Speculation for RE: Addressing Unanticipated Consequence

Andrew Darby[a], Peter Sawyer[b], Nelly Bencomo[b,*]

[a]Lancaster University, ImaginationLancaster, Lancaster, UK
[b]Durham University, Department of Computer Science, Durham, UK

Abstract

Context: Software innovations frequently lead to unanticipated consequences with significant impacts, a problem exacerbated by emerging technologies and inherent uncertainty in problem domains. Traditional Requirements Engineering (RE) practices are often inadequate for addressing this complexity. Our work is motivated by the need to prefigure emergent properties early in the development process to mitigate such risks.
Objectives: This paper introduces the Consequences and Futures Model, a conceptual framework designed to help software development teams better understand the potential impacts of software innovations at the pre-design phase. The overarching goal is to equip developers with techniques to facilitate proactive risk evaluation and mitigation planning.
Methods: The model draws on three real-world case studies of software systems with unintended consequences: the UK Post Office's Horizon system, Apple's AirTags, and DJI drones in Ukraine. The model draws on existing research outwith the Software Engineering domain, which was evaluated through three qualitative workshops using speculative strategies to provoke forward-thinking analysis.
Results: The workshops successfully demonstrated the model's potential to help participants understand technology's potential impacts and identify areas for further exploration. The model encouraged development teams to consider design decision consequences from multiple perspectives: the organization, users and non-users, the technology itself, and the wider world. It

*Corresponding author.
Email address: nelly.bencomo@durham.ac.uk (Nelly Bencomo)

prompted teams to explore both probable and plausible futures using speculative strategies. The model stimulated creativity, aiding participants in aligning their understanding of subject technologies and broadening their perspective on social implications.
Conclusions: The Consequences and Futures Model enhances RE by integrating Speculative Design techniques to address the risks of innovation. It offers structured exercises for exploring a problem space beyond traditional requirements gathering. While not predictive, it supports imaginative scenario-building to uncover hidden risks. The model's operationalisation is preliminary and we outline a lightweight protocol for practitioner use. Integration into RE processes is left to future work.

Keywords: emergent properties, creative strategies, consequences, requirements engineering, design fiction, futures

---

1. Introduction

We live in a rapidly-changing world in which many of the primary motors of transformation are innovations in software. Some of these innovations change how we do things in an incremental way, while others radically disrupt our personal, social and business lives. Change agents usually want their innovations to have beneficial consequences, but innovations can lead to *unanticipated consequences*, "changes that are neither intended nor recognized by the members of a system" [1]. Unanticipated real-world impacts present a problem that is becoming increasingly urgent as more and more decisions are devolved to the opaque workings of machine learning technologies.

At the heart of the problems lies uncertainty; uncertainty about the nature of the problem domain, uncertainty about the behaviour of the software and uncertainty about the impact of the software on the various environments (business, user, physical, social, etc.) in which it will operate. Understanding and resolving these uncertainties is traditionally the responsibility of requirements engineering (RE). Conventional RE practice involves developing a detailed understanding of the problem at hand and, once this had been achieved, specifying a solution. This can be done in a single upstream phase using a plan-based development model, or more commonly nowadays (and arguably more effectively), in increments using an agile process. This understand-then-specify model is conceptually appealing, but the history of

RE is full of examples where the scale and complexity of the problem have proven beyond the ability of the developers to identify and build the software that its stakeholders desire [2].

Worse, the challenges of scale and complexity are sometimes compounded by volatility in the problem domain and in the technology landscape. Legislation, social norms, user expectations, standards, tools, infrastructure and fashion may all be in flux. New technology, perhaps even the software solution itself, may be the primary agent of change in the problem domain. Mutation of the problem space is sometimes intentional, especially where a new technology, as happened with ride sharing platforms, is purposely disruptive [3]. However, the risk is that the effects go beyond the scope of the problem originally envisaged leading to malign unanticipated consequences. The amplification of extreme views and the facilitation of user surveillance afforded by social networking, the reinforcement of social injustices caused by facial recognition systems, and the contribution to global heating by cryptocurrency mining are all examples of malign consequences arising from software innovations.

The scope of stakeholders' goals, resource constraints and failures to envision possible scenarios all help to inhibit our ability to identify and implement effective mitigations to uncertainties associated with software innovations. Software Engineering, and Requirements Engineering in particular, is focused on delivering stakeholder value. Requirements engineers, product managers and stakeholders have few tools beyond (e.g.) encouragement to think about exception scenarios or obstacles to goal satisfaction with which to evaluate degrees of uncertainty or risks that may seem intangible.

Despite the risks and the limitations of RE practice, the demand for solutions to ever more complex and poorly defined problems continues apace. This is what motivates the work described in this paper, our investigation of whether there are ways to prefigure emergent properties early in the development process. Our vision is that developers and stakeholders can better evaluate risks and plan for their mitigation by building controls into their software, thereby reducing the likelihood of unanticipated consequences, or at least improving preparedness when unanticipated consequences do emerge.

We have developed a model and associated process that has been designed to stimulate users, designers and requirements engineers to explore the problem space and the consequences of potential solutions with a greater focus on possible impacts than is conventionally the case in RE. The paper addresses the problem of unanticipated consequences stemming from software innovations and explores a novel approach to tackling this challenge using speculative design techniques. It introduces the Consequences and Futures Model, a conceptual framework designed to help development teams explore and understand the potential consequences of their design decisions in, and through, a rapidly changing world.

The paper's contributions are as follows:

- A conceptual Consequences and Futures Model that helps teams surface plausible impacts for stakeholders, the software's users and the wider world.
- A workshop protocol with speculative strategies to widen consideration of impacts beyond near-term probable futures.
- An initial evaluation through three workshops highlighting benefits and limitations.

Detailed operationalisation and lifecycle integration are outside this paper's scope.

The rest of the paper is structured as follows. Section 2 reviews related work, Section 3 describes our methodology, Section 4 discusses three well-known exemplars of innovations that have had unintended consequences, Section 5 introduces our Consequences and Futures Model, Section 6 evaluates the model, Section 7 discusses the results of the evaluation and finally, Section 8 concludes the paper and points to future work.

2. Related Work

Consequences, explained in relation to diffusion theory, are "the changes that occur to an individual or to a social system as a result of the adoption or rejection of an innovation" [1]. Rogers describes the dimensions of consequences as being; direct, indirect, anticipated, unanticipated, desirable and undesirable. Direct consequences are an immediate response to an innovation, while indirect consequences occur as a result of direct consequences, as "consequences of consequences". Consequences may occur closer to or further away from their origin; the design action. Their immediacy can be described, in terms of both the speed with which they occur and the directness of their connectivity. Anticipated consequences are "recognized and intended by the members of a system", while unanticipated consequences are not. By their nature anticipated consequences are to varying degrees of certainty predictable, while unanticipated consequences may be either unpredictable or unforeseen. The "functional" and "dysfunctional effects" of an innovation are experienced as an issue of desirability, where the functional are desirable consequences and the dysfunctional undesirable. The desirability of consequences are a matter of perspective and, as such, contestable [1, p. 462-463].

To summarise, the dimensions of consequence may be understood in relation to the immediacy of consequences to a design action (direct vs indirect), the predictability of those consequences (anticipated vs. unanticipated) and the desirability of any consequences arising from the design action (desirable vs. undesirable) [1].

2.1. Addressing unanticipated consequences

Rogers work on consequence within diffusion theory [1, p.441] is informed by Merton's [5, 4] work on the consequences of planned actions. Merton's work highlights *assessment errors*, *myopia*, *fundamentalism*, *lack of foreknowledge* and *self-defeating prophecy* as five factors that limit an actor's ability to anticipate direct and indirect consequences [5].

Sveiby et al. highlight *myopia* as the most common of Merton's limiting factors followed by *assessment errors* and *fundamentalism* [6]. From their innovation research perspective, Sveiby et al. cite pro-innovation bias, informed by either myopia or fundamentalism, as a key issue limiting attention and research into unanticipated consequences [6].

De Zwart offers a critical analysis of Merton [5, 4] that warns against a conflation in the later work [4], regarding the synonymous use of *unintended consequences* and *unanticipated consequences* [7]. The Law of Unintended Consequences states that deliberate actions have unforeseen effects. While, eliding unanticipated and unintended may be commonplace in many fora, De Zwart argues that the term *unanticipated consequences* has disappeared from the literature and that the conflation of terms has limited academic attention on the category of "unintended but anticipated consequences". His argument makes clear that intention (intended vs. unintended) should be added to the dimensions of anticipation, direction and desirability when considering the nature of the consequence of actions. Figure 1 visualises De Zwart's argument, showing the dimensions of consequence (direction, anticipation and desirability) as discussed in diffusion theory [5, 1], highlighting the importance of the dimension of *intention* and the potential concurrence of side-effects and trade-offs with planned actions [7], and demonstrating the scope of unintended and unanticipated consequences. Within Figure 1, the continuous grey lines represent the potential consequences of planned actions, several of these could occur simultaneously. The dashed grey lines represent a pathway that can't be followed, as a consequence cannot be both intended and unanticipated. We also highlight the potential response to each line of potential consequence as being either acceptance, mitigation, amplification or exploration.

It is useful to consider the diagram as classifying potential consequences according to whether they are '*unintended and anticipated*', '*intended and anticipated*' or '*unintended and unanticipated*'. This last classification is our main focus in this paper because it is the one most neglected by existing software development processes.

From an engineering design perspective, Walsh et al. define unintended consequences as "behaviors that are not intentionally designed-into an engineered system yet occur even when a system is operating nominally" and argue that such consequences may arise as a result of either "the bounded rationality of human designers" or those designers willfully ignoring anticipated consequence [8]. The concept of bounded rationality recognises that designers have "limited time and mental resources and therefore may not have sufficient information to be able to make perfectly rational design decisions", while designers' willful ignorance may be attributable to the significance of the anticipated consequence not being known or being considered a dissenting opinion among decision makers, and more routinely they may simply arise as side effects or design trade-offs [8].

Walsh et al. note other strategies that attempt to address unintended consequences such as premortems and redteaming [8]. Premortems aim to benefit from a kind of prospective hindsight, "In a premortem, team members assume that the project they are planning has just failed—as so many do—and then generate plausible reasons for its demise. Those with reservations may speak freely at the outset, so that the project can be improved rather than autopsied" [9]. Redteaming takes an adversarial approach and may be performed using manual or automated techniques with the aim of improving organisational decision-making by incorporating alternative analysis and perspectives [10] or as a goal-based approach to specific client issues. Pieters and van Cleeff argue that the complexity of factors relating to the foresight, assessment and attribution of consequence mean that we "can no longer rely on an ethics of consequences" and should instead turn to the precautionary principle to guide our design actions [11], even though doing so may, among other potential negative consequences, stifle innovation [12].

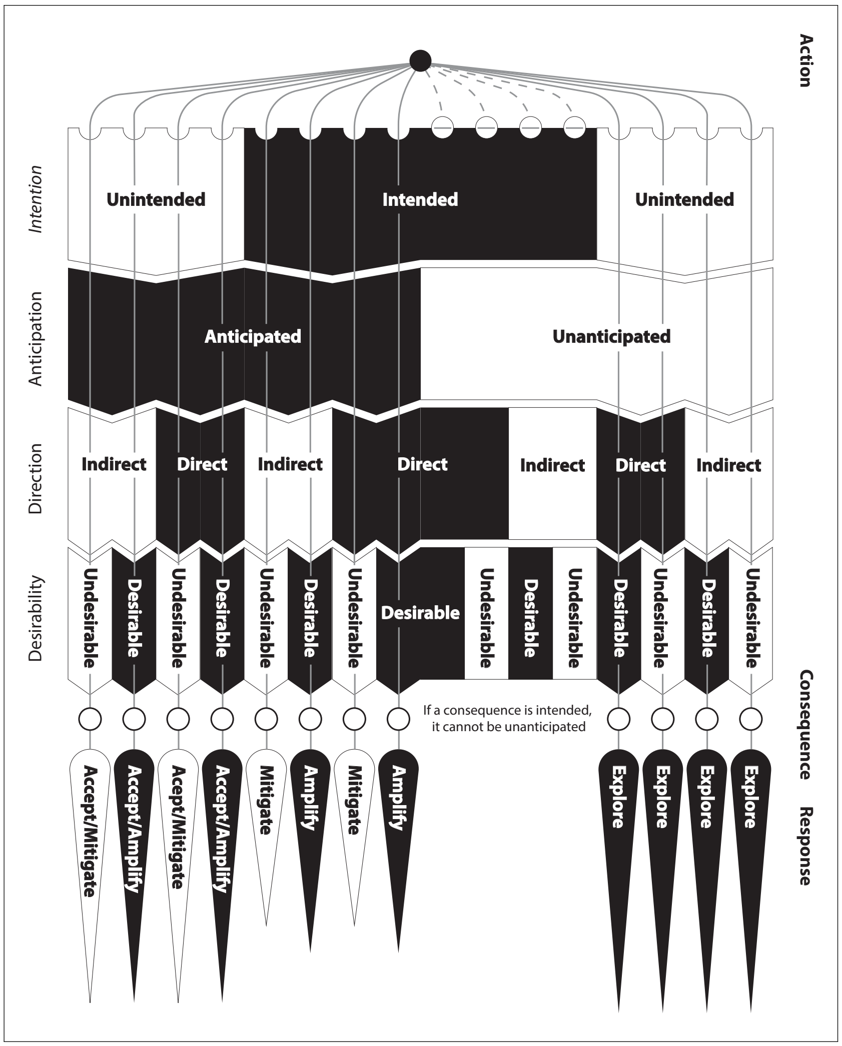

Figure 1: A diagram presenting the potential consequences of planned actions, drawing on Merton 1936 and De Zwart 2015

The goal of the work is to aid innovation by reducing potential *unintended and unanticipated* consequences, and through exploration recognise and identify potential consequences that can then be considered within the two *anticipated* sections in Figure 1. Our ultimate aim is to equip software developers with techniques to help them identify where their software might have impacts on customers, users or the wider societal context in which the software will operate that, using existing RE practices, they might otherwise fail to recognise.

2.2. Addressing the future

The future is a temporal horizon of the present and as such it can never begin [13, p. 140]. It is best understood as being continuously emergent. It results from two drivers. Firstly, the available knowledge and resources that make up the circumstances of the present time. Secondly, the discourses, visualizations and enactments that form our varied and often contested visions of the future. There is constant interplay between circumstance and vision, our contexts and our intents, that continuously generates 'mutually exclusive' futures within an *open future* [13]. The future is fiercely contested, no two people make sense of the world in exactly the same way and "different worldviews and values disclose different truths" [14] leading to myriad futures. To help futurists disambiguate and simplify the complexity of futures, Voros, drawing on [16], helpfully classifies them as to the possibility and desirability of their coming to fruition — labeling them probable, plausible, possible and preferable futures [15]. This taxonomy entered design discourse via [17], then [18]. The scenarios and wildcards visualised in some versions of the futures cone [19, 20] are omitted for the sake of clarity in others [18, 21]. However, it is important to note that these scenarios and wildcards represent a range of alternative possible visions within open futures.

The futures field makes explicit the link between consequences and futures in its adoption of techniques like the futures wheel [22]. Several variants of the futures wheel have been developed [23] that help users map expanding dendritic networks arising from the potential consequences of trends or specific events[1]. Though logical this move is problematic, as the relation between the "thing or circumstance which follows as an effect or result from something preceding" (OED) reduces in certainty over time due to the speed of change

[1]The expansion is often limited to a fixed number of orders.

and the complexity of society. This being the case, addressing consequences over different time periods necessitates an acknowledgment and acceptance of degrees of uncertainty, and an engagement with a range of future possibilities. "To anticipate a future is in some way to 'predict' its possibility, not its certainty" [24] and so, generating alternative futures is a way of outlining the possibility space around specific actions. As a whole the futures field tries to "provide policy-makers and others with views, images, alternatives etc. about futures in order to inform the present" [14] and there is a particular focus on plausible futures in service to organisational preparedness thereby enabling anticipatory action.

As a discipline, Requirements Engineering approaches the future pragmatically. RE's purpose is to identify the properties of an effective solution to a business problem with sufficient clarity such that software engineers are able to devise the means by which these properties can be achieved in software. Collectively, these required properties (hence "requirements") have intended consequences; to make a business process more efficient, to make a product appealing to users, etc. [25]. Its work is fundamentally prospective in that the futures it specifies are informed by commercial imperatives and are expected to occur in the near-term. A barrier to effective RE is the complexity of the problem domain such that understanding of competing priorities and consensus among stakeholders proves hard to achieve. This challenge is inevitably compounded as the problem evolves over time as, for example, business competitors innovate themselves. These factors have contributed to the widespread adoption of Agile methods where the need to constantly reevaluate requirements and design decisions drives software development as a series of short and narrowly-scoped understand-specify-implement-test cycles. This incremental delivery of even large systems is accepted as the price of resolving lack of clarity and responding to change. As the nature of the problem to be solved comes into increasing focus, aided by users' experience with partial solutions to the problem, the development team is better able to adapt their solution accordingly. A side-effect of this, however, is an intensified focus on the near-term [26], with developers able to look little further forward than the date of the next release.

Human Computer Interaction (HCI) explores the interrelationships between humans and computers across various contexts. As an interdiscipline, preoccupied with understanding the interconnections between humans and evolving technologies, it has much of its focus on the novel and the new, yet grounds itself by continuously pressing the boundaries of technological feasibility. It deals with the probable and plausible near-future, and while it has often prospectively considered commercial viability it is also unafraid to speculate. The assumed futures that inform HCI practice often pass undeclared, despite areas of HCI making explicit use of envisioning to set research direction [27]. Ubicomp is one example where this is not the case. Ubicomp is an area within HCI, in which a remotely plausible vision guided the subsequent development of a new paradigm in computing [28], despite its initial commercial potential being slight.

As a discipline, much of Design practice is orientated towards prospective offerings to market and a commercially-viable probable future. However, alternative design practices, like Speculative Design, set aside the constraints of commercial imperatives in order to critically explore the conjunction of desirability and potentiality [29, 18] querying ideologies and values on the way. One method, design fiction [30], borrows from design thinking as well as a range of creative practices, and is "the practice of creating tangible and evocative prototypes from possible near futures, to help discover and represent the consequences of decision making" [31].

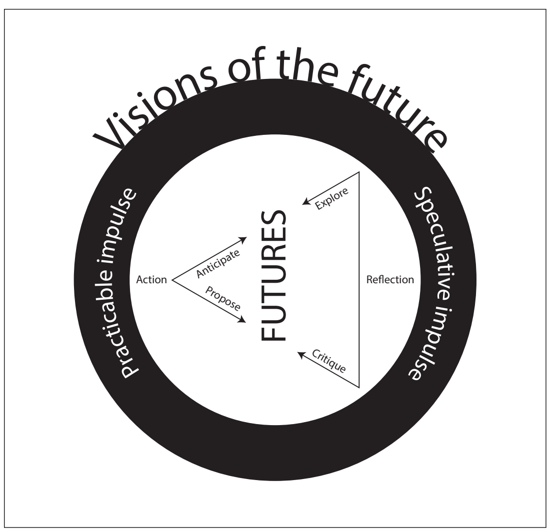

Figure 2: Visions of the future

Visions of the future are driven by a Utopian impulse, the hope for betterment. That impulse may be immediate or more distant, practicable or more speculative, and can be informed by different epistemological stances. Figure 2, presents, the authors' characterisation of these motivations in relation to several key terms highlighted in italics below. RE, and much of Design, are driven by a *practicable impulse* to act in the present, their visions are are motivated by useful practical *action*. And as such, they are framed prospectively, they seek to *anticipate* users' needs and desires, and to *propose* offerings to market. Some Design practices have a *speculative impulse* creating alternative futures, their visions are orientated toward *reflection*. And as such, they seek to question and explore users' needs, and in doing so educate their desires. They are framed as investigations that are broadly in support of the market and the status quo and whose purpose is to *explore*, or alternatively as challenges to that status quo which aim to offer *critique*. In HCI, Ubicomp's visions, demonstrate that these impulses do not have to work in opposition they can operate through oscillation. Shifting between the pursuit of near-future technological possibility driven by both the speculative and the practicable, Ubicomp, has also operated under the umbrella of a larger speculative vision of a more radical distant future. In response to the speed of recent innovation and the attendant societal challenges, in areas such as Artificial Intelligence, a new field has developed. Responsible Innovation is aimed at "aligning research and innovation to the values, needs and expectations of society" [32] and it encourages a robust querying of innovation through anticipation, reflexivity, inclusion and responsiveness [33]. While commercially orientated, Responsible Innovation research engages deeply with plausible futures.

2.3. Getting creative

RE's practices have been usefully reframed as creative problem-solving to support RE practitioners to draw on other bodies of work [34]. Maiden et al. argue that "By framing requirements engineering as creative problem solving we can gain new insights into it, and recruit new knowledge from other disciplines to understand it better and support it more effectively"[34]. Part of the challenge addressed by Maiden et al. is that developing 'good' software requirements cannot be done without a detailed understanding of the (business/market/social/physical/political/etc.) world in which the software will operate and on which it will act. Sometimes the world is just too complex and the world is always changing in ways and at a pace that cannot be fully known. Requirements engineers, product managers, software designers, usability experts, security red teamers, and everyone involved in the conception and development of software make decisions informed in part by assumed futures; their conception of what the world will look like during the lifetime of their software. Making the assumed future explicit is rare in software development practice, but explicitly envisioning the future has found utility in HCI and Ubiquitous Computing research [27]. This has included experimentation with design fiction as a method to explore consequence, in HCI [35] and RE [36].

Our research draws on various aspects of these approaches. We seek a practical means for requirements engineers and others to reason about the future beyond the probable, assumed futures that drive many design decisions to enable more robust decision-making when designs do not encounter their assumed worlds.

3. Research Design

Here we articulate the research questions addressed and outline the make-up of the research team, our methodological approach and the phasing of our inquiry. We go on to describe our methods, data collection and analysis.

Our hypothesis, based on the insights developed in section 4, is that there are immediate economic and longer-term societal benefits to be gained by development teams from their deeper engagement with the potential impacts of software innovations. To address this we sought to develop a practice-orientated technique for software designers to engage with both the consequences and futures that their work may give rise to in order to alleviate negative impacts, such as reputational harm and correction costs. These considerations motivated us to investigate what could help development teams to better understand the possibilities created by a new software product or service. This led us to formulate the following research questions:

- RQ1 Can we design a technique or tool to help development teams surface and structure potential consequences of software innovations beyond assumed probable futures?
- RQ2 To what extent can such a technique or tool be made compatible with existing practices?

This exploratory study adopts an interpretivist approach and uses qualitative methods to address the research questions inductively. The outcomes sought relate to basic, or pure, knowledge in response to RQ1 and applied knowledge in answer to RQ2. Our methodological approach, Research through Design is informed by both Design Science [37, 38] and Design [39]. Taking a multi method approach we used three qualitative methods to address our research questions. The methods included, case studies, a contextual review and workshops. To establish the currency, relevance and scope of our research we used several contemporary real-world cases as examples. Also, as the consequences of technological innovation can be far-reaching, we opted to use a contextual review [40, 41] in our engagement with the literature to resist a singular discipline's interests. Finally, we used workshops as a means to address the domain-specific issues of consequences in RE and SE [42].

The research team comprised the Principal Investigator (an Associate Professor in Computer Science) with expertise in design at runtime, a Co-Investigator (a Professor of Software Systems Engineering) with expertise in requirements engineering, and a Research Associate (a doctoral candidate), whose research focused on design, speculation and participation, and who had extensive experience in workshop facilitation.

The research study was conducted in three phases. The first phase developed reflections on three real world socio-technological systems. These reflections alongside concepts and theories, noted in section 2, informed the development of a conceptual model. Then, in the second phase, the conceptual model was used to guide participant inquiry in three workshops. Two workshops were conducted with a group of industry professionals (with a focus on a prospective technological deployment) and one with academics (with a focus on speculation arising from the further development of a technology). Each topic represented a pervasive, incrementally evolving technology that continues to pose development challenges and is likely to do so into the future, and, as such, provides a suitable subject to support the model's evaluation. Finally, in the third phase the utility of the conceptual model in-use was considered.

In the first phase, we engaged in an extended dialogue around three socio-technological systems that had been widely reported in national UK news media. We purposively selected the UK Post Office's Horizon electronic point-of-sale system (see subsection 4.1), Apple's AirTag tracker technology (see subsection 4.2), and DJI's commercial drones in the context of the war in Ukraine (see subsection 4.3). These cases were chosen because they were well documented, of high public salience, and exemplified real-world software or software-intensive innovations that had produced unanticipated consequences. Collectively, they span contrasting domains—public infrastructure, consumer technology, and dual-use commercial systems—offering diverse yet comparable contexts for examining how software design decisions and organisational values can interact with wider social and geopolitical conditions to produce unintended outcomes. These instances allowed us to explore unanticipated consequence in the actual world as part of three different software or software-intensive products.

Thereafter, we conducted a contextual review [41] to give us an understanding of the evolution of work on consequences related to innovation and software. Rather than applying strict inclusion or exclusion criteria, the selection of material for the contextual review was guided by the need for interdisciplinarity, historical and conceptual breadth, socio-cultural relevance, diverse forms of evidence, and clear alignment with issues emerging from the real-world cases, ensuring that the review captured the multifaceted and evolving nature of technological consequences. Accordingly, the review was shaped by relevance, conceptual richness, and the ability of material to speak to evolving debates across multiple domains, rather than by formalised criteria. The review drew on academic literature [1, 4, 5, 6, 7, 8, 9, 10, 11, 12], court judgments [43, 44] and press reports [45, 46, 47, 48, 49, 50, 51] to illuminate the multifaceted nature of the topic area within a broad socio-cultural context. Drawing on the literature and insights developed studying our three example cases we developed concept maps, in line with [41], and synthesised a conceptual model, which was realised visually for use in workshop settings.

The research team purposively selected industry and academic groups to participate in the workshops which were held between 28th July 2022 and the 11th October 2022. The workshops took place in large meeting rooms at participant groups' organisational bases in England, UK.

The participants in the academic workshop (WA n=5: M=3, F=2) included academics and postgraduates drawn from a university Computer Science department. Participants in the first (WI1 n=8: M=7, F=1) and second (WI2 n=5: M=4, F=1) industry workshop were practitioners from our industry partner, a large telecoms provider. There was complete overlap and continuity between participants in the two industry workshops. In all three workshops individuals' participation was voluntary. The Principal Investigator had a preexisting relationship with both the industry partner and the Computer Science department participating in the study.

No data was collected with regard to non-participation. Each of the workshops lasted two hours, the practitioners and salaried academics received no remuneration, while the postgraduate participants were offered a nominal remuneration –a £25 Amazon voucher– to thank them for their participation.

The lead author facilitated all three workshops in-person, while the second author acted as an in-person observer for the first and third workshop and observed using Microsoft Teams in the second workshop. The research team, made field notes to preempt any privacy concerns our industry partner might harbour regarding digital collection techniques, and collected data during, and after, each workshop to provide a rich description of events for subsequent thematic analysis.

Our reflections on three example cases of software-related innovations becoming newsworthy, section 4, alongside an engagement with the literature on unanticipated consequences and futures, in section 2, allowed us to identify significant areas of interest for development teams to address and debate. In response, we developed a conceptual framework, discussed in section 5, to guide the workshop activities described later, in section 6.

4. Example cases

In the following three subsections, we discuss examples of recent incremental innovations of high salience with well-documented unintended impacts across different domains. In each case insufficient attention appears to have been paid to plausible consequences and which we believe illustrate the need for developers to have better tools with which to explore the plausible futures around technological deployments.

We highlight the UK Post Office's Horizon electronic point-of-sale system, see 4.1, DJI's commercial drone technology in Ukraine, see 4.3, and Apple's tracker technology, AirTags, see 4.2. As we discuss these instances we take the consequences of purposive actions, as discussed in [5] and [1], to be synonymous with the consequences of design actions made in software development and we highlight these concepts, in italics.

4.1. The Post Office Horizon system

The Horizon electronic point of sale (PoS) and accounting system, 'Europe's largest non-military IT contract' [44, Technical Appendix 15], was developed by Fujitsu for the UK Post Office. Subpostmasters, self-employed agents running post offices, were contractually required to use Horizon. Inevitably, Horizon had defects, some of which are now known to have caused accounting errors. The service contracts between Fujitsu and the Post Office disincentivised both parties from believing or acting on subpostmasters' reports of Horizon's failures. Instead, the Post Office's policy was to prosecute any subpostmasters who did not make good any shortfall from their own funds, deeming them liable for the errors [43]. No-one in the Post Office, Fujitsu or the courts seemed to consider it suspicious that a large body of people of previously good character had suddenly taken to embezzlement, and that this conversion to criminality coincided exactly with the introduction of a new PoS and accounting system. As a result, hundreds of Subpostmasters suffered financial penalties. Some were jailed. Some took their own lives. Eventually, a group action by 550 subpostmasters exposed Horizon's defects and the egregiousness of the Post Office and Fujitsu's corporate strategies [47, 48].

To highlight the values embedded in the design of the Horizon system, we used a values taxonomy [52]. We highlight the Post Office's view of the Horizon system's infallibility and note how that speaks to the values of *Wealth* and *Honesty* being prioritised. In the view of the Post Office, Horizon ensured *Honesty* within the socio-technical system, by delivering *Security* with a high degree of *Competence*. The pursuit of Subpostmasters through the courts was necessitated by a fiduciary requirement to protect *Wealth*.

Horizon was not developed as a high integrity system, and as such from the developer perspective software defects were accepted as inevitable. Every non-trivial software system has defects and these sometimes result in failures, such as producing the accounting errors that Horizon generated. However, these need not have produced the impacts that they did; those on subpostmasters' livelihoods, reputations and health. The wider system that Horizon operated within was at fault. The Subpostmasters, as the software's users, were directly effected by the software defects' failures, while the Post Office as the client was indirectly effected. Within the Post Office, a myopic perception of the Horizon system's infallibility held significant sway. Enough so that its mitigation strategies, to defend itself against financial loss rendered it immune to the effects created by the software defects' failures. After all, if the software is faultless then its users must be at fault. So, the Post Office could recoup any losses via the Subpostmasters (who were contractually obliged to regularly make good discrepancies), or write them off (following Subpostmaster prosecutions). This situation left the Subpostmasters exposed to the effects of the software defects' failures which were then amplified by the wider socio-technological system. Software development contracts gave Subpostmasters no access to key data and without that data to aid their defence the Post Office's prosecutory policies became essentially predatory. As such, the Subpostmasters suffered the *intended anticipated indirect undesirable consequences* that were a feature of the wider system design, particularly the contracting process and prosecution practice, in which the Horizon software was used.

4.2. The Apple AirTag

Marketed as a way to keep track of your things the Apple AirTag, released in 2021, was predictably and quickly misused, e.g. for stalking and car theft. Non-Apple, primarily Android, phone users are particularly vulnerable to AirTag misuse. Apple eventually responded to Android phone users' concerns with the release of the Tracker Detect app. Even now, however, while the Apple ecosystem automatically alerts iPhone users to Airtags in their vicinity, Android users have to actively run scans for AirTags. This is despite TU Darmstadt's AirGuard app [46] demonstrating that generating automatic alerts for Android users was possible.

Focusing on values embedded in the design, we noted that in addition to being a part of an ecosystem that embodies *Wealth*, Apple's AirTags selectively highlight other values, such as *Accomplishment* as expressed through design aesthetics, *Helpfulness* as demonstrated by their ease-of-use, and user *Security* as shown by core software features. However, values are also expressed by what actions are not taken, and with the AirTag Apple chose not to provide for the *Security* of non-users, thereby giving considerable scope for potentially subversive, malicious or even criminal intent by AirTag's users.

The potential indirect collateral damage to Non-Apple users arising from Apple's primary design decisions has either been deemed acceptable as an *unintended anticipated indirect undesirable consequence* or it was an *unintended unanticipated indirect undesirable consequence*. Given the seemingly obvious potential for harm, the former would be an example of worrying corporate priorities (and expression of values) and the latter an example of a somewhat mystifying myopia. As only by focusing without exception on the motivating scenario might other use-cases fail to be considered.

4.3. The war in Ukraine and DJI Drones

The use of commercial drone technology in the Ukraine War has been of particular interest to US policymakers in terms of near-term plausible futures [45] and to the Ukrainian government in terms of the direct consequences surrounding their use in the war. Both Ukraine and Russia have used commercial drones to carry out reconnaissance. Ukraine has also modified commercial drones to drop munitions. Russia has targeted drone pilots with artillery bombardments using DJI's *Aeroscope* drone detection product. These repurposings of the drones and Aeroscope were *unintended anticipated undesirable consequences* that DJI accepted might happen and attempted to mitigate through their terms and conditions which precluded the technologies use in warfare.

The values embedded in the design of DJI's Aeroscope and drones are highlighted by the company as being *Creativity*, *Competence*, and, *Freedom*. However, in war, their users have subverted these values in pursuit of advantage in *Security*.

Ukraine called on the commercial drone manafacturer, DJI, to block Russia's use of their drones in the region [49]. DJI's CEO responded stating that the company could not change the product for a number of technical reasons [50]. Without appropriate technical solutions to navigate the demands of the two sales territories, the company stopped sales in Russia and Ukraine [51]. The company's drones still operate in both regions. DJI's inability to respond to the technical challenges presented by these governmental requests suggests either that an *unintended anticipated direct* and, potentially (at least, from a governmental perspective), *desirable consequence* had been considered in some way disadvantageous by the company and was not therefore pursued through development, or, that an *unintended unanticipated direct and desirable consequence* had been left unexplored. DJI were ill-prepared for the urgent new use-cases generated by the outbreak of war and the subsequent demands created by the shift in the context of use of their drones.

When considered through the lens of consequences, as visualised in Figure 1, the above cases highlight some of the ways in which the consequences of design actions operate concurrently or sequentially and also how they could be addressed differently. Focusing on the top two tiers of Figure 1, the visualisation presents three dyads; *intended and anticipated*, *unintended and anticipated*, and *unintended and unanticipated*. The first of these, the *intended and anticipated*, speaks to product or service development with the consequences of design actions necessitating either mitigation or amplification to appropriately shape the features of the product or service. The second, the *unintended and anticipated*, speaks to broader issues of acceptability, where the mitigation or amplification of the consequences of design actions fall beyond the scope of the software development. The third, the *unintended and* *unanticipated*, speaks to the need to widen scope and pursue further exploration of less likely possibilities in near-future states.

These insights drawn from the cases described in sections 4.1, 4.2 and 4.3 were used to drive the development of a conceptual framework taking into account, the following four areas; the technological product or service itself–*Technology*, the context of its development in terms of the company–*Organization*, the context of the world–*World* it will operate in and the range of possibilities that its use suggests–*Use*. This conceptual framework is explicated below.

5. Conceptual Framework

The *Consequences and Futures Model*, see Figure 3, was iteratively developed. We synthesised insights from the three case analyses (Section 4) and the contextual review (Section 2) through concept mapping. We developed practical insights through an early exploratory session that drew on values [54], the adoption cycle [1] and the PESTLE framework [19][2]. Further concept mapping was organised visually as a guide to participant speculation and, finally, speculative strategies (Section 5.5) were incorporated to support participants in operationalising the conceptual relationships. From the case studies, we observed how organisational biases, selective value prioritisation, and external disruptions amplified unintended consequences. From the literature, we identified key themes such as the multidimensional nature of consequences, limiting factors for organisation's consequence recognition, the importance of values in shaping design, and the roles of futures thinking and creative strategies for speculative exploration. These insights converged into five guiding principles for addressing unanticipated consequences: (1) broaden the consequence space; (2) integrate organisational and societal perspectives; (3) expand temporal scope; (4) operationalise exploration through structured speculative strategies; and (5) adopt multi-domain analysis. These principles directly shaped the structure and content of the concept map, to ensure it would provide a practical framework for reasoning about the wider impacts of software innovations. Iterations were reviewed collaboratively within the research team and refined to ensure internal coherence and usability.

[2]This session explored aviation futures with an international consultancy firm.

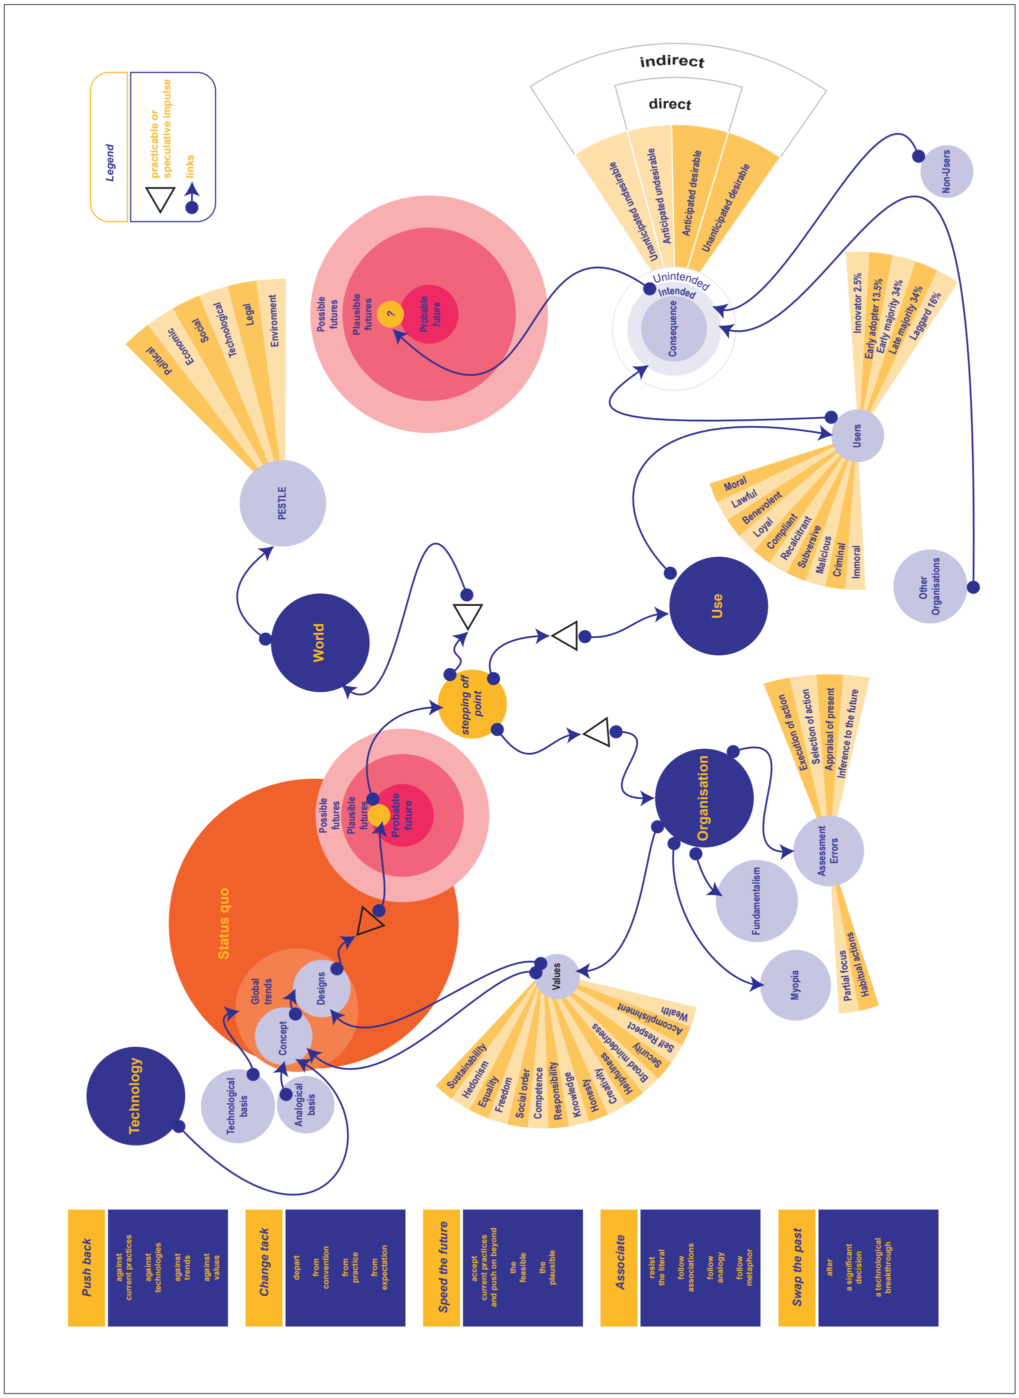

Figure 3: The Consequences and Futures Model. The creative strategies described in the panels on the left are informed by Pierce [53].

Below, we describe the concepts operationalised through the model and connect them back to the literature, in order to demonstrate how the factors were derived based on relevancy to our objectives. The *Consequences and Futures Model* uses concepts related to the creation and use of organisations' technologies in a changing world, with key terms highlighted in italics as follows: *Organisation*, *Use*, *Technology*, and *World*; and two areas related to temporality, namely *probable* and *plausible* futures.

In using the *Consequences and Futures Model*, we take anyone participating in a design or requirements exercise, software practitioner, customer, user, etc. to be acting as a part of the development team. The model provides a means for the development team to communicate among themselves and a process map to help guide the work.

5.1. Organisation

This section, see Figure 4, of the model focuses on the *Organisation* behind the planning of a design, it invites the development team to consider the impact of the organisation, and its units, on the development. Following Sveiby et al. we draw on Merton to highlight the 'factors that limit an actor's possibility to anticipate both direct and indirect consequences' [6]. Prioritising prevalence, as demonstrated by [6], we focus on the three most common limiting factors *Myopia*, *Fundamentalism* and *Assessment errors* and set aside *Lack of foreknowledge* and *Self-defeating prophecy*.

We briefly address each factor. *Myopia* is informed by values, it occurs when the desirability of beneficial consequences becomes the overriding focus of attention, such as the Post Office's unshakeable belief in the infallability of the Horizon system. *Fundamentalism* happens when an organisation's dominant values drive actions because of a 'felt necessity' [5, p. 903], such as the Post Office prosecutory policy. So to consider *Myopia* and *Fundamentalism* development teams must critique the organisation by addressing its values, these are 'the criteria people use to select and justify actions and to evaluate people (including the self) and events' [54]. Sutcliffe et al. introduced a values taxonomy relevant to software development [52], which we mapped onto the Schwartz values framework [54] to demonstrate the advantage of the streamlined taxonomy as a more readily usable element within the conceptual framework[3].

[3]Sutcliffe et al.'s Wealth, Self Respect, Broad mindedness, Helpfulness, Creativity,

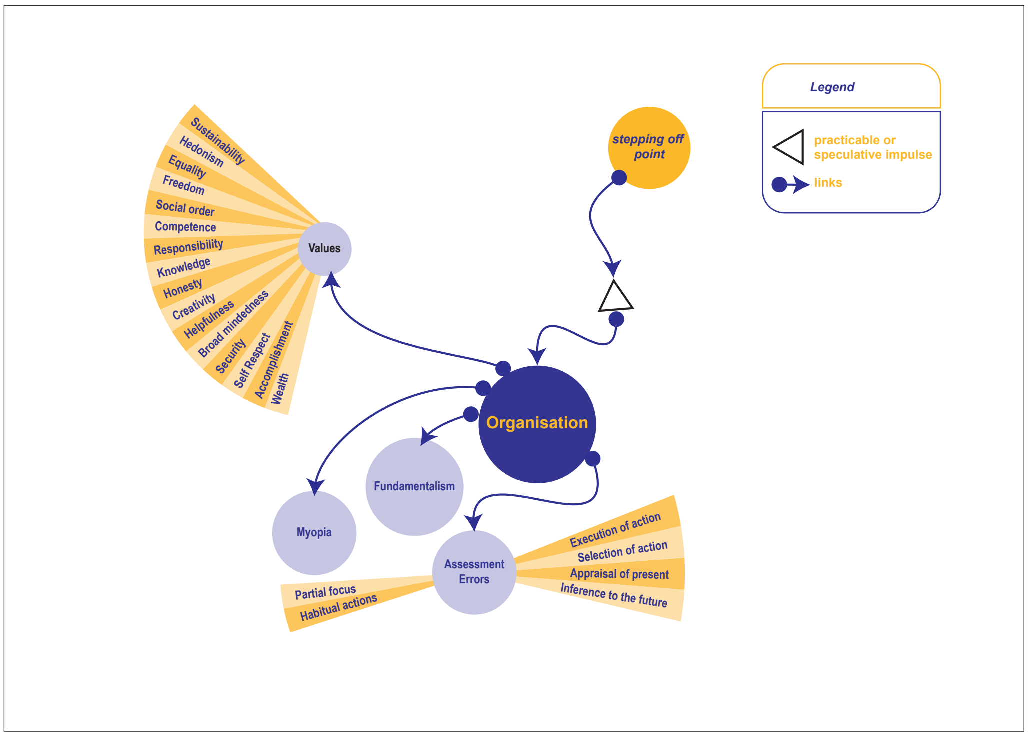

Figure 4: The Consequences and Futures Model - detail of the Organisation section. The Assessment Errors, Myopia, and Fundamentalism parts of the model draws on Merton [5]. The values part of the model draws on Sutcliffe et al. [52] which is itself informed by Schwartz [54].

This part of the model invites the development team to critically reflect on the quality of the organisation's assessment. Merton highlights how *Assessment Errors* occur, noting that 'we may err in our appraisal of the present situation, in our inference from this to the future objective situation, in our selection of a course of action, or finally in the execution of the action chosen' [5, p. 902]. Additionally, he draws attention to the fallacious assumption that past successful actions will continue to prove effective, noting that the belief becomes fixed in 'the mechanism of habit' [5]. He also points out errors of partial focus, ranging from 'simple neglect' to 'pathological obsession' [5]. Apple's development of the AirTag, and its neglect (if accidental) of Android users' safety, appears to be an example of *Partial focus*. Adopting Merton's considerations, markers for *Partial focus* and *Habitual action* fan out to the left of the *Assessment Errors* marker, while fanning out to the right are markers for *Execution of action*, *Selection of Action*, *Appraisal of present*, and *Inference to the future*. The Post Office's pursuit of Subpostmasters over assuring the software was fit for purpose could be described as a problem of *Selection of Action* in the Horizon socio-technical system and DJI's ability to register risks associated with responding to operating in the context of a war suggests issues with their making of *Inference to the future*.

5.2. Use

This section, see Figure 5, focuses on the design in *use* and invites the development team to consider how proximity relates consequences directly, or indirectly, to *Other Organisations*, *Non-Users* and *Users*, as well as the *Organisation* itself. Merton, Rogers, and De Zwart [5, 7, 1] underpin the potential consequence space. This is intended to invite consideration as to how expected, acceptable or probable a consequence of the design decision in use might be.

Honesty, Responsibility, Social order, Freedom and Equality directly map to individual Schwartz values. Competence is encompassed by a synonym of an individual Schwartz value, Capability. Accomplishment is encompassed by a similar concept, Success. Security broadens the scope of two Schwartz values, Family security and National security by dropping the modifiers. Knowledge is most closely aligned with two Schwartz values, Intelligence and Wisdom. Hedonism directly maps to a class level Schwartz values, which incorporates the individual Schwartz values; Pleasure, Enjoying life and Self-indulgence. Sustainability may be aligned under the Universalism class level Schwartz values, which incorporates the individual Schwartz values; Protecting the environment, and Being part of nature.

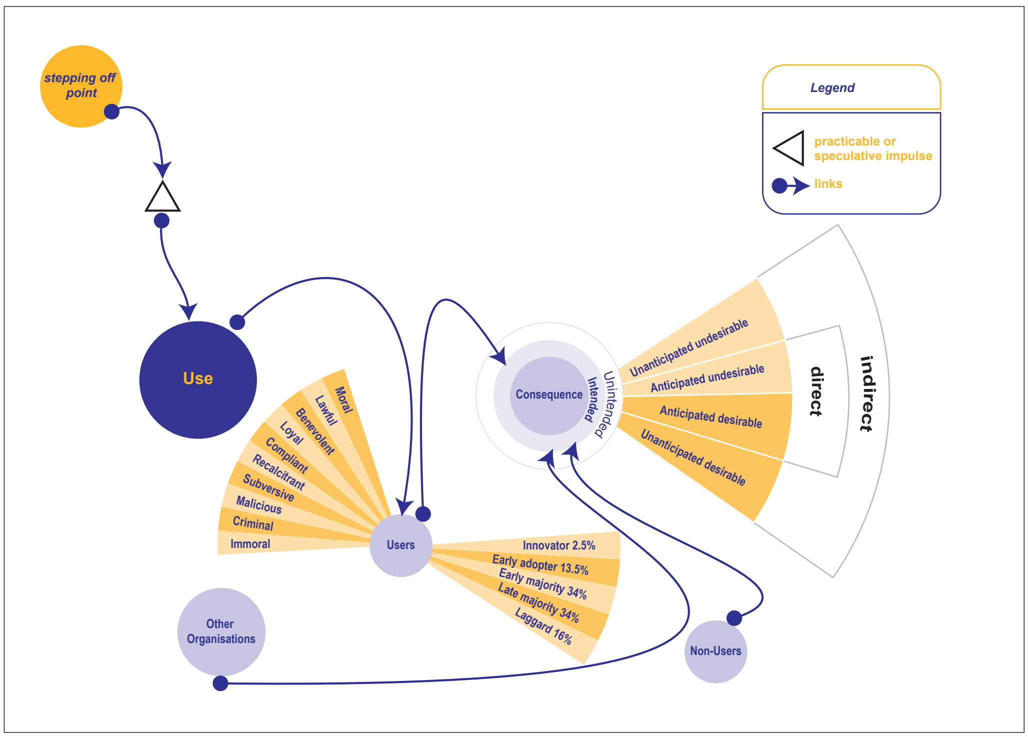

Figure 5: The Consequences and Futures Model - detail of the Use section. The users part of the model draws on Darby [55] and Rogers [1], while the consequences part of the model draws on Merton [5], Rogers [1] and De Zwart [7].

We highlight the *Users* attitude to use and the position in the adoption cycle as key points of interest. The range of development teams' attitudes is given, to support a consideration of user impact and perspective, and is an extension of the spectrum presented in [55, p. 151]. The adoption cycle [1] is used to help development teams to explore the different impacts that the point of adoption may have on usage, especially as that intersects with user attitude. The significance of this is illustrated by DJI's seeming ambivalence to the (mis)use of their drones in the war in Ukraine, see 4.3.

5.3. Technology

This part of the model, see Figure 6, invites the development team to consider technology as things are in the present, in the context of the status quo from a number of viewpoints. Firstly, it asks the development team to consider the *Analogical* [34], and *Technological basis* of the design *Concept* and to further consider how that *Concept* and the *Design*(s) that springs from it form part of *Global trends*. Then, it asks the development team to consider how *Values* are embedded in a design, how a design acts prospectively to describe a probable future, and how the conditions of the wider world support that possibility. Applying a values taxonomy [52] to the Apple AirTags example case, see section 4.2, highlighted the design decision-making within the product development. The values taxonomy [52] is included to support speculation regarding values embedded in the design. These values terms can be employed to query organisational *Myopia* and *Fundamentalism*, in 5.1, as well as supporting exploration of development teams' attitudes, in 5.2.

As they develop a technological design into an innovation companies plan for the market to be configured in a particular manner based on high degrees of likelihood. After launch an innovation's success is reliant, to some degree, on the stability of its context of use. However, in the real world contexts of operation are unpredictable and subject to change, which may have radical implications for an innovation.

5.4. World

This section, see Figure 7, invites the development team to consider potential external factors from the wider world that may act on the *Status Quo* changing the foundations for the context of use of the design decision irrevocably. Though minor events, especially those directly related to a design's supply chains and sales territories, may play a more significant role in many

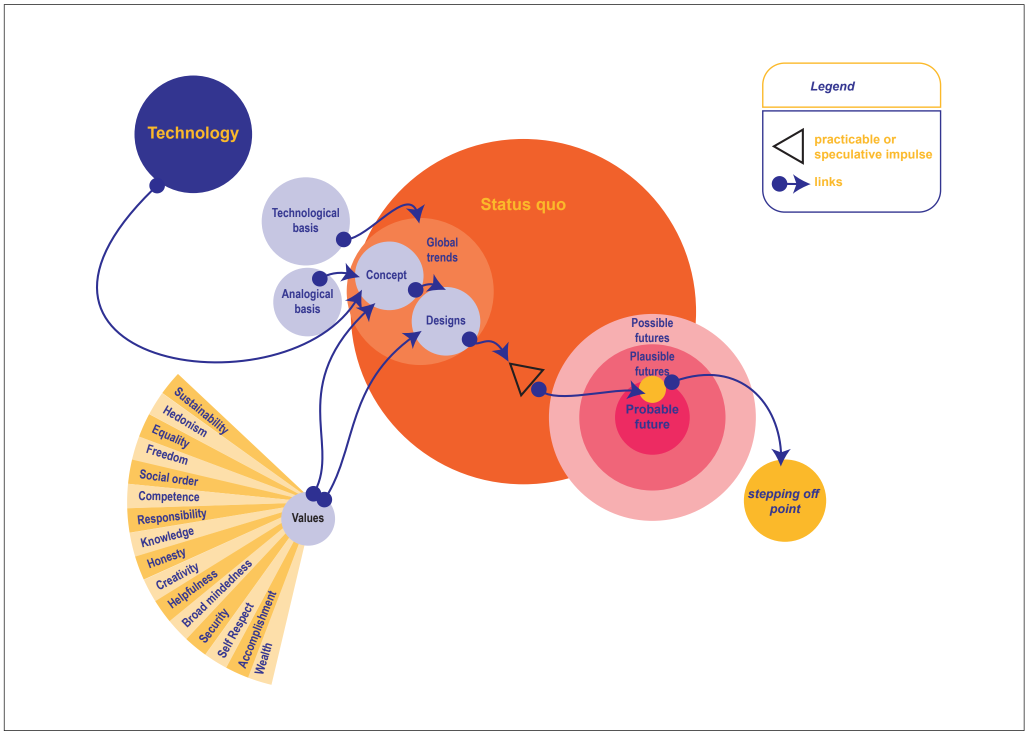

Figure 6: The Consequences and Futures Model - detail of the Technology section. The values part of the model draws on Sutcliffe et al. [52] which is itself by informed Schwartz [54]. The futures part draws on Dunne and Raby [18] and other authors detailed in the text. The analogical basis is informed by Maiden et al. [34] and the status quo relates to Taylor [56].

designs' operational contexts it is important to also consider major events, especially those that may be disruptive, aberrant, catastrophic or anomalous [19]. A development of Aguilar's EPTS [57], the *PESTLE* (*Political*, *Economic*, *Social or Socio-cultural*, *Technological*, *Legal* and *Environmental*) tool is helpful in querying macro pictures of the industry environment. *PESTLE*'s use here is inspired by its incorporation into Taylor's futures work [19], which specifically focuses on plausible futures. The impact of the war in the Ukraine on the governmental and customer expectations of DJI drones, see 4.3, recommends *PESTLE*'s inclusion.

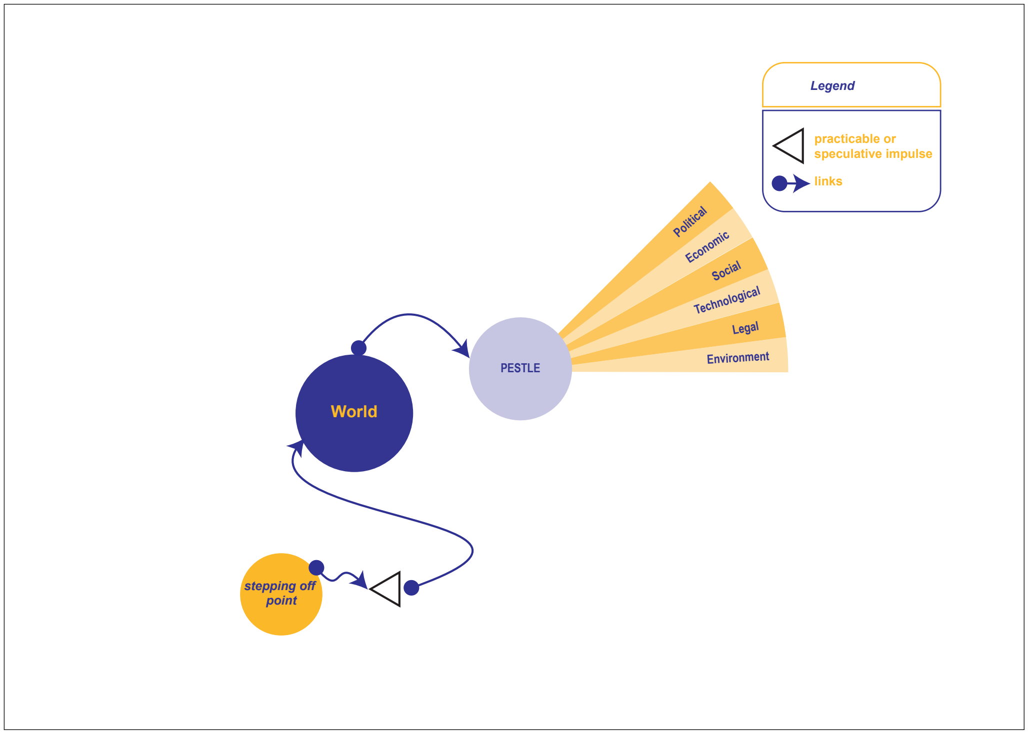

Figure 7: The Consequences and Futures Model - detail of the World section. This part of the model draws on Taylor [56] which is itself informed by Aguilar [57].

5.5. Speculative Strategies

Speculation about the future is not untethered imagination. Speculative, like prospective, design practices work within the laws of physics and focus on plausible and probable futures [18] and both, while differently orientated, are underpinned by the concept of prefiguration, the need to imagine beforehand [53]. Five speculative strategies, presented in panels on the left side of Figure 3, were drawn from Pierce's observation's on frictional tendencies [53] and presented to support development teams' engagement with the conceptual framework. Pierce's frictional tendencies; analogical, divergent, oppositional, accelerational and counterfactual, were originally conceived as 'a tool for teaching design students how to concretely compose and construct speculative, critical, and conceptual designs' [53]. These were more accessibly termed *Associate*, *Change tack*, *Push back*, *Speed the future*, and *Swap the past*, respectively. They are common creative strategies some of which are seen in Requirements Engineering [34], and others in Design [58], Speculative Design [18], and Design Fiction [30]. These speculative strategies are oriented toward the development of probable and plausible futures, as described by [15, 16, 18]. The probable future invites the development team to generate a practicable assessment of the technology and the plausible future invites a speculative exploration, as described in Figure 2.

The yellow circle at the model's centre provides a *Stepping off point* for speculation that is informed by the probable world that the technology under consideration is designed to be used in. Along each of the paths towards the *World*, *Organisation* and *Use* parts of the model, a triangle indicates the option to pursue either a practicable or more speculative impulse, as described in Section 2.2.

It should also be noted that we do not claim that the conceptual framework is a tool for prediction or analysis. Rather, it provides a speculative jumping off point for development teams' creative exploration. By tracing various pathways across the *Consequences and Futures Model*, we aim to support the creation of broadly plausible futures formed in productive tension with participant understandings of probable and possible futures.

6. Evaluation

The evaluation was conducted by piloting the conceptual framework in a series of three workshops. Across these workshops, one with academics WA and two with industry professionals WI1 & WI2, the research team aimed to better understand how the *Consequences and Futures Model* might be employed in workshop settings to usefully address over-the-horizon change for the purposes of understanding requirements and informing design decisions.

The workshops allowed the research team to gather feedback on the use of the conceptual model in explorations of possible futures related to incremental technological innovation in two domains, internet search in WA and digital twins in WI. The research team selected or agreed workshop topics that exemplified pervasive, incrementally evolving technologies that continue to pose development challenges and are likely to do so into the future.

As is consistent with our methodology evaluative criteria emerged through observation, thematic analysis, and participant feedback. Our evaluation of the *Consequences and Futures Model* considered the engagement and creativity of participants and the scope of speculative exploration, as well as participant satisfaction, perceived efficacy and critique. These indicators were were derived from participant's spontaneous comments, tone and engagement, in line with [59]'s qualitative approach to evaluating informal conversation.

Following an introductory section outlining the workshop protocol, these evaluative criteria provide useful sub-headings for our evaluation.

6.1. Workshop protocol

The WA and WI1 workshop sessions were designed to run as stand-alone two-hour sessions. These sessions were structured in three parts; groundwork, speculation, and discussion. The WI2 session, was added after WI1's completion, and was planned as an iterative extension of the speculative and discussion parts of WI1. Participants were seated at a large table around a large-format poster of the *Consequences and Futures Model*. The lead author facilitated the process, and Post-it notes™, Sharpies™and a whiteboard were used to aid in the generation, sharing and organisation of ideas.

Next, we outline each of the three parts; groundwork, speculation, and discussion:

Part 1: Groundwork: This initial phase was intended introduce the model and to establish the topic area and create a shared starting point for the session. Drawing on the technology section 5.3 of the framework, we aimed to draw out working definitions of the technological innovation that formed the focus of the workshop. Thereafter explorations were made of the concepts, designs and values underpinning that innovation, the analogies that explicate it, and applicable trends within the status quo. Activities included developing a brief working definition of the core technology, mapping underlying technologies and their relations, detailing value propositions, identifying underpinning values, and discussing trends and factors in the wider world affecting their development.

Part 2: Speculation: This section focused on generating speculations about the *World*, the *Organisation*, and the *Technology* in *Use*. Participants were introduced to five creative strategies, see 5.5, and asked to develop potential futures using each one. Working with the prompts – *Speed the future*, *Associate*, *Change tack*, *Push back*, and *Swap the past* – participants generated speculations under a two-minute time constraint for each strategy. The time allocated to each activity was limited in order to avoid fixation [60] and encourage a light-touch approach to idea generation. Participants were encouraged to speculate about both unanticipated and anticipated consequences and to generate ideas by intersecting different elements of the model as they considered the different creative strategies.

Part 3: Discussion: The session concluded with a group discussion, where the goal was to have participants describe the space between what they believed would happen and what could happen, and highlight any insights gained during the process of arriving at their various speculations.

6.2. Engagement and creativity of participants

The groundwork section of WI1 was covered quickly and coherently, participants agreed a working definition of Digital Twins as a "data-based and logical representation of the physical world with bidirectional communication" and identified a wide range of underpinning technologies, many of which are themselves new and developing, including Mixed Reality and IoT. They described the organisation's current Digital Twins projects based on Building Information Modelling and immersive training, and identified internal threats to the development of future Digital Twins applications. Participants' considerations of internal threats focused on weak organisational thinking that might lead to a failure to recognise the potential of Digital Twins for the organisation.

When participants responded to our values query about their current projects they articulated the corporate values the company espouses and aspires to rather than offering a critical assessment. Also, when asked to describe the analogies and metaphors at play in their Digital Twins work the responses showed the participants' uncertainty, with abstract concepts, questions, and corporate visions offered up alongside analogies and metaphors. Responding to questions about the trends and external factors exerting an influence on future Digital Twins applications they were more certain. Participants highlighted key trends such as; Data-driven decision making, Business automation, Industry 4.0, and Metaverse virtual world. They also highlighted external factors, such as; energy efficient and clean energy informed by climate crisis and remote working informed by Covid pandemic acceleration.

Overall, the WI1 participants responded well to the speculative strategies they were invited to use, they particularly favoured the analogical, accelerational and oppositional ones. However, they found it difficult to develop ideas using the counterfactual strategy in the available time, and the deviational and divergent strategies also generated fewer ideas than might be expected, though this may be the result of a rushed prompt by the facilitator. Throughout WI1 and WI2, participants were able to frame their evolving understanding of the Digital Twins concept by developing an extensive set of speculations, including predictive extrapolations, plausible and possible futures. Of the workshop sessions, the observer noted that "It must have been tiring but everyone appeared to enjoy it and everyone contributed. The degree of creativity and the participants' openness to speculation was impressive."

The WA participants responded well to our prompts to use speculative strategies in combination with other elements of the model. Below, we showcase the creativity of participants and their degree of engagement with some illustrations of how they responded. Prompts to consider probable and plausible futures, intersected with current and other possible values, resulted in a flurry of speculations: for example,

- search results were imagined that provided traditional measures of accuracy, regional specificity, and, inferred users emotional responses;
- search engines were imagined that made explicit use of AI chatbots leading to suggestions of anthropomorphic search engines or where hedonism provided an alternative guiding value.

Applying PESTLE, WA participants developed further speculations: for example,

- explicitly costing searches in terms of their environmental impact;
- applying legal constraints on deepfake content;
- prioritising requirements for equality of access, etc.

The WA participants then focused on attitudes to use to speculate about the effects of various phenomena, including:

- shifts in users' values and preferences (attitudes to ads, value of ease of access to information);
- increased state control c.f. Weibo;
- the emergence of purposely malign search engines c.f. the dark web.

Engaging with the prompts and the subsequent discussion the participants broadened their understanding of search engines as a class of technology.

6.3. Scope of speculative exploration

Reflecting the WI participants' roles as engineers and technologists, a key theme through which the group explored the development of Digital Twins was through automation. A line of speculation built on the technological trend to assert a probable future based on the status quo. The human-in-the-loop was conceived as failing which lead to the human being replaced in decision-making through AI systems and in physical capacities by robotic or nano-robotic systems. Significant assumptions underpinning this move toward automation included, forms of virtual reality being presumed to be freeing, forms of physical reality being considered constraining and that AI technologies could adequately provide the prediction and interpretation necessary to address complex needs.

Countering this line of reasoning, other speculations of plausible futures by WI1 and WI2 participants –prompted by the PESTLE framework– explored the social impact of these probable technological trends, they speculated about the impact of job losses and the challenges posed by non-working economies. They noted the emergence of the notion of Universal Basic Incomes proposed in contemporary economics as a potential panacea for the problems that their plausible technological developments would create. They also addressed the impacts on their organisation, recognizing that the process of moving toward non-working economies through automation could lead to deficits in organizational memory, skills loss, and a disaffected workforce.

Drawing insights from these speculations together, in the discussion the WI participants focused on the definition of Digital Twins and how they might improve interactions between physical and virtual entities, as well as virtual and virtual entities. In doing so they eschewed the requirement for a physical object to have a direct relation to a digital twin.

Similarly, WA participants broadened their understanding of search engines as a class of technology. The exercise was carried out at pace with participants sharing ideas and insights into potential opportunities for the evolution of search engine technology, as well as some of the potential pitfalls. However, there was little time or energy left for the group to converge around a particular design or set of designs for next-generation search engines, or, more importantly, to draw out insights and requirements based on the relation between the plausible futures they posited and nearer-term probable realities.

6.4. Participant satisfaction, perceived efficacy and critique

In discussion, WA participants reported that they enjoyed the format of the workshop. For the postgraduate students particularly, the idea of speculation as a means to reason about technology was novel and rewarding. After the workshop, one of the students showed how much they valued the model by emailing a request for a copy of it to enable further study.

In discursive feedback at the end of the workshop, WI1 participants were positive about having the space to think differently about the Digital Twins and being encouraged to think beyond the enterprise. Participants identified a need to develop a more nuanced understanding of thresholds between humans and physical and virtual realities as a result of Digital Twin technologies. They enjoyed the explicit prompting to think about what might go wrong and what the future might look like.

Following the workshops, the WI participant group provided written feedback. They were broadly positive about the experience of participating in the workshops, and they were pleased with the result, however they were critical of the iterative nature of the process. Interestingly, participants felt that a greater diversity in the participant group might have produced different, better or more creative results. However, they were not specific as to the nature of that diversity.

While, WA participants were under no obligation to demonstrate benefit to an employer, individual WI participants held at minimum a tacit expectation that the time taken out of their day-to-day work would not be misspent and preferably that it would demonstrate explicit benefit. Indeed, the R&D team leader hoped that our workshops would help the team develop a better understanding of Digital Twin technologies. The concept had emerged as a point of interest for several of the company's business units and the R&D team was responsible for developing a consistent definition and evaluating their potential utility to the company. Through our workshops using the model, WI participants were able to frame their evolving understanding of the Digital Twins concept by developing an extensive set of speculations, including predictive extrapolations, plausible and possible futures. By the end of the sessions, the R&D team planned to present their redefinition of Digital Twins to upper management, it would also, assumedly, inform their subsequent research and development agenda.

Our evaluation demonstrates the workshop participants' creative engagement with *The Futures and Consequences Model* resulting in positive participant feedback and an acknowledged degree of utility.

7. Discussion

We return to our research questions to guide our discussion, which are formed, as follows:

RQ1 Can we design a technique or tool to help development teams surface and structure potential consequences of software innovations beyond assumed probable futures?

RQ2 To what extent can such a technique or tool be made compatible with existing practices?

Our evaluation shows that participants were able to surface and structure potential consequences of software innovations beyond assumed probable futures answering RQ1 by proving a degree of utility for our partner organisation. However, the response to RQ2 and the extent to which *The Futures and Consequences Model* is compatible with existing practices requires further research.

Our aim in this work is to investigate how Requirements Engineers can identify and evaluate potential impacts of software even where the software is highly innovative and the wider (e.g. business, social) environment is volatile, subject to change and hard to predict. To do this we have drawn on work on consequence and futures. Discussion of consequence and futures share key dimensions, the predictability and desirability of consequence as discussed in communication studies [1, p. 60] are articulated in terms of probability and preferability in futures and speculative design discourse [15, 18]. The consequences of a software product on business or wider society follow from decisions made by the designer. In most cases, the designer considers these consequences as they would apply if the future was the same as the present - a static world. The future is difficult to predict and is therefore both harder for the designer to reason about and has less immediate value; why let the future influence a design decision when it might never happen? Pragmatically, and understandably, practitioners keep the scope of their design decisions tight and their consideration of futures even tighter. While speculating about the social consequences of design decisions is uncommon in industry practice–and considering that aspects of the workshop challenged some WI participants' worldviews, particularly regarding the neutrality of technologies–participants were open to and engaged well with the exercises undertaken. The WA academics, not having actual design decisions to make, found it even easier to engage with futures exploration and, as a result, perhaps gained less value. It should be noted in relation to RQ2, that any discomfort created for participants' in encountering unfamiliar concepts in the workshop setting is not in itself antithetical to compatibility with existing practices, as some discomfort may promote productive speculative frictions.

In software and systems development, to be manageable, consequences need to be as predictable, desirable and direct as possible, with any negative impacts mitigated against in development, and through supported release, using (e.g.) agile feedback loops. As demonstrated in Section 4 this approach can have very poor social outcomes. Our approach is different. We seek to make critical consideration of less predictable, desirable and direct consequences manageable so as to reduce the incidence of design decisions that turn out to be bad. The cause and effect relation implied by consequences is messy. This is due not only to design decisions in isolation but how these decisions interact with a host of other elements; in the business domain, in social norms, in organisational culture, and so on. When considering possible consequences, we do not need to consider only those that result in the poorest of outcomes (even if we could predict what those might be). Rather, there is value in any widening of the set of possible consequences of design decisions. Exploring the possibility space encourages a long and broad view of consequence which may be used to frame the context for the direct views required by RE.

As WI1's call for greater diversity in the workshop process attests, it is not only the chronological distance from a present that defines a future, it is the viewer's belief in a futures' possibility and probability, which is itself informed by their own position. Any perception as to its desirability is informed by the viewers' ability to conceive its impacts on their own, and others', future lives. In this way acts of anticipation are conditioned by one's perspective. Across all of the workshops, the various switches between present and future presented a significant but stimulating challenge for participants. The iteration of creative strategies was an intentional part of the process. However, some participants questioned the approach taken with WI1 & WI2, reportedly feeling they were retreading territory. However, WA participants were more positive with one going so far as to request the model for further study. For them the value resided almost tacitly in the act of determining the design, their perceived benefit was more experiential. The participants' call for greater diversity in this kind of workshop, e.g. beyond requirement engineers and the wider software team, also addresses a potential myopia in our approach. There is a strong argument that key stakeholders including decision makers and the people driving the vision for the software also need to buy in and be involved in the process.

We argue that the iteration of creative strategies is a feature rather than a bug and is essential to the development of new perspectives, though we recognise that the value of the approach may be questioned by participants resistant to perceived repetition. Iteration's value is evidenced by the development of the WI team's understanding of Digital Twins technology, and its role within the organisation, across the workshops. Nevertheless, we accept that we need to manage expectations and demonstrate value. Moving toward a diversity of participants that iterate through creative strategies should serve to broaden the perspectives on consequence produced through the exercise. That said, there are a number of issues that could impinge on the successful use of *The Futures and Consequences Model* in industry settings. The model directly encourages critique of organisational weakness and as a result it could be challenging to operate in companies with more hierarchical organisational cultures, particularly in terms of psychological safety [61]. The same point might be made in terms of hierarchical expectations within national cultures with high expectations of deference. The model also encourages a critique of technology that may sit at odds with more techno-optimistic assessments from potential participants.

Time is a precious commodity and accessing industry practitioners was especially difficult. Our first conception of the workshop design was heavily modified to run at two hours instead of the 1.5 days we had originally envisioned due to this very reason. Piloting the model challenged the authors presumption that 1.5 days was required for a successful workshop and, with further work, a single 3.5 hour session now appears achievable for industry use. With time-limited access to participants, the pacing of content was, at times, stressful for participants, with some becoming visibly tired by the demands of the workshop. A number of the activities required participants to develop an understanding of unfamiliar worldviews, creative strategies, and conceptual devices, that each demand significant attention and time. For example, participants' familiarity and facility with analogies and metaphor is likely to be variable, and allowing time to introduce these devices, and ensure they are understood, is essential to support full engagement with the framework. That said, the perceived value of the model ultimately dictated time allocation, and WI1 demonstrated sufficient value to the industry partner to justify investment in WI2. This kind of work requires championing from within the organisation to convince organisational gatekeepers of its value, both as a complement to current practice and risk mitigation and as a developing element of best practice.

Finally, RE research requires specialised approaches, often demands strict experimental controls, and often imposes constraints that conflict with the inherent characteristics of speculative approaches. Addressing the potential social consequences of innovation is not a common approach within software or systems engineering. Although the workshop participants were alert to social change and the fact that technological innovation plays a significant role in this, the WI R&D team, in particular, viewed technological innovation as inherently neutral, in line with instrumental theories of technology. As the conceptual framework assumes the relation between values and technology is more complicated, in line with a critical theory approach [62], the framing of the workshop can prove challenging to participant worldviews.

8. Limitations and Future work

Feedback from the workshop participants and our own observations of how the participants' understanding evolved provides evidence that our model can help stakeholders gain a wider understanding of a technology and its impact. This has been most strongly demonstrated where our participants were operating in exploratory mode, seeking to better understand a technology and its potential for impact.

The evidence for our model's potential for people working in prospective mode, i.e seeking to envision new products and services, is weaker. Our industry participants, who had an R&D rather than a product development role proved adept at using the model to find commonality among the multiple and sometimes poorly-aligned company-wide conceptualisations of digital twin technologies and at postulating new product features. The model stimulated the team's creativeness in a way that was possibly useful for the company. However, when discussion led the participants into potentially negative impacts of envisioned developments of digital twins, e.g. deskilling, they didn't pursue the line of reasoning that led them there, instead tending to gloss over the negatives. This limitation may be a reflection of the fact that digital twins were already a de-facto technology within the company, even if their maturity of adoption varied considerable across the company's different divisions and business functions. For example, it is possible, in line with work on psychological safety [61], that there was an unacknowledged nervousness that postulating negative side-effects of digital twins was to implicitly criticise management decisions approving its adoption within the company. We postulate that had the team been tasked with evaluating an emerging technology, not yet adopted by the company, there might have been more appetite for a more balanced investigation of the potential for both +ve and -ve impacts.

Our ultimate goal is that the enriched insights into technology impacts that our model offers should feed into design decisions in the form of product requirements. For this to happen, we would need to develop a practice, based on our model, that practitioners see as clearly beneficial; that can help designers avoid unanticipated or unintended negative consequences and which can be integrated with existing requirements and design practices. To better understand the extent to which our model represents an advance towards this goal, we will need to transform such insights into product requirements in a live project and trace how requirements are handled throughout a full project life-cycle. Such a longitudinal study would be the goal of the next stage in our research.

However successful we and others might ultimately be in developing speculative techniques that integrate into established RE processes, we do not expect that the time and human resources needed for their effective application would be justified in all cases. Our motivation for this work was sparked by the observation that some innovations have widespread societal impact, and our conviction that a means was needed to understand what these might be early enough to avoid or mitigate the negative impacts. The question then is, how to recognise where there is a risk of such impacts, without needlessly investing resources where the risk is low. How to characterise the applications where such a risk exists and identify those that would benefit from speculative RE methods is a second strand of work that we plan for the future.

9. Conclusion

In this paper, we have proposed augmenting established Requirements Engineering practices with techniques drawn from Speculative Design. Our motivation is to make the development of innovative software less prone to unanticipated consequences. We argue that the need to sensitise developers and other stakeholders to the likelihood that their software can have consequences beyond those envisioned is amply demonstrated by recent events. This need is critically urgent as new technologies, perhaps most notably Generative AI, increasingly disrupt how businesses operate and even how wider society operates.

The work introduces and evaluates an initial adaptation of speculative design techniques for RE, providing insights into their efficacy and practical applicability. The principal contribution of our work is a *Consequences and Futures Model* that organises elements of Speculative Design into a structured set of exercises that a development team can use to surface and structure potential consequences and develop a broad understanding of a problem space. The model is designed to broaden awareness of potential consequences rather than exhaustively predict them, recognizing that many future factors remain inherently uncertain. Such prediction is impossible, as many factors that determine consequence (e.g. future developments in politics, society, technology, etc.) are impossible to predict with confidence. Rather, the application of the model should help the team gain an understanding of the potential for their software to have impacts to an extent that existing RE practice does not support. Our other contributions are a lightweight workshop protocol intended to support the operationalisation of the model, and an initial evaluation of the model's use in three workshops that highlights its benefits and limitations.

More generally, we aim to demonstrate the value of speculative strategies in making explicit the wider consequences often overlooked in software development.

Our findings suggest that the model can serve as a valuable addition to RE practices, particularly for projects with significant societal or organizational impact. However, further work is required. Since many software projects are relatively narrow in scope, guidance is needed to help practitioners identify those in which the potential impact is sufficiently broad in scope and costly enough to justify the cost of applying the *Consequences and Futures Model*. We anticipate this will entail refining workshop protocols with templates, digital tools, and clearer role guidance to improve scalability and usability. This will also help us integrate the model with existing requirements engineering practices to support seamless adoption. Accordingly, and subject to funding, we will apply the model across the entire product development lifecycle. This work is planned for a domain in which generative AI is being introduced in order to deepen our understanding of its sustained impact on decision-making. It will further help us evaluate whether the model improves the identification of unintended consequences compared with traditional approaches.

10. Acknowledgements

This paper results from the Twenty20Insight project funded by the UK's Engineering and Physical Sciences Research Council (Grant reference number EP/T017627/1). The research study had ethical approval from the Science Faculty Ethics Committee at Durham University (Reference: COMP-2022-05-04T16_15_46-hdcl45). Many thanks to our participants for generously sharing their time, expertise and enthusiasm. The first author was further supported by the Wellcome Trust under Grant reference number [222858/Z/21/Z].

References

[1] E. M. Rogers, Diffusion of innovations, 4 ed., Simon and Schuster, New York, USA, 1995.

[2] T. Justinia, The uk's national programme for it: Why was it dismantled?, Health Services Management Research 30 (2017) 2–9. URL: https://doi.org/10.1177/0951484816662492. doi:10.1177/0951484816662492, pMID: 28166675.

[3] A. Komanduri, Z. Wafa, K. Proussaloglou, S. Jacobs, Assessing the impact of app-based ride share systems in an urban context: Findings from austin, Transportation Research Record 2672 (2018)

34–46. URL: https://doi.org/10.1177/0361198118796025. doi:10.1177/0361198118796025.

[4] R. K. Merton, Social theory and social structure, Simon and Schuster, New York, USA, 1968.

[5] R. K. Merton, The unanticipated consequences of purposive social action, American sociological review 1 (1936) 894–904.

[6] K.-E. Sveiby, P. Gripenberg, B. Segercrantz, A. Eriksson, A. Aminoff, Unintended and undesirable consequences of innovation, in: XX ISPIM conference, The Future of Innovation., ISPIM, Vienna, 2009, pp. 1–16.

[7] F. De Zwart, Unintended but not unanticipated consequences, Theory and Society 44 (2015) 283–297.

[8] H. Walsh, A. Dong, I. Tumer, Towards a theory for unintended consequences in engineering design, Proceedings of the Design Society: International Conference on Engineering Design 1 (2019) 3411–3420. doi:10.1017/dsi.2019.348.

[9] G. Klein, Performing a project premortem, 2014. URL: https://hbr.org/2007/09/performing-a-project-premortem, last accessed: 29-08-2023.

[10] D. F. Longbine, Red teaming: Past and present, 2008. URL: https://apps.dtic.mil/sti/citations/ADA485514, accession Number: ADA485514.

[11] W. Pieters, A. van Cleeff, The precautionary principle in a world of digital dependencies, Computer 42 (2009) 50–56.

[12] H. E. Miller, K. J. Engemann, The precautionary principle and unintended consequences, Kybernetes 48 (2019) 265–286.

[13] N. Luhmann, The future cannot begin: temporal structures in modern society, Social Research 43 (1976) 130–152.

[14] R. A. Slaughter, Sense making, futures work and the global emergency, Foresight 14 (2012) 418–431. URL: https://doi.org/10.1108/14636681211269897.

[15] J. Voros, A Primer on Futures Studies, Prospect, the Foresight Bulletin 6 (2001) 1–7.

[16] N. Henchey, Making Sense of Future Studies, Alternatives 7 (1978) 24–28.

[17] S. Candy, The Futures of Everyday Life: Politics and the Design of Experiential Scenarios, Ph.D. thesis, University of Hawai'i at Manoa, 2010.

[18] A. Dunne, F. Raby, Speculative Everything, The MIT Press, Cambridge, Massachusetts, 2013.

[19] C. W. Taylor, Creating Strategic Visions, Strategic Studies Institute, US Army War College, Carlisle Barracks, PA, 1990.

[20] T. Hancock, C. Bezold, Possible futures, preferable futures., The Healthcare Forum Journal 37 (1994) 23–29.

[21] P. Coulton, J. Lindley, Vapourworlds and Design Fiction: The Role of Intentionality, The Design Journal 20 (2017) S4632–S4642. URL: https://doi.org/10.1080/14606925.2017.1352960. doi:10.1080/14606925.2017.1352960. arXiv:https://doi.org/10.1080/14606925.2017.1352960.

[22] J. Glenn, Futurizing teaching vs. futures courses, Social Science Record 9 (1972) 26–29.

[23] J. C. Glenn, The futures wheel, in: J. C. Glenn, T. J. Gordon (Eds.), Futures research methodology-version 3-0, The Millennium Project, Washington DC, USA, 2021, p. Chapter 6.

[24] T. Fuller, Will small become beautiful?, Futures 32 (2000) 79–89. URL: https://www.sciencedirect.com/science/article/pii/S0016328799000683. doi:https://doi.org/10.1016/S0016-3287(99)00068-3.

[25] Z. Jin, Chapter 2 - requirements engineering methodologies, in: Z. Jin (Ed.), Environment Modeling-Based Requirements Engineering for Software Intensive Systems, Morgan Kaufmann, Oxford, 2018, pp. 13–27. URL: https://www.sciencedirect.com/science/article/pii/

B9780128019542000029. doi:https://doi.org/10.1016/B978-0-12-801954-2.00002-9.

[26] R. Kasauli, E. Knauss, J. Horkoff, G. Liebel, F. G. de Oliveira Neto, Requirements engineering challenges and practices in large-scale agile system development, Journal of Systems and Software 172 (2021) 110851. URL: https://www.sciencedirect.com/science/article/pii/S0164121220302417. doi:https://doi.org/10.1016/j.jss.2020.110851.

[27] S. Reeves, Envisioning ubiquitous computing, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12, Association for Computing Machinery, New York, NY, USA, 2012, p. 1573–1582. URL: https://doi.org/10.1145/2207676.2208278. doi:10.1145/2207676.2208278.

[28] G. Bell, P. Dourish, Yesterday's tomorrows: notes on ubiquitous computing's dominant vision, Personal and ubiquitous computing 11 (2006) 133–143.

[29] J. Auger, Speculative design: crafting the speculation, Digital Creativity 24 (2013) 11–35.

[30] J. Bleecker, Design Fiction: A short essay on design, science, fact and fiction, Near Future Laboratory, San Francisco, USA, 2009.

[31] J. Bleecker, N. Foster, F. Girardin, N. Nova, The Manual of Design Fiction, Near Future Laboratory, San Francisco, USA, 2022.

[32] B. C. Stahl, C. Aicardi, L. Brooks, P. J. Craigon, M. Cunden, S. D. Burton, M. D. Heaver, S. D. Saille, S. Dolby, L. Dowthwaite, D. Eke, S. Hughes, P. Keene, V. Kuh, V. Portillo, D. Shanley, M. Smallman, M. Smith, J. Stilgoe, I. Ulnicane, C. Wagner, H. Webb, Assessing responsible innovation training, Journal of Responsible Technology 16 (2023) 100063. URL: https://www.sciencedirect.com/science/article/pii/S2666659623000069. doi:https://doi.org/10.1016/j.jrt.2023.100063.

[33] R. Owen, J. Stilgoe, P. Macnaghten, M. Gorman, E. Fisher, D. Guston, A Framework for Responsible Innovation, John Wiley & Sons,

Ltd, New York, USA, 2013, pp. 27–50. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118551424.ch2. doi:https://doi.org/10.1002/9781118551424.ch2.

[34] N. Maiden, S. Jones, K. Karlsen, R. Neill, K. Zachos, A. Milne, Requirements engineering as creative problem solving: A research agenda for idea finding, in: 2010 18th IEEE International Requirements Engineering Conference, IEEE Computer Society, 10662 Los Vaqueros Circle, P.O. Box 3014, Los Alamitos, CA 90720-1314, 2010, pp. 57–66. doi:10.1109/RE.2010.16.

[35] A. Rapp, Design fictions for learning: A method for supporting students in reflecting on technology in human-computer interaction courses, Computers & Education 145 (2020) 103725. URL: https://www.sciencedirect.com/science/article/pii/S0360131519302787. doi:https://doi.org/10.1016/j.compedu.2019.103725.

[36] A. Darby, E. Tsekleves, P. Sawyer, Speculative requirements: design fiction and RE, in: 2018 IEEE 26th International Requirements Engineering Conference (RE), IEEE, Banff, AB, Canada, 2018, pp. 388–393.

[37] J. Zimmerman, J. Forlizzi, S. Evenson, Research through design as a method for interaction design research in hci, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '07, Association for Computing Machinery, New York, NY, USA, 2007, p. 493–502. URL: https://doi.org/10.1145/1240624.1240704. doi:10.1145/1240624.1240704.

[38] J. Zimmerman, J. Forlizzi, The Role of Design Artifacts in Design Theory Construction, Artifact 2 (2008) 41–45. doi:10.1080/17493460802276893.

[39] W. Gaver, What should we expect from research through design?, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '12, Association for Computing Machinery, New York, NY, USA, 2012, p. 937–946.

[40] C. Barnes, G. Melles, Managing interdisciplinarity: a discussion of the contextual review in design research, in: International Association of Societies of Design Research (IASDR) Conference, School of Design,

The Hong Kong Polytechnic University, Kowloon, Hong Kong, 2007, pp. 1–14.

[41] C. Gray, J. Malins, Visualizing Research: A Guide to the Research Process in Art and Design, Ashgate, Aldershot, UK, 2004.

[42] R. Ørngreen, K. T. Levinsen, Workshops as a research methodology, Electronic Journal of E-learning 15 (2017) 70–81.

[43] T. H. M. J. Fraser, Ewhc 606 (qb), [2019]. URL: https://www.bailii.org/ew/cases/EWHC/QB/2019/606.html.

[44] T. H. M. J. Fraser, Ewhc 3408 (qb), [2019]. URL: https://www.bailii.org/ew/cases/EWHC/QB/2019/3408.html.

[45] H. Moumen, Commercial drone tech proliferates in ukraine, 2022. URL: https://www.nationaldefensemagazine.org/articles/2022/9/7/commercial-drone-tech-proliferates-in-ukraine.

[46] A. Heinrich, N. Bittner, M. Hollick, Airguard - protecting android users from stalking attacks by apple find my devices, in: Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec '22, Association for Computing Machinery, New York, NY, USA, 2022, p. 26–38. URL: https://doi.org/10.1145/3507657.3528546. doi:10.1145/3507657.3528546.

[47] BBC News, Post office scandal: Public inquiry to examine wrongful convictions, 2022. URL: https://www.bbc.co.uk/news/business-60369875.

[48] N. Wallis, The Great Post Office Scandal, Bath Publishing, Bath, UK, 2021.

[49] DroneXL, Video shows risk of flying a dji drone in ukraine right now, 2022. URL: https://dronexl.co/2022/03/16/vice-prime-minister-dji-block-dji-drones, last accessed: 05-04-2022.

[50] H. Kesteloo, Ukrainian vice prime minister urges dji ceo frank wang to block dji drones used by russians immediately, 2022. URL: https://dronexl.co/2022/03/16/vice-prime-minister-dji-block-dji-drones/.

[51] A. Wong, Chinese drone maker dji halts sales in russia and ukraine, 2022. URL: https://www.vice.com/en/article/v7dapb/dji-halts-drones-sales-russia-ukraine-war.

[52] A. Sutcliffe, P. Sawyer, W. Liu, N. Bencomo, Investigating the potential impact of values on requirements and software engineering, in: Proceedings - 2021 IEEE/ACM 43rd International Conference on Software Engineering, Proceedings - International Conference on Software Engineering, IEEE, United States, 2021, pp. 39–47. doi:10.1109/ICSE-SEIS52602.2021.00013, funding: This research was partially supported by EPSRC (EP/T017627/1) Twenty-20 Insight. ; 43rd IEEE/ACM International Conference on Software Engineering: Software Engineering in Society, ICSE-SEIS 2021 ; Conference date: 25-05-2021 Through 28-05-2021.

[53] J. Pierce, In tension with progression: Grasping the frictional tendencies of speculative, critical, and other alternative designs, in: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, Association for Computing Machinery, New York, NY, USA, 2021, pp. 1–19. URL: https://doi.org/10.1145/3411764.3445406.

[54] S. H. Schwartz, Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries, in: M. P. Zanna (Ed.), Advances in Experimental Social Psychology, volume 25, Academic Press, San Diego; London, 1992, pp. 1–65. URL: https://www.sciencedirect.com/science/article/pii/S0065260108602816. doi:https://doi.org/10.1016/S0065-2601(08)60281-6.

[55] A. Darby, Design Fiction and Participation: From Social Dreaming to Speculative Heterotopia, Ph.D. thesis, Lancaster University, 2023.

[56] C. W. Taylor, Creating Strategic Visions, Strategic Studies Institute, US Army War College, Carlisle Barracks, PA, 1990. URL: https://apps.dtic.mil/sti/pdfs/ADA231618.pdf.

[57] F. J. Aguilar, Scanning the business environment, Macmillan, New York, NY, USA, 1967.

[58] T. Brown, et al., Design thinking, Harvard business review 86 (2008) 84.

[59] J. Swain, B. King, Using informal conversations in qualitative research, International Journal of Qualitative Methods 21 (2022) 16094069221085056. URL: https://doi.org/10.1177/16094069221085056. doi:10.1177/16094069221085056. arXiv:https://doi.org/10.1177/16094069221085056.

[60] D. G. Jansson, S. M. Smith, Design fixation, Design studies 12 (1991) 3–11.

[61] A. Edmondson, Psychological safety and learning behavior in work teams, Administrative Science Quarterly 44 (1999) 350–383. URL: https://journals.sagepub.com/doi/abs/10.2307/2666999. doi:10.2307/2666999. arXiv:https://journals.sagepub.com/doi/pdf/10.2307/2666999.

[62] A. Feenberg, Transforming technology: A critical theory revisited, Oxford University Press, New York, USA, 2002.