Beyond Authorial Burden

JOEY JONES, Electronics and Computer Science, University of Southampton, Southampton, United Kingdom of Great Britain and Northern Ireland
DAVID MILLARD, Electronics and Computer Science, University of Southampton, Southampton, United Kingdom of Great Britain and Northern Ireland

Interactive Digital Narrative (IDN), a primarily web-first hypertextual medium, creates an overhead of writing for authors. There are many methods and tools that have been proposed for writing IDNs with the goal of reducing this burden, but because there is no comprehensive model of the Authoring Burden assessing the impact and appropriateness of these approaches is difficult. We have undertaken interviews with IDN authors (n=14) to understand how they manage the authoring burden within their own projects. Based on these interviews, and drawing on the existing literature, we propose a model of the Authoring Burden comprised of three parts: Content Creation, Dynamic Authoring, and Programming/Tool Creation. The initial size of this burden is set by the Author's Goals informed by their Capability and the Audience/Publishing Context. We also find 29 strategies employed by authors to manage the burden. There are five distinct types. Embracing and Reducing strategies impact the overall scale of the challenge, whereas Generative, Reuse, and Decoupling strategies move work between the three parts of the model. We validate our model with focus groups comprising different sets of experts (n=8). Our model shows that many strategies for managing the burden transform rather than reduce work, and that the most appropriate strategy for a given author will be highly dependent on their personal goals and capabilities. It shows claims to alleviate the 'burden' of authoring labour may often shift the nature of the labour itself, or the design of a given IDN, into unwanted forms.[1]

CCS Concepts: • **Applied computing** → **Media arts**; • **Human-centered computing** → **Hypertext / hypermedia**.

Additional Key Words and Phrases: Authoring, Hypertext, Hypertext Fiction, Literary Hypertext, Interactive Fiction, Interactive Digital Narrative, Game Design, Design Patterns

1 Introduction

Interactive digital narrative (IDN) is a medium with roots in early hypertext and mainframe computing [23].[2] IDNs are digital game-like narratives where the player typically engages with the experience through agential play, by making choices. Contemporary IDNs are primarily a web-first phenomenon. They are often published online using web architecture (for example with Twine and JavaScript, or Inform and web interpreters like Parchment). Community development of these works, including programming help and running game jams, happens through online forums, and through long-running community fixtures such as the Interactive Fiction Competition which has been running online since 1995.

The nature of IDN imposes a workload on the author: whether it is a hypertext, a simulated environment, a set of scenes or even a database, every plot branch, alternative description or dialogue choice creates extra

[1]Some of the text of this paper appeared as Experiencing the Authorial Burden, 2024 [18].
[2]A shorter version of this paper appeared as Experiencing the Authorial Burden [18]

Authors' Contact Information: Joey Jones, Electronics and Computer Science, University of Southampton, Southampton, United Kingdom of Great Britain and Northern Ireland; e-mail: J.D.Jones@soton.ac.uk; David Millard, Electronics and Computer Science, University of Southampton, Southampton, United Kingdom of Great Britain and Northern Ireland; e-mail: dem@soton.ac.uk.

ACM 1559-114X/2025/8-ART
https://doi.org/10.1145/3757746

writing labour over and above the labour inherent in writing narratives. A digital work, where it isn't as limited by material restrictions such as page counts, can also end up much larger than non-digital interactive narratives, such as printed gamebooks or plays with alternative endings. This overhead when writing for an interactive system is referred to as the 'authorial burden' [17].

Frequently in the literature, there are claims that a new tool or approach will help alleviate the authorial burden. To pick some examples: this claim is made in support of Ryan et al.'s dialogue combination annotation scheme [28]; Talbot & Youngblood's behaviour markup of play scripts [32]; and Dobrovosky *et al*'s argument for using machine learning to create artificial intelligence of non-player characters in Serious Games [9]. In each case, the authors present an approach to authorship that focuses on the use of automation with the justification that it will save on author time. These arguments invite the question whether the aesthetic trade-off for the author is worthwhile — pouring all of the paint directly from the tin onto the canvas would save a painter time, but usually wouldn't cohere with the painter's vision for their work. There is also the question of whether the labour of annotating scripts, editing machine-generated output and programming behaviours is more interesting and worthwhile to the authors of interactive narrative than the creative labour of making bespoke authored content. How do IDN authors currently experience and mitigate the authorial burden? With the intention of exploring this question, we undertook a series of fourteen qualitative interviews with authors of narrative-based game forms, including text-adventures, visual novels, and choice-based fiction. Out of these interviews, a model was developed of the authorial burden and the strategies used to overcome or transform it, as well as capturing the context in which authoring is occurring and its influence on those decisions. This model was then discussed with three further groups of authors, who acted as expert panels to validate the model and refine it further.

In Section 2 of this paper, we present an overview of interactive narrative authoring, the authorial burden, and previous conceptual work in this space. In Section 3, we present the methodology for the interview series and expert panels. In Section 4 we provide an overview of the results of the interviews in the form of our conceptual model of the authorial burden. In Section 5 we discuss the outcomes of the expert panel validation exercise. In Section 6 we discuss core outcomes of the model. Finally, Section 7 concludes with our key findings and suggestions for future work.

2 Background

Interactive digital narratives can take many different forms, from choice-based hypertext and Twine stories like *With Those We Love Alive* (Porpentine, 2013), parser-based text adventures such as *Counterfeit Monkey* (Short, 2012), online storylet-based works like *Fallen London*, interactive films such as *Bandersnatch*, procedurally managed dramas like *Façade*, and more. Following Koenitz, we can say interactive digital narratives are 'expressive narrative forms implemented as a computational system [...] and experienced through a participatory process' [22]). This echoes Crawford's earlier conceptualisation of an 'interactive' narrative as requiring a system which listens (is participatory), thinks (is computational), and speaks (is expressive) [8]. In a similar way to how a playwright or composer creates the means for a potential performance, the author of interactive narrative creates a system for generating possible narrative experiences [21]. What is dissimilar is that the eventual experience is partially co-created through reader/player decisions. The interactive narrative that is experienced can be considered as an event, and the program which allows that experience as a system of variations [7].

A traditional novelist may hold in mind different visions of a work or produce multiple drafts or variants, in what has been referred to as fluid text [6]. However, the labour of producing a program that can dynamically reveal alternative traversals of a work produces additional labour for the author of interactive digital narrative who faces combinatorial problems arising from the interactive structure of their works. An often cited early exploration of the authorial burden is found in Bruckman's 1990 paper on the 'Combinatorics of Storytelling' [5], where she describes the burden of exponential growth of branches in the context of interactive film. If you have

one scene which can lead to two other possible scenes, and those second set of scenes can lead to two different scenes, the content needed to be written doubles at each step. This is referred to by Garbe as the *Authoring Wall*, an exponential demand for written content (though it can apply to any other kind of content as well, such as art and animation) [14]. Strategies that seek to manage this branching through structural patterns, or state tracking, will instead tend to grow in complexity: the author has to keep in mind a 'possibility space' when writing the narrative, and more possible combinations of story-state increase the mental overhead [15]. This may often eventually slow the production of new content, as the number of possible states needed to account for becomes hard to conceptualise, in what Garbe refers to as the *Complexity Ceiling* [14].

The burden is often referred to as a problem for new tools or approaches to solve. For instance, Fendt et al propose the strategy of using fake choices for overcoming the authorial burden [11], an approach later taken up by Dominguez [10] in the context of narrative roleplaying games. Later the term is used two separate papers by Ryan, with combinatorial dialogue authoring [28], natural language generation Ryan et al. propose as strategies for overcoming this burden [27]. Most recently Sun et al.'s work creating an LLM-powered storylet framework with the intention of reducing the 'implementation burden' [31]. Several of these claims are explored and challenged in Section 6.

Kitromili's interview series explored interactive narrative authors' experiences of the authoring process as a whole, and her approach informed our own investigation [20]. Within her schema of authoring issues, the authorial burden can be considered one form of *scalability* issue. Recent conceptual work in this area has been done by Jones [17], who presents a taxonomy of strategies that authors can use to approach the authorial burden. This was used to organise the strategy codes in the interviews presented in this paper and will be discussed below.

Our objective is to develop these ideas on limits and approaches into a more comprehensive model of the choices that authors make to manage the authorial burden. Our intention is to take into account both the scale of content (the authoring wall) and the complexity of state management (the complexity ceiling). With this goal in mind we interviewed a number of interactive narrative authors to understand how they approached the problem.

3 Methodology

Interactive digital narrative writers and designers are experts in their own craft, and grapple with design and implementation questions at every stage of a work's development. Between April and December 2022, the first author undertook a series of semi-structured interviews with creators of IDN about their writing process in relation to issues of scope management. The study was approved by the University's ethics committee (study number: ERGO/FEPS/70996).

A qualitative, inductive approach was chosen as most appropriate for gathering genuine first-hand insights from authors without pre-judging their difficulties or priorities. The IDN authors were chosen through a mix of opportunistic use of personal contacts and snowballing. As the first author is a published IDN author, and thus embedded in interactive narrative creation, there were unique opportunities to speak with people who may otherwise have been more unknown or less reachable. Moreover, the specific advantage of peer-to-peer interviewing of experts, over the model of the expert and the outsider, is that a greater store of knowledge is initially shared and thereby more knowledge can be revealed during the interview [4]. An initial roster of participants was drawn from existing connections, and at the end of each interview they were asked who they would recommend speaking with, leading to four more interview subjects that were successfully contacted and interviewed. This snowballing allowed for more variety in the participants, though with much smaller success rate in requests. People are more likely to give a significant chunk of their time to someone they are familiar with than to someone who contacts them coldly.

Participant Number	Main interactive narrative types	Snowballed
P1	Branching choice-based	P7
P2	Parser-based text adventure, procedurally generative text	
P3	Visual Novel	
P4	Branching choice-based, parser-based text adventure	
P5	Parser-based text adventure	P11
P6	Homebrew choice-based engine	P14
P7	Drama engine, agent-based	
P8	Parser-based text adventure	P9
P9	Parser-based text adventure	
P10	Choice-based, narrative generation	
P11	Parser-based text adventure, storylet based	
P12	Choice-based, Storylet-based	
P13	Choice-based	
P14	Choice-based, narrative generation	

Table 1. Table of Interviewees

The aim of this process was not to achieve a demographic representation, (which for this form of long-form interview would not be feasible), but rather to hear from people working across a range of core interactive narrative forms (see Table 1) in both hobbyist and commercial contexts. In total 14 people were interviewed, 8 of whom had written works for a commercial market, with the rest primarily producing freely available games. Every author (or often co-author) of commercial games had also developed non-commercial works (a common practice for game developers). All of the authors had published at least one IDN that was hypertextual in structure or playable within a web browser. Many of the participants took part in the interactive fiction community, submitting freely-playable works to online game jams and competitions. All of the interviewees had worked with media and story structures beyond the specific work for which they were interviewed, which provided additional insight.

Before each interview, authors were given a participant information sheet to read the details of how the study would be performed and its intended focus. Each signed a consent form. The interviews took place over a recorded video call. They were asked a series of questions on how they came to write one of their published works, the challenges related to managing scope and authorial burden, and the approaches they used to overcome these challenges. The interviews were semi-structured around a series of 'experience questions' designed to encourage rich responses about each author's specific design practices [16]. Often follow-up questions would be asked for additional clarity, and authors were usually encouraged to speak at length when they came to matters of particular interest to their craft. Each interview lasted approximately an hour.

The recorded interviews were later inductively coded, primarily by the first author. From what the participants said, common themes relating to the authorial burden were identified, as well as the wider context in which they came to author their works (which in turn shaped the workload). This process was done iteratively with the assistance of the second author, with earlier interviews re-coded with relevant codes that were added after later interviews.

Synthesising the insights of the interviews, a conceptual model of the authorial burden, based around the main themes, was created. This model in turn was presented back to a different group of IDN authors in a series of expert panels for validation. For the panels, the participants were recruited through posts on social media and interactive fiction web forums, ensuring a different mix of people. In total, eight authors of interactive

	Codes	P	M	Definition
Authorial Burden	Dynamic Authoring	12	26	Authoring stateful game narrative.
	Feature Creep	6	11	Expansion of scope during a project.
	Programming Labour	6	9	Producing computer-facing code.
	Scopability	11	27	How possible it is to judge workload at the outset.
	Writing Burden	11	19	Narrative text workload.
Capabilities	Motivation To Write	11	34	Reasons for writing interactive narrative.
	Technical Limitation	10	24	Drawbacks of existing authoring tools.
	Tool Mastery	12	25	Author ability with authoring tools.
	Writing Ability	6	20	Author ability with writing.
Design Goals	Agency	5	7	Implementing gameplay with a feeling of freedom of choice.
	Evenness	5	7	Maintaining quality between branches.
	Genre	10	21	Impact of genre on game structure.
	Legibility To Player	9	28	Making the impact of player's choices clear.
	Narrative Quality	8	18	Good prose, dialogue, story, use of language etc.
	Player Affordances	8	19	Accessibility, hinting, and other player aids.
	Puzzle Design	8	22	Creation of puzzles.
	Replayability	7	7	Designing for a repeatable, varied experience.
	Types of Choices	8	14	Different forms of player choices.
Release Environment	Audience Feedback	10	13	Actual and expected response from players.
	Publishing Context	12	22	How and where a work is distributed.

Table 2. List of themes and codes (excluding strategies), with number of times mentioned and the number of participants who mentioned the code, and a brief definition.

narrative took part in three panels (two with three people, one of two). All panelists had made at least one released parser-based or hypertext work, with most having made several.

In each case, the conceptual model of the authorial burden was introduced to the panelists both before the interview in preparatory documents and at the start of each discussion session. The panels were structured as more open conversations than in the original interviews, with questions around the model discussed both directly and implicitly, and — in line with embracing strength of co-expert interviewing — the participants were encouraged to ask one another follow-up questions. The goal of the panels was to test the conceptual model, with experts in the field best placed to offer verification, challenges, and caveats.

4 Results

In the interviews a total of 19 codes were identified across four themes: *burdens, capabilities, authorial goals, and release environment*. These are shown in Table 2 with code counts and brief definitions.

In addition we identified 29 additional codes across a fifth theme: *strategies*. Each of these codes represented a specific approach to managing the authorial burden. These are shown in Table 3 organised into the five sub-categories identified by Jones [17]: strategies that reduce what needs to be created, reuse content, decouple content, generate content, or embrace the burden. Unlike our main coding activity, discovering and ascribing the strategies to these sub-categories was deductive insofar as the categories were pre-theorised.

In the following sections we will explore each of these themes in turn. Participants are referenced as P1 - P14.

Strategy Themes	Codes	P	M	Definition
Decoupling	Episodification	3	4	Splitting narrative into self-contained segments.
	Modularisation	4	6	Repeatable or expandable segments of gameplay.
	Storylets & Plot Clouds	3	4	Dynamic selection of narrative chunks.
Embracing	Automation	4	8	Programming labour to save future rote work.
	Collaboration	8	16	Increasing the number of authors.
	Flow	9	25	Enhancing creative capacity of an author.
	Iterative Writing	5	9	Implementing ideas had during production.
	Knowledge Management	1	1	Using tools to manage epistemic state.
	More Time	7	9	Expanding production time.
	No Rewrites	3	3	Avoiding editing and redrafting.
	Planning & Prototyping	8	12	Planning the design before writing.
	Tool Creation	5	9	Developing tools to aid authoring.
	Visualisation	4	7	Diagramming design structures.
	Writing Out Of Sequence	3	3	Authoring later parts of a narrative earlier.
Generating	Procedural Content	9	21	Programming rules to automate creation.
	Story Sifting	1	2	Surfacing implicit narrative.
Reducing	Abstraction	2	2	Using non-realist elements to simplify.
	Cutting	10	20	Removing planned or written content.
	Reducing Actions	5	23	Limiting player verbs or choice range.
	Reducing Objects	8	17	Limiting items, characters, or locations.
	Reducing Play-Time	2	2	Shortening the experienced length of the work.
	Reducing Tracking	2	2	Limiting the statefulness of the narrative.
	Words vs Assets	2	5	Using text and dialogue over visual assets.
Reusing	Cumulative Variables	3	5	Accretive agency rather than immediate branching
	House Style	1	3	Established design patterns for multiple authors.
	Recycling Work	4	5	Keeping and repurposing unused ideas and work.
	Shared Assets	2	3	Reusing game elements in different contexts.
	Structural Patterns	8	31	Merging, looping, and other design structures.
	World Model	4	8	Implementing repeatable default game behaviour.

Table 3. List of strategy codes, with number of participants (P) who mentioned the code, total number of unique mentions (M), and a short definition.

4.1 Burdens

"It's always more work than anticipated, right? And the question is whether you keep ploughing. A lot of game concepts start out as something very small, right? They start as a joke. They start as a shitpost. They start out as a vibe. And it never seems to be as big an idea. Sometimes you have some clue, but you never know all the challenges when you first have the idea." - P2

Most of the participants directly spoke about the **Writing Burden** of creating content for an interactive system. Although as P11 expressed, 'the word count or the branch number or the play time or all these standard metrics that people use to measure the size and scope of games' could be deceptive and not 'map on really with the workload for the writing'. Participants described how the design itself would entail a concomitant writing workload. There was the amount of work involved in writing any significant creation made from text

(a workload shared by novelists and playwrights): P4 relates that, 'even writing the throughline is a significant challenge because you just need so much plot content'. And then there was the work that arose specifically out of the interactive nature of IDN. This workload could emerge through 'combinatorial complexity', which P4 explained: 'You can take two reasonable looking things and combine them. You have a hammer and a nail. You try hammering the nail'. The amount of writing needed increases as more combinable elements are added to the design.

Dynamic Authoring involves writing conditional or variable content that can require different game-states to instantiate. It could also involve the 'hard work' of 'creating grammars: taking sentences that you would like the work to be able to output and generalizing them... while still keeping underlying structure' (P2). As P6 describes, dynamic authoring is a combined activity of authoring narrative and designing state engines, 'you're writing the puzzle and you're writing the words at the same time... the writing and the programming feel more like the same job'.

Programming Labour — the creation of new game-behaviour, or developing new tools — was identified by some as a 'frustration' (P5), for others it was just a different kind of work that allows authors to use 'different parts of [their] brain' (P6). Where the game design required a large programming overhead for new chunks of content, this could hinder development. P14 said that 'the biggest problem with [a game] was that writing those specific plotty scenes was really difficult' because of the need to program a lot of bespoke behaviour, with the consequence being that the developer 'wrote far fewer of those than... originally intended', which resulted in the game feeling 'a little bit thin'.

Scopability, the possibility of judging the amount of work required for a project, was usually low at the start of making games. As P2 explained, 'the idea in your head at the start of a project is never what comes out at the end', because 'in order to make a project work, you have to nail down certain structural things which may change everything, and that includes what is the proper size of the thing'. This difficulty was pervasive but not absolute, as some chunks of content became easier to judge. P11 described gaining a sense that '100 objects are going to take this long' to make. Still, in the working out of a design, and the process of writing it, scope inevitably changes. This means that 'it is not practical to scope a project once and then be done with it' (P2).

Feature Creep would increase workload as new ideas emerged during creation or testing. What could seem at the start like 'a quick, easy project that wouldn't take that long', could soon 'balloon out of control' (P14). Feature creep was often seen as inevitable ('we ran over, obviously, like you do in every game' —P10), but this was not always a bad thing. Sometimes the author finds a better way to implement their work, often from early play testing, as P9 experienced: 'eventually I saw how I could do things better... I started seeing these possibilities really as I started playing the game and started getting transcripts'. Often the narrative or setting as-presented would suggest possibilities to players during testing that the authors hadn't considered, and to improve the responsiveness and reasonableness of the IDN, the author would want to add in some response (whether an alternative puzzle solution or just a specific reason the player's idea wouldn't work), with P7 saying that if someone during a 'play test makes a move' that makes sense but is currently disallowed by the system, then 'I think it is our duty in to write a custom response'.

4.2 Capabilities

"Just because I know what good writing looks like doesn't mean I can produce it." - P4

Authors of interactive narrative have different abilities and interests. Authors can and do increase their technical and creative abilities, but at any given moment in a project there is a limit to what they can achieve on their own which varies substantially from author to author. They would come to each project with their own prior **Writing Ability**, or 'design knowledge, something that the author builds up themselves in relation to what they're working on and the environment in which they're creating' (P10). Participants were keenly aware of

their own strengths and weaknesses in their writing ability, 'knowing your abilities, what you're capable of' (P9). Authors described playing to their strengths, as different game forms would tax different skills. For this reason, P4 maintained that 'a bad author can make a great parser game' but not a great choice-based game.

Digital narrative is created with programming tools including both general and domain specific languages, integrated development environments, and game engines. Authors varied in their **Tool Mastery**. 'You work with the tools that you know how to use', and in turn the tools impact the vision of the game: 'what's ready to hand shapes your imagination a lot' (P2). Once familiar with one tool, authors experience a considerable 'cost in becoming familiar with something else' (P1). This opportunity cost of time and effort in switching to more appropriate tools meant authors would often make do with what they were already familiar with.

The tools themselves each have their own **Technical Limitations** to what can be achieved with them and sometimes 'even just getting things to work' could be 'slow' (P14). Technical limitations could also be imposed from without from publishing platforms, such 'super constraining' restrictions on the App Store (as noted by P10), or the file limitation on the Interactive Fiction Competition website, which led P11 to include less audio in one project submitted to the competition.

4.3 Authorial Goals

"I wanted to have a sense that people could have quite a different experience... I wanted to have people able to express themselves or have the characters responding in different ways." —P13

An author of interactive narrative may have any number of goals for their work which will have concomitant impacts on the burden of work. The interactive element of IDN is found in the player making choices— whether that be explicit choices over the direction of the plot, tactical game-related decisions, choices to express character, or just moving the protagonist around a virtual space. As per Murray [24], we can distinguish agency as an enhanced form of interactivity: the player not only takes action, but the results of that action can be seen. Player agency is the feeling that 'choices matter'. Several of the participants spoke directly on trying to increase player **Agency** as a design goal, where a player 'could actually make a real decision that had a major effect on the story' (P4). Attempting to implement player agency can be a great driver of workload, but is seen as creating a rewarding experience. P4 thought that implementing 'enough choice' would 'draw reader engagement significantly'. Even IDNs that purposefully restrict player agency as part of their authorial goal often 'require a certain amount of work' (P11) such as in redirecting or blocking player actions with unique messages.

For some authors, **Replayability** was a core design goal. They wanted to give players repeatable, varied experiences. Replayability for a narrative-centred work involves genuinely writing different content to experience rather than just giving the appearance of alternative content. As P1 expressed, 'whenever there's a lot of content' experienced on repeated play, 'if anybody goes through and reads all of it, I want them to feel rewarded rather than to be cheated'. That's not to say this always has to branch the plot absolutely, but it does entail writing different responses. 'Making your way through the story in a pretty varied way', in P13's experience, could be facilitated by giving players opportunities to 'express themselves', seeing 'the characters responding in different ways'. Replayability wasn't a universal goal among the participants. Some authors expressed a desire to give limited time, non-compulsive experiences. P9 contrasted 'casino'-like videogames which can be 'played too much' with what 'interactive fiction can do: it can almost say OK, you know you've had enough. You can move on.... We're not gonna overstimulate'.

If the author wants to include choices that have an impact on the game or story, this can involve wanting players to have some understanding of the effectance of their decisions. Having the effects of choices be legible to players is to increase responsiveness, and writing responses takes work. A lot of work creating textual variation or deeper game mechanics could be done without it necessarily being 'visible' (P4, P13), 'transparent' (P8), 'exposed' (P7), or 'surfaced' (P11) to the player. Others mentioned occasions when their game design efforts

weren't sufficiently noticed. P9 mentioned a game feature that 'people didn't pay attention to... so it had less punch than I thought it would'. If the author intends the player to make choices and have their choices reflected back at them in the text, then, as P11 pointed out, if 'there's no indication' that the player is experiencing a 'different branch... then some players will just skim right over it... You might implement all this stuff and they won't even notice it'. Thus, an additional challenge in writing interactively is writing the interaction so that the choices have **Legibility to the Player**.

A common feature of interactive digital narrative is puzzles, especially (but not exclusively) in works in the adventure game tradition. **Puzzle Design** is its own implementation challenge, but the creation of puzzles can also involve writing narrative responses and framing around the puzzles. The various authoring challenges of designing puzzles include judging the difficulty, creating the desired player experience, and iterating new puzzle content. When it came to designing puzzles with the right level of difficulty for the player to solve, P3 admited 'I don't have a good sense for how difficult puzzles are. So it's a little bit easier just to make something where there aren't conceptually complex puzzles'. In contrast, P11 had a distinct vision for how puzzles should feel to the player: each should be 'an event in the narrative', rather than a more abstract experience of 'standing in front of some machine where you have to push all the buttons and manipulate stuff'. Where one set of design lead to greater simplicity, the other (in service of the vision for the game) tended towards more complexity through greater narrative integration.

Authors would also strive for high **Narrative Quality**. Authors expressed a desire not to 'bore the player' (P5), and to be 'engaging and enjoyable' (P14), but they also attempted to afford the player experiences of 'mystery' (P7), 'fun' (P14), 'surprise and delight' (P14), and a 'density of experience' (P2). As P7 expressed, 'graphical fidelity and gameplay mechanics have gotten polished to a point that really narrative is the way you differentiate your IP'.

When there were parallel paths through a narrative, maintaining **Evenness** of authoring effort could be an issue. If there are, say, two paths through a section of an IDN, authors typically want the player to have a satisfying experience regardless of which path is taken. They want the player to 'always be getting something juicy' (P11). A problem here is that as a writer, coming up with one idea for a character in a scene or an outcome of a choice may be hard enough, but coming up with a second or 'especially... a third' of comparative weight and quality is 'inspirationally, more difficult' (P8). Evenness of author effort spent on writing in relation to how likely it was for a player to experience an optional section of the narrative could also be a challenge. 'What felt the hardest was to have was to maintain roughly equal discoverability of each path', P4 maintained, not wanting 'to spend a significant amount of time on a path that only 0.1% of players would do'.

The **Genre** itself often directly shaped the strategies authors could employ to keep content production manageable: a comedy game with less consistency of tone might be easier to write for when the author 'couldn't really come up with a viable choice of two actual serious things' a player might do (P1). A horror game could afford to 'brutally' end stray plot branches so 'only a handful of paths move forward' (P3). While opening up some possibilities for reducing workload, genres can also create new difficulties. For example, 'combinatorial explosion has always been a major issue with detective games' (P4): if the game involves collecting and combining clues to solve a mystery, then the author has created a situation of 'quadratic complexity' if the player can attempt to combine each clue with each other clue.

Player Affordances were added by authors, often based on actual or expected audience feedback, in order to enhance the player experience beyond the narrative or ludic elements. Affordances include features such as an 'automap' and 'mode for screen readers' (P8), 'hint' system (P9, P11), 'giving an alternative way to the solution' (P9), and 'prompts to save your game at different points' (P3). Making an 'effort with accessibility' (P3) would enrich a work, though at the expense of feature creep.

4.4 Release Environment

"Events and competitions are useful in that they bring a sense of community around, they bring deadlines... a lot of the projects I've actually finished have been for some kind of contest or some kind of event." - P2

Interactive narratives aren't created in a vacuum. A work that isn't just a personal creation is made for an intended audience, which may or may not be commercial. The authorial goals are impacted by the feedback from non-authors related to the project: testers, publishers, and other collaborators. This **Audience Feedback** can impact design throughout the creation process, but also between works. In discussions about influencers the concept of 'audience capture' is often found, referring to 'the processes by which influencers are influenced by their audiences' [19]. Creators without their own overriding vision or principles, can allow themselves to increasingly tailor their output to an audience, producing only what they think their audience wants to see. In creative writing, this is sometimes called 'writing for the market' (for instance, [2]), and shouldn't be seen as wholly captured or uncaptured: authors are always in some relationship with the audience expectations and their own creative vision. While participants often received 'good feedback' (P6, P8) and 'messages' (P7) about their work, they weren't necessarily captured by their audience. Often authors would be writing foremost for their own tastes, especially in hobbyist works. As P5 said, 'I work all day at a boring job, nine to five, and I don't do that so that I come home and write for somebody else... I want to write something that I will find cool'. Similarly, the small stakes of writing interaction fiction could be creatively freeing, with P9 saying, 'I don't have a huge audience that I'm going to let down if I mess up'.

The **Publishing Context**— whether the game was made for an online game jam or competition, or for a commercial market (and what that market was) — impacted the author's drive to finish a work, as well as shaping the content and size of the project. One author was motivated to finish a work for a 'book deal' (P1) while another 'enjoy[ed] having the deadlines' (P9) that competitions would bring. For most of the participants, as is the case across other creative fields [13, 33], authoring interactive works wasn't their main source of income. For those for whom it was, choosing to make money this way was most often about the desire to work in the medium, though pecuniary motivations could still play a role. P4 reports beginning a commercial project 'under duress', saying, 'I had a cat that my ex-wife wanted and then to have this we took him to the emergency vet and it was $4000 in debt... and so my main goal was to pay off cat debt and that was my intent on starting [the game]'. In this case, the commercial work was achieved off of a portfolio of successful non-commercial works. This pattern was seen among most of the authors working in the commercial sphere: they were able to transition (sometimes gradually, often incompletely) from the hobbyist sphere, building on craft experience and personal reputation built making works for free. Again, this transition as a process, through portfolio work (working on short-term projects, sometimes simultaneously) is a pattern seen in other creative industries such as music, and the wider videogame industry [1]. For authors working in a completely non-commercial way, the publishing context is quite different. Game jams and competitions are frequently spurs to publish. P2 reflects a common sentiment saying, online 'events and competitions and are useful in that they bring a sense of community around you, they bring deadlines'. The deadlines ensure the game is brought to a completion, and the community ensures an audience, to the extent that most of the hobbyist participants tended to favour releasing in the context of a web-hosted event.

4.5 Strategies

Our largest theme, detailed in Table 3 was a list of strategies used by authors to manage the authorial burden. The design elements of interactive narrative (their underlying structure, and the ways the player can interact with them) aren't chosen solely due to the impact on writing, but rather there is an interplay with the author's goals for the game, and what they have to do to make it happen. Some strategies are implemented on a project-wide scale, others on a scene-by-scene basis. Some appear in the initial design, others, such as cutting-content, are

implemented in an ad hoc way during creation as required. We found that the categories can and do overlap with one another.

4.5.1 Decoupling.
"The content model that we had... allowed scenes to dynamically reshuffle themselves in ways that made sense, and yet did not require us to make as much new content in order for it to work." — P7

Decoupling strategies are those in which content is written that doesn't depend on most of the earlier content happening before it, and so can be written much the same regardless of existing state. These strategies were often implemented on a project-wide scale. Content could be broken up in to distinct chapters with **Episodification**, 'trying to keep [the chapters] separate mostly' (P8).

Works could be written with **Modularisation** where a story segment could 'happen at different points... flexibly' (P13). Authors writing modular interactive narrative experience a tension between trying to engender a sense of agency and expressiveness through the continuity of the story, and managing complexity. One of the challenges with separation is that continuity can require referring to previously established facts in a story: if those facts can be different then full separation of sections is harder. P13 experienced this writing optional romance arcs in one story: 'for some of the romances... they're fairly separated, but they're responsive to the things that are going on and there are a couple of characters who are more tied in with the plot'.

Modular content could be delivered in the form of semi-independent story segments, sometimes described as **Storylets or Plot Clouds** (a metaphor for a segment of plot that moves around like a cloud). These are self-contained scenes that may appear in different orders or not at all depending on the player's traversal. The benefit of this approach P14 explains, 'is that the story is completely unset, but it's kind of waiting to be used. It's not really decided what anything's gonna be until the last possible moment so that we can make sure at that moment we can give you the most interesting thing that we've got'.

4.5.2 Embracing. "I really did almost all the writing in probably four or five very long sittings. Five to ten hours of just sitting down and doing the writing." —P3

Rather than seek to change the work itself, the authorial burden could be embraced (much as Stern encouraged designers to do with combinatorial explosion, one element of authorial burden) [30]. This can be done on the technical side, through **Automation** of rote processes, or **Tool Creation**, making more appropriate languages, engines, generators and the like for the particular project. This ranged from 'script validation' (P7) to a 'custom domain specific language' (P10). These approaches are the sort often advocated in the academic literature. Sometimes tools might be made that didn't end up as useful as first anticipated. P7 reports the team making 'a number of tools... catered towards the game' that they were making, however 'for speed' they ended up 'editing JSON files' directly. Time spent programming automation tools is time not spent authoring, and so the returns on this strategy are by no means certain. P9 experienced this directly after making one tool for automation that generated too much possible content, 'I just procrastinated and crashed out because there was just so many potential things based on my script and then when I went through it all there just wasn't enough' The automated script had created '5050 pages' which was too much to go through and when P9 finally did make a start, there still wasn't enough viable in these pages for their game concept. 'It should have helped me fail forward fast', P9 said, but it didn't.

Existing tools for **Visualisation** or **Knowledge Management** could be employed to better conceptualise the work being done. Often less technical means were used for this, such as 'hand drawn flow charts' (P3). The writing process itself could be adjusted, though authors differed considerably on what they found useful. Some authors reported success with a period of extensive **Planning and Prototyping** beforehand, including 'engineering of the system... from a paper prototype' (P10). During creation, **Writing Out Of Sequence** could be used, producing, for example, 'chunks of the dialogue trees' (P8) in whatever order inspiration struck them. A bold strategy was omitting self-editing and focusing on getting as much done with **No Rewrites**: P14 advised

that 'every idea you have, write it down... Don't editorialize yourself too much. Don't be too fussy about choosing which path you take... write through it.'

The capacity of the authors themselves could be extended. This could mean bringing in more team members with **Collaboration**. This might split the workload, but also help with the complexity overall, with P2 maintaining that 'it is much easier to work in a group than it is to work… for sustained periods on complicated projects alone'. Where commercial pressures such as wages and release windows weren't a factor, simply allowing **More Time** for the project was common: 'the difference between a personal project and a studio project,' P10 described, was that with a 'passion project... if you're blowing your scope, you can say this is going to be a five year project instead of a two year project'.

Authors could attempt to induce a greater experience of **Flow** to write for longer with fewer distractions. When that experience of flow was missing, the results could be very difficult to bear for authors. P4's experience with one long commercial project was of 'spending four hours where three and a half hours was overcoming self loathing and thirty minutes is writing'. Flow can be interrupted in multiple ways. Sometimes authors are more attuned to one kind of work than another at any given moment: 'when I sit down my brain can be in a coding or a writing mode and it can be hard to switch' (P5). This can lead to discontinuities of flow: 'if I sit down thinking I'm going to write a bunch of descriptions and then I realize, oh wait, there's a bug in this thing... It feels like my brain struggles to change gears'. Sometimes the sheer size and complexity of a project can arrest flow, 'this thing is so massive, I just have to concentrate on it' (P8).

4.5.3 Generating.

"ProcGen... makes for very good snippets of narrative, but not for very good overarching narrative" - P2

There is a common belief among non-authors that content could be generated very quickly and cheaply from Artificial Intelligence or **Procedural Content** generation more broadly (e.g. in Dobrovsky et al [9]). However, the participants didn't typically conceptualise generation as a means to save authoring time: 'sometimes this is phrased as [procedural generation] gives you twice the content for only five times the work' (P2). Authors often had a more nuanced view of the power and expressiveness of procedural approaches. This was particularly explored by P14, who said, 'I don't believe in procedural generation of content. I believe in procedural assembly of authored content. So you have a pool of interesting things which have been written by human because humans are interesting and computers aren't. And the computer is just gluing them, so we can maximize the usage of all the content that we write.'

A different generative approach was mentioned by P10: using **Story Sifting** techniques to create narratives out of generated content or emergent gameplay. 'The idea is that you have a series of events that are presented to the system... And what the system does is say from here, which sequences am I interested in? Which one of those are good stories?' (P10)

Many authors are at least 'a little skeptical' (P7) about the use cases for procedural generation in interactive digital narrative. This was especially the case among those with experience creating things with it. P2 had worked extensively with text generation and had several insights into its limitations. 'One of the big problems with ProcGen,' P2 maintained 'which continues to be a problem is that it makes very good snippets of narrative but not for very good overarching narrative'. While P2 admits that some of the 'interesting things' within procedural generation include ways to try 'get around the problem', procedural generation in general is considered 'very much not a complete story, it is hints of story. Its narrative hooks'. A narrative hook is the start of a plot, or a prompt for an author to develop one, and these, P2 says are 'quite easy for a machine to develop... and quite difficult for it to explore fully'. This was P14's experience as well, finding that when authoring procedurally, the content was 'repetitive and not very interesting and a missed opportunity' so in the end its use was mainly relegated to minor repeating interstitial scenes as 'it wasn't generally very useful' as the author 'spent an enormous amount of time fighting with a system just to make it stop being dumb'.

At worse, this shallow procedural creation of narrative becomes readily apparent to the player. As P5 says, it 'probably feels like walking around like a Wild West set at universal, where you know that it's all just a façade'. This becomes a problem 'when you get shown a piece of text and you can see where the placeholders are in it. You know you can see the bit that is essentially saying "You are in a blank; there is a blank here; you can blank or blank".' This obviousness of the template is the weakness of shallow implementation of procedural generation. As P5 continues, when 'you see that ten times and it's only mildly better than seeing literally the same piece of text ten times, I feel like for the same effort you could have actually just crafted ten different things and I would be exactly as as bored with seeing it... I would sooner take the advantage of text and and craft like 100 short sentences that are very different than procedurally generate 1000 words, but it will feel kind of like a a robot wrote it'.

To make the generated content seem non-robotic, the designer must use a considerable amount of discernment. As P2 jokes, 'the joy of procedural generation is that it gives you twice the content for only twice the work... sometimes this is phrased as it gives you twice the content for only five times the work'. Behind this framing is a serious point about craft goals, P2 continues, 'You can save some kinds of work and introduce some kinds of variety, the question is whether those kinds of variety are valuable to you'. It's an expressive medium in itself with its own craft considerations, not primarily a means to reduce labour.

P10 explored two further issues with procedural generation as a strategy. One finding was that it could require a lot of content upfront before it seemed interesting: 'you can build it out so that it's super unimpressive until it's really impressive or you can have incremental gains'. Another problem was that it was hard to guarantee quality. P10 says, 'this is also what I found in industry, because... game studios want guarantees. You do procedural narrative stuff and they want to know that it will always be good and the more exciting and procedural you make things, the harder it is to say "Nah man, don't worry. It's all good"'. This could particularly be a problem if working on an Intellectual Property (IP) where story possibilities have to be signed off by the IP holder. P10 says, 'if we had been [working on] some wild IP like Batman... and then the IP holder for Batman was like, "Send us over the story so we can approve it"', they would be forced to reply to the effect of, "'I can't do that. Here's a million different ways a story could play out"'. It wouldn't 'humanly possible' for anyone to see all the possibilities.

Considering the power of using a computer to select content and the various drawbacks presented here, P14 offers a way forward that had worked on several projects: 'Procedural generation is largely a waste of time and is something that people who don't know how to write want because they don't know how to write... There's one good thing about it though, which I really like, which is the ability to bring content to bear on the player. What procedural stuff lets us do is say wherever you are, whenever you are, I'm going to give you the most interesting thing that we have in our collection of things. And I'm going to make sure that it's relevant to you'. As a strategy to distribute narrative chunks, these kinds of approaches leverage the power of digital systems, while maintaining the desired richness of hand-authored content.

4.5.4 Reducing.

"Can I actually justify to myself that I'm subjecting the player to this much text?... I can probably say the same funny thing in one sentence and just delete everything before... And then of course, in doing that, you delete a bunch of nouns and there's a bunch of other stuff that you don't need to implement and just by saying less you make less to implement." - P5

Perhaps the most straightforward way of limiting the authorial burden in a project after it has commenced is just to limit ambitions through **Cutting** content. In the initial design and beyond, the amount of varying content can be limited by **Reducing Objects**, the 'extraneous things' (P13) such as items, places, or characters. It can be limited by **Reducing Actions**, limiting the choices, verbs or other interaction methods available to the players. Implementing actions, especially where they create combinatorial complexity, 'can cause a lot of headaches' (P11)

which can be ameliorated by limited the range of ways the player can interact with the work. Complexity can be reduced through greater decoupling of content segments by **Reducing Tracking** of game-state, limiting the number of 'eternal flags' (P8) being checked. Overall workload can be reduced by limiting through **Reducing Play Time**. This latter approach was explored by P2: 'most of the work was about packing in an awful lot of experience into a very small time and that takes a lot of work and you can't have that kind of density of experience over a multi-hour project'.

The medium of the work itself can allow reduction of both static and dynamic authoring. **Abstraction**, such as cutting elements extraneous to the narrative that might add to the realism of the situation but would be more work to implement such as checking 'who's in the room when you have a conversation' (P7), or adding game-like elements that don't necessarily make sense in-fiction, including arbitrary limitations on action. Interactive narratives that are primarily based in text can leverage the benefits of using **Words vs Assets**, in that the author can declare things about the world much quicker and cheaply in text: 'you can write a huge game and you're not having to worry about making the 3D models or finding all the voice actors... because you're just writing sentences' (P5).

4.5.5 Reusing.

"I had been keeping a collection of different story ideas and it just so happened that I figured out a way to put several together." —P3

Within a work, content can be reused. This can be through using **Shared Assets**, repurposing the same creation or 'programming structure' (P8) in different contexts. Another common way is using **Structural Patterns** such as 'branch and bottleneck' (P1), or using game loops, so that the same content can appear again (this can overlap with Moduluarisation and other Decoupling strategies, though content can be Decoupled without Reusing elements, and content can be Reused without necessarily Decoupling sections).

The author can use a **World Model**— 'a sensible default with interesting overrides' (P14). Having a set of rules for default behaviour allows locations, objects, and characters that can reappear (this is a core feature of text adventures). States can be tracked as variables that accumulate, such as relationship 'affinity scores' (P3), or in-game resources like money, reusing the same **Cumulative Variables** that can be changed in multiple places in the work. This can allow players to make choices that 'would add to stats' allowing agential play without immediately 'branching the plot' (P12).

Ideas and content can be reused between works, **Recycling Work** between projects. Institutional knowledge can be built up in a studio or publisher in sets of guidance on **House Style**, thereby 'institutionalizing design knowledge' (P10).

5 Conceptual Model of the Authorial Burden

Figure 1 illustrates how we have brought the insights above together into a conceptual model of the authorial burden. Authors have a vision for the games they want to create which determines their **Authorial Goals**. These goals are shaped by their **Capabilities** (their skills in writing, the tools they know, their imagination etc.), and the **Audience and Publishing Context** (including competitions, game jams, commercial launches, publishers). Often the vision is incomplete, and they write to find out more!

The design emerging from these goals sets the initial **Authorial Burden** (which may or may not be known before they start writing), the relative size of which is also influenced by the author's capability (a given project would be easier for an experienced author with relevant skills than an inexperienced one without those skills). There are three different kinds of work involved:

(1) **Static Authoring**: creating more unchanging content. (e.g. writing lots of descriptions)

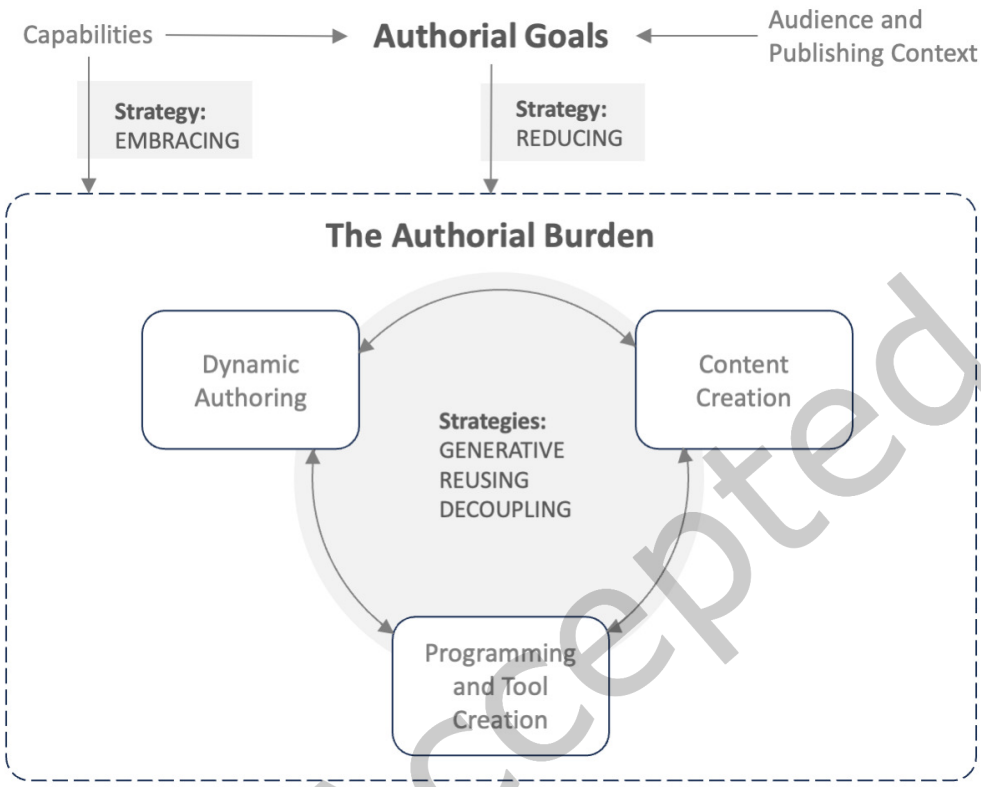

Fig. 1. A conceptual Model of the Authorial Burden

(2) **Dynamic Authoring**: creating content which might appear based on some conditions, or appear in different contexts. (e.g. writing lists of objects and descriptions that can be recombined to create many different objects)
(3) **Programming/Tool Creation**: expanding system capacities (e.g. creating a new game engine or scripting language to do something that couldn't easily be done in existing tools)

As part of an authors' initial design, or later when they realise there's too much work, they implement strategies for managing the burden. These strategies can:

(1) REDUCE the overall workload (e.g. cutting content) effectively adjusting the goals
(2) EMBRACE the project as it is, and try to increase capacity instead (more writers, more time taken, crunch, improving skills)
(3) REUSE elements (writing, code, assets)
(4) DECOUPLE content, separating it into self-contained sections to manage complexity
(5) GENERATE content (e.g. through procedural generation)

The first two strategies impact the relative size of the burden by either adjusting the goals (REDUCE) or increasing the capability of the author (EMBRACE).

The last three strategies (GENERATE, REUSE, DECOUPLE) do not necessarily change the size of the burden but instead primarily change the nature of the work. For example, a *description generator* might reduce the amount of unique descriptions you need to write (Content Creation is reduced), but it is still work to make the generator to begin with (Tool Creation is increased), and to populate it with enough content so that it feels varied and interesting enough (Dynamic Authoring is increased).

6 Validation with Experts

As described in Section 3 we presented the model back to a different set of experts through three panels in order to validate the model. In this section we will explore the outcomes of those discussions and present the implications of the conceptual model on how the authorial burden, and scope issues more broadly, are talked about in the context of interactive narrative; look at questions and caveats raised by the experts; and sketch out potential avenues for future research in this area.

6.1 Implications

Rarely can creative work be sped up without changing the process of creation and the form that the output takes, and often alleviating one aspect left more time for other work rather than reducing labour overall. As one panelist put it, after interrogating the design and removing content, 'what I found was I didn't write less as a result... I was picking what I wanted to write about'. Our model of the authorial burden emphasises that every proposed strategy to alleviate authorial burden changes the kind of creative labour that is performed (and thereby not always reducing work absolutely), and that these changes may not align with author visions, interests or skill-sets, or be appropriate for the intended audience of the work. This was repeatedly born out in the interviews and confirmed in the expert panels: often authors will use a strategy to move work into a domain they find more interesting or are more comfortable creating within. For instance, game designers who are also tool-makers are often seen spending time programming a specific engine to meet their requirements, rather than trying to amend their vision to fit with an existing tool, even though that would usually be quicker. Conversely, less technically-minded authors will often stay with a tool long after it has ceased meeting all their requirements, due to the high opportunity cost of switching to something better suited or attempting to make a tool themselves.

Authorial 'burden' is often not experienced as burdensome, and one common feedback from the panelists was about the negative connotations of the term itself. Creative labour is the part that's enjoyable for the author, and there is little desire to abstract it away. As one panelist said, 'sometimes the strategy doesn't reduce the workload, it frames it in a way that you actually enjoy doing'. While commercial games are made under context of time and financial constraints, hobbyists are more likely to approach the craft in a similar way to many novelists: as an enjoyable but not particularly profitable past-time.

One panelist who had worked with machine-generated prose pointed out that it was a lot of work to get it to output desired content, and the content created needed a lot of editing to 'force disparate content to make sense', but in taking a 'bottom up' approach, it could be used as a source of creative inspiration. Others were even less enthusiastic: 'It's never "how I'm using it in my own work" and always "here's how it's going to disrupt that other industry that I know nothing about" ', one said while another reflected that 'there's no part of [my writing] where I'm like, "I wish an AI would do this for me". If I didn't find it interesting I just wouldn't do it.'

The aesthetic texture of hand-crafted content is very different to the 'procedural oatmeal' that can arise through machine generation. So even when the designers do want to avoid authoring work (e.g., to meet commercial constraints), the techniques touted to alleviate authorial burden often typically would not meet the desired quality and specificity of output. Where generative approaches are sought by authors, they are picked for their specific affordances. One participant stressed the distinctiveness of generation as its own artistic medium: 'you are not

going to save work overall if you want the same kind of product', but, 'you can save some kinds of work and introduce some kinds of variety: the question is whether those kinds of variety are valuable to you.'

This same dynamic is present in non-generative strategies as well. For example, the use of illusory agency has been proposed as a way to avoid authorial burden by offering choices but always redirecting to the same branch regardless [11]. While merging branches has long been recognised as a fundamental technique of hypertext authoring [3], committing to only or primarily using illusory agency in a game that is positioned to players as being high agency is in itself a specific vision for a work, and one which would become obvious on replay or review of that work: the illusion cannot hold indefinitely. Authors make games that according to a vision, and that vision is rarely 'pretend to be some other game-type entirely'. If having high player agency is the design goal that is creating the authorial burden, removing that player agency doesn't solve the burden, rather it creates some other game entirely.

6.2 Challenges

During the panel some explanatory challenges arose for the model. The 5-part schema of strategies [17] was a clarifying lens to view authorial design choices but it elides the temporal aspects of design. For instance, one expert in the first panel discussed how she would 'collapse continuities' in her hypertext fiction as part of an iterative process of shrinking the initial workplan by merging different split branches. While this is both a reducing and reusing strategy (specifically, cutting some proposed branches by way of reusing others), it is also a strategy for a specific time in the design process: when an initial hypertextual pattern has been planned but has yet to be fully implemented.

Another way of splitting the strategies would be to consider which take place at which stage of authoring, and the scope of each strategy, from the level of a single scene or room, all the way to game-wide changes. So many of the strategies for managing the authorial burden happen at the outset of the project, with expectations of how much work it will be shaping the initial vision for a game. In Pinchbeck's examination of the design of Dear Esther, he remarks that the whole setting and structure of the game was made with 'the economics of building' in mind: setting the game on a remote island was explicitly a strategy to diegetically limit objects to implement, especially other people [26]. As per Kitromili [20], authoring interactive narrative is made up of different activities from ideation and planning to writing and editing, and strategies for managing work are appropriate to different stages.

One panelist offered the challenge: instead of the 'authorial burden', what about the 'authorial responsibility'? 'Burden' has the connotation of something that must be shouldered or shirked. As a term, it leaves something to be desired. Still, the term does leave the correct impression that authorial choices each leave a weight of work to be done, and that the structure of interactive narratives, especially those with plot branches or variable narrative combinations, itself creates more of this work, whether or not it is felt as a burden.

Another panelist felt that author work in presentation, in user interface, and in marketing, is its own kind of labour, which has concomitant strategies that may or may not work (he described snake-oil salesmen promising '500 clicks with this one tool'). Indeed, care should be taken not to present 'authorial burden' as a synonym for all authorial labour, including the creation of paratextual elements, but rather as a description of the work that arises specifically out of the interactive nature of the story. The author might take care over how their hypertext is presented to a competition, but that care is typically independent of the labour that was needed to flesh out the different plot branches they decided to implement.

Finally, a panelist who was foremost a game engine designer and secondarily a game author, was keen to highlight the specific burdens in tool creation. Even before someone makes a game with a tool, the designer-programmer must wrestle with the work that emerges out of their programming design decisions. While this programming labour can be an activity undertaken by some authors, and many of the work management strategies

are similar (such as modularisation or loose-coupling), an author of a tool for making games has a different kind of aim than the creator of a game. A hypertext system or text adventure engine might embody some assumptions about possible stories (especially if they implement a world-model) but it isn't itself a story. The combination of variable state management and the management of narrative continuity is what makes interactive narrative authoring so thorny. They have all the problems that face the programmer and the novelist as well as the unique problems from arise from the intersection of the different crafts.

7 Discussion

In the process of developing the themes out of nearly 20 hours of interviews and panels, a range of core tensions emerged between the strategies discussed, and the authors conceptualisation of their craft and the claims commonly made about tools that could reduce their authorial labour.

7.1 Craft Knowledge

Some of the strategies were approaches a person could select from a list and attempt to implement in a project. **Cutting** content is an obvious strategy along these lines that, while requiring some discernment to use effectively, is one that most authors would naturally use in the course of things regardless of their craft background. Other strategies were more like cultivated practices, or authorial virtues: ways of living and working that enabled small continuous gains. An example of this is the practice of **Recycling Work**: it is not usually feasible for an author to retroactively recycle previous ideas, code snippets, orphaned lines of dialogue and the like, unless the author has cultivated a habit of keeping their notes and saving these things. Similarly, an author might seek **Collaboration**, but there are skills to being an effective collaborator. In this way, we can distinguish between short-term strategies that one can seek to apply and the long-term strategy of growing capabilities. The growing of authorial capabilities is the active development of craft knowledge. In the context of writing, [25] describe the craft knowledge of writers as "not simply a body of content knowledge but substantive knowledge of how to craft and shape a text". It is operational knowledge that can be used in service of creative goals. Kitromili characterises this craft knowledge as interactive narrative "authoring activities" — such as writing, ideation, and structuring [20]. These authoring activities are things that authors do, of course, but also things that the author *learns how to do*. Authors have various capabilities including mastery of specific tools, and practical design knowledge of structural patterns. Authoring interactively involves using a varied set of skills that can be cultivated. Spierling argues, access and education for authors is a way to encourage authoring overall, and being able to think and write interactively is a skill that writers of other media don't automatically possess [29]. The more acquainted authors are with common interactive patterns, and familiar they are with their tools, the more ambitious projects they will be able to complete. While **House Style** is a way that companies attempt to teach authors institutionally valuable craft knowledge, and invariably authors learn in the process of making, specifically developing design know-how remains a perhaps under-explored **Embracing** strategy.

7.2 Tool Mastery vs Tool Creation

There are many interactive narrative tools continually being developed, but tool creation as a strategy is in tension with tool mastery. Creating bespoke engines is programming labour which can have a huge time cost, and so using existing narrative engines, programming languages, visualising tools etc. can save a lot of possible work, with the compromise that what can be made will be limited to what is possible within those tools. Many of the participant authors had strong preferences for tools they were already familiar with. Tool mastery is an aspect of their growth as authors, and deep knowledge of a tool can be hampered by continual switching: often authors will use a deeply unsuited tool that they are already familiar with rather than pay the opportunity cost of learning something else. As such, extending existing tools may sometimes be a more promising strategy for

improving the capacity of existing interactive narrative authors than the creation of entirely new tools: there may be a compromise point between building a better thing, and building a thing that will be used.

7.3 Conceptual Distinctions

The **Reusing** and **Decoupling** themes have considerable overlap. Often in practice, in modular design, and when sharing assets, or procedurally selecting content, the sections are decoupled precisely so that they can be reused, recombining in multiple possible places. However, there is a distinction. A work made up entirely of disconnected vignettes would be fully episodic, its sections decoupled from one another, without any content being reused. Conversely, a loop design structure reuses story segments without necessarily decoupling them from the narrative.

Another possible conflation was in the similarities of storylets with strategies of modularisation. A storylet is a modular form of content: a more-or-less self-contained section of story that can be deployed on a modular basis. However, there are other ways to design with modules that don't involve storylets at all (for example, game levels), and storylets can be deployed entirely calligraphically, placed in a fixed, interlocked sequence of passages. So again, while in practice there is a great deal of overlap in design thinking with these two strategies, they're not always co-extensive.

7.4 Supplementary and Incompatible Strategies

Some of the strategies, on their own, do not directly contribute to a work being authored. Indeed, some forms of tool creation can actively detract from the time one has available to actually write. Other strategies such as visualisation, planning and protoyping, or knowledge management might help an author understand their work and aid in testing, but no amount of looking at the work will directly translate into more scenes written without writing those scenes. These strategies can help gains through a project, but they're supplementary to the authoring. To put it another way, visualising may show an author whether their project is massively out-of-scope, but it isn't sufficient on its own to do anything about the problem. Several of the strategies have clearly incompatible approaches. For example, the process of cutting excessive content is in conflict with the approach of avoiding rewriting wherever possible. Indeed, these approaches appear to be set up for two very different kinds of game experience: the short-form, intense, closely written work, in contrast to the long-form, extensive, more loosely written epic. The approach that is appropriate for authors to take depends on what is congruent with their goals. Some strategies flow directly from certain authorial goals and some would be in conflict. If an author wants to foster a sense of agency in the player, a feeling that their choices are impacting the story, then excessive episodification or reduction might conflict with that goal. Conversely, if a player is creating an episodic story, then the fewer states they have to track, the more distinctly episodic they can make the work. In this way, some strategies might be said to synergise.

7.5 Analytical Claims of the Model

In the background section above, four papers were identified which make claims that some design or tool approach can relieve the authorial burden. Using the conceptual model, we can now properly assess these claims. Almost no one has to author for interactive narrative, it isn't a hardship imposed on people, but rather a skilled practice that authors enjoy doing. Simply put, the authorial burden isn't always experienced as burdensome, and strategies that aim to remove the author's skilled presence in the process are likely to be resisted by those very same authors. **Generating strategies** especially can be fraught for exactly this reason. When a new form of generation is suggested that promises to relieve the authorial burden, we should ask what form of labour it leaves to the author and whether this is something they are asking for.

7.5.1 Considering Illusion of Agency. Fendt et al. argue that the 'illusion of agency' can be used as a 'step in reducing the authorial burden of games' [11]. Their claim is that players report a similar sense of agency when their choices are reflected back at them in a linear game story, as when they really do make choices which branch the plot. They ran an experiment, seeking to test this claim with two groups of players experiencing very similar games, one with 'illusionary' branching choices (choices that merged the plot immediately after being made), and another with the same choices leading to alternative plot branches. Their finding was that reported agency was similar between both groups. They only ran the experiment with the participants once, and so they didn't record how players would feel about their sense of agency once they know that their choices had no impact on the wider plot. They conclude that it is immediate feedback to the players that allows them to experience a sense of agency. Applying the conceptual model of the authorial burden here, illusory agency is a **Reducing** strategy, and one which may or may not be aligned with any given author's **Authorial Goals**. If an author is interested in facilitating agential play, they probably want the agency to be non-illusory. While some authors report being pleased when players read more into the story than has been implemented, most authors of branching narrative really do want to write and offer players multiple branches. However, if an author wishes for their work to be a replayable *and* agential experience, then clearly an illusory approach will be insufficient, at not least if used predominantly throughout a work: the player will quickly see on replay that their choices have no wider impact. Similarly, if the game is intended to be a longer experience than Fendt et al.'s very short experiment games [11], the lack of story continuity will quickly become apparent to players unless there is very heavy **Episodification**. The model of the authorial burden leads us to consider how appropriate a strategy is in light of the author's goals (for agency, replayability, length, genre etc.), giving us reason to pause before accepting claims that an approach will reduce labour.

7.5.2 Considering Combinatorial Dialogue Authoring. In their paper *Combinatorial Dialogue Authoring*, Ryan et al. present an experimental approach for annotating snippets of dialogue which then can be used to generate conversations [28]. Authors must not only create dialogue but also annotate it, allowing the dialogue elements to be greatly recombined. Applying the conceptual model of the authorial burden, this is a strategy that shifts the burden of work from static to dynamic authoring, situating the author in a role of writer, annotator, and polisher of 'small nuances lost on the computer'. If authors prefer this form of work, and if their vision for their work involves a great deal of dialogue, then indeed this approach could be promising for those kinds of projects. However, as a general approach it might not find broad appeal if increasing authorial leverage eliminates desired qualities in the interactive narrative.

7.5.3 Considering Natural Language Generation. In their next paper, *Toward Natural Language Generation by Humans* [27], Ryan et al make a point consonant with the discussion above, that 'computer-generated dialogue, here and now, is so impoverished relative to the human-authored equivalent that its huge alleviation of authorial burden is not worth the attendant drop in content quality'. To avoid this issue, but still gain some of the massive generative affordances of natural language generation (NLG), they develop the preliminary design of an authoring tool for conversation variation, using probabilistic context-free grammar. As with text variation in Inform or Twine, or context-free grammar tools like Tracery, the author has control over the range of possibilities, while allowing considerable generative power. For authors this shifts the burden of work from static authoring of dialogue to dynamic production of grammars. This form of approach could appeal to authors who already use textual variation (as several in the study did). In alleviating one form of the burden, it changes the nature of the work to a form of labour that some authors would find interesting and worthwhile while others may find onerous.

7.5.4 Considering Generative Artificial Intelligence. In contrast to the skepticism the interactive narrative authors in the study expressed towards generative artificial intelligence approaches to authoring (GenAI), Fisher takes the

position that GenAI is currently, or will soon, reduce the authorial burden in interactive narrative [12]. Applying the conceptual model of the authorial burden to this claim we have reason to pause. There is no grounds to suppose that current GenAI techniques produce content approaching the narrative quality and consistency of hand authored material, as Ryan et al. in the preceding section recognise clearly [27], even despite the massive improvements the technology has seen in recent years. But moreover, even if the narrative quality was equivalent to a skilled human author, writing a prompt isn't the same as authoring content, or even constructing bespoke dynamic systems for generating or delivering authored content. This kind of GenAI strategy shifts authorial labour from all forms of authoring, to the generation and curation of generated material that is largely outside of the author's fine control. That is to say, it removes much of the authoring from the process. While some people may prefer the labour of designing prompts to induce a GenAI to output something useful, this doesn't appear to be something desired by many already-existing interactive narrative authors who have more discerning authorial goals, and are often willing to develop the capabilities —the craft knowledge— to achieve these goals.

8 Conclusion

In the field of interactive digital narrative, alleviating an authorial burden is the implicit motivation (and often explicit motivation) of much tool development, but the concept of the authorial burden itself was undertheorised. This paper develops the conceptual model of the authorial burden. Crucially, it recognises that not all authoring labour is experienced as burdensome. Going beyond 'authorial burden', we should consider that much authorial labour is something authors choose to partake in as a preferable activity. This model was developed out of lengthy, rich, qualitative interviews with authors themselves (with wide experience of a range of interactive narrative sub-types). One of the core insights of the model is that the strategies authors use to manage labour do not necessarily reduce it overall, but rather move it between forms that might better match their authorial goals and skills. Researchers proposing new tools and design approaches should be mindful about whether or not they are seeking to shift the burden of work away from the areas authors actually find creatively fulfilling. Strategies for managing authorial labour often change the nature of what is being produced, and so in the creation of any given interactive narrative, which strategies are appropriate will vary.

The model of the authorial burden describes and situates various kinds of strategies that authors use. This framework allows better understanding of whether a design strategy reduces labour overall, or shifts it to a different form of burden, and whether it does so by compromising the initial vision of the work. Furthermore, categories that REDUCE, EMBRACE, REUSE, DECOUPLE or GENERATE are themselves theoretically generative: designers can use these as spurs to identifying novel solutions to thorny authorial burden problems on any given project.

One limitation of the study is that although many specific strategies were identified, many more could be described through discussions with even more authors. The study itself was focused on the creation of works with a strong authorial vision, usually written by one or two people, or at most a small team. Approaches and constraints are likely to be different for larger projects, although it should still be possible to conceptualise the strategies used by big studios within the model.

Engaging with the craft knowledge that interactive narrative authors possess is invaluable to understanding the process of authoring. Our model of the authorial burden helps to capture some of this craft knowledge and provides a theoretical foundation for future work on IDN authoring. We also hope that it inspires others to undertake field-work that engages with authors, ultimately leading to a better understanding of the real challenges that authors face in their creative practice.

References

[1] Daniel Ashton. 2011. PATHWAYS TO CREATIVITY: Self-learning and customising in/for the creative economy. *Journal of Cultural Economy* 4, 2 (May 2011), 189–203. doi:10.1080/17530350.2010.506325

[2] David Baker. 2010. *On Demand: Writing for the Market in Early Modern England.* Stanford University Press, Stanford.

[3] Mark Bernstein. 1998. Patterns of hypertext. In *Proceedings of the ninth ACM conference on Hypertext and hypermedia : links, objects, time and space—structure in hypermedia systems: links, objects, time and space—structure in hypermedia systems (HYPERTEXT '98).* Association for Computing Machinery, New York, NY, USA, 21–29. doi:10.1145/276627.276630

[4] Alexander Bogner and Wolfgang Menz. 2009. The Theory-Generating Expert Interview: Epistemological Interest, Forms of Knowledge, Interaction. In *Interviewing Experts*, Alexander Bogner, Beate Littig, and Wolfgang Menz (Eds.). Palgrave Macmillan UK, London, 43–80. doi:10.1057/9780230244276_3

[5] Amy Bruckman. 1990. The Combinatorics of Storytelling: "Mystery Train Interactive". *The MIT Laboratory* (April 1990).

[6] John Bryant. 2002. *The Fluid Text: A Theory of Revision and Editing for Book and Screen.* University of Michigan Press. Google-Books-ID: 1w4wpOdPbu4C.

[7] JR Carpenter. 2017. In the event of a variable text. *Convergence* 23, 1 (Feb. 2017), 98–114. doi:10.1177/1354856516678357 Publisher: SAGE Publications Ltd.

[8] Chris Crawford. 2020. Fundamentals of Interactivity | Interactive Storytelling Tools for Writers | Chris Crawford. http://www.erasmatazz.com/library/the-journal-of-computer/jcgd-volume-7/fundamentals-of-interactivi.html

[9] Aline Dobrovsky, Uwe M. Borghoff, and Marko Hofmann. 2017. Applying and Augmenting Deep Reinforcement Learning in Serious Games through Interaction. *Periodica Polytechnica Electrical Engineering and Computer Science* 61, 2 (May 2017), 198–208. doi:10.3311/PPee.10313

[10] Ignacio X. Domínguez, Rogelio E. Cardona-Rivera, James K. Vance, and David L. Roberts. 2016. The Mimesis Effect: The Effect of Roles on Player Choice in Interactive Narrative Role-Playing Games. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16).* Association for Computing Machinery, New York, NY, USA, 3438–3449. doi:10.1145/2858036.2858141

[11] Matthew William Fendt, Brent Harrison, Stephen G. Ware, Rogelio E. Cardona-Rivera, and David L. Roberts. 2012. Achieving the Illusion of Agency. In *Interactive Storytelling*, David Oyarzun, Federico Peinado, R. Michael Young, Ane Elizalde, and Gonzalo Méndez (Eds.). Springer, Berlin, Heidelberg, 114–125. doi:10.1007/978-3-642-34851-8_11

[12] Joshua A. Fisher. 2023. Centering the Human: Digital Humanism and the Practice of Using Generative AI in the Authoring of Interactive Digital Narratives. In *Interactive Storytelling*, Lissa Holloway-Attaway and John T. Murray (Eds.). Springer Nature Switzerland, Cham, 73–88. doi:10.1007/978-3-031-47655-6_5

[13] Henrik Fürst. 2023. Continuing in a creative career: Claiming an artistic identity and aligning trajectories among early career novelists. *Poetics* 100 (Oct. 2023), 101818. doi:10.1016/j.poetic.2023.101818

[14] Jacob Garbe. 2020. *Increasing Authorial Leverage in Generative Narrative Systems.* Ph. D. Dissertation. UC Santa Cruz.

[15] Jacob Garbe, Max Kreminski, Ben Samuel, Noah Wardrip-Fruin, and Michael Mateas. 2019. StoryAssembler: an engine for generating dynamic choice-driven narratives. In *14th International Conference on the Foundations of Digital Games (FDG '19).* Association for Computing Machinery, New York, NY, USA, 1–10. doi:10.1145/3337722.3337732

[16] S.E. Hove and B. Anda. 2005. Experiences from Conducting Semi-structured Interviews in Empirical Software Engineering Research. In *11th IEEE International Software Metrics Symposium (METRICS'05).* IEEE, Como, Italy, 23–23. doi:10.1109/METRICS.2005.24

[17] Joey Donald Jones. 2022. Authorial Burden. In *The Authoring Problem: Challenges in Supporting Authoring for Interactive Digital Narratives*, Charlie Hargood, David E. Millard, Alex Mitchell, and Ulrike Spierling (Eds.). Springer International Publishing, Cham, 47–63. doi:10.1007/978-3-031-05214-9_4

[18] Joey Donald Jones and David Millard. 2024. Experiencing The Authorial Burden. In *Proceedings of the 35th ACM Conference on Hypertext and Social Media (HT '24).* Association for Computing Machinery, New York, NY, USA, 78–87. doi:10.1145/3648188.3675134

[19] Daniel Jurg, Marc Tuters, and Ike Picone. 2024. "Alex, DO NOT BACKPEDAL ON SANDY HOOK!": Reactionary Fandom, Cancel Culture, and the Possibility of 'Audience Capture' on YouTube. *Television & New Media* (Sept. 2024), 15274764241277473. doi:10.1177/15274764241277473 Publisher: SAGE Publications.

[20] Sofia Kitromili, James Jordan, and David E. Millard. 2020. What Authors Think about Hypertext Authoring. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media (HT '20).* Association for Computing Machinery, New York, NY, USA, 9–16. https://doi.org/10.1145/3372923.3404798

[21] Hartmut Koenitz. 2015. Towards a Specific Theory of Interactive Digital Narrative. In *Interactive Digital Narrative.* Routledge. Num Pages: 15.

[22] Hartmut Koenitz, Mirjam Palosaari Eladhari, Sandy Louchart, and Frank Nack. 2020. INDCOR white paper 1: A shared vocabulary for IDN (Interactive Digital Narratives). doi:10.48550/ARXIV.2010.10135 Version Number: 2.

[23] David E. Millard and Charlie Hargood. 2021. Hypertext as a Lens into Interactive Digital Narrative. In *Interactive Storytelling: 14th International Conference on Interactive Digital Storytelling, ICIDS 2021, Tallinn, Estonia, December 7–10, 2021, Proceedings.* Springer-Verlag, Berlin, Heidelberg, 509–524. doi:10.1007/978-3-030-92300-6_51

[24] Janet H. Murray. 1998. *Hamlet on the Holodeck: The Future of Narrative in Cyberspace.* MIT Press, Cambridge, MA, USA.

[25] Debra Myhill, Teresa Cremin, and Lucy Oliver. 2021. Writing as a craft: Re-considering teacher subject content knowledge for teaching writing. *Research Papers in Education* (Sept. 2021), 1–23. doi:10.1080/02671522.2021.1977376

[26] Dan Pinchbeck. 2008. Dear Esther: An Interactive Ghost Story Built Using the Source Engine. In *Interactive Storytelling (Lecture Notes in Computer Science)*, Ulrike Spierling and Nicolas Szilas (Eds.). Springer, Berlin, Heidelberg, 51–54. doi:10.1007/978-3-540-89454-4_9

[27] James Ryan, Andrew Fisher, Taylor Owen-Milner, Michael Mateas, and Noah Wardrip-Fruin. 2021. Toward Natural Language Generation by Humans. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 11, 4 (June 2021), 53–56. doi:10.1609/aiide.v11i4.12840

[28] James Owen Ryan, Casey Barackman, Nicholas Kontje, Taylor Owen-Milner, Marilyn A. Walker, Michael Mateas, and Noah Wardrip-Fruin. 2014. Combinatorial Dialogue Authoring. In *Interactive Storytelling (Lecture Notes in Computer Science)*, Alex Mitchell, Clara Fernández-Vara, and David Thue (Eds.). Springer International Publishing, Cham, 13–24. doi:10.1007/978-3-319-12337-0_2

[29] Ulrike Spierling and Nicolas Szilas. 2009. Authoring Issues beyond Tools. In *Interactive Storytelling (Lecture Notes in Computer Science)*, Ido A. Iurgel, Nelson Zagalo, and Paolo Petta (Eds.). Springer, Berlin, Heidelberg, 50–61.

[30] Andrew Stern. 2008. Embracing the Combinatorial Explosion: A Brief Prescription for Interactive Story R&D. In *Interactive Storytelling (Lecture Notes in Computer Science)*, Ulrike Spierling and Nicolas Szilas (Eds.). Springer, Berlin, Heidelberg, 1–5. doi:10.1007/978-3-540-89454-4_1

[31] Yuqian Sun, Phoebe J. Wang, John Joon Young Chung, Melissa Roemmele, Taewook Kim, and Max Kreminski. 2025. Drama Llama: An LLM-Powered Storylets Framework for Authorable Responsiveness in Interactive Narrative. doi:10.48550/ARXIV.2501.09099 Version Number: 1.

[32] Christine Talbot and G. Michael Youngblood. 2012. Spatial Cues in Hamlet. In *Intelligent Virtual Agents (Lecture Notes in Computer Science)*, Yukiko Nakano, Michael Neff, Ana Paiva, and Marilyn Walker (Eds.). Springer, Berlin, Heidelberg, 252–259. doi:10.1007/978-3-642-33197-8_26

[33] David Throsby and Anita Zednik. 2011. Multiple job-holding and artistic careers: some empirical evidence. *Cultural Trends* 20, 1 (March 2011), 9–24. doi:10.1080/09548963.2011.540809 Publisher: Routledge _eprint: https://doi.org/10.1080/09548963.2011.540809.

Received 16 April 2025; revised 16 April 2025; accepted 28 July 2025