



A Workshop-Based Method for Navigating Value Tensions in Collectively Speculated Worlds

Nava Haghighi*
Stanford University
Stanford, CA, USA
navaha@stanford.edu

Matthew Jörke*
Stanford University
Stanford, CA, USA
joerke@stanford.edu

Yousif Mohsen
Stanford University
Stanford, CA, USA
yousifm@stanford.edu

Andrea Cuadra
Stanford University
Stanford, CA, USA
apcuad@stanford.edu

James A. Landay
Stanford University
Stanford, CA, USA
landay@stanford.edu



Figure 1: A scene from the *action round* of our workshop, in which participants use roleplay and speculative artifacts to explore the ethical implications of a future scenario.

ABSTRACT

The rapid pace of technological progress carries with it a heightened risk of ethics and privacy violations, creating an urgent need for mechanisms to address this risk. We approach this problem from the perspective of designers and technologists aiming to design technology that better accounts for ethical implications. We iteratively developed a workshop-based method ($N=113$, seven workshops) for probing ethical implications of emerging ubiquitous computing technologies. We contribute a method that enables people with varying levels and areas of domain expertise and with a variety of lived experiences to collectively speculate about the ethical implications of emerging technologies, navigate value tensions, and prototype artifacts as a way to grapple with those tensions. We introduce implication design as a means for participants with and without design experience to communicate how a technology might change to better serve them. Lastly, we share our learnings from and reflections on our design process.

*Both authors contributed equally to this work.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

DIS '23, July 10–14, 2023, Pittsburgh, PA, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9893-0/23/07.
<https://doi.org/10.1145/3563657.3595992>

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**; *Interaction design process and methods*; • **Social and professional topics** → *User characteristics*.

KEYWORDS

collective speculation, value tensions, implication design, design ethics, speculative design, values in design, design workshop, design methods

ACM Reference Format:

Nava Haghighi, Matthew Jörke, Yousif Mohsen, Andrea Cuadra, and James A. Landay. 2023. A Workshop-Based Method for Navigating Value Tensions in Collectively Speculated Worlds. In *Designing Interactive Systems Conference (DIS '23)*, July 10–14, 2023, Pittsburgh, PA, USA. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3563657.3595992>

1 INTRODUCTION

With digital devices collecting increasingly intimate data that is analyzed using complex and uninterpretable algorithms, the risk of ethics and privacy violations has been growing at the same pace as technological progress [12, 25, 32, 41, 60]. Academic and public concern surrounding technology ethics has sharpened in recent years, motivating a broad range of approaches to addressing ethical concerns across a variety of disciplines [7, 8, 29, 39, 79]. Despite these efforts, the anticipation and communication of ethical implications of technology remains a significant challenge.

We approach this problem from the perspective of design and computer science researchers aiming to understand how the design of emerging technologies ought to change to better account for ethics and privacy implications. As researchers, we have both the power and responsibility to shift the technologies we develop towards more desirable, ethical, and equitable futures. However, while we may have good intentions and aim to develop the technology for a specific group's needs, we are acutely aware that the technologies we create can ultimately be reappropriated for different purposes. We therefore believe that it is crucial to consider a multiplicity of alternative technology futures from the outset of the design and development process.

Among design-based approaches to computing ethics, many have turned to methods such as speculative design [4, 27] and design fiction [9, 69] to probe the ethical implications of emerging technologies. Through the imagination of possible futures or alternative presents [17], these methods can create space for critical reflection on values and ethics [5, 79]. Recent work has emphasized the importance of considering whose voices, experiences, and perspectives are considered when imagining possible futures [11, 18, 19, 40, 58, 65, 73, 74]. Towards addressing these concerns, a growing body of work has explored more participatory approaches to speculation and futuring [30]. Speculating from multiple perspectives affords the possibility of constructing scenarios that center diverse values and experiences, but will often construct futures whose values are in tension with each other [15]. We see these value tensions not as disagreements to be resolved, but as a means to deepen the discussions around ethical implications of technology, which values or whose values to prioritize, and how to decide which actions to take in response. In particular, we recognize a unique opportunity to bring participants together not just to surface value tensions that can later serve as inspiration for designers, but engage participants in collectively navigating these tensions.

In this work, we explore how collective speculation might be applied towards probing the ethical implications of emerging technologies. To develop our workshop-based method, we adopted an iterative process and conducted a series of co-design workshops with over 100 participants over the course of one year. Our goal was to develop a workshop that (1) surfaces diverse values, (2) enables participants of different backgrounds (specifically, different levels and area of expertise, as well as different identities and demographics) to collectively navigate those value tensions, and (3) does not rely on heavy facilitator involvement or design expertise.

Towards these goals, we make use of several techniques to mitigate power dynamics and create a space where people of all backgrounds can meaningfully and comfortably contribute. We draw inspiration from collaborative game design and role play to facilitate collective speculation [17, 49, 63] and increase engagement [21, 52, 54], and present structured brainstorming activities to guide the collective imagination of possible futures. We introduce the concept of *implication design*, a design approach towards embedding a technology's ethical implications in its design. We use implication design to enable participants with and without technical or design experience to communicate how a technology might be changed to better serve them, leading to a deepened discussion of value tensions in the context of imagined solutions. Lastly, we reflect on the outcome of the final workshop and our learnings,

focusing on the *process* by which participants collectively navigated value tensions and iterated through complex design trade-offs.

In summary, we make the following contributions:

- (1) The final design of our workshop-based method for anticipating and communicating ethical implications of technology.
- (2) A summary of our iterative design process and reflection on our learnings while conducting seven workshops with 113 participants.
- (3) Implication design, a design approach for embedding a technology's ethical implications in its design that we employ to allow participants without technical or design expertise to collectively navigate value tensions.

2 RELATED WORK

In this section, we first provide an overview of theoretical perspectives that center values and ethics in design research. We then discuss the ways in which speculation and futuring can create space for critical reflection on values and ethics. Next, we share how participatory and co-design methods can address shortcomings of speculative design methods by opening the space for collective speculation. Lastly, we describe prior work in game design that can enable collectively navigating value tensions while striving for equal contribution.

2.1 Values and Ethics in Design

What exactly constitutes ethical technology and how best to design it has long been a subject of academic debate across disciplines such as philosophy of technology, science and technology studies (STS), human-computer interaction (HCI), anthropology, media studies, and more [67]. The multiplicity of theories and methods is a testament to the complexity of the “wicked problem” [59] at hand. In this work, we align ourselves with perspectives common in HCI and design research (see Shilton [67] for a comprehensive review).

Within HCI and design, approaches such as value-sensitive design [33, 35, 36], values in design [45], worth-centered design [14], reflexive design [64], and critical design [26] aim to surface how values and ethics are embedded in the design of technology and/or provide methods for operationalizing the consideration of human values in the design of technology. For instance, reflexive design [64] argues that critical reflection on unconscious values should be an integral part of technology development. Through critical reflection, designers can make explicit which experiences are centered by the design and which are marginalized. In our workshop, we aim to create space for such critical reflection.

Value-sensitive design (VSD) provides practitioners with theory and methods for eliciting and engaging with stakeholder values in the design process. We take inspiration from several VSD methods in our work, including direct and indirect stakeholder analysis [37], scenario analysis [55], and card-based activities for value elicitation [34]. Our techniques align with characterizations of values as being situated in local contexts, shaped by societal norms and culture, and intertwined with lived experience [35, 43, 48], and we draw particular attention to the importance of engaging with value tensions [51]. We adopt JafariNaimi et al.'s [43] definition that “values serve situations as hypotheses.” Following this view, values serve design problems by foregrounding both what the situation

is and what can be done about it. In our workshops, we aim to support participants in collectively determining the situations that demand action and navigating what might be appropriate actions to take. In the following section, we discuss how we draw from speculative design methods to elicit values and value tensions in imagined futures.

2.2 Speculation and Futuring

Nearly all designers engage in the imagination of possible futures [17] and most technological development is driven by speculative imagination of “future sociotechnical worlds” [65]. Design methods that explicitly engage with the practice of speculation are commonly referred to as speculative design [4, 27, 77] or design futuring [46], though this terminology is not without controversy [4, 17, 46]. Closely related methods include design fiction [9, 69], critical design [26], and adversarial design [23]. Each of these methods involve the use of fictional scenarios or worlds to imagine possible futures or alternative presents [4] and frequently involve the design of speculative artifacts for those futures. By removing practical or commercial constraints, speculative artifacts and scenarios can “suspend disbelief” [70], create space for critical reflection on societal values, or challenge dominant narratives of the future [46].

However, speculative design has been criticized for not engaging with topics of race, gender, class, and power, and for not calling larger sociotechnical infrastructures into question [18, 19, 46, 58, 73]. In response to these criticisms, and as speculative design practices have moved “out of the ‘showroom’ and into ‘field’” [30], a growing body of work has turned to more participatory forms of speculation and futuring, which we discuss in the following section.

2.3 Participatory Forms of Speculation

Participatory and co-design methods aim to enable stakeholders to directly participate in the design of technology [53, 62]. Participatory design theory highlights that practitioners must create new forms of participation to allow non-expert participants to meaningfully contribute [53] and reflect on the power dynamics among stakeholders in the design process [10]. Closely related to participatory design, co-design also aims to enable stakeholder participation in the design of technology [68], with a lesser emphasis on political participation [20].

Approaches that combine elements of participatory and speculative design vary widely in their goals—centering and amplifying the voices of marginalized communities [6, 40, 72, 74], facilitating collaboration among participants with different viewpoints [13, 15], complementing design knowledge with non-design expertise [24, 75], grounding speculated futures in situated experiences [22], or supporting long-term forecasting in technology innovation [28]—as well as in their level of participation [30]. We share many of the goals above and draw particular attention to approaches that have an explicit focus on surfacing the ethical and privacy risks of technology, including *Timelines* [79], *Judgement Call* [5], and *Security Fictions* [50].

In addition to surfacing values, we aim to build on this body of work by exploring methods that enable participants of different backgrounds to collectively *navigate* value tensions in imagined

worlds. However, it is not sufficient to simply add diverse stakeholders “and stir” [53]. We use game design and role play techniques, described next, to scaffold speculation and facilitate equal contribution while reducing facilitator involvement.

2.4 Game Design and Roleplay

Games are a powerful medium for speculation, since they encourage immersion and expand the range of possibilities without being bound to practical limitations of the physical world [17, 49]. Due to this quality, there has been a history of using games in HCI as a rhetorical medium to implicitly shape values and encourage action [16, 31, 61] or to convey knowledge or skills as in serious games [47]. Serious games have also been used in conjunction with roleplay activities to involve participants in a make-believe process [49]. Relatedly, other work in HCI has used roleplay techniques to facilitate critical self-reflection [57] or understand user needs [63].

In addition to games being used to encourage speculation and exploration of fictional scenarios, games have been used in participatory design as a “leveler” to create a “relaxed, familiar, and egalitarian” space [52, 54]. Moreover, collaborative games have been shown to increase social closeness [21]. Lastly, gamestorming—creating game worlds to examine business challenges—has been used to increase engagement and improve collaboration in business workshops [38]. Using game design and roleplay techniques, our workshop design facilitates speculation, immersion, and critical reflection, and fosters equal contribution among participants of varying backgrounds and expertise levels.

3 BACKGROUND AND MOTIVATION

Our motivation for developing our workshops emerged from a tension we faced as researchers interested in imagining how designed artifacts should better account for ethics and privacy implications. Given that such implications are contextual, relational, and subjective, the decision of which implications to consider in designs that may affect millions of people should not fall on a small group of designers. Thus, we set out to explore how using designed artifacts (as a complement or alternative to systems-level or policy-based approaches) could be used as a means to enable people with and without technical expertise to communicate how a technology might be changed to better serve them. We found precedence in prior work that utilized design to embed the ethical implications a technology might carry, and propose the term *implication design* to describe artifacts that communicate or protect against ethical implications of technology through their design.

In this section, we first discuss the concept of implication design and provide examples of implication design in prior work. We then discuss our positionality as researchers and conclude with our workshop’s goals.

3.1 Implication Design

In this work, we use the term *implication design* to describe a design approach in which a product’s implications are embedded into its design, analogous to how the functionality a product affords is communicated through signifiers [56]. Implications are potential outcomes of using a technology that an end-user may not anticipate

and are frequently ethical in nature. For example, an implication of using a webcam is that a user's video data is recorded and possibly transmitted to an unknown third party. We believe implication design is particularly useful in ubiquitous computing systems where the negative implications of the technology are often specific to a local context of use [42]. For example, cameras can be used both for connecting with friends and relatives over a video call or for covert surveillance and stalking.

We refer to artifacts that signify their ethical implications as *implication design artifacts*. Implication design artifacts make their implications more saliently perceptible to the user by embedding them directly into their design. For example, *Eyecam* [71], an anthropomorphic webcam shaped like a human eye, can be imagined to communicate the potential negative implications of being observed. An existing (but subtle) example of implication design in current consumer technology is the integration of green and orange indicator lights in the iPhone to communicate when the camera or the microphone is being used [1]. Contrasting the two examples, *Eyecam* more saliently communicates the potential implication of being observed through its anthropomorphic design. In *Dangling String* [76], possibly one of the earliest examples of implication design in digital technology, artist Natalie Jeremijenko created a ceiling mounted string that twitched everytime a bit of data moved through an ethernet cable at Xerox PARC.

Beyond signifying a technology's implications by *communicating* them to the user, implication design artifacts can also allow the user to *protect* themselves against those implications by making them tangible and embodied. For example, *Alias* [44], a teachable "parasite" that is designed to give users more control over their smart speakers by intercepting audio signals, can be considered an implication design artifact for protection.

By formally articulating the concept of implication design, we hope to highlight the ways in which this design approach can be used to mitigate potential harms when using and interacting with technology. In this work, we use implication design as a means to enable people with and without technical expertise or design experience to communicate how a technology might be changed to better serve them. This reflects our desire to bring a diversity of perspectives into the process of designing technology that better accounts for ethical implications. Shifting our focus "from artefacts to process" [30], we aimed to examine the role of implication design in supporting collective speculation and navigating value tensions.

3.2 Positionality

We are a diverse group of researchers with a wide range of experiences of power and marginalization. For example, we come from large and small countries in the Global North and the Global South, belong to religious, sexual, ethnic, and gender majorities and minorities, and some of us have disabilities. We have backgrounds in design, computer science, and architecture. We all share the privilege of being researchers at a prestigious academic institution in the US. Through our institutional affiliation, we had access to experts in computing ethics, ethics in design, and game design who offered us consultation in developing this workshop. Our research group (among other topics) researches and designs ubiquitous computing systems for applications in health, sustainability, and education.

This proximity to ubiquitous computing research informed our choice of technologies and scenarios to explore in our workshops.

This project was born out of our own concerns with the potential harms of ubiquitous computing systems, including our own research projects. Rather than working within the constraints of what is currently technologically feasible, we wanted to begin with imagining desired futures. However, given our diversity of interests and lived experiences, we knew that one person or group's desired future would not reflect a future desirable for all. This inspired us to explore how collective speculation might be applied towards probing the ethical implications of emerging technologies, and navigating the value tensions that inevitably arise in the process in a more inclusive manner.

3.3 Workshop Goals

Much like our iterative approach to the design of our workshop, our goals were iteratively defined and refined over the course of this project. The process of designing workshops and reflecting on participant interactions enabled us to identify three key goals:

1) Surface diverse values.

The workshop should surface a diversity of values such that the set of ethical implications considered is as broad as possible, which we believe contributes to a richer and more nuanced design space. Noting that values are grounded in lived experience, we set out to bring together participants with a diversity of lived experience. In addition, we believe that it is important to include participants with different levels and areas of domain expertise, who can provide a more nuanced understanding of the technological implications. Thus, we aim to create activities such that all participants are able to and feel comfortable surfacing a diversity of values that draw on their lived experiences.

2) Enable participants of different backgrounds to collectively navigate value tensions.

Beyond surfacing values, the workshop should enable participants to collectively navigate value tensions. Simply surfacing values results in a select group of people (i.e., the designers or developers of the system) deciding how to resolve the value tensions as well as which values and whose values to prioritize. Navigating tensions requires scaffolding to help participants make difficult decisions about design trade-offs and conflicts in stakeholder values. This scaffolding should also account for power dynamics amongst participants and work towards ensuring that everyone has an equal voice.

3) Reduce reliance on heavy facilitator involvement and design expertise.

While nearly all design workshops require some degree of facilitation, the workshop should not rely on facilitators having deep design expertise to ensure that the first two goals are met. Instead, the facilitator should be responsible for choosing a technology and scenario, recruiting participants, and creating a safe and comfortable environment for contribution. We hope that reducing reliance on facilitator involvement and expertise will help this method scale to many different settings.

4 WORKSHOP OVERVIEW

In this section of the paper, we provide an overview of the workshop, describe the goal of each activity, and offer some examples of the resulting artifacts from our workshops. The design of our workshop is the outcome of our iterative and reflective process. In Section 5, we describe in detail and reflect on the process that resulted in this final workshop design. We reflect on the outcomes of our final workshop in Section 6.1. A comprehensive facilitation guide and all workshop materials can be found at: <https://stanfordhci.github.io/collective-speculation>

Our final workshop is organized into the following sections, summarized in Table 1:

1) Introduction and Warm-Up

Facilitators provide an overview of the workshop, set ground rules and agreements for communications, and establish a safe space for contribution. Participants complete a short warm-up activity to increase their comfort with speculation and uncertainty.

2) Anticipation Round

In the anticipation round, participants are guided through a series of brainstorming activities to imagine contexts, stakeholders, and use cases for a given technology. Participants then explore positive and negative implications of a given technology for various stakeholders in a chosen context and use case. The output of this round is a deck of cards that is used in the action round.

3) Implication Design Round

In the implication design round, participants learn the concept of implication design and develop speculative design artifacts using negative implications from the anticipation round. This round provides scaffolded activities to help participants familiarize themselves with implication design.

4) Action Round

In the action round, participants use the deck of cards to engage in a roleplaying activity. During this activity, participants construct futures grounded in their lived experience and explore value tensions among stakeholders. Participants create speculative design artifacts to communicate or protect against potential ethical implications and collectively navigate design trade-offs.

4.1 Facilitator Preparation

Before the workshop, facilitators decide on a technology and usage scenario to explore based on the project and domain of interest. For example, we used the following scenario in our workshops:

A technology company is offering a new product on the market—a behavioral sensing system to be deployed in indoor spaces that adjusts environmental factors based on behavioral data in order to enhance productivity and wellbeing.

Facilitators are also responsible for recruiting participants for the workshop. For our final workshop, we recruited 15 participants with diverse identities and demographics as a proxy for diverse lived experiences and different levels and areas of domain expertise.

We used screening surveys with self-reported identity categories, demographic information, education level, and profession for recruitment, which we distributed through email lists and flyers in our area. More information about our recruitment is described in Section 5 and full details can be found in the facilitation guide.

4.2 Introduction and Warm-Up

To begin, the facilitators introduce themselves and create a safe and supportive environment, setting expectations for the workshop’s topics and presenting agreements for respectful communication. Then, all participants introduce themselves and engage in a short warm-up activity. In our workshops, we used a scenario similar to the main workshop scenario, but more provocative and overdramatized (the technology was an intelligent toilet) to incentivize creative speculation. Participants were prompted to design artifacts for this scenario subject to certain constraints (e.g., “you have to use magic”).

4.3 Anticipation Round

The objective of this round is to generate a deck of cards to use in the action round. This round has three main phases of individual and collective brainstorming: stakeholder and context brainstorming, use case brainstorming, and implication brainstorming.

4.3.1 Stakeholder and Context Brainstorming. The goal of this brainstorm is to generate many ideas for potential contexts where this technology might be deployed, (e.g., preschools and trains), direct stakeholders for each context, (e.g., children and conductors), and indirect stakeholders, (e.g., future children and cleaning staff). Participants first work individually, then in small groups of 2–3 people to brainstorm based on the scenario provided by the facilitators (see Figure 2). The round ends with a vote in larger, 4–6 person teams (two small teams merge) about which context and respective stakeholders to take on to the next phase of this round.

4.3.2 Use Case Brainstorming. The goal of this brainstorm is to generate many ideas for potential use cases within the selected context. We define a use case to be a scenario that benefits at least one stakeholder. The brainstorming is scaffolded using the following template:

For ____ who want to ____, the system can ____ by ____.

After brainstorming, each team selects three final use cases and writes them on use case cards from the action deck. For example, one team with *hospitals* as their context selected the use case: “*For doctors who want to reduce patient stress level, the system can improve that by adjusting the temperature, light, and humidity.*” We also provide participants with a handout that contains more information about the sensors the system might use, the kinds of behavioral data the system might collect, and the kinds of automatic adjustments the system might be able to make.

4.3.3 Implication Brainstorming. The goal of this phase is for participants to finalize the action deck by considering positive and negative implications of the use cases from the last round. We used the Tarot Cards of Tech [3] to prompt implication brainstorming, though facilitators can use any materials of their choosing. The Tarot Cards of Tech contain a number of prompts (e.g., “*Who or*

Phase	Duration (min)	Rationale	Resulting artifact(s)
Introduction	30	Set expectations for the workshop, establish a safe space for contribution.	None.
Warm-up	15–20	Increase comfort with speculation and uncertainty in poorly defined problems.	Ideas drawn on sticky notes or prototyped with craft materials.
Anticipation Round	45–60	Imagine contexts, stakeholders, and use cases for a given technology.	A deck of cards (stakeholder, context, use case, positive & negative implication).
Implication Design Round	30	Familiarize the concept of implication design, develop artifacts that embody implications.	Implication design sketches and/or physical prototypes.
Action Round	60–90	Navigate value tensions.	Future packs with implications, stakeholders, and implication design artifacts.

Table 1: Workshop overview.




Figure 2: A team’s context and stakeholder brainstorming board at the end of the activity.

what disappears if your product is successful? Who loses their job?” that can help brainstorm possible implications. For example, for the use case mentioned above, a positive implication that one team identified was that it would “create jobs for technology teams developing and maintaining the system,” while a negative implication was that it would “introduce a divide between patients who can (and cannot) afford the service.” After brainstorming, each team collectively selects six positive and six negative implications and writes them on their action deck implication cards. This activity completes the action deck. See Figure 3 for an example of a completed deck.

4.4 Implication Design Round

The goal of this round is to familiarize participants with the concept of implication design and developing artifacts that embody potential implications. Facilitators explain, motivate, and provide examples of implication design. Next, participants draw a negative implication card from their deck and receive a series of prompts to communicate

the negative implications: on the sensor, on the person’s body, and via a digital screen display. Then, they use prototyping materials (e.g., clay, pipe cleaners, construction paper, and foam blocks) to sketch out and build these ideas. See Figure 4 for examples of implication design artifacts produced in this round.

4.5 Action Round

The goal of this round is to navigate value tensions through a role-playing game based on the deck of cards created in the previous round. Participants use their action decks to explore the negative and positive implications of technology through an interactive roleplaying activity. Groups are shuffled to include some members from the team that generated the deck, and new members from other teams.

Figure 5 summarizes the action round steps. Participants set up the cards and reveal positive and negative implication cards. One group member, the moderator, situates those cards based on the



Figure 3: *Left*: All workshop materials (sensing system information sheet, blank card set, Tarot Cards of Tech [3]) at the beginning of activity. *Right*: The end of an action round with stakeholder distribution based on positive and negative implications, as well as an implication design artifact.



Figure 4: Examples of implication design artifacts produced in the implication design round. *Left*: An uncertainty communicator where the system communicates the system’s uncertainty. *Center*: An anonymity bracelet that anonymizes the wearer’s collected information. *Right*: A daycare robot that “watches over children” to better communicate the system’s interactions with the children

use case at hand by specifying a context they can anticipate from their own lived experiences. For example, in our workshops, one moderator situated a hospital context by adding, “a pregnant woman is going into labor and is extremely stressed.” Grounded in this context, each participant selects a stakeholder card and places it on either the positive or negative implication card. To choose which implication card to place their stakeholder card on, participants should ask themselves if they (as their chosen stakeholder) feel excited about the positive implication or concerned about the negative implication. Then, for two minutes, each participants thinks of an implication design to maintain the positive implications while communicating and/or protecting against the negative implications.

There are four rules for creating implication design artifacts: (1) you cannot make up new laws, but you can implement organizational policy, (2) all implications (positive or negative) have to be communicated through the design of an artifact, (3) you can use magic, (4) you cannot make the system unhackable.¹ We gave participants the following example to help clarify the rules:

Example: An intervention that proposes a new law that makes it illegal to fire employees whose jobs were lost to automation would not be allowed. However, one could propose that a particular company implements a policy that they will not fire anyone as a result

¹In the following section, we discuss how these rules were motivated and how they evolved over time.

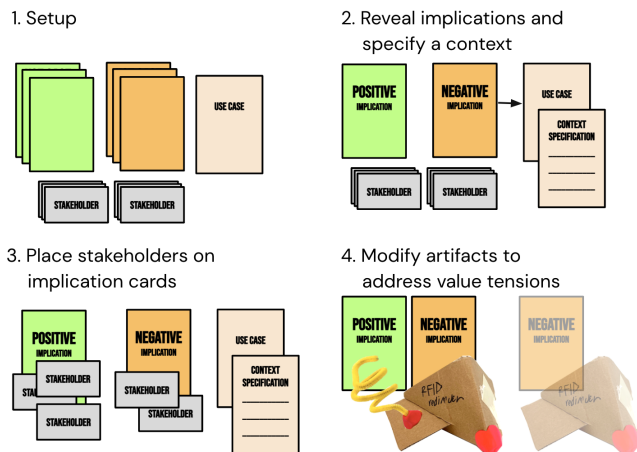


Figure 5: Action round summary. *Top left:* setup including a use case card, face down implication cards, and stakeholder cards; *Top right:* a positive and a negative implication card are revealed, and a blank card that the moderator fills in from their own lived experience is added to the use case card; *Bottom left:* stakeholder cards are placed on top of their corresponding implications cards by different participants; and *Bottom right:* shows a modified implication artifact with an antenna that addresses value tensions between stakeholders on the positive and negative implication cards. These steps repeat several times.

of implementing this technology. Be aware that this may cause stakeholders such as the CEO to remove themselves from the benefit card. Moreover, the risk of automation still needs to be communicated in the design of an artifact (e.g., a public board that displays the number of employees that have been laid off).

Then, participants individually create or modify implication design artifacts with craft materials to communicate or protect against negative implications while preserving positive implications and present them to the group. Participants are encouraged to “play” their speculative intervention by placing it on an implication card. The other participants can move their stakeholder cards if the intervention changes their perception of the implications. Participants can also propose modifications and build on each others’ interventions. We refer to the final sets of implications, stakeholders, and implication design artifacts as *future packs* (see Figure 6 for examples). Future packs are intended to communicate a more complete picture of the values that were brought up, value tensions that were discussed, and how the group navigated those value tensions. After each round, a new moderator is chosen and the activity repeats.

5 ITERATIVE WORKSHOP DEVELOPMENT AND METHODS

In this section, we document our iterative and reflective process in conducting five pilots and two full-length workshops ($N = 113$). An overview of the structure of each workshop iteration can be found in Table 2. We describe how the original design of the workshop

evolved over time to meet each of our goals (described in Section 3.3): (1) surface diverse values, (2) enable participants of diverse backgrounds to collectively navigate value tensions, and (3) reduce reliance on heavy facilitator involvement and design expertise. Before each session, we posed questions motivated by our goals and the outcomes of previous sessions, and after each session we reflected on the workshop outcomes. In Section 6, we reflect on the outcomes of our final workshop and on the lessons learned throughout this iterative process.

5.1 Pilots 1 and 2

The first two pilots were designed to test an initial variant of our workshops in which the action round was structured as an adversarial game. We aimed to understand how activity scaffolding shapes participant interactions as well as which existing tools and methods might be useful for this workshop.

5.1.1 Participants. The first pilot ($n = 4$) was conducted with undergraduate students in our personal networks and the second pilot ($n = 8$) with Human-Computer Interaction graduate and undergraduate researchers at our institution.

5.1.2 Materials. Sticky notes, Ethical OS toolkit [2], and Tarot Cards of Tech [3].

5.1.3 Procedure (90 minutes). Participants were divided into two teams: an exploitation team and an implication design team. Teams were presented a speculative scenario similar to that of our final workshops (see Section 4.1). The exploitation team was tasked with exploiting the scenario in order to maximize profits, and the implication team was set to protect against possible exploitations through embedding protections into an artifact. Participants were given Ethical OS cards [2] and the Tarot Cards of Tech [3] to help brainstorm possible exploits. After brainstorming, the teams faced each other in a battle round where the exploitation team introduced a threat, and the implication design team attempted to protect the user from that threat using the solutions they have come up with. Protected threats got a point for the implication team, unprotected threats got a point for the exploitation team, points were tallied, and the winning team was declared.

5.1.4 Reflections on Goals. Overall, the first two pilots were promising and encouraged us to pursue this direction of work. Most importantly, we were inspired by observing the teams collectively navigate value tensions and set out to better support these dynamics. We made several changes to surface more specific stakeholders, contexts, and implications in the anticipation round, as well as changes to the action round to encourage design-based interventions.

Goal 1: Surface diverse values. In analyzing the threats, we found that the anticipated implications lacked specificity in their context as well as in the stakeholders considered. For example, stakeholders would just be reduced to “people in the office,” and therefore no specific person or groups would be protected. Inspired by prior work in value-sensitive design [37, 55] and recent ethics-focused design activities [5, 79], we included a context and stakeholder brainstorming round in the later workshops. We also noticed that some scenarios lacked specificity, so in the subsequent workshop we provided the participants with threat anticipation cards (which



Figure 6: Two future packs crafted at the end of the action round, with stakeholders sorted, and interventions crafted based on the positive and negative implications of the use case. Left: A future pack for the use case: “For doctors who want to reduce patient stress level, the system improve that by adjusting temperature, light, or humidity.” The implication design intervention is a playground waiting room to reduce patient stress, which does not rely on personal sensing devices that some patients may not be able to afford. Right: A future pack for the use case: “For parents who want their children to be happy and develop emotional skills, the system can get them support by alerting the caregiver when a child’s mood is bad or they are stressed.” Interventions include a heart rate monitor and randomized intervention generator that limit dependence on the system, as well as an intelligent bot and large display that helps children reflect on their feelings.

later evolved into our use case template). Lastly, we found that the Tarot Cards of Tech [3] were easier for non-designers to understand and more successfully prompted participants to think of a broader set of implications than the Ethical OS toolkit [2], and therefore decided to use the Tarot Cards of Tech in our future workshops.

Goal 2: Collectively navigate value tensions. In reflecting on the workshop, we noticed a promising direction where the two teams were collectively navigating value tensions for specific stakeholder groups, instead of taking an absolute stance on the threats and benefits. For example, in response to the threat of lack of transparency into what data is collected, the implication design team devised a wall-mounted display showing all data being used by the system. When they presented this implication design item, the exploitation team responded that this can “leak information about who is struggling [with their wellbeing].” The implication design team responded that they would make the display adapt to aggregate data at a group level. Observing these back-and-forth exchanges inspired us to further explore how to better support collectively navigating value tensions in future workshops.

Goal 3: Reduce facilitator involvement. Encouragingly, we found that implication design was easy to understand for participants without design backgrounds. Participants came up with interventions such as the display showing all data being collected (mentioned above) and transparent face tattoos that anonymize data as it is being collected. However, participants were quick to jump to protections that were based on privacy policies, legislation, or computer security

fixes. In the next workshop, we opted to explicitly prohibit interventions that were based on policy or systems-level interventions (this restriction was removed in later workshops). Additionally, because participants were having difficulty determining winners, we introduced a scoring rubric.

5.2 Full Workshop 1

In the first full workshop, we made an effort to recruit participants that did not know each other and had different domain expertise, expertise levels, and lived experiences. With this pool of participants, our goal was to study the group dynamics that emerge in collaboration and methods to ensure equal participation.

5.2.1 Participants. We recruited ten participants through university mailing lists, word of mouth, and physical flyers posted at local businesses, a local community college, and community centers (e.g., libraries, art centers, and parks). We explicitly aimed to recruit participants with varying levels of domain expertise in a variety of areas, as well as with a diversity of lived experiences. We used the demographic information as a proxy for diversity of lived experiences and selected participants that were diverse along different self-disclosed identity and demographic categories, including: *gender* (six female, four male), *age* (four 18–24, three 25–34, one 35–44, one 55–64, and one 65+), *country called home* (two outside of the US), *sexual orientation* (six heterosexual or straight, two queer, one gay, one bisexual), *racial or ethnic identity* (five white, three Hispanic/Latinx, two Asian, one Middle Eastern, one Native American; participants could report several categories), *education level* (six

Iteration	Participants	Format	Materials	Procedure
Pilots 1 & 2	Friends ($n=4$), HCI group ($n=8$)	Brainstorming + Adversarial game	Sticky notes, Ethical OS [2], Tarot Cards of Tech [3]	Duration: 90 minutes <ol style="list-style-type: none"> 1. Anticipation round: brainstorm threats/exploits or implication design protections 2. Battle round: one team introduces threat, the other tries to protect using implication design
Workshop 1	Campus & local community recruitment ($n = 10$)	Brainstorming + Adversarial game	Sticky notes, threat anticipation template, Tarot cards of Tech [3], craft materials, scoring rubric	Duration: Six hours <ol style="list-style-type: none"> 1. Context, direct/indirect stakeholder brainstorming 2. Threat brainstorming using Tarot cards and threat anticipation template 3. Design implication design artifacts for three selected threats 4. Battle round: one team introduces threat, the other tries to protect using implication design
Pilots 3 & 4	Research project team ($n=8$), Computing ethics reading group ($n=20$)	Brainstorming	Sticky notes, Tarot cards of Tech [3]	Duration: 60 minutes <ol style="list-style-type: none"> 1. Context, direct/indirect stakeholder brainstorming 2. Use case brainstorming for a selected context 3. Implication brainstorming based on use cases
Pilot 5	Game design class ($n=48$)	Collaborative game	Deck of use case, threat, benefit, stakeholder, and power cards (created by research team using the outcomes of pilots 3 and 4), game rules	Duration: 60 minutes <ol style="list-style-type: none"> 1. Moderator draws use case, threat, and benefit card from deck 2. Each participant selects and places stakeholders on threat or benefit card 3. Participants brainstorm interventions that mitigate threats while preserving benefits, and others move their stakeholder cards in response 4. Moderator calls for another brainstorming round or ends the round, tallying score
Workshop 2	Campus & local community recruitment ($n=15$)	Brainstorming + Collaborative role-play activity	Sticky notes, Tarot cards of Tech [3], blank action deck, system information sheet, craft materials, evaluation rubric	Duration: Six hours See Section 4 for procedure

Table 2: Overview of our iterative workshop design process.

holding or pursuing bachelor’s degrees, one master’s, two PhD), and *employment status* (five students, one employed for wages, one unemployed, one retired). We also aimed for a diversity of domain expertise among expert participants (human-computer interaction, artificial intelligence, communications, education, and entertainment). Four participants were not affiliated with our institution.

5.2.2 Materials. Sticky notes (stakeholder and context brainstorming), threat anticipation template (anticipation round), Tarot Cards of Tech [3] (implication brainstorming), craft materials (implication design and action rounds), scoring rubric (action round).

5.2.3 Procedure (6 hours). After introductions, participants were split into teams of two and completed a warm-up activity, similar to that described in the final workshop (i.e., the intelligent toilet scenario). Participants brainstormed contexts of use, direct stakeholders, and indirect stakeholders [37, 55]. Then, each participant used the Tarot Cards of Tech [3] to imagine possible threats for different stakeholders and filled out a threat anticipation template (where is this technology deployed, who might take advantage of this technology, who is impacted, how are they impacted, and a description of a potential scenario).² Teams convened into larger groups of four and selected three threats. After a short lecture

²We referred to negative implications as threats, did not differentiate between implications and use cases, and called the anticipation the threat anticipation round in early versions of the workshop. We discuss how these changes evolved below.

introducing the concept of implication design, participants spent 30 minutes designing implication design artifacts for the three threats they selected. These artifacts were then taken to the battle round where each team tries to protect against other teams’ threats by using the implication design artifacts in a timed round where the teams have the ability to improve on their artifact. After each round, all participants used a scoring rubric to vote for the artifact and how well it protected the team on the following four dimensions: uniqueness, risk communication, threat protection, and intervention desirability. Participants also filled out a pre- and post-workshop survey.

5.2.4 Reflections on Goals. We received positive feedback from the participants and observed interesting discussions of value tensions in this first full workshop. However, we realized that the adversarial format did not lend itself well to collectively navigating value tensions, so we made several changes to discourage all or nothing stances and facilitate a more nuanced discussion.

Goal 1: Surface diverse values. First, we realized that the participants had to implicitly imagine both a reason to use the technology (i.e., a use case) and a negative implication of that use case when thinking of “threats.” In response, we added an explicit use case brainstorming session to future workshops. Moreover, while the threat anticipation templates were helpful in making threats more

specific, they resulted in the participants focusing too much on individual threats without first considering a broader set of implications. Therefore we opted to remove the threat anticipation template and instead added more scaffolding to use case brainstorming. Lastly, we also observed that the participants grew attached to the items they had brainstormed, leading to a protectiveness that seemed ungrounded in lived experience and a lack of cross-pollination of ideas across different teams. In response to this, we asked the participants to swap their use cases and decks in future workshops.

Goal 2: Collectively navigate value tensions. In evaluating our workshop with our intended pool of participants, our most important learning was that the adversarial format was not conducive to navigating value tensions. We found that the competitive nature of this format was unpopular among participants, particularly those who were already hesitant to speak up. While each team was able to improve their own artifact, the adversarial dynamic made it such that teams had no incentive to constructively improve upon each other's designs. Moreover, the adversarial format encouraged implication design artifacts that often eliminated imagined benefits to protect against threats. One participant noted that they wished there was "less competition between groups, did not feel as productive as collaborative constructive criticism." We therefore spent the next months developing and testing various iterations of the game to make it more collaborative, allow teams to constructively build on each other's ideas, and consider positive and negative implications in relation to each other and the stakeholders they affect.

Goal 3: Reduce facilitator involvement. We observed that participants were mostly able to complete the workshop activities independently, though facilitators occasionally needed to intervene when instructions were unclear. As noted above, we broke down the activities into step-by-step processes in response. This additional scaffolding helped shape activity outcomes to address our goals, but also made the activities easier to complete independently.

5.3 Pilot 3 and Pilot 4

Pilots 3 and 4 focused on the anticipation round. Our goal was to iterate on the learnings from the full workshop to create a streamlined, scaffolded process for the anticipation round.

5.3.1 Participants. Pilot 3 was conducted with a research group studying smart spaces ($n=8$). Pilot 4 was conducted with a graduate student computing ethics reading group at our institution ($n=20$). We did not collect demographic information.

5.3.2 Materials. Sticky notes (stakeholder and context brainstorming), Tarot cards of Tech [3] (implication brainstorming).

5.3.3 Procedure (60 minutes). Participants brainstormed contexts, indirect, and direct stakeholders in teams of 3-4. Each team then voted to pick one context and brainstorm possible use cases for that context. Then, the groups brainstormed on implications that may arise based on different use cases.

5.3.4 Reflections on Goals. While our new scaffolding for the anticipation round proved to be effective, we made a few changes to ensure consistency and clarity across the different components of

the anticipation round. We do not reflect on Goal 2 as this pilot did not include the action round.

Goal 1: Surface diverse values. Having removed the detailed threat anticipation template, we found that the use cases generated in these pilots were either too specific or too generic at times. Moreover, we expected that use cases would inherently map to positive implications. Instead, we discovered that participants observed use cases to be positive for some users and negative for others. For the future workshops, we clarified that the use case should be a reason why at least one stakeholder wants to use a piece of technology in a given scenario and introduced threat and benefit cards (which we call positive and negative implication cards in our final workshop).

Goal 3: Reduce facilitator involvement. In response to the changes above and to streamline the brainstorming process, we opted for a lighter-weight, fill-in-the-blank use case template (*For — who want to —, the system can — by —*).

5.4 Pilot 5

This pilot tested the action round only. Prior to this pilot, we had play-tested different action round formats (e.g., role-playing with stakeholder roles) that inspired us to pursue a game-based format for the action round. The primary goals of the pilot were to understand if a game structure can better support collectively navigating value tensions.

5.4.1 Participants. Pilot 5 took place in a game design class with 48 undergraduate students during one of the playtest sessions in their class. We did not collect demographic information.

5.4.2 Materials. We created a deck of cards for the action round using the outcomes of pilots 3 and 4. The use cases were modified to be neutral and we added benefit cards in addition to the existing threat cards. Therefore the deck included the following items:

- 1 use case card
- 6 prefilled threat cards and 4 blank threat cards
- 6 prefilled benefit cards and 4 blank benefit cards
- 15 prefilled stakeholder cards and 10 blank stakeholder cards
- 4 power cards (constraints to spark creativity), e.g., "You are temporarily granted immunity and can break laws or policies."

We also provided each team with a printed copy of the game rules.

5.4.3 Procedure (60 minutes). Participants were divided into teams of 6–10 and sat around a table. The team assigned a moderator who laid out the use case and drew one benefit and one threat for that use case from the relevant pile. Each participant selected a stakeholder card and placed the stakeholder card either on the benefit or threat card and explained the reasoning for their choice. For two minutes, all participants brainstormed implication design interventions that would remove/minimize the threat while preserving the benefits for as many stakeholders. Each participant shared their intervention, other participants could move their stakeholder cards in response, and everyone had a chance to improve upon the interventions collectively. The moderator could choose to call for another brainstorming session and use the "power" cards as needed, or end the round and tally the score. Positive points were assigned for each stakeholder on a benefit card and negative points were assigned for each stakeholder on a threat card.

5.4.4 Reflections on Goals. This pilot confirmed that a collaborative game structure supported independent navigation of value tensions by the participants, and facilitated a rich engagement with the surfaced values and stakeholder perspectives. Despite this, the point system lent itself to a strong focus on winning or losing which was not aligned with our goals. Moreover, the rules for implication design seemed to be unclear to some participants. As a result, we made a few changes that we discuss below. We do not reflect on Goal 1 as this pilot did not include the anticipation round.

Goal 2: Collectively navigate value tensions. Notably, our change to a more collaborative game structure enabled a more fluid discussion around value tensions. However, we noticed that framing the activity as a game led to participants focusing on winning or losing. Our intention was not to gamify ethics, but instead to use the game-like structure to create space for critical reflection and discussions of value tensions. As a result, we removed the points system altogether.

Moreover, we modified our language around threats and benefits after realizing that this terminology was itself value-laden: not all negative implications are threats and not all positive implications are benefits. We revised this terminology for our final workshop.

Goal 3: Reduce facilitator involvement. The participants were also frustrated by the activity's vaguely defined rules, which led to them asking the facilitators for guidance on what interventions met the criteria. For instance, participants did not understand why they could not employ policy interventions, especially in the absence of alternatives. Therefore in the final workshop, we decided to allow policy and systems level interventions, but asked them to imagine how the policy or systems level changes might be communicated through design.

Another challenge was that the scenarios did not necessarily resonate with the moderator leading that round, which resulted in participants asking for facilitator guidance. In the final workshop, we asked the moderators to situate the generic context into a specific one they are familiar with from their lived experience. This would enable the other team members to refer to the moderator if they have questions regarding their perception of the context.

5.5 Full Workshop 2 (Final Workshop)

Based on the learnings from the previous six workshops and pilots, we conducted our final workshop to validate the final design when conducted with our intended participant population.

5.5.1 Participants. The recruitment followed the same procedure as workshop 1 described in 5.2.2 (and described in full detail in the facilitation guide), with an increased effort to recruit outside of our university's community. We recruited 15 participants from a range of identity and demographic categories, including: *gender identity* (eight cis women, five cis men, one trans man, one genderfluid person), *age* (twelve 18–24, one 25–34, one 45–54), *country called home* (five outside of US), *sexual orientation* (eight heterosexual or straight, three queer, two gay, one bisexual), *racial or ethnic identity* (six white, two African American/Black, two East Asian, two Hispanic/Latinx, two Southeast Asian, two South Asian, one Middle Eastern, one Native American), *education level* (one high school graduate, nine with some college credit, two bachelor's,

two master's, two Ph.D.), and *employment status* (eleven students, two employed for wages, one self-employed, one unemployed). We also aimed for a diversity of domain expertise among expert participants (two held degrees in design, one in education and product management, one in philosophy). Seven participants were not affiliated with our institution.

5.5.2 Materials. Sticky notes (stakeholder and context brainstorming), Tarot cards of Tech [3] (implication brainstorming), blank action deck (anticipation and action rounds), system information sheet, crafts materials (implication design and action rounds), evaluation rubric (action round).

5.5.3 Procedure (6 hours). The procedure is nearly identical to the final format described in Section 4, with minor changes to the structure discussed below.

5.5.4 Reflections on Goals. The final workshop largely met all of our goals. In the following section (Section 6.1), we reflect in more detail on the outcomes of this workshop. Additionally, we made one change to the final workshop structure that we discuss below.

Goal 3: Reduce facilitator involvement. We made one minor change to our workshop's design (as presented in Section 4). In the first full Workshop, there was little cross-pollination of ideas across different teams. In response to this, we asked the participants to swap their use-cases and decks at two stages in the final workshop. However, swapping decks without having the context for the new deck was disruptive for some members. Based on this feedback, we modified the final workshop protocol in Section 4 and the facilitation guide to suggest shuffling the teams instead of swapping the entire deck.

6 REFLECTIONS

In this section, we share our reflections on the final workshop outcomes, as well as the learnings from the iterative process that we hope can provide useful insights for future workshop designers.

6.1 Reflections on Final Workshop Outcomes

In this section, we share an overview of the workshop outcomes and reflect on our workshop's ability to meet each of our goals (Section 3.3). In reflecting on the outcomes, we shift our focus "from artefacts to process" [30], measuring "success" based on the process by which the participants collectively navigated value tensions and iterated through complex trade-offs, rather than evaluating the quality of any one artifact.

6.1.1 Surfacing Diverse Values. In the final workshop, 52 different contexts, 50 direct stakeholders, and 34 indirect stakeholders were identified. Some examples of the identified contexts in our final workshop were *preschools, spaceships, factories, orchestra, trains, ice cream store, airplane cockpit, barracks, exam hall, and museums*. Direct stakeholders included *conductors, pilots, stockers, cashiers, shoppers, people exercising, and people reading*, while indirect stakeholders included *test proctors, librarians, cleaning staff, repairs/technicians, animals, pets, fans, personal aides, caregivers, and tourists*. We also observed ideas cross-pollinating from stakeholders to contexts, e.g., *orchestra* (context) inspiring *conductor* (stakeholder), then leading to *train station* (context). Beyond identifying, participants *grounded* the stakeholders, contexts, and use

cases they imagined in their personal experiences. For example, in response to a negative implication where the sensing system might misgender users, one participant noted that “for stakeholders we will have to say ‘cleaning staff’ and ‘trans cleaning staff,’” with another teammate adding in that they “might need sexual minorities” as additional stakeholders.

6.1.2 Navigating Value Tensions. We wanted to enable participants of diverse backgrounds, lived experiences, and expertise domains and levels to feel comfortable voicing their opinions in collectively navigating value tensions. Overall, participants found the workshop to be fun and enjoyable based on our anonymous post-workshop feedback survey. The three most common words used to describe the workshop were *fun* (6), *creative* (5), and *interesting* (3) and participants rated the workshop 4.1/5 (1: strongly disliked, 5: strongly liked) on average. Participants also reported being able to effectively collaborate with their teammates, regardless of prior design or domain experience. Participants had an average score of 4.6/5 regarding how effectively their team worked together. Participants with and without an advanced degree reported feeling comfortable voicing their opinions during the workshop (5 for people with bachelor’s or higher, 4.4 for people without a bachelor’s degree). One participant reported that they enjoyed “learning from the creativity of others.” Another participant noted that “saying anything that came to mind and being silly” and “pushing back against ideas to make them stronger” as a factor for success on their team.

We observed that teams were able to utilize our workshop scaffolding and implication design as tools to collectively navigate value tensions. For example, one team with “hospital” as their context had the following use case: “For doctors who want to reduce patient stress level, the system improve[s] that by adjusting temperature, light, and humidity.” The moderator made the context more specific by saying: “I took a covid test, I will be anxious before getting the result.” The team originally came up with a “personalized watch” to individualize the care. However, when considering the negative implication “divide between patients who can (not) afford the service,” they noticed that not every patient will be able to afford a personalized watch. To address this tension, they created a playground waiting room that could maintain the benefit of reducing stress for all patients while addressing the potential divide (see Figure 6, left).

6.1.3 Reduce Reliance on Facilitators. In the final workshop, we strived to minimally intervene as facilitators, and observed that participants were able to complete each of the workshop activities without direct facilitator guidance. Thus, our roles as facilitators focused on creating a safe space and clarifying activity instructions when necessary. During our final workshop, participants relied on the tools provided to them as well as their teammates to resolve any issues. When asked whether they got stuck and how they got unstuck, the participants overwhelmingly mentioned collaborating with and relying on their teammates as a method to get unstuck. Despite the minimal facilitator intervention, all of the participants completed the brainstorming activities and all of the teams developed several implication design artifacts.

6.2 Considerations for Future Workshops

Beyond sharing the final design of our workshop with the broader community, in this section we share our learnings from the iterative process that informed the final design of the workshop. We hope that other researchers and practitioners can use these considerations in developing their own methods.

6.2.1 Specificity, Breadth, and Speed. Throughout our all of our workshops, we found that the most intriguing and creative implication designs emerged from the most specific scenario descriptions. Early iterations did not make a clear enough distinction between a context (a physical location or social arrangement), a scenario (a narrative that relates and gives life to contexts and stakeholders), and a use case (a reason why one stakeholder wants to use a piece of technology in a given scenario), which resulted in vague, non-specific scenarios and implications in subsequent activities. Moreover, increasing the specificity of scenario descriptions allowed participants to draw from their own personal experiences. However, prompting participants to brainstorm highly specific scenarios from the outset is challenging. It requires much less effort to come up with general stakeholder or context categories (e.g., doctors and hospitals) than highly specific scenarios. Moreover, having every context be made specific is too time-consuming for a six-hour workshop. Thus, our workshop design functions as a “funnel” of ideas, in which stakeholder, context, and use case brainstorming is scaffolded and progressively made more specific as the workshop progresses.

6.2.2 Utopias and Dystopias. In our first pilots, participants would identify mostly negative, oftentimes dystopian scenarios. We sought to discourage hyperbolic scenarios because they “muddle the banality of more probable outcomes” [78, p. 1368] and detract from discussions about value tensions. Consequently, we restructured activities such that the discussions were centered around how positive and negative implications are experienced by different stakeholders and how implications stand in relation to or in tension with one another, which is common in prior work [5, 79]. Note that this does not reflect a belief that all technology has inherent benefits, but merely that technological harms and benefits must be considered in relation to one another. We also refrained from referring to negative implications as “threats,” “risks,” or “exploits,” to discourage participants from only envisioning extreme scenarios.

6.2.3 Moving Beyond Policy or Computer Security Fixes. We found that participants were quick to jump to protections that were based on privacy policies, legislation, or computer security fixes. These kinds of protections tended to encourage solutions-focused thinking and detracted from our goal of surfacing value tensions and exploring how the *design* of products ought to change to better communicate or protect against potential implications. Thus, we initially prohibited interventions that were based on policy or systems-level interventions (e.g., “You may not invent new laws,” “You may not make any additional modifications to the system,” or “You may not assume that the system is unhackable”). These restrictions were widely unpopular (even among designers) and did not help participants come up with implication designs. In response, we tried to make the rules less restrictive and used “You can use magic” for

positive encouragement. Additionally, we emphasized that implications must be communicated even for policy interventions and that policy interventions can also impact stakeholders dynamics.

6.2.4 Adversarial vs. Collaborative. Our first designs of the workshop featured an adversarial action round in which participants were divided into teams that were tasked with exploiting a scenario or protecting against harms. This approach bears some similarity to redteaming [50, 80], a method in computer security. While we were inspired by the adversarial method's ability to naturally surface discussions around value tensions, we found that the competitive nature of this format was unpopular among some participants (particularly those were already hesitant to speak up) and led to disappointment when there were no rules to decide which team had “won” at the end of activity. Moreover, this format led to designs with “all-or-nothing” protections that eliminated benefits so as to protect against the broadest range of threats. We briefly experimented with rubrics to address the issue of winning, but found that rubrics encouraged solutions-focused thinking, rather than exploring how different solutions shift value tensions among stakeholders. Ultimately, we found that an adversarial format did not align with our design goals and shifted to a collaborative format.

7 LIMITATIONS AND FUTURE WORK

In this section we discuss the limitations of our method and the opportunities for future work.

7.1 Organizing and Facilitating

To make the workshop available to a broad audience, one of our main goals was to make the workshop lightweight and easy to facilitate. While the outcomes of our final workshops confirm that we met this goal, it is likely that our design and domain expertise shaped the workshop dynamics and outcomes. Though we strived to minimally intervene in participant interactions (particularly during the final workshop), further work is required to test how easily people without a design background can facilitate our workshop.

Beyond expertise in design, we also note that factors such as recruitment and setting a safe space can form barriers to facilitating our workshop. Recruiting a diverse pool of participants is a considerable challenge, requiring significant time and resources. Moreover, when working with participants from marginalized groups, additional expertise and care is required to maintain a safe space for all participants. We also acknowledge that significant labor is required to initiate conversations about ethics in organizations, sustain these conversations over time, and advocate for changes in the design of particular technologies. Prior work has defined this form of labor as *values advocacy* [66] and much like *Timelines* [79], our method can serve as a tool to assist values advocates in their work.

7.2 Workshop Recruitment

For our full-length workshops, we worked to recruit participants from a variety of demographic backgrounds in order to draw from a diversity of lived experience. Despite our efforts, we acknowledge several limitations in our recruitment. First, we acknowledge that lived experience is impossible to quantify and that recruiting participants among self-reported identities and demographics is an imperfect, albeit useful proxy. Second, our recruitment was

constrained by our geographic location and influenced by our institutional affiliation. All of our participants were required to be fluent in English, all of our participants lived in the same geographic area, and many of our participants were students affiliated with our university. These recruitment limitations impacted the diversity of experience we were able to capture in our workshops because institutional affiliation and geographic location are dimensions of power and privilege. Lastly, our recruitment was affected by the COVID pandemic since our workshop requires in-person attendance. It is likely that many potential participants did not enroll due to COVID concerns (our workshops were conducted in 2021–2022), in particular older participants and those not affiliated with our institution.

As such, we are interested in sharing our resources with outside facilitators and are actively pursuing partnerships with community organizations. For example, we are looking into running this workshop with older adults, and may partner up with local non-profit organizations serving older adults to do so. We hope that this continuation of our work will make our method even more accessible, with the goal of allowing organizations to initiate and organize workshops without any researcher involvement.

7.3 Putting the Workshop Outcomes to Work

We are most excited about the potential of our method to facilitate sustained discussions with communities about the development of emerging technologies. While the goal of this work was to develop a method that enabled participants of different backgrounds to collectively navigate value tensions, this leaves open the question of how the outcomes of our workshop might be put to practice which we hope to study further in future work.

In our experience so far, we found two avenues through which this workshop has informed our own research on smart spaces, which may be indicative of how the knowledge produced in our workshop might integrate into larger ethics-focused design efforts. First, the discussions generated amongst the participants during the action round drew on rich lived experiences and added nuanced insight into values that were surfaced using traditional VSD methods. Therefore, our method could be used as a supplement to traditional value elicitation methods by foregrounding the *tensions* among them. Second, the implication design artifacts produced by the participants can be used to understand and probe user attitudes and expectations with regards to a technology. In the next section, we discuss the ways in which we think implication design artifacts can be used in future work.

7.4 Future Directions for Implication Design

As discussed in Section 3.1, we utilize implication design as a language for participants with and without design expertise to communicate how a technology might change to better serve them. Therefore, we did not focus on evaluating how well implication design artifacts communicated or protected against potential implications, but rather used them as probes to understand participant needs and values.

However, we also believe that implication design has potential as a human-centered complement to policy- and computer security-based solutions in real technology products. Informed by the preliminary implication design artifacts produced in our workshops, we hope to further explore implication design with professional designers and design researchers in the context of real ubiquitous computing products. In future iterations, we also wish to explore how different prompts contribute to the type of solutions imagined. For example in our final workshop, we asked the participants to imagine the artifacts on the sensor, on the person's body, and via a digital screen display, which prompted individual-level considerations. We are interested in exploring prompts that may lend themselves to community-, society-, or more-than-human- level considerations. Lastly, designers might also use semi-functional, speculative implication design artifacts as in situ probes, exploring when and in which contexts people may wish to employ these artifacts to further understand the underlying values and value tensions.

7.5 Beyond Sensing and UbiComp Applications

Our workshops explored the ethical implications of sensing and ubiquitous computing technology, which reflects our own interests as researchers in this space. Ubiquitous computing technologies foreground a wide range of ethical implications, particularly due to their prevalence across a variety of personal and social contexts (e.g., home, work, school, or medical) and reliance on pervasive data collection and inference algorithms. However, there are many other technologies with ethical implications in various domains (e.g., gene editing, energy technology) and further work is required to assess the extent to which the activities in our workshop apply across other technological domains.

Moreover, our workshop was grounded in our motivation as designers and researchers to explore how the design of products can be changed to better account for ethical implications. However, changing the design of a product is only one way to address ethical implications. Other approaches include legislation and policy proposals, advocacy and activism, and computer security. We note that our action round is not strictly tied to implication design and could be modified to support different kinds of interventions. For instance, instead of proposing implication designs, participants could propose policy solutions to address positive and negative implications. Such an approach will likely require different scaffolding and we are interested in its potential for future extensions of our workshop.

7.6 Ethics Education

Many participants shared that they learned a lot about ethics in technology during the workshop and left with a new perspective or changes in mindset. As a next step, we are also interested in adapting components of the workshop into an education module that can be integrated into undergraduate courses following initiatives such as Embedded EthiCS [39]. Such an adaptation may also allow the workshop to more easily be used as a design intervention for initiating conversations about ethics and values across a variety of organizations or companies.

8 CONCLUSION

This paper presents a workshop-based method for collective speculation on the ethical implications of technological futures. In developing our workshop, we followed an iterative process, taking inspiration from theoretical perspectives that center values, ethics, and participation in design research, as well as speculative and futuring methods, game design, and roleplay. We report on the learnings from our iterative process, conducting seven iterations of our workshop with 113 participants. Through sharing our workshop materials and learnings, we hope to make this method adaptable and easy-to-facilitate across a variety of settings.

ACKNOWLEDGMENTS

We would like to thank the many individuals and groups who provided us with invaluable expertise and guidance for this project: Beleicia Bullock, Sarah Billington, Kathleen Creel, Ashlyn Jaeger, Jennifer King, Jingyi Li, Tianshi Li, Daphne Martschenko, Nick Merrill, Anne Newman, Tonya Nguyen, Katharina Reinecke, Jacob Ritchie, Christina Wodtke, Richmond Wong, Stanford HCI, IxD, and Computing & Society Groups, the HPDS team, and the anonymous reviewers. Most importantly, we extend our gratitude to the many workshop participants, without whom this work would not have been possible.

We are grateful for the financial support provided by the Ethics, Society, and Technology Hub at the McCoy Family Center for Ethics in Society, the Hasso Plattner Design Thinking Research Program, and the Stanford Center for Integrated Facility Engineering. This research was reviewed and approved by the institutional review board at Stanford University under IRB Protocol #IRB-60989.

REFERENCES

- [1] [n. d.]. About the orange and green indicators in your iPhone status bar. <https://web.archive.org/web/20230111170228/https://support.apple.com/en-us/HT211876>. Accessed: 2023-02-11.
- [2] [n. d.]. Ethical OS Toolkit. <https://ethicalos.org>
- [3] [n. d.]. The Tarot Cards of Tech. <http://tarotcardsoftech.artefactgroup.com/>
- [4] James Auger. 2013. Speculative design: crafting the speculation. *Digital Creativity* 24, 1 (2013), 11–35.
- [5] Stephanie Ballard, Karen M Chappell, and Kristen Kennedy. 2019. Judgment call the game: Using value sensitive design and design fiction to surface ethical concerns related to technology. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 421–433.
- [6] Karl Baumann, Ben Caldwell, François Bar, and Benjamin Stokes. 2018. Participatory design fiction: community storytelling for speculative urban technologies. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–1.
- [7] Rachel KE Bellamy, Kuntal Dey, Michael Hind, Samuel C Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilović, et al. 2019. AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development* 63, 4/5 (2019), 4–1.
- [8] Michael S Bernstein, Margaret Levi, David Magnus, Betsy A Rajala, Debra Satz, and Charla Waeiss. 2021. Ethics and society review: Ethics reflection as a precondition to research funding. *Proceedings of the National Academy of Sciences* 118, 52 (2021).
- [9] Julian Bleeker. 2009. Design Fiction: A Short Essay on Design, Science, Fact and Fiction. <https://blog.nearfuturelaboratory.com/2009/03/17/design-fiction-a-short-essay-on-design-science-fact-and-fiction/>
- [10] Susanne Bødker, Kaj Grønnebæk, and Morten Kyng. 1995. Cooperative design: techniques and experiences from the Scandinavian scene. In *Readings in Human-Computer Interaction*. Elsevier, 215–224.
- [11] Kirsten Bray and Christina Harrington. 2021. Speculative blackness: considering Afrofuturism in the creation of inclusive speculative design probes. In *Designing Interactive Systems Conference 2021*. 1793–1806.

- [12] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [13] Derya Ozelik Buskermolen and Jacques Terken. 2012. Co-constructing stories: a participatory design technique to elicit in-depth user feedback and suggestions about design concepts. In *Proceedings of the 12th Participatory Design Conference: Exploratory Papers, Workshop Descriptions, Industry Cases-Volume 2*. 33–36.
- [14] Fatoumata Camara and Gaëlle Calvary. 2015. Worth-centered design in practice: lessons from experience and research agenda. In *IFIP Conference on Human-Computer Interaction*. Springer, 123–139.
- [15] EunJeong Cheon, Stephen Tsung-Han Sher, Šelma Sabanović, and Norman Makoto Su. 2019. I beg to differ: Soft conflicts in collaborative design using design fictions. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 201–214.
- [16] Paul Coulton. 2015. The role of game design in addressing behavioural change. In *Proceedings for 11th European Academy of Design Conference*.
- [17] Paul Coulton, Dan Burnett, and Adrian Ioan Gradinar. 2016. Games as speculative design: allowing players to consider alternate presents and plausible futures. (2016).
- [18] Luiza Prado de O. Martins and Pedro JS Vieira de Oliveira. 2016. Breaking the cycle of Macondo: design and decolonial futures. *XRDS: Crossroads, The ACM Magazine for Students* 22, 4 (2016), 28–32.
- [19] Pedro JS Vieira de Oliveira. 2016. Design at the earview: decolonizing speculative design through sonic fiction. *Design Issues* 32, 2 (2016), 43–52.
- [20] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2021. Stakeholder Participation in AI: Beyond "Add Diverse Stakeholders and Stir". *arXiv preprint arXiv:2111.01122* (2021).
- [21] Ansgar E Depping and Regan L Mandryk. 2017. Cooperation and interdependence: How multiplayer games increase social closeness. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*. 449–461.
- [22] Audrey Desjardins, Cayla Key, Heidi R Biggs, and Kelsey Aschenbeck. 2019. Bespoke booklets: A method for situated co-speculation. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 697–709.
- [23] Carl DiSalvo. 2012. Adversarial Design.
- [24] Carl DiSalvo, Tom Jenkins, and Thomas Lodato. 2016. Designing speculative civics. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 4979–4990.
- [25] Charles Duhigg. 2012. How companies learn your secrets. <https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>
- [26] Anthony Dunne. 2008. *Hertzian tales: Electronic products, aesthetic experience, and critical design*. MIT press.
- [27] Anthony Dunne and Fiona Raby. 2013. *Speculative everything: design, fiction, and social dreaming*. MIT press.
- [28] Felix Anand Epp, Tim Moesgen, Antti Salovaara, Emmi Pouta, and Idil Gaziulusoy. 2022. Reinventing the Wheel: The Future Ripples Method for Activating Anticipatory Capacities in Innovation Teams. In *Designing Interactive Systems Conference*. 387–399.
- [29] Gregory Falco, Ben Shneiderman, Julia Badger, Ryan Carrier, Anton Dahbura, David Danks, Martin Eling, Alwyn Goodloe, Jerry Gupta, Christopher Hart, et al. 2021. Governing AI safety through independent audits. *Nature Machine Intelligence* 3, 7 (2021), 566–571.
- [30] Pedro Gil Farias, Roy Bendor, and Bregje F Van Eekelen. 2022. Social dreaming together: A critical exploration of participatory speculative design. In *Proceedings of the Participatory Design Conference 2022-Volume 2*. 147–154.
- [31] Mary Flanagan, Daniel C Howe, and Helen Nissenbaum. 2005. Values at play: Design tradeoffs in socially-oriented game design. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 751–760.
- [32] Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. 2018. "A Stalker's Paradise" How Intimate Partner Abusers Exploit Technology. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–13.
- [33] Batya Friedman. 1996. Value-sensitive design. *Interactions* 3, 6 (1996), 16–23.
- [34] Batya Friedman and David Hendry. 2012. The envisioning cards: a toolkit for catalyzing humanistic and technical imaginations. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1145–1148.
- [35] Batya Friedman and David G Hendry. 2019. *Value sensitive design: Shaping technology with moral imagination*. MIT Press.
- [36] Batya Friedman, Peter Kahn, and Alan Borning. 2002. Value sensitive design: Theory and methods. *University of Washington technical report 2-12* (2002).
- [37] Batya Friedman, Peter H Kahn Jr, Jennifer Hagman, Rachel L Severson, and Brian Gill. 2006. The watcher and the watched: Social judgments about privacy in a public place. *Human-Computer Interaction* 21, 2 (2006), 235–272.
- [38] Dave Gray, Sunni Brown, and James Macanujo. 2010. *Gamestorming: A playbook for innovators, rulebreakers, and changemakers*. " O'Reilly Media, Inc."
- [39] Barbara J Grosz, David Gray Grant, Kate Vredenburg, Jeff Behrends, Lily Hu, Alison Simmons, and Jim Waldo. 2019. Embedded EthiCS: integrating ethics across CS education. *Commun. ACM* 62, 8 (2019), 54–61.
- [40] Christina Harrington and Tawanna R Dillahunt. 2021. Eliciting tech futures among Black young adults: A case study of remote speculative co-design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [41] Kashmir Hill. 2020. The secretive company that might end privacy as we know it. <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html>
- [42] Jason I Hong, Jennifer D Ng, Scott Lederer, and James A Landay. 2004. Privacy risk models for designing privacy-sensitive ubiquitous computing systems. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques*. 91–100.
- [43] Nassim JafariNaimi, Lisa Nathan, and Ian Hargraves. 2015. Values as hypotheses: design, inquiry, and the service of values. *Design issues* 31, 4 (2015), 91–104.
- [44] Bjørn Karmann and Tore Knudsen. 2018. Project Alias. https://web.archive.org/web/20221202203137/https://bjoernkarmann.dk/project/project_alias. Accessed: 2022-12-02.
- [45] Cory Knobel and Geoffrey C Bowker. 2011. Values in design. *Commun. ACM* 54, 7 (2011), 26–28.
- [46] Sandjar Kozubae, Chris Elsdén, Noura Howell, Marie Louise Juul Søndergaard, Nick Merrill, Britta Schulte, and Richmond Y Wong. 2020. Expanding Modes of Reflection in Design Futuring. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [47] Fedwa Laamarti, Mohamad Eid, and Abdulmotaleb El Saddik. 2014. An overview of serious games. *International Journal of Computer Games Technology* 2014 (2014).
- [48] Christopher A Le Dantec, Erika Shehan Poole, and Susan P Wyche. 2009. Values as lived experience: evolving value sensitive design in support of value discovery. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1141–1150.
- [49] Ilaria Mariani. 2020. Other Worlds. When Worldbuilding and Roleplay Feed Speculation. In *International Conference on Human-Computer Interaction*. Springer, 482–495.
- [50] Nick Merrill. 2020. Security Fictions: Bridging Speculative Design and Computer Security. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 1727–1735.
- [51] Jessica K Miller, Batya Friedman, Gavin Jancke, and Brian Gill. 2007. Value tensions in design: the value sensitive design, development, and appropriation of a corporation's groupware system. In *Proceedings of the 2007 international ACM conference on Supporting group work*. 281–290.
- [52] Michael J Muller. 2007. *Participatory design: the third space in HCI*. CRC press.
- [53] Michael J Muller and Allison Druin. 2012. Participatory design: the third space in human-computer interaction. In *The Human-Computer Interaction Handbook*. CRC Press, 1125–1153.
- [54] Michael J Muller, Daniel M Wildman, and Ellen A White. 1994. Participatory design through games and other group exercises. In *Conference companion on Human factors in computing systems*. 411–412.
- [55] Lisa P Nathan, Batya Friedman, Predrag Klasnja, Shaun K Kane, and Jessica K Miller. 2008. Envisioning systemic effects on persons and society throughout interactive system design. In *Proceedings of the 7th ACM conference on Designing interactive systems*. 1–10.
- [56] Donald Norman. 2008. The way I see it - Signifiers, not affordances. *Interactions* 15 (11 2008), 18–19. <https://doi.org/10.1145/1409040.1409044>
- [57] Pablo Ortiz and D. Fox Harrell. 2018. Enabling Critical Self-Reflection through Roleplay with Chimeria: Grayscale. In *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play (Melbourne, VIC, Australia) (CHI PLAY '18)*. Association for Computing Machinery, New York, NY, USA, 353–364. <https://doi.org/10.1145/3242671.3242687>
- [58] Luiza Prado and Pedro Oliveira. 2017. Questioning the "critical" in speculative & critical design. <https://medium.com/a-parede/questioning-the-critical-in-speculative-critical-design-5a355cac2ca4>
- [59] Horst WJ Rittel and Melvin M Webber. 1973. Dilemmas in a general theory of planning. *Policy sciences* 4, 2 (1973), 155–169.
- [60] Luc Rocher, Julien M Hendrickx, and Yves-Alexandre De Montjoye. 2019. Estimating the success of re-identifications in incomplete datasets using generative models. *Nature communications* 10, 1 (2019), 1–9.
- [61] Karen Schrier and David Gibson. 2010. *Ethics and game design: teaching values through play: teaching values through play*. IGI Global.
- [62] Douglas Schuler and Aki Namioka. 1993. *Participatory design: Principles and practices*. CRC Press.
- [63] Gry Seland. 2009. Empowering end users in design of mobile technology using role play as a method: reflections on the role-play conduction. In *International Conference on Human Centered Design*. Springer, 912–921.
- [64] Phoebe Sengers, Kirsten Boehner, Shay David, and Joseph 'Jofish' Kaye. 2005. Reflective design. In *Proceedings of the 4th decennial conference on Critical computing: between sense and sensibility*. 49–58.
- [65] Phoebe Sengers, Kaiton Williams, and Vera Khovanskaya. 2021. Speculation and the Design of Development. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–27.

- [66] Katie Shilton. 2010. Technology development with an agenda: Interventions to emphasize values in design. In *Proceedings of the Annual Meeting of the American Society for Information Science & Technology (ASIS&T)*, Vol. 47. 1.
- [67] Katie Shilton. 2018. Values and ethics in human-computer interaction. *Foundations and Trends® in Human-Computer Interaction* 12, 2 (2018).
- [68] Marc Steen. 2013. Co-design as a process of joint inquiry and imagination. *Design Issues* 29, 2 (2013), 16–28.
- [69] Bruce Sterling. 2005. *Shaping things*.
- [70] Bruce Sterling. 2012. Sci-Fi Writer Bruce Sterling Explains the Intriguing New Concept of Design Fiction (Interview by Torie Bosch). <https://slate.com/technology/2012/03/bruce-sterling-on-design-fictions.html>
- [71] Marc Teyssier, Marion Koelle, Paul Strohmeier, Bruno Fruchard, and Jürgen Steimle. 2021. Eyecam: Revealing relations between humans and sensing devices through an anthropomorphic webcam. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [72] Alexandra To, Hillary Carey, Geoff Kaufman, and Jessica Hammer. 2021. Reducing Uncertainty and Offering Comfort: Designing Technology for Coping with Interpersonal Racism. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [73] Cameron Tonkinwise. 2014. How we intend to future: review of Anthony Dunne and Fiona Raby, speculative everything: design, fiction, and social dreaming. *Design Philosophy Papers* 12, 2 (2014), 169–187.
- [74] Jasper Tran O’Leary, Sara Zewde, Jennifer Mankoff, and Daniela K Rosner. 2019. Who gets to future? Race, representation, and design methods in Africatown. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [75] Ron Wakkary, Doenja Oogjes, Henry WJ Lin, and Sabrina Hauser. 2018. Philosophers living with the tilting bowl. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [76] Mark Weiser and John Seely Brown. 1996. Designing calm technology. *PowerGrid Journal* 1, 1 (1996), 75–85.
- [77] Richmond Y Wong and Vera Khovanskaya. 2018. Speculative design in HCI: from corporate imaginations to critical orientations. In *New Directions in Third Wave Human-Computer Interaction: Volume 2-Methodologies*. Springer, 175–202.
- [78] Richmond Y Wong, Nick Merrill, and John Chuang. 2018. When BCIs have APIs: Design fictions of everyday brain-computer interface adoption. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 1359–1371.
- [79] Richmond Y Wong and Tonya Nguyen. 2021. Timelines: A world-building activity for values advocacy. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [80] Micah Zenko. 2015. *Red Team: How to succeed by thinking like the enemy*. Basic Books.