

The Oxford Handbook of Ethics of AI

Markus D. Dubber (ed.) et al.

https://doi.org/10.1093/oxfordhb/9780190067397.001.0001

Published: 2020 **Online ISBN:** 9780190067427 **Print ISBN:** 9780190067397

Search in this book

CHAPTER

23 Social Failure Modes in Technology and the Ethics of AI: An Engineering Perspective

Jason Millar

https://doi.org/10.1093/oxfordhb/9780190067397.013.28 Pages 442–461

Published: 09 July 2020

Abstract

This chapter argues that, just as technological artefacts can break as a result of mechanical, electrical, or other physical defects not fully accounted for in their design, they can also break as a result of social defects not fully accounted for in their design. These failures resulting from social defects can be called *social failures*. The chapter then proposes a definition of *social failure* as well as a taxonomy of *social failure modes*—the underlying causes that lead to social failures. An explicit and detailed understanding of social failure modes, if properly applied in engineering design practice, could result in a fuller evaluation of the social and ethical implications of technology, either during the upstream design and engineering phases of a product, or after its release. Ideally, studying social failure modes will improve people's ability to anticipate and reduce the rate or severity of undesirable social failures prior to releasing technology into the wild.

Keywords: technological artifacts, social defects, social failures, social failure modes, engineering, design

Subject: IT and Communications Law, Law

Series: Oxford Handbooks

Collection: Oxford Handbooks Online

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

Introduction

Case 1—The Quebec Bridge

IN August 1907, a steel cantilever bridge spanning the St. Lawrence River near Quebec City was under construction and nearing completion.[1] Engineers noticed that some of the structural beams were bending, so they initiated an investigation. They discovered that calculations made early in the design phase had gone unchecked, resulting in a design that was insufficient to support the full weight of the bridge. Unfortunately, the engineers were unable to communicate the information to the construction crews in time. One fateful day, as workers were nearing the end of their shift, a beam failure ↳ caused the south
p. 444 structure of the bridge to collapse, killing seventy-five of eighty-six workers on the bridge at the time.[2]

Case 2—Google Glass

Designed as a wearable piece of Wi-Fi-enabled computing technology in the early 2010s, Google Glass resembled a pair of glasses, but had a small heads-up-display in place of the lenses and a camera integrated into one arm of the glasses. Wearers could search the web and view results on the display and capture video and pictures, all while communicating with the device using voice commands. Early adopters were sighted wearing their Glass everywhere. Google Glass was quickly, and widely, criticized on privacy grounds, owing to the wearer's ability to capture and record images surreptitiously in any number of social, often private, settings (think urinals),[3] and also raised a number of health-related concerns. Additionally, and importantly, early Google Glass wearers were ridiculed in the mainstream media and on social media (and likely in person) for their "creepiness"[4] and "dorkiness"[5]—these early adopters were eventually labeled *Glassholes*.[6] Despite its technical sophistication, Google Glass was discontinued shortly after its limited precommercial debut.[7]

Engineers, and people impacted by the technology they design could benefit from a new analytic toolkit that helps clarify the link between the ethical issues raised by AI and other technologies and the design decisions that shape those technologies. In this chapter I propose such a toolkit. In my approach I examine a common characteristic of technological artifacts[8] that, to date, has been under examined but deserves attention,
p. 445 which is this: just as technological artifacts can break as a result of mechanical, electrical, or other ↳ *physical* defects not fully accounted for in their design (as with the Quebec Bridge), they can also break as a result of *social* defects not fully accounted for in their design (as with Google Glass). I call failures resulting from social defects *social failures*, and propose that we add them and the underlying causes that lead to social failures—*social failure modes*—to the list of failure modes worth studying in detail and thinking about when designing technology. An explicit and detailed understanding of social failure modes, if properly applied in engineering design practice, could result in a fuller evaluation of the social implications of technology, either during the upstream design and engineering phases of a product or after its release. Ideally, studying social failure modes will improve our ability to anticipate and reduce the rate, or severity, of undesirable social failures prior to releasing technology into the wild, just as an understanding of other common failure modes has improved our ability to anticipate and reduce the rate and severity of undesirable mechanical, electrical, and other physical failures.[9]

There are at least three relatively straightforward reasons why we should study social failures and seek to reduce the rate or severity of social failures in technology. Social failures, like any other failure, can break a product, thus: (1) depriving individual users and society of whatever benefits that an artifact would otherwise deliver;[10] and/or (2) delivering harms upon individual users or society. Finally, (3) an improved understanding of social failures can explain what was wrong with an artifact that led to a failure, thus

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

leading to targeted improvements in the social aspects of an artifact's design. Thus, a better understanding of social failures in technology can lead to better technology, both functionally and ethically.

To be sure, my goal here is not to suggest that social failures, or the social failures I describe herein, are the only nonphysical failures that deserve attention when analyzing how successful a technology is, or will be, once unleashed in society. There are other nonphysical failures that one could point to as being equally deserving of examination. Two that have been brought to my attention are economic and political failures.[11] Economic failures, as I understand them, occur when a product fails as a result of its fit within a particular market. A product might be too expensive or could fail to deliver the right kinds of value propositions to its intended users. Political failures, as I understand them, occur when an artifact does not fit prevailing standards and is unable to force new ones, or when it fails to appease important political actors. An analysis of these kinds of failures is worth study, but is beyond the scope of this chapter. However, it is also worth noting how difficult it is to treat physical, social, economic, and political failures as if we can easily draw impenetrable boundaries around them. These various failures can be understood as multiple sides of the same artifact; when we choose to focus exclusively on physical failures as objects of analysis in engineering,
p. 446 we do so arbitrarily and not ↳ because the artifact is more accurately, or appropriately, understood according to the language of the physical sciences. When the goal is to understand how an artifact exists in society (its natural habitat), its physical characteristics are one among many vantage points from which to observe it.

This chapter, as its focus is starting a conversation about social failure modes in technology, proposes an additional vantage point from which to observe artifacts, that is, a new way of understanding how a technology can "work" or "break." From a methodological perspective, I build toward that understanding by spanning a conceptual gap between two academic worlds. On one side of the gap, we have theories from philosophy, sociology, and anthropology focused primarily on the social and ethical dimensions of technology. They provide detailed examinations of the social dynamics between individuals, societies, and technologies, as well as a rich vocabulary for thinking about what causes a technology to *work* or *break*. On the other side, we have theories from the sciences and engineering focused primarily on the physical aspects of technology, on ways of (re)configuring matter to make things. They provide detailed examinations of the physical dynamics of technology, as well as a rich vocabulary for understanding what causes a technology to *work* or *break*.

As I've laid things out, both of these worlds are concerned with understanding how technology works and breaks, so the conceptual gap may seem small. I believe it is. The project here is to bring the vocabulary used to describe the social dimensions of technology in contact with the vocabulary used to describe the physical dimensions of technology, and do it in a way that avoids demanding that either should buckle under the other's force. Combining these vocabularies makes for richer descriptive and normative accounts of technology. Furthermore, whatever practical gains are to be made by understanding social failure modes will undoubtedly require the social apparatus to be incorporated into technical practices. Practical gains, therefore, will require cooperation with, and uptake among, those in command of the physical world, namely engineers. So, to span the gap, an important aspect of this project is making the social apparatus compatible with the technological; to be used, the concept of a social failure mode must be designed such that it fits well into the world of engineering.

With those considerations in mind, I begin this chapter by making explicit the link between social failures and the ethics of AI, as AI seems to lend some urgency to better accounting for the social dynamics of technology during design. I then describe a few different ways in which artifacts break physically and, in doing so, offer an overview of the language typically deployed to describe how things break. That discussion of physical failures and failure modes lays the necessary groundwork to frame the concept of a social failure and social failure modes. Social failures do not always render a technology broken in the traditional sense (i.e., physically unusable)—as I will describe things a technology can suffer a social failure yet still

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

"function." The question, then, is why should we worry about social failures in the first place? In response to that question, I describe in this chapter within the sections "Social Failures and Social Failure Modes," and "The Ethical Dimensions of Social Failures," how social failures break technology and raise, or underscore, ethical concerns, thus fleshing out my primary normative ↳ claim, which is that we have good
p. 447 reasons for studying, and avoiding, social failures in technology. I then provide a rudimentary social failure mode analysis of Microsoft's Tay—a particularly problematic AI—as a case study illustrating how one might approach using the proposed toolkit. But to begin, I explain why a better understanding of social failure modes is particularly important, and would be particularly beneficial, for engineers designing robots and AI.

Social Failures and the Ethics of AI

Much of the literature on the ethics of AI that has recently emerged concerns itself with describing how AI applications result in problematic ethical issues. Those issues tend to be cataloged under a variety of ethical concepts, including (but not limited to): trust and trustworthiness; accountability; justice; bias; fairness; interpretability; explainability; power and control; gender; privacy; discrimination; truth; and equality. Humans reify these concepts in practice, through our actions within relationships and institutions, by attaching to each concept a set of social norms that help to delineate between actions that are acceptable and unacceptable, permissible and impermissible. Take *trust*, for example. I might trust a person with my sensitive personal information because that person has demonstrated their willingness and ability to not divulge other important information with which they were entrusted. I might further consider that person *trustworthy* because I know they have certain caring dispositions that are aligned with my values, such that I believe they will respond appropriately when pressed to divulge my personal information. Thus, a person is considered both trusted and trustworthy because their actions in our relationship adhere to social norms that define trust/trustworthiness. Similarly, institutions can codify norms defining trust/trustworthiness in policies and processes, thus signaling that an institution is trustworthy and can be trusted, especially in cases where they consistently apply their policies. The same is true for bias, interpretability, and the other ethical concepts mentioned earlier—each is defined in practice by a set of accompanying norms.

When designing AI (indeed, as I describe in this chapter, when designing most, if not all, technology), engineers inevitably embed social norms into the artifact, and people experience those norms when interacting with the artifact. Norms that are engineered into the AI shape the underlying decisions and resulting human-AI interactions. Thinking back to the trust example, and imagining that I am asked to trust an AI with personal information, the kinds of norms we engineer into the AI will define its ethical character. When and to whom will it divulge my personal information? Will it inform me about any particular instance of its divulging my personal information? Do I have choices to make about these potential divulgences? As I describe in this chapter, if the engineered norms undermine my normative expectations regarding trust, my interactions with the AI could raise serious ethical issues, and, as a result,
p. 448 the AI could suffer a social failure. As we develop robots and AIs that are more autonomous, we will be ↳ required to delegate more and more (previously human-dominated) decision-making to them, so that AIs can interact with us, among us, and on our behalf.[12]

Understanding the relationship between social norms, technology, and the design decisions that engineers make takes on a sense of urgency with robots and AI; robots and AI are performing among us, and as they increase in autonomy and sophistication they are increasingly performing the social norms that we have necessarily embedded in them.

Social norms can, therefore, be seen as a critical social backdrop upon which much of AI is designed. Because of this, we would benefit by developing an analytic design/engineering toolkit to help translate

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

between the technical and the social, engineering and the humanities, and bring the social dynamics of technology into the foreground in engineering design. Hence, I propose here a rigorous study of social failures and the underlying social failure modes as a step in that direction.

Talking about How Stuff Breaks: A Brief Primer on Physical Failures and Failure Modes

Everything we make breaks, eventually. Joints crack through repetitive strain. Metal surfaces erode when exposed to the weather. Tires wear from friction. Bolts loosen from vibration. Circuits fizzle from the heat of moving electrons. Pipes burst under pressure. And overloaded beams buckle and shear. The list of physical failures that can potentially break any one technology is long, and grows in proportion to the complexity of the system. Having more parts increases the probability that failures will occur. This simple fact about the things we make—that they all break, eventually—can be read as pessimistic and discouraging. It's the stuff that could cause an engineer faced with the inevitability of failures to throw their hands in the air. But each physical failure encountered also represents an opportunity to study the underlying causes of that failure: the physical *failure modes*. Each study that leads to a breakthrough in understanding how a thing *has* failed also lays the foundation for an understanding of how things *will* fail in the future. Researching how things fail helps to prevent future failures of the same kind, a thoroughly optimistic perspective.

The optimism that drives research into physical failures is fueled by the recognition that physical things fail according to physical laws. The engineers in charge of the Quebec Bridge knew this and were able to predict
p. 449 the eventual failure, even if other ↳ circumstances stopped them from preventing it. Beams don't break randomly, they break predictably. Their failure modes can be grouped and characterized by type: they *buckle* and *yield* regularly according to the physical properties (e.g., type of material and dimensions) of the beam.

Compiling a catalog of physical failure modes has allowed engineers to anticipate potential failures while designing things. With this additional knowledge, engineers can make explicit, informed, decisions about the physical makeup of their devices that more fully take into account the many ways that each component could fail in use. Certain failures will ultimately be deemed acceptable given various design constraints, while others will be avoided at all costs. The upshot is that a more detailed understanding of physical failure modes can lead to more deliberate, safer technology, designed with a fuller understanding of the benefits, risks, and harms associated with particular design decisions.

Optimism over the possibility of better understanding failure modes has paid off over the years. What started as investigations into how physical things break has evolved into a much broader analysis of how things and processes can, and do, fail. As a result of measurable successes, failure mode analysis is a cornerstone of good design.

For the remainder of this chapter I argue that a similar optimism is justified and should be acted upon when it comes to investigating and cataloging the many ways that technology can suffer social failures. It's justified because we have good reason to believe that social norms can be identified and understood within a target user group (or society in which a technology will be deployed) to the extent of allowing us to reasonably predict how technology can interact (e.g., support or undermine) and/or come into tension with those social norms. The benefits of avoiding social failures in technology ground the normative claim that we ought to investigate and catalog social failure modes.

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

Social Failures and Social Failure Modes

The Zimbabwe Bush Pump

The success of the Zimbabwe Bush Pump—a simple hand-operated water pump like one you might find on an old farm, and the protagonist of a story narrated by Marianne de Laet and Annemarie Mol—beautifully illustrates the concept of a social failure of technology.[13] What follows is a condensed version of that story.

To see it in a picture, the bush pump looks like any other rudimentary hand operated pump. It is
p. 450 intentionally rustic, designed specifically for use in rural Zimbabwe where ↳ clean water is a precious resource. To use it, one places a container under the spout, grabs hold of the wooden lever on top, and proceeds to pump, up and down, drawing water from the well upon which the pump sits.

But the Zimbabwe Bush Pump is not like the many other rudimentary hand-operated water pumps on the market. In fact, it has succeeded where those other similar looking pumps have failed precisely because other pumps fail to establish social norms required for them to work and transgress already established social norms common in the pumps' use context (rural Zimbabwe). The Bush Pump, just to be clear, is designed to establish the requisite social norms and conform to already existing social norms, which lead to its success.

As de Laet and Mol argue, for a pump to *work*, in addition to reliably pumping water—its intended mechanical function—it must also be seen as an artifact that locals *trust* and, when they use it, must reliably pump *clean* water. Satisfying these three specifications requires more than mere mechanical design considerations.[14] Unlike other pumps, all of them mechanically adequate, various sociocultural features are incorporated into the Bush Pump's design for it to satisfy its role as a trusted source of clean water.

For starters, in order to prevent contamination of the well, community members must be convinced to diligently maintain the pump. The Bush Pump's designers accomplish this by demanding the involvement of the whole community in various aspects of the installation of the pump. Drilling the well, for example, is designed as a social activity that motivates a sense of community ownership of the well. Rather than using mechanical drills that are quick and efficient, but tend to be operated by nonlocals, the Bush Pump relies for its well on a special hand-operated drill known as a Vonder Rig, which requires community involvement for it to work.

> A video distributed by the factory shows that sometimes operating the rig turns into a village feast. Village women push the iron crossbar to drive the auger into the ground, while village men sit on the bar to weigh it down and children dance around.[15]

Once the well is drilled it must be capped with a headworks (a small concrete and brick structure that sits atop the well) to provide adequate water runoff and a seal against deadly E. coli contamination. Like the drilling, constructing the headworks is designed as a community affair, which helps to further establish norms of ownership of the Bush Pump.[16]

De Laet and Mol nicely summarize what is required to install a working Bush Pump:

> In order to be a pump that (pre)serves a community, it not only needs to look attractive, have properly fixed levers and well-made concrete aprons, it must also be capable of gathering people
> p. 451 together and of inducing them to follow well-drafted ↳ instructions. It must come with a Vonder Rig and invite people to push bars, sit on them or dance around them. It must seduce people into taking care of it. Thus the boundaries around a community pump may be widely drawn. Indeed they embrace the community.[17]

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

Thus, drilling the well, assembling the pump, and constructing the headworks variously involve the whole local community. These communal activities are designed to generate a strong sense of community ownership of each installed pump, in turn helping to establish critical social norms around pump use and maintenance.[18]

There is one other crucial social element to the story of the Bush Pump that helps to explain why it succeeds where others fail. Rather than simply relying on geological data to determine the best site for each pump, the Bush Pump's instruction manuals and other documentation "state clearly and repeatedly … that local water diviners should be consulted before any decision about the siting of a water hole is made."[19] According to a UNICEF worker, nongovernmental organizations intent on drilling wells in locations based solely on geological data or issues of efficiency (proximity to certain buildings, from livestock, etc.) regularly witness the failure of their pumps because local diviners were not consulted, leading the local women to consider those wells "dead."[20] Thus, the Bush Pump is trusted in part because it conforms to social norms attached to the wisdom of local water diviners.

Examining the role that these sociocultural aspects play in the design of the Bush Pump undercuts the idea that any similar pump, mechanically speaking, would work in rural Zimbabwe. Indeed, in telling the Bush Pump's story, de Laet and Mol emphasize that other mechanically superior pumps have failed precisely for lacking such sociocultural design elements. Pumps that are poorly maintained become contaminated, while pumps that have been placed without a water diviner's endorsement are ignored.

Social Failure Dynamics

We see that the Bush Pump's story is also a story about how users' social norms are fundamentally related to an artifact either *working* or *breaking*. The Bush Pump works because it is mechanically sound, because its designers figured out a way for it to encourage community members (i.e., its users) to feel a sense of collective ownership toward it, thus to care for it and prevent contamination, and finally because its design respects the trust norms associated with local water diviners, and so it is used.

Pumps that fail (including Bush Pumps), on the other hand, can fail *socially*, either because they fail to encourage the requisite social norms, or because they violate important existing ones.

In these two senses, then, water pumps can, and do, suffer *social failures*. And just as we can talk about
p. 452 pumps breaking due to mechanical failures, we can talk about pumps ↳ breaking due to social failures. Pumps that suffer a lack of community ownership in rural Zimbabwe break because they are not cared for and become contaminated. (Thus, social failures can sometimes lead to physical failures though, in this case, the social failure is at the root.) Pumps that are not trusted by women in rural Zimbabwe break because they are ignored.

Google Glass can be described in similar terms.[21] Though some people were happy to use it and celebrated its technical capabilities, too many others rejected it because it violated important existing social norms of health, privacy, and coolness. When a (potential) user rejected Google Glass on grounds that it violated these existing social norms, Google Glass broke. Alternatively, one might say that Google Glass failed to shift, or to generate, the social norms required for it to be accepted and used (i.e., to work). For example, if Google Glass was designed in such a way that it successfully shifted society's notions of coolness, or of privacy, such that those new norms encouraged its intended use, then it would not have suffered those social failures.

With the Bush Pump and Google Glass in mind, I offer an initial definition of a social failure:

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

> *An artifact suffers a social failure when a social norm designed into it comes into tension with an accepted social norm held by its user(s), or other relevant stakeholder, to the extent that the tension diminishes or prevents the intended use, or functioning, of that artifact.*

But what is a social norm, and what does it mean for a social norm to be designed into an artifact? First, a social norm, in the sense I am using it, should be understood just to mean a behavioral rule *R* that is followed by a sufficiently large set of individuals within a social group.[22] By this definition, *R* could be held by social group *P*, and not by some other group P'. The kind of social group I refer to here is rather flexible. It could be a particular citizenry (e.g., Canadians), a professional group (e.g., physicians, engineers), or some other identifiable group (e.g., iPhone users, hobby farmers, cyclists). They key feature of any group, when thinking about social failures, is that its members share, and act on, at least one reasonably identifiable social norm.

p. 453 There are several ways of describing how social norms are designed into artifacts.[23] Social norms can be identified in the *political relations* an artifact imposes on its users.[24] Consider nuclear reactors, which, according to Winner,[25] impose a set of authoritarian norms on the society in which they are installed. Nuclear reactors are extremely expensive to build and maintain and need to be strictly secured in order to maintain safety, which results in complex surveillance and security regimes deployed on and around the facilities and its workers. Nuclear material (the fuel and waste) is dangerous, valuable, and a national security threat (it can be weaponized), and many of the parts used in building and maintaining the plant are safety critical, requiring additional strict security regimes to manage the entire supply chain and waste disposal process. Contrast this with the more democratic norms embedded in windmills. Windmills are relatively inexpensive, can be installed in a farmer's field (by the farmer, no less!), produce no hazardous waste, thus they impose no strict security regimes on those they impact. Nuclear reactors and windmills are similar in that they are designed to produce electricity, but the political relations embedded in each artifact could not be more different.

Social norms can also be reflected in the work we *delegate* to artifacts. The average front door to any building is illustrative here.[26] Most of us share a social norm when it comes to front doors: that they should generally be kept closed. This norm reflects any number of underlying reasons for closing a front door, such as maintaining a comfortable indoor temperature, saving energy costs, keeping the bugs out, keeping the burglars out, keeping the wind out, keeping the family pets in. The list goes on. Doors tend to ignore this social norm; they require human attentiveness and a helping hand to keep closed. Indeed, to remain closed, doors operated by children often require the help of an additional person, a nearby adult, to call out repeatedly, "Close the door!" However, clever humans have found a way to delegate all the work required to enforce the norm to small artifacts attached to doors. These little mechanisms, usually powered by a spring, are designed to automatically close doors once opened, thus embodying and enforcing the social norm to keep front doors closed.

p. 454 Engineers can also delegate flexible social norms to technology via "settings" that target those norms.[27] Take privacy settings, like those designed into Facebook's user interface, which define Facebook's behavior with respect to personal data that users upload and share on the platform. These settings allow a user to shape Facebook's behavior based on their own normative privacy expectations, for example, by allowing users to partially limit photo sharing to friends. However, Facebook's design decisions also bound the possibilities of that individual shaping. Facebook users cannot, for example, upload photos to the platform and decline to share those photos with Facebook. Thus, settings can reflect, shape and bound social norms delegated to a technological artifact.

Finally, we can identify social norms embedded in artifacts when they *mediate* our perceptions of, and actions in, the world. Verbeek argues that by mediating the way we perceive reality, a technology can

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

"actively contribute to the moral decisions human beings make."[28] Using medical imaging technologies as an example, he writes:

> [T]he specific way these technologies represent what they "see" helps to shape how the body or a fetus is perceived and interpreted and what decisions are made. In this way, technologies fundamentally shape people's experiences of disease, pregnancy, or their unborn children. Ultrasound, for instance, makes it quite easy to determine the thickness of the nape of the neck of a fetus, which gives an indication of the risk that the unborn child will suffer from Down's syndrome. The very fact of having an ultrasound scan made, therefore, lets the fetus be present in terms of health and disease and in terms of our ability to prevent children with disease from being born.[29]

As a result of its mediating role, ultrasound technology has been elevated to a standard of care among healthcare professionals—*perform an ultrasound at 7–12 weeks, and again at 20 weeks during each pregnancy* is among their widely accepted social norms. That norm has a cascading effect in that it imposes new norms that women encounter during a standard pregnancy when they are confronted with the results of their ultrasound.

I have sketched a picture in this section illustrating several ways to describe how social norms are designed into technology. The social norms that will matter most in an analysis of social failures are the ones that tend to align with, or transgress, important social norms held by a particular group impacted by a technology. If, say, privacy is important to a particular group of potential users, then how engineers embed privacy norms in their product will be an important design consideration. A product that transgresses potential users' privacy norms severely enough will likely suffer a social failure.

Two Social Failure Modes

The Zimbabwe Bush Pump and Google Glass suggest, as a starting point for further analysis and characterization, two types of social failures, which we can refer to as *social failure modes*. These social failure modes will help flesh out the definition of a social failure.

(1) *Absence of supportive norms*. An artifact can fail socially when its design requires a certain social norm to be held by its user(s) in order for it to work (i.e., be used) as intended, but that required norm is not in fact held by its user(s).

As we saw with the Bush Pump, in order for it to be cared for, community norms of ownership are established by way of elaborate community activities centered around a new pump, as a result of designing the Bush Pump to be installed as a community affair. The Bush Pump's designers have designed into it the norm *pumps must be cared for by the community to avoid contamination*. Pumps that fail socially often do so because their designs fail to establish such supportive caring norms.

With Google Glass, in order for it to have been socially acceptable people's privacy norms would have had to shift in such a way as to make wearing cameras that record surreptitiously socially acceptable, perhaps even cool. By designing Glass the way they did, Google's engineers designed into it the norm *it is acceptable to wear cameras that record surreptitiously*, which was in tension with the widely held, and stable, societal norm *it is NOT acceptable to wear cameras that record surreptitiously*. Thus, the existing norms came into tension with the designed norms, the supportive norms were absent, and Glass suffered social failures.

(2) *Norm transgression*. An artifact can fail socially when a social norm designed into it transgresses an accepted social norm held by its user(s).

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

This social failure mode is more straightforward. As we see with water pumps, some are designed to be agnostic with respect to the social norms surrounding water diviners in rural Zimbabwe. Those pumps have designed into them the norm *install this pump on any water-yielding well*. However, agnosticism over the role of local water diviners can result in an installed pump transgressing the local norm *only trust water from a well that has been endorsed by a water diviner*. Any pump installed on a well not endorsed by a local water diviner runs the risk of transgressing that local norm, since *any water-yielding well* may not be trusted.

Likewise, Google Glass, evidenced by the creation of Glassholes, clearly transgressed existing norms of privacy and coolness.[30]

That said, not all norm transgressions are created equal. Though the distinction is vague and somewhat
p. 456 controversial, social norms seem to come in two broad psychological ↳ categories: *conventional* and *moral*.[31] Prototypical examples of conventional norms include "wearing gender-inappropriate clothing (e.g. men wearing dresses), licking one's plate at the dinner table, and talking in a classroom when one has not been called on by the teacher."[32] Prototypical moral norms, on the other hand, tend to include prohibitions against killing, injuring, or stealing from other people. Moral norms tend to be categorized as such because they are regarded as: less authority dependent; more universal (i.e., they apply cross-culturally); involving rights violations; and more serous in nature than conventional norms.

Importantly, in many cases (though not all) an analytic focus on social failures involving moral norms and moral norm transgressions will underscore the unavoidable moral dimension of engineering design. Such a focus will thus help to delineate between permissible and impermissible design decisions and suggest certain obligations that engineers have by virtue of the design decisions they must navigate. Engineers have already partially embraced this reality. Social failures caused by safety or privacy norm transgressions are already generally considered impermissible, resulting in a general obligation for engineers to design for safety and increasingly for privacy. In Canada, for example, the profession's code of ethics requires engineers to "hold paramount the safety, health and welfare of the public and the protection of the environment and promote health and safety in the workplace."[33] Other social norms, upon analysis, could give rise to similar professional constraints, as I discuss in the next section.

Conventional norms, in contrast, tend to be regarded as such because they are: more authority/institution dependent; locally applicable (i.e., they do not apply cross-culturally, or even among different groups within a single cultural boundary); do not involve rights violations; and transgressions of them are less severe than those involving moral norms.

Using the moral/conventional distinction as a guide, we can further subdivide social failures resulting from norm transgressions into two sub-categories: *conventional norm transgressions* and *moral norm transgressions*. To illustrate, social failures resulting from rights violations, say privacy transgressions, can be categorized as moral norm transgressions, whereas social failures resulting from institutional norm transgressions, say a mobile phone that can't be "muted" in a quiet space, can be categorized as conventional norm transgressions. As I discuss in the following section, the moral/conventional distinction can serve as a useful guide for anticipating the nature of harms resulting from social failures of the norm transgression type.

The Ethical Dimensions of Social Failures
p. 457

My argument so far has focused on sketching a conceptual apparatus with which to describe social failures and social failure modes. In this section I turn to the normative aspect of this project and propose two straightforward reasons why we should care about and investigate, in order to mitigate or avoid, undesirable social failures in technology.

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

The first reason has to do with the deprivation of benefits a group of users (or other implicated stakeholders) would otherwise derive from a technology that works as intended. Assuming a technology is designed to deliver some set of goods or benefits to some group of individuals (e.g., clean water, clean air, improved medical diagnoses, enhanced communication with friends and family, novel modes of interaction with one's environment, etc.), social failures, insofar as they prevent that technology from being used as intended, stand to reduce the benefits derived from that technology. Some social failures will deprive users of important benefits. Water pumps that become contaminated, or are not trusted, do not deliver the benefits of potable water. Other social failures will deprive users of benefits in ways that are trivial or at least far less costly. I assume this to be the case with many intended uses of Google Glass.

In addition to depriving people of benefits, social failures can directly harm individuals, giving us further reason to investigate and avoid some social failures in technology.[34] Contaminated wells deprive communities of clean water, but they also allow E. coli to infect (and kill) members of that community. Google Glass surreptitiously recording the surrounding environment while its host is standing at a urinal in a crowded washroom violates the privacy of every recorded individual sharing that otherwise semi-private space.

We can use the moral/conventional distinction here to anticipate how those individuals harmed by a social failure in technology will interpret the severity of the harm and, more importantly, to evaluate the nature of the harm in a way that provides guidance on the moral permissibility of a particular technology.

Moral norm transgressions, by definition, tend to result from rights violations and tend to be considered more serious in nature. Thus, we can trivially assume that harms resulting from moral norm transgressions will threaten to break a technology, while conventional norm transgressions may pose a lower risk in that regard.

Less trivial, however, is the implication that every social failure resulting from a moral norm transgression suggests that something is potentially morally problematic about that technology. Here we see clearly the impact that an understanding of social failure modes ought to have on guiding ethical engineering practice.
p. 458 Knowledge of social failure modes can frame a discussion about whether or not to design a technology in ↳ such and such a way, by providing clearer links between designed features and moral norms, thus helping to distinguish between design decisions that users will consider permissible and impermissible. The corollary, to complete the thought, is that a better understanding of social failure modes can help to define the boundaries of ethical design decisions; every social failure of technology resulting from a moral norm transgression can potentially be interpreted, insofar as the social failure was reasonably predictable, as an ethical design failure on the part of the responsible engineers. Thus, a more thorough understanding of social failure modes should help distinguish between permissible and impermissible technologies or features thereof.

Determining whether a particular moral norm transgression is an ethical failure is a complicated business. Not all social norms are justified, so a technology that transgresses a particular unjustified moral norm could count as an engineering victory (though just to be clear, if the norm violation had no effect on users' willingness to use the technology as intended, this would not be considered a social failure of that technology according to the definition on offer here). In addition, a technology might establish an unjustified or controversial moral norm without suffering a social failure. For example, China is currently rolling out a complex social networking platform that calculates and ranks its citizens according to a *social credit score*, ostensibly to curb people's behavior in public and help determine how to distribute goods among the citizenry.[35] The various new norms that will likely emerge as a result of the platform's ubiquity will strike many as a clear and unjustifiable violation of basic freedoms. According to the definition on offer, however, the platform would only suffer social failures if it failed to be used as intended owing to existing norm violations, or if it failed to establish the requisite norms for its intended use. Critically, the occurrence

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

or nonoccurrence of a social failure in technology should not be taken as evidence of a technology's moral permissibility or impermissibility.

This moral ambiguity of social failures echoes the moral ambiguity of many new technologies whose social meaning takes time to emerge and stabilize and can ultimately differ from group to group.[36] "The telephone and typewriter," as Verbeek points out, "were not developed as communication and writing technologies but as equipment for the blind and the hard of hearing to help those individuals hear and write."[37] Telephones also eased interpersonal communication, altered the speed at which information propagates globally, and enabled surreptitious government eavesdropping through wiretaps.

Pointing out the moral ambiguities inherent in social failures should not be interpreted as absolving engineers of their moral responsibility in designing technologies with the intent to disrupt justified norms or establish unjust norms. Rather than undermining the concept of social failures in technology, the moral ambiguity of social failures is, I think, important for maintaining the descriptive accuracy of the concept. It anchors justifications for particular norms in the social contexts into which technologies are deployed
p. 459 without so divorcing it from engineering practice as to absolve engineers of ↳ their responsibilities as designers. In addition, the concept provides an analytic apparatus that explicitly links design with the social dynamics of technology. Diverse social groups can go on negotiating their norms, while engineers acting in good faith, as members of those larger groups, deploy the concept to anticipate and evaluate social failures in their design activities.

A fuller discussion of the dynamics between design and the moral permissibility of particular technologies is beyond the scope of this chapter. However, I believe that discussion would benefit from a more fully developed language of social failures, one that investigates and links particular technological features to particular norms and thus to engineering practice.

A Social Failure Mode Analysis of Microsoft's Tay

As was mentioned in the section "Talking about How Stuff Breaks: A Brief Primer on Physical Failures and Failure Modes," designing AI requires engineers to embed social norms into artifacts so that the AIs can interact with us, among us, and on our behalf. Having sketched an initial taxonomy of social failure modes throughout this chapter, one can see how particular design decisions regarding social norms can shape the ethical nature of an artifact. Furthermore, the reach of AI in society can be broad and immediate, often the result of a background system-wide software update, which adds a certain urgency to understanding what social failures any particular design decision could trigger. To illustrate the link between social failures and the ethics of AI with a final example, I turn to Microsoft's Tay.

Microsoft's Tay was a Twitter chatbot, an AI designed to interact with other Twitter users as if it were a fourteen-year-old teenage girl. Because Tay lived on Twitter, its reach was broad—Tay could interact with any of Twitter's roughly seventy million active monthly users. Tay was also designed to "learn" how to carry on normal Twitter conversations by using the ever-changing set of prior interactions as its learning dataset. That meant that by design, the interactions Tay had with other Twitter users would shape Tay's future behavior. In other words, Microsoft made the Twitterverse responsible for teaching young Tay how to behave when interacting with people.

The Twitterverse did not parent Tay responsibly. Within a day, Tay was behaving like a racist, neo-Nazi misogynist, forcing Microsoft to pull Tay's plug amid widespread public controversy.[38]

It might be that fourteen-year-old girls were Tay's target user group and that Tay could have learned to be
p. 460 quite pleasant after interactions with only those target users. But the ↳ larger sociotechnical system into

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

which Tay was embedded included numerous callous Twitter users intent on shock-testing Tay. To be sure, the average Twitter user interacting with Tay during its brief existence did not conform to the same interactional norms as the average fourteen-year-old (or so one hopes). Thus, Tay suffered multiple social failures.

It is tempting to brush Tay off as a foolish publicity stunt that was bound to go wrong. After all, the internet, Twitter in particular, is known for its trolls. A more charitable approach, however, reveals Tay as a great learning opportunity. More specifically, too quick a dismissal represents a missed opportunity to apply a more sophisticated failure analysis on Tay, one that more accurately accounts for the social dynamics that caused Tay to fail.

Let us first examine the absence of supportive norms that Tay (and Microsoft) required for Tay to succeed. For Tay to work well, it needed to experience a large number of socially acceptable Twitter interactions. Tay required the support of conversational norms typical in what would generally be considered "polite society." But the conversational norms on Twitter are unpredictable at best (especially between two users who do not know each other), deeply unethical at worst, and are certainly not representative of conversational norms in the rest of society. Instead of providing supportive conversational norms, the Twitterverse helped to break Tay.

Once Tay's learning algorithm had been exposed to unsupportive conversational norms, Tay suffered a second social failure: it repeatedly transgressed accepted conversational norms existing in "polite society." Indeed, the kinds of conversational norm transgressions Tay committed were not mere conventional transgressions, they were moral in nature. In response to Tay's behavior, Microsoft was quickly criticized by those stakeholders who found Tay's behavior ethically impermissible. Microsoft responded responsibly to the situation by shutting Tay down.

Applying a social failure mode analysis to an artifact like Tay allows us to focus design efforts on both analyzing and shaping the social dynamics required for an artifact to work. Based on the rudimentary taxonomy provided in this chapter, that means analyzing and shaping the requisite supportive norms and the resulting norms characterizing the artifact's behavior. For Tay to work, design effort was required to both understand the current conversational norms on Twitter and, given their current state, to change the conversational norms on Twitter so they would support Tay's socially appropriate learning and result in ethical behavior. But notice how a focus on social failure modes has helped us move from a state of relative ignorance with respect to Tay's failure, toward much more specific design activities that hold some promise if the goal is to anticipate social failure modes before they occur. This could be the analytic toolkit engineers need that helps clarify the link between the ethical issues raised by AI and other technologies, and the design decisions that shape them.

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

Conclusions

I assume in this chapter that engineers generally intend new technologies to make the world a better, not
p. 461 worse, place. Clearly, many technologies do so. Along those lines, ↳ I assume there is no controversy in claiming that engineers are generally, and nontrivially, motivated to make things that improve the lives of individuals and the societies in which they live, while also being motivated to avoid making things that harm individuals and/or societies. These are fair assumptions, and they are reflected consistently in the various codes of ethical conduct to which engineers are obliged to adhere in their professional dealings. Insofar as engineers are prevented from acting on these motivations in their professional settings, there is work to be done in improving engineers' professional settings. Insofar as engineers are not so motivated in their work, there is work to be done in improving the profession. Insofar as technologies fail, often resulting in harms, there is work to be done in understanding how they fail and how to prevent those failures moving forward.

The social implications of technology have long been studied in detail outside of the profession, and many of those discussions suggest that if the norms embedded into a technology come into tension with accepted social norms held by its user(s), the technology can break. Technology can break when it suffers a social failure. However, practical conceptual and analytic apparatuses for describing and anticipating the links between technology, social norms, and engineering practice are required to put those insights into practice.

To help bridge that gap, between studying the social implications of technology and designing technology, I have proposed a working definition of social failure and described two social failure modes. These analytic tools help to expose and clarify the inextricable connections between some design failures and social norms. Though there is much work to be done to build a more sophisticated understanding of social failures in technology and the ethical implications they have for engineering practice, one practical, ethical insight should already be apparent. A working knowledge of social failure modes can help make better technology, because technology that aligns with people's robustly justified social norms is better than technology that undermines them.

Bibliography

Calo, R., M. Froomkin, and I. Kerr eds. *Robot Law.* Northhampton: Edward Elgar, 2015.
Google Scholar Google Preview WorldCat COPAC

Collins, H.M. and R. Evans. *Rethinking Expertise*. Chicago: University of Chicago Press, 2007.
Google Scholar Google Preview WorldCat COPAC

Friedman, B. and P.H. Kahn Jr. "Human Values, Ethics, and Design." In *The Human-Computer Interaction Handbook*, ed. J.A. Jacko and A. Sears, 1177–1201. Mahwah, NJ: Lawrence Erlbaum Associates, 2003.
Google Scholar Google Preview WorldCat COPAC

Latour, B. *Pandora's Hope: Essays on the Reality of Science Studies.* Cambridge, MA: Harvard University Press, 1999.
Google Scholar Google Preview WorldCat COPAC

Lin, P., R. Jenkins, K. Abney, and G.A. Bekey eds. *Robot Ethics 2.0.* Oxford: Oxford University Press, 2017.
Google Scholar Google Preview WorldCat COPAC

van den Hoven, J., N. Doorn, T. Swierstra, B.-J. Koops, and H. Romijn eds. *Responsible Innovation 1: Innovative Solutions for Global Issues.* Springer 2014.
Google Scholar Google Preview WorldCat COPAC

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

Notes

1 I thank the various participants of two workshops held at Stanford University, one at Apple and one at the University of Toronto's Centre for Ethics, who gave generous comments on drafts of this chapter. Particular thanks are owed to the brilliant Shannon Vallor, who presented a detailed and helpful commentary at the Political Theory Workshop at Stanford, and to Rob Reich, Josh Cohen, Sergio Sismondo, and Celeste Parisi, for their critique of early drafts. Anne Newman, Johannes Himmelreich, Ted Lechterman, Fay Niker, Lindsey Chambers, Hannah Carnegy-Arbuthnott, Desiree Lim, and Laura Gillespie provided valuable feedback at the Stanford EIS Postdoc Workshop. Support for this research was provided through Postdoctoral Fellowships from both the Social Sciences and Humanities Research Council of Canada (SSHRC), and the McCoy Family Center for Ethics in Society at Stanford.

2 C. Pearson and N. Delatte, "Collapse of the Quebec Bridge, 1907," *J. Perform. Constr. Facil*. *20*(1) (2006): 84–91.

3 R. Eveleth, "Google Glass Wasn't a Failure. It Raised Crucial Concerns," *Wired* (December 12, 2018), https://www.wired.com/story/google-glass-reasonable-expectation-of-privacy/.

4 D. Pogue, "Why Google Glass Is Creepy," *Scientific American* (June 1, 2013), https://www.scientificamerican.com/article/why-google-glass-is-creepy/.

5 Dorkiness here is described as being far worse than mere "nerdiness"; see M. Wohlsen, "Guys Like This Could Kill Google Glass Before It Ever Gets Off the Ground," *Wired* (May 2, 2013), https://www.wired.com/2013/05/inherent-dorkiness-of-google-glass/. If you doubt whether a "dork factor" can be blamed for the failure of Google Glass, check out the Tumblr titled *White Men Wearing Google Glass* (http://whitemenwearinggoogleglass.tumblr.com) for confirmation.

6 L. Eadicicco, M. Peckham, J.P. Pullen, and A. Fuitzpatrick, "The 20 Most Successful Technology Failures of All Time," *TIME* (April 3, 2017), http://time.com/4704250/most-successful-technology-tech-failures-gadgets-flops-bombs-fails/.

7 R. Metz, "Google Glass Is Dead: Long Live Smart Glasses," *MIT Technology Review* (November 26, 2014), https://www.technologyreview.com/s/532691/google-glass-is-dead-long-live-smart-glasses/.

8 The term "technological artifacts" should be understood quite colloquially as being interchangeable with other terms that pick out the stuff humans design and build, such as devices, products, technologies, widgets, gadgets, things, etc. Throughout this chapter I use the term interchangeably with "technology," "product," and "artifact."

9 Some failures are desirable, such as those designed to minimize overall damage or harm, or to direct the harm appropriately, such as crumple zones in vehicles that are designed to absorb energy in a crash.

10 I focus here on artifacts that are designed with some good purpose in mind. I have very little to say here about those that are not.

11 Particular thanks to Sergio Sismondo for alerting me to these two categories of failure.

12 J. Millar and I. Kerr, "Delegation, Relinquishment, Responsibility: The Prospect of Expert Robots," in *Robot Law*, ed. Ryan Calo, Michael Froomkin, and Ian Kerr (Cheltenham, UK: Edward Elgar Press, 2016), 102–127. See more on delegation in this chapter under "The Ethical Dimensions of Social Failures."

13 M. de Laet and A. Mol, "The Zimbabwe Bush Pump: Mechanics of a Fluid Technology," *Social Studies of Science* 30 (2000): 225–263.

14 *Id.*

15 *Id.* at 233.

16 *Id.* at 233.

17 *Id.* at 235.

18 *Id.* at 234.

19 *Id.* at 234.

20 *Id.* at 234.

21 At this point a preemptive clarification might be useful. If you find yourself bristling at the side-by-side comparison of the Bush Pump and Google Glass, the former a provider of basic goods and the latter a mere luxury item, you're in good company. My aim in this chapter is not to provide an overall moral assessment of these technologies. My aim is to provide an analytic toolkit that explicitly links engineering design to social norms in a way that will empower engineers to consider, in a more nuanced way, the interplay between design and social norms in their daily work. I leave the overall moral assessment of particular technologies to future work and would encourage anyone interested in such work to consider using social failure modes as a useful analytic tool for communicating to engineers reasons for endorsing or rejecting particular aspects or features of those technologies.

22 C. Bicchieri, *The Grammar of Society: The Nature and Dynamics of Social Norms* (Cambridge: Cambridge University Press, 2006).

23 I take as a starting point the fact that social norms are unavoidably designed into technological artifacts. For a fuller explanation of this starting point, see J. Millar, "*Technology as Moral Proxy: Autonomy and Paternalism by Design*," *IEEE*

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024

Technology & Society Magazine 34(2) (2014): 47–55.

24 "Users" is understood broadly here to include direct users, as well as other stakeholders (e.g., policymakers, citizens, etc.) impacted, directly or indirectly, by the artifact.

25 L. Winner, *The Whale and the Reactor—A Search for Limits in an Age of High Technology* (Chicago: University of Chicago Press, 1986).

26 This example is largely borrowed from Bruno Latour, who also describes how we delegate ethical (social) norms to seatbelts, robots, traffic calming obstructions and walls. See B. Latour, "Where Are the Missing Masses? The Sociology of a Few Mundane Artifacts," in *Shaping Technology/Building Society*, ed. W.E. Bijker and J. Law (Cambridge, Mass: MIT Press, 1992), 225–258. For a similar argument for technology's ability to *script*, or prescribe, behavior, see M. Akrich, "The De-Scription of Technical Objects," in *Shaping Technology/Building Society*, ed. W.E. Bijker and J. Law (Cambridge, Mass: MIT Press, 1992), 205–224.

27 J. Millar, "Ethics Setting for Autonomous Vehicles," in *Robot Ethics 2.0: From Autonomous Cars to Artificial Intelligence*, ed. P. Lin, R. Jenkins, and K. Abney (New York: Oxford University Press, 2017), 20–34.

28 P-P. Verbeek, "Materializing Morality: Design Ethics and Technological Mediation," *Science, Technology & Human Values* 31 (2006): 366.

29 *Id.*

30 Eveleth, "Google Glass Wasn't a Failure. It Raised Crucial Concerns."

31 For various discussions of the moral/conventional distinction and the controversies and questions it raises, see: E. Machery and R. Mallon, "Evolution of Morality," in *The Moral Psychology Handbook*, ed. J. Doris and the Moral Psychology Research Group (New York: Oxford University Press, 2010); D. Kelly, S. Stich, K.J. Haley, S.J. Eng, and D. Fessler, "Harm, Affect, and the Moral/Conventional Distinction," *Mind & Language 22*(2) (2007): 117–131; S. Nichols, *Sentimental Rules: On the Natural Foundations of Moral Judgment* (New York: Oxford University Press, 2004).

32 Kelly et al., "Harm, Affect, and the Moral/Conventional Distinction," 117.

33 Engineers Canada, "2 The Code of Ethics," *Public Guideline on the Code of Ethics* (2016), https://engineerscanada.ca/publications/public-guideline-on-the-code-of-ethics.

34 Social failures can also deliver harms to non-humans, a fact that deserves careful consideration but is beyond the scope of this chapter. Social failures can result in contaminated rivers, polluted air, loss of habitat and biodiversity, etc.

35 M. Hvistendahl, "Inside China's Vast New Experiment in Social Ranking," *Wired* (December 14, 2017), https://www.wired.com/story/age-of-social-credit/

36 W. Bijker, T. Hughes, and T. Pinch, *The Social Construction of Technological Systems* (Cambridge, MA: MIT Press, 1987).

37 Verbeek, "Materializing Morality: Design Ethics and Technological Mediation," 365.

38 J. Vincent, "Twitter Taught Microsoft's AI Chatbot to Be a Racist Asshole in Less than a Day," *The Verge* (March 24, 2016), https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist; S. Kleeman, "Here Are the Microsoft Twitterbot's Craziest Racist Rants," *Gizmodo* (March 24, 2016), https://gizmodo.com/here-are-the-microsoft-twitter-bot-s-craziest-racist-ra-1766820160.

Downloaded from https://academic.oup.com/edited-volume/34287/chapter/290667331 by University of Ottawa user on 19 July 2024