A long game can work if it builds narrative energy, which apparently my game did not.

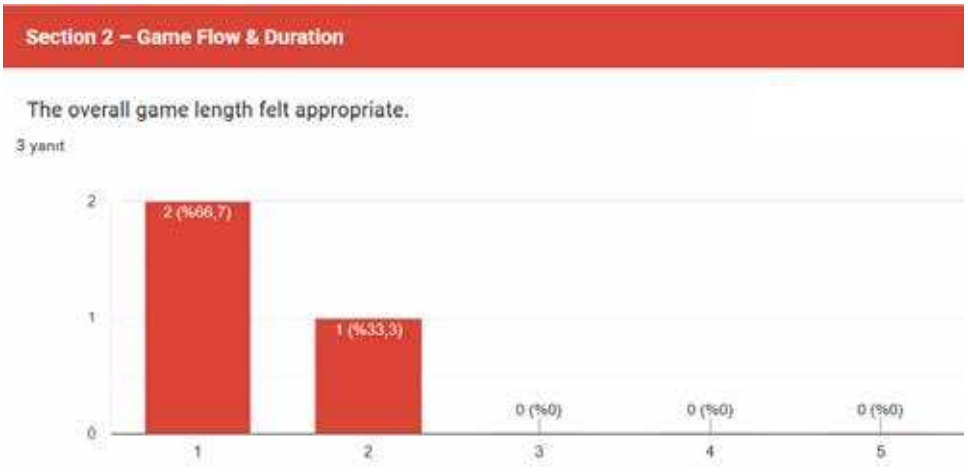

The ratings for pacing and downtime sat mostly between 2 and 3 which means that the players were disengaged. Players waited for turns rather than anticipating them. The score for balance between luck and strategy was slightly more forgiving (around the 3.5 range). This shows that the underlying distribution of events, tiles and cards didn't feel unfair. Instead, the problem was that strategy didn't matter because the systems were too bloated to give players meaningful

agency. In MDA terms, the mechanics existed, but the dynamics never formed, so the aesthetics collapsed (Hunicke et al., 2004).

Fun, Agency and Interaction: A Game That Didn't Let People Play

The most important numbers in the survey were those tied to interaction. Multiple questions related to player–player engagement scored below the midpoint. People simply weren't interacting because the prototype wasn't giving them reasons to.

In the early version, no mechanics required cooperation, negotiation, interference or even attention to what others were doing. Everyone working on their own stat sheets, their own movement paths, their own preparation for Boss fights felt like the game was a sort of multiplayer solitaire.

The surveys confirmed this bluntly:
"I felt involved in what other players were doing" score: low
"I had opportunities to affect other players' outcomes" score: low
"The game encouraged interaction": below neutral

These numbers mattered more than any mechanical feedback because The Final Exam is fundamentally a social roguelite. The surveys revealed that the prototype failed at one of the most important goals of any multiplayer game: giving players a reason to care about each other's actions.

Theme and Experience: Strong Fantasy, Weak Execution

Interestingly, the Story & Experience Survey showed something different. Players responded positively to the theme. They liked the university concept, the tile-based campus, the ridiculousness of academic obstacles and the student archetypes. It was an indicator that I could achieve a thematically-relatable and fun-to-play board game. Even without illustrations finished, my friends found something by themselves. The fantasy and the world were intact but the systems were problematic.

This distinction became crucial later when revising the game because I only had to change how the mechanics expressed this fantasy.

Frustration and Flow: The Energy System as a Design Failure

One survey question asked whether players understood the flow of the game. Surprisingly, this scored slightly above neutral. Players did understand what they were supposed to do. They just didn't enjoy doing it.

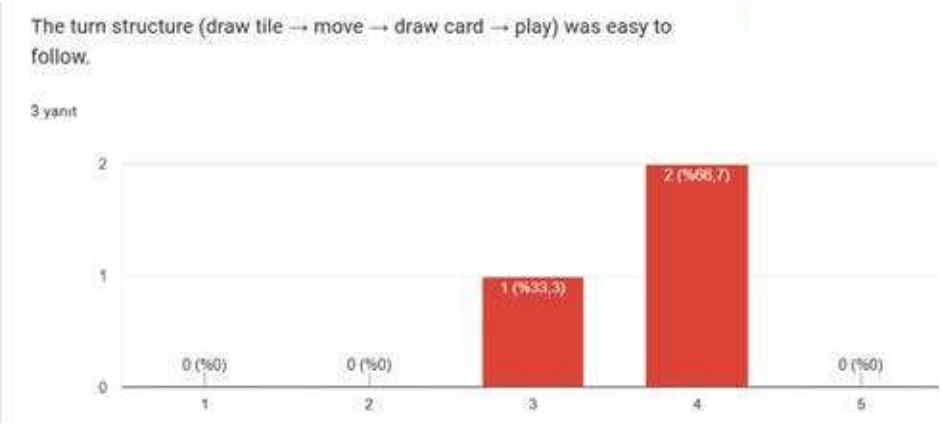

Most open comments revolved around the Energy system. Players wrote about confusion, unnecessary tracking and long pauses. One comment explicitly said the system made the game feel like "homework," which is ironically the opposite tone I wanted.

The Energy system was supposed to be the game's core resource management system. But soon I realized that I was creating a solution for an unexisting problem. Tracking energy manually was only a task and I could find another, more interesting way of resource management.

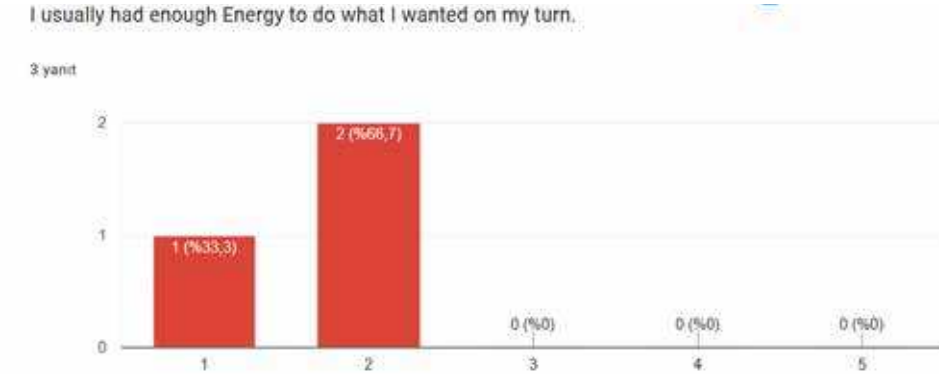

Winning, Losing and Emotional Pacing

The survey question "Would you like to play again?" scored around 3. It was humbling and expected at the same time. Even the professor said: "Ma non si dice così!" ("That's not the way to say it!"). This score revealed the emotional flatness of the prototype. No one was angry at the game but no one was excited either. The outcome of the game felt predetermined. The progression toward the ending felt slow.

In other words, the prototype did not produce the "one more run" instinct that defines the roguelite genre.

The surveys made that failure measurable.

What the Surveys Made Impossible to Ignore

The two surveys together told me five things:

1. The game was far too long for the amount of fun it produced.
2. Interaction was severely lacking, killing any chance of emergent narrative.
3. The Energy + stat system forced players to think like accountants.
4. The theme was strong enough to survive a redesign.
5. The experience of playing was flat.

The next section, therefore, is about designing a new identity for The Final Exam based on what the data demanded.

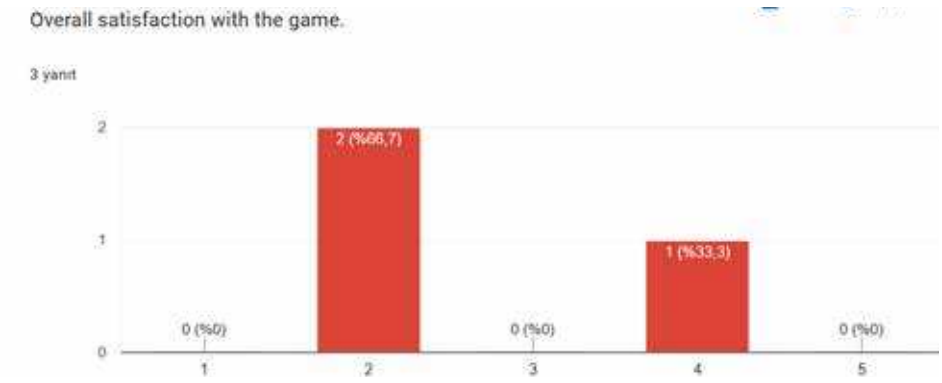

5.5.4 Redesign: Less is More

As we were doing the first playtest, even before the surveys and the interview, it became quite clear that the game needed subtraction before considering adding anything. I had to strip away every mechanic that did not directly serve the experience I wanted to create.

The process began with a post-playtest conversation where we listed eleven problems and potential solutions. At first, they looked like adjustments. But once we unpacked each one, their implications became much deeper. They pointed toward a different game entirely.

1- no need for d20 die

2- instead of counting energy, every player has a predefined energy each turn (for example 5 energy each turn)

3- every player starts with cards in their hands (lets say 5)

4- items can be drawn as cards from the decks

5- more mobility cards

6- more interaction with players

7- other ways to obtain credits other then boss fights

8- one global stat instead of four

9- tiles have unique triggers (for example: when you step on a cafeteria tile, gain 1 energy for that turn)

Eleven problems and potential solutions

5.5.4.1 Removing the Weight: Systems That Had to Go

a. Eliminating the Energy System:

The Energy mechanic was the most obvious casualty. It was meant to be an elegant form of resource management similar to some of my favorite deck builder games: Hearthstone and Slay the Spire. Although these two belong to two different genres, they share some similarities. All deck builders and card games use some sort of “energy” or “mana” system for their resource management system. Every card has a mana or energy value. Players can play their cards by paying from their energy or mana pool. This simple system adds strategic depth with-

out overcomplicating it. Players must consider which cards to hold for the next turn or how much energy they need to play their cards back to back in one turn. As discussed previously, this implementation could work if The Final Exam was a digital card game. By trying to count Energy manually caused friction, frustration and downtime. Every turn required recalculations, memory checks and micro-adjustments.

Mana value of 2 in Hearthstone

Energy value of 0 in Slay the Spire

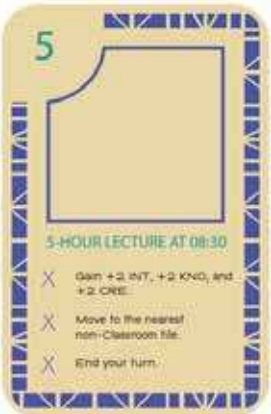

Energy value of 5 in first version of The Final Exam

Our first idea was simplifying it to a fixed-per-turn system (5 Energy each turn). But as I was redesigning the game and rewriting the rule book, I realized that Energy was malfunctioning with other systems. I was looking forward to adding more social interaction and tension by allowing players to intervene in Encounters. If players had to play cards to intervene, how much would those cards cost? 0 Energy? Or would they be a different type of card? And players would be pleased calculating their Energy manually again, even though it was fixed-per-turn?

The eventual solution was removing Energy entirely. But once it was gone, the game instantly breathed better. The “3 cards per turn” limit replaced it, maintaining the sense of resource management without the overhead. Playing cards still required strategic thinking.

New cards without Energy mechanic

b. Collapsing Four Stats into One:

In the first version of the game, the Knowledge–Charisma–Intelligence–Creativity model looked meaningful in theory. It gave characters flavor. Every character had different sets of stats corresponding to their character. Also every Boss fight would require different stats, meaning that players had to collect these stats somewhat equally in order to obtain the most Credits possible. Furthermore; cards, items and passive abilities revolve around this system. Just like Energy, I thought it would add depth and planning ahead. The problem was that the system didn't have a cap. Every character started at reasonable levels: 2-5 points of each stat. However, some combinations of different character passives, items and cards could skyrocket the stats, making it really difficult to design an interface for tracking. Video games essentially allow, even want, maximizing stats. In Hades, killing the final boss is quite difficult since he has a great amount of max HP (maximum hitpoint). With the right combinations and a little bit of luck from randomly generated outcomes, it is possible to kill it in a blink of an eye. However, implementing an interface for such a system was impossible in a board game.

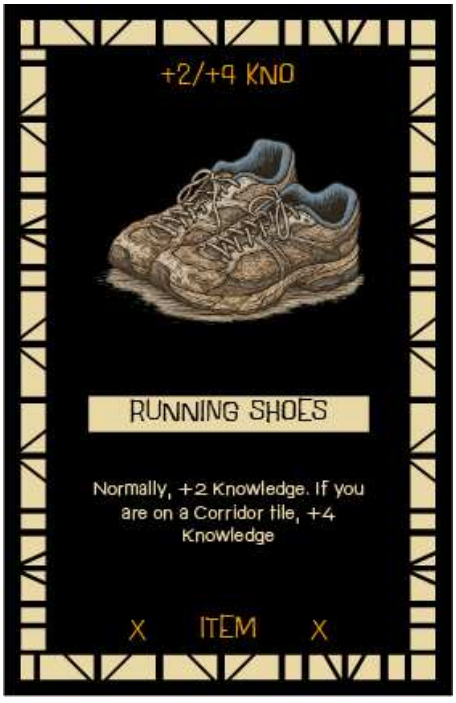

New Items with the renewed Knowledge stat

Condensing everything into one global stat: Knowledge solved multiple problems at once:

•It removed all bookkeeping from Student Boards.

•It clarified Item design: every bon became readable at a glance.

•It simplified Encounter resolution to a single formula.

•It prevented analysis paralysis.

More importantly, it created space for players to focus on the choices rather than the numbers.

c. Cutting an Entire Environment

The original game had five environments, including an “Outside” area. On a 8x8 board it worked flawlessly but reducing the board size required a revision of Environment count too. Cutting it tightened the game world and concentrated the emergent stories around the core identity of the university.

5.5.4.2 Rebuilding the Skeleton: Systems That Needed to Be Added

Eliminating mechanics was only half of the transformation. I also had to add systems that created the missing structure, tension, and interaction.

a. Introducing Tile Passives

Originally, tiles existed only to determine which deck players would draw from. They were functional but inert. The new prototype introduced one-sentence tile passives that immediately gave environments purpose:

•Entering a Classroom gives you +5 Knowledge for this turn.

•Entering a Cafeteria increases your card limit to 4 cards this turn.

•Entering a Corridor allows you to move 1 additional tile without drawing a card or triggering another tile passive.

• Entering a Bathroom lets you draw 1 Item, and at the end of your turn you must discard down to 3 Items if you exceed the limit.

These passives transformed movement from “walking to a card source” into “navigating a landscape with opportunities and tradeoffs.” The map finally became a living system rather than a decorative board.

b. Rebuilding Decks

The first version’s decks contained only “playable cards” and “traps.” The bosses were added to the Classroom deck after every Classroom tile was placed on the board. The bosses were the only way to obtain Credits, just like Encounters now. Initially, it was plausible and coherent with the game’s theme. First players –as students– would explore the campus, take classes, go through various situations, and then would take their exams. Just like in real life. However adding

them later caused a lack of pacing, surprise and narrative rhythm. The redesign created a much more dynamic structure:

•Encounters (now the primary source of Credits)

•Events

•Actions

•Traps

•Items

This change did two things at once:

- It made every draw meaningful, dangerous or surprising.
- It created the “run structure” typical of roguelites: bursts of ad-

vantage, moments of loss, occasional chaos and steady progression through risk.

Crucially, Encounters became the core loop, not a late-game add-on. They created the tension and flow that the first prototype completely lacked.

c. Designing for Social Interaction

The surveys made it clear that the biggest missing piece was player–player interaction. A roguelite without tension is weak; a multiplayer roguelite without negotiation, sabotage, or alliances is simply not a game.

The redesign addressed this directly:

- Players can now help each other during Encounters (with negotiated rewards).
- They can intervene to sabotage Encounters.
- They can trade Items.
- They can coordinate to stop a leader.

These systems created the "friction" that emergent narrative theory identifies as essential for story formation (Webb & Cesar, 2019; Kybartas et al., 2020). Every rule that encouraged silence was removed. Every rule that encouraged tension, diplomacy, or betrayal was added.

The result was a shift from multiplayer solitaire to tabletop politics – exactly the tone the game needed.

d. Shortening the Board and Tightening Movement

Shrinking the grid from 8×8 to 6×6 was a mechanical necessity. With a smaller space:

- players naturally collide more often,
- movement stops feeling like a commute,
- tile passives become relevant,
- and map layout becomes a shared story with more chaos.

It was one of the first changes that sprung to mind but maybe will be the most impactful.

5.5.4.3 The Redesign as Identity Formation

With this redesign process the game became a clearer, patched version of the first prototype. The new system has:

- a clear core loop,
- clean stats,
- purposeful tile interactions,
- meaningful randomness,
- real social tension.

Its identity is more consolidated around these changes: a comically hostile campus filled with academic dangers and petty rivalries. Now it has a structure where students race, bargain, block and occasionally ruin each other's plans. The Final Exam is still in the development phase as I am writing this but it definitely started becoming a real game.

Hollow Knight, 2017

CHAPTER 6

CONCLU-SION

6.1 Revisiting the Purpose of the Thesis

This thesis set out to understand whether the distinctive qualities of roguelike and roguelite games; such as procedural variation, systemic storytelling, unpredictability and iterative progression could meaningfully translate into a physical tabletop format. Digital roguelikes rely on computational processes that constantly reshape the player's experience. The central question guiding this work was whether these principles could remain coherent, functional and expressive once removed from a digital environment and rebuilt through analog components, human interpretation and social interaction.

The aim was to explore how roguelike games' design philosophy can be adapted to create the emergent narrative experience in a tabletop environment. This required examining how stories emerge from systems rather than scripts, how randomness can structure play rather than disrupt it and how player decision-making changes when the game becomes a shared, physical and social space. Through a combination of theoretical research and hands-on design, the project sought to identify which elements of the genre survive translation, which require reinterpretation and which are fundamentally transformed by the shift from screen to table.

6.2 What the Thesis Demonstrates

Across the theoretical chapters and the practical development of The Final Exam, this thesis demonstrates that roguelike design principles can be translated into a physical medium. This requires careful reinterpretation. Direct replication doesn't work. Analog games operate under completely different conditions compared to their digital cousins. They are slower, social, physically constrained and reliant on players to manage the system. The research showed that emergent storytelling does not come from complexity itself, but from how players interact with simple, repeatable systems.

The thesis also makes clear that randomness behaves differently across formats. In digital spaces, procedural generation can produce sharp difficulty spikes or intricate layouts without affecting flow. On a tabletop, randomness must be structured, legible and easy to resolve or it risks becoming noise. The work in Chapter 4 and 5 illustrates that physical randomness needs to be paired with visible decision-points so that players always feel agency.

A second key insight is that social dynamics replace computational depth in analog emergent storytelling. Tabletop games' shared interpretation and social interaction substitute digital roguelites' generating variety through algorithms. This shift places players at the center of narrative creation. The system becomes the players. The first prototype failed precisely because it ignored this: it behaved like an automated simulation rather than a social game. The redesign ought to succeed because it embraces human behavior as the primary narrative engine.
Ultimately, the project demonstrates that the core of roguelike identity can survive outside the digital environment. To achieve this, it must be rebuilt using tools that belong to the tabletop.

6.3 Reflection on the Design Process

Working on The Final Exam made the difference between theoretical knowledge and practical design unmistakably clear. At the beginning, I approached the project with the confidence that a strong conceptual foundation would naturally lead to a strong game. I quickly learned that systems that look coherent in diagrams or academic texts often behave very differently once players touch them. The first prototype forced me to confront this gap directly.

I entered the project with a tendency to add mechanics, believing that complexity would create depth. Through prototyping and feedback, I learned the opposite. Depth comes from clarity. The strongest parts of the game emerged only after I removed the systems that were getting in the way. Cutting the Energy system, collapsing the stats and shrinking the board reshaped the identity of the entire project. Subtraction became the creative act that allowed the real game to surface.

The process also changed my understanding of emergent storytelling. In video games, emergence is associated with procedural generation and algorithmic variation. Through prototyping, I realized that in analog games, emergence comes from people. Their interaction and improvisation emerges from the story. Designing this game taught me that emergent narrative in physical play is less about simulating the unexpected and more about enabling it.

Finally, this project reshaped the way I see my own role as a future designer. I started with a rigid idea of how a roguelite “should” work. I ended with a better understanding of how to adapt a genre’s philosophy to a new medium. The difficulties of the first prototype were necessary experiences that sharpened my ability to recognize what serves a good game and what distracts from it. The Final Exam is still evolving, but the process documented in this thesis marks the moment it stopped being an imagined system and became a real, playable foundation to build on.

References

Akcaoglu, M., & Green, L. S. (2018). Teaching systems thinking through game design. Educational Technology Research and Development, 66(6), 1505–1523.

Bernhard, P., & Deschamps, J.-P. (2021). Dynamic games with randomly arriving players.

Bosman, F. G., & van Wieringen, A. L. H. M. (2018). I have faith in Thee, Lord: Criticism of religion and child abuse in the video game The Binding of Isaac. Religions, 9(133).

Bowman, S. L. (2014). Social conflict in role-playing communities: An exploratory qualitative study. Analog Game Studies / WyrdCon Companion.

Burgess, J., & Jones, C. (2021). Player co-creation and emergent narratives in the Total War series.

Cambridge University Press. (n.d.). Entropy. In the Cambridge Dictionary. https://dictionary.cambridge.org/dictionary/english/entropy

Cavazza, M., Charles, F., & Mead, S. J. (2002). Interacting with virtual characters in interactive storytelling. In Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (pp. 318–325). ACM.

Chauvin, A., Levieux, G., Donnart, J.-Y., & Natkin, S. (2015). Making sense of emergent narratives: An architecture supporting player-triggered narrative processes. In 2015 IEEE Conference on Computational Intelligence and Games (CIG) (pp. 1–8). IEEE.

Chauvin, A., Mottin, E., Buemi, A., & Buemi, E. (2015). Narrative processes: Reifying and managing emergent narratives in persistent game worlds.

Chen, Y. (2023). Cyber discipline: How roguelite games engage players – A perspective from Foucault's Discipline and Punish. Proceedings of the International Conference on Interdisciplinary Humanities and Communication Studies, 6, 504–511.

Clark, D. B., Tanner-Smith, E. E., & Killingsworth, S. S. (2016). Digital games, design, and learning: A systematic review and meta-analysis. Review of Educational Research, 86(1), 79–122.

De Cooman, G., & De Bock, J. (2021). Possibilistic equilibria under imprecision: A revision.

De Kegel, B., & Haahr, M. (2020). Procedural puzzle generation: A survey. IEEE Transactions on Games, 12(1), 21–40.
de Lima, E. S., Feijó, B., & Furtado, A. L. (2022). Procedural generation of branching quests for games. Entertainment Computing, 43, 100491.
Dubbelman, T. (2016). Narrative game mechanics. In F. Nack & A. S. Gordon (Eds.), Interactive Digital Storytelling: ICIDS 2016 (pp. 39–50). Springer.
Fullerton, T. (2018). Game Design Workshop: A Playcentric Approach to Creating Innovative Games (4th ed.). CRC Press.
Gellel, A., & Sweetser, P. (2020). A hybrid approach to procedural generation of roguelike video game levels. In Proceedings of the 15th International Conference on the Foundations of Digital Games (FDG '20) (pp. 1–8). ACM.
Grey, D. (2013, May 14). Screw the Berlin Interpretation! Games of Grey. https://www.gamesofgrey.com/blog
Guo, S., Peterson, B., & Tran, L. (2022). Homogenization of random walks in balanced random environments. Electronic Journal of Probability.
Gustafsson, V., Holme, B., & Mackay, W. E. (2020). Narrative substrates: Reifying and managing emergent narratives in persistent game worlds. In Proceedings of the 15th International Conference on the Foundations of Digital Games (FDG '20) (pp. 1–6). ACM.
Guzsvinecz, T. (2022). The correlation between positive reviews, playtime, design and game mechanics in souls-like role-playing video games. Multimedia Tools and Applications, 82, 4641–4670.
Ho, M. J., & Carter, M. (2019). Roguelike ancestry network visualisation. In 2019 IEEE Conference on Games (CoG) (pp. 1–8). IEEE.
Hunicke, R., LeBlanc, M., & Zubek, R. (2004). MDA: A formal approach to game design and game research. In Proceedings of the AAAI Workshop on Challenges in Game AI.
Ip, B. (2011a). Narrative structures in computer and video games: Part 1. Games and Culture, 6(2), 103–134.
Ip, B. (2011b). Narrative structures in computer and video games: Part 2. Games and Culture, 6(3), 203–244.
Johnson, M. R. (2015). The use of ASCII graphics in roguelikes: Interface and aesthetic. Games and Culture, 10(6), 547–574.

Kosof, J. (2021). Building creativity support tools for dungeon design in tabletop games: A participatory design study. Northeastern University.
Kybartas, B., Verbrugge, C., & Lessard, J. (2020). Tension space analysis for emergent narrative. IEEE Transactions on Games, 12(3), 235–246.
Martinho, C., & Sousa, M. (2023). CSSII: A player motivation model for tabletop games. In Proceedings of the 18th International Conference on the Foundations of Digital Games (FDG 2023) (pp. 1–3).
Mathys, J. (2020, October 8). Hades has a larger word count than The Iliad and The Odyssey combined. Game Rant. https://gamerant.com/hades-word-count-illiad-odyssey/
Peinado, F., & Gervás, P. (2004). Transferring game master's knowledge to a computer role-play storytelling system.
Pereira, Y. H., Ueda, R., Galhardi, L. B., & Brancher, J. D. (2019). Using procedural content generation for storytelling in a serious game called Orange Care. In Proceedings of SBGames 2019 (pp. 192–194).
Pereira de Araujo, R., & Souto, V. T. (2017). Game worlds and creativity: The challenges of procedural content generation. In A. Marcus & W. Wang (Eds.), Design, User Experience, and Usability (pp. 443–455). Springer.
Peng, X., Quaye, J., Rao, S., Xu, W., Botchway, P., Brockett, C., ... Dolan, B. (2024). Player-driven emergence in LLM-driven game narrative. arXiv preprint arXiv:2404.17027.
Pozzi, F., Ceregini, A., Ivanov, S., Passarelli, M., Persico, D., & Volta, E. (2024). Digital vs. hybrid: Comparing two versions of a board game for teacher training. Education Sciences, 14(318), 1–14.
Prins, J. P., Valls-Vargas, J., Ontañón, S., & Bidarra, R. (2023). StoryWorld: Generating believable quests from simulated worlds with character memory and desire. In Proceedings of FDG 2023 (pp. 1–4).
RandomisedTV. (2025). [FWR] Balatro set seed glitched – 1m25s speedrun [Video]. YouTube. https://www.youtube.com/watch?v=lUdjBSwX5SQ
Risi, S., & Togelius, J. (2020). Increasing generality in machine learning through procedural content generation. arXiv preprint arXiv:1911.13071v2.

Ruiz, V. S., & Mota, S. R. da. (2024). Use of the roguelike genre as a narrative tool: Memories and child development in The Binding of Isaac. UFMG Proceedings on Game Studies.
Ruiz, V., & Mota, S. (2022). Use of the roguelike genre as a narrative tool: Memories and child development in The Binding of Isaac. Journal of Gaming & Virtual Worlds, 14(2), 173–185).
Salen, K., & Zimmerman, E. (2004). Rules of Play: Game Design Fundamentals. MIT Press.
Sim, Y. T., & Mitchell, A. (2017). Wordless games: Gameplay as narrative technique. In ICIDS 2017 (pp. 137–149). Springer.
Smith Nicholls, F., & Cook, M. (2023). "That Darned Sandstorm": A study of procedural generation through archaeological storytelling. In FDG 2023 (pp. 1–4).
Tapscott, A., Thompson, B., Carenini, G., & Tate, R. (2018). Modeling conversation as a game: Turn-taking and narrative structure.
Webb, A. M., & Cesar, P. (2019). Narratives in multiplayer tabletop play: Friction, coordination, and story emergence.
Wee, C. (2024). "If it's held dear, it'll get pushed through": Transmedia narratives, play cultures, and soft canon in tabletop role-play.
Wong, Z. S. Y., Wang, P., & Fang, J. (2021). On human generation and perception of randomness: Sequence complexity and streaks.

Ludography

Angband developers. (1990–). Angband [Video game]. Open-source community project.
Balatro / LocalThunk. (2024). Balatro [Video game]. Playstack.
Biskup, T. (1994–). ADOM (Ancient Domains of Mystery) [Video game].
Blizzard Entertainment. (2012–). Diablo III [Video game series]. Blizzard Entertainment.
Blizzard Entertainment. (2014). Hearthstone [Video game]. Blizzard Entertainment.
CD Projekt Red. (2015). The Witcher 3: Wild Hunt [Video game]. CD Projekt.
Dead Cells / Motion Twin & Evil Empire. (2018). Dead Cells [Video game]. Motion Twin.
FromSoftware. (2011–2016). Dark Souls [Video game series]. Bandai Namco Entertainment.
id Software. (1993). Doom [Video game]. id Software.
Jackson, S., & Livingstone, I. (1982–). Fighting Fantasy [Gamebook series]. Puffin Books.
Kojima Productions. (1998). Metal Gear Solid [Video game]. Konami.
McMillen, E. (2018). The Binding of Isaac: Four Souls [Card game]. Studio71.
McMillen, E., & Nicalis. (2011). The Binding of Isaac [Video game]. Nicalis.
MegaCrit. (2017). Slay the Spire [Video game]. Humble Bundle.
Mossmouth, LLC. (2008). Spelunky [Video game]. Mossmouth.
NetHack DevTeam. (1987–). NetHack [Video game]. NetHack DevTeam.
Shiny Shoe. (2020). Monster Train [Video game]. Good Shepherd Entertainment.
Subset Games. (2012). FTL: Faster Than Light [Video game]. Subset Games.
Supergiant Games. (2020). Hades [Video game]. Supergiant Games.
Toy, M., Wichman, G., Arnold, K., & Lane, J. (1980). Rogue [Video game]. Epyx.
TSR. (1974). Dungeons & Dragons [Tabletop role-playing game]. TSR.

This thesis examines how emergent storytelling works in roguelike and roguelite games, and how their narrative logic can shift from digital systems to physical tabletop play. It traces the genre's evolution and shows how procedural generation, permadeath, repetition and systemic design act as narrative engines rather than simple mechanics. Moving to tabletop formats changes everything: without automated updates, players must manage variation, state and negotiation themselves, creating new constraints but also new opportunities for shared, emergent stories.

The Final Exam, the design project at the centre of this thesis, tests how roguelike principles—procedural variety, failure loops, risk and small narrative beats—can be adapted to a physical game. Playtests and redesigns reveal the challenges of balancing randomness and difficulty without digital support, but also how systemic thinking still drives meaningful play. The thesis argues that roguelikes aren't defined by code or platform; they work wherever rules interact dynamically. When systems invite discovery and surprise, players naturally generate their own stories—on a screen or around a table.

Candidate: Can Uygurcetin
Supervisor: Andrea di Salvo
The Final Exam's Art Director: Eda Bolver