

Information processing and societal threat

Andrew Perfors¹

¹University of Melbourne

Human information processing and decision making is an extremely complex topic that we can only scratch the surface of here. In general, it is shaped by all aspects of our cognition, from the lowest level (what we remember and pay attention to) to the highest level (what our values are). And, because people don't exist in a vacuum from each other, these things are all heavily influenced by other people: the societal and cultural context we live in. This chapter will give an overview of the most important factors relevant to the topic of threat – how individuals behave in response to threats. This can be split into two subtopics: first, what shapes how we *assess* threats, and second, what shapes how we make decisions once we decide a threat is present.

HOW DO PEOPLE ASSESS THREATS?

Threat assessment varies considerably depending on whether the threats are immediate or more long-term. For immediate threats one of the main issues is making good enough decisions quickly enough, whereas for long-term threats the more difficult problems often revolve around how to accurately assess risks, outcomes, and costs that are often highly uncertain or lie far in the future. Immediate threats tend to be personal, for the simple reason that societies cannot communicate and mobilise (as a cohesive unit) on a rapid time scale; any immediate societal action is mostly a combination of institutional action and the emergent phenomenon built out of the immediate threat responses of many individuals.

Longer-term threats can be both personal and societal (although of course every societal threat is interpreted and evaluated individually by each person). A significant challenge that individuals have when evaluating personal long-term threats is that they very often do not have all of the information they need, or are susceptible to biases that make them susceptible to poor decision making. That said, individuals avoid some of the largest issues that face societies over the long-term, which is the problem of integrating information and mobilising action on a collective level.

This chapter focuses on information processing in response to threat on both individual and societal level, and thus mainly over the long term. However, since some of the psychological processes are rooted in how we respond to threats in the immediate term, we begin by considering that issue.

Immediate threat assessment

The most salient part of immediate threats is their rapidity. Whether they come from a person or the environment, speed is of the essence: it is usually better to have a fast response, even if it's not perfect, than it is to not act at all because of the time taken to figure out what to do (Gilbert, 1993). Similarly, long-term costs are much less relevant: it is better to live to see another day and deal with those costs then.

Because of these factors, for immediate threats, assessment and decision-making are closely intertwined, and often not conscious (Flannelly et al., 2007). Assessment is often based on subtle cues that a person might not even be aware of, much less able to verbalise: a person's body language might seem "off" or hard to reconcile with their words; you might hear a sound or see something that seems out of place. As such, one's sensitivity to those cues is primarily emotional and instinctual rather than rational (Panksepp, 1998). It is heavily influenced by prior experience, which means that people who have been through traumatic experiences have elevated threat assessment in situations that resemble them; this can be an advantage if the threat is really there, but has negative effects when it is not (Bryant and Harvey, 1997; Steimer, 2002). Regardless, emotion is dominant, and the most relevant emotion is fear (Loewenstein et al., 2001). Physiologically, fear acts to optimise fast decision-making: it spikes adrenaline and cortisol, increases heart rate, and pulls resources out of long-term needs like digestion and into short-term capabilities like the muscles and brain (LeDoux, 2000). Mentally, attention gets very focused, and time feels slower: although our responses are mediated by cognitive beliefs to some extent (MacLean, 1990), people under immediate threat don't have the sensation of considering lots of different options, but instead quickly decide on one. The array of options is usually limited and instinctual, being some variation on fight, flee, or freeze (Panksepp, 1998; Flannelly et al., 2007).

This was of necessity an extremely brief overview, but the important takeaway points are that to a large extent, our immediate threat assessment is instinctual, unconscious, emotion-driven (especially fear driving), and optimised for speed. These traits bear the hallmark of evolutionary selection in response to immediate threats. Unfortunately, as we'll see in the rest of the chapter, they are often deleterious or counterproductive when it comes to long-term or societal threats.

Societal threat assessment

Societal threats are unique because the information, effects, and costs are usually very uncertain as well as highly distributed amongst people. This makes it far more difficult to not only identify when there is a threat, but also agree on what it is and what to do about it. As a result, we are extraordinarily vulnerable to threats where the information about them is dispersed and inaccessible to many, that are hard to communicate, and where the costs are uncertain and delayed. Threats that have all of these characteristics and *also* must be acted on relatively quickly in order to avoid the worst outcomes are thus our Achilles heel, our kryptonite. This section explains why each of these factors are so challenging.

Communication is the common element that underlies almost all of the things that make societal threat processing difficult. Consider the problem of simply noticing that there is a threat. For societal-level threats, there is never a single individual – and often not even a single institution – that has access to the data that would indicate that the threat exists. As an example, the recognition of the threat of climate change took many decades. The first realisation of the role that atmospheric CO₂ plays in temperature dates back to the 19th century (Arrhenius, 1896), but at the time this was of only theoretical interest because the rate of CO₂ production was so low. It wasn't until the 1950s that concern about global warming began to emerge, but it initially existed in the form of small clues distributed across different fields, from ocean chemistry (Revelle and Suess, 1957) to atmospheric analysis (Keeling, 1960) to mathematical modelling (Manabe and Wetherald, 1967), among others. Even after that point, it took several more decades for the scientific community to come to a rough consensus, simply because the problem was so complex. And all the while, individual citizens perceived nothing directly: any tiny effects were swamped by the enormous daily and yearly weather variation that occurs naturally.

For even a small part of society (scientists) to register climate change as a threat, then, required extensive and ongoing communication as well as considerable expertise and time. Luckily, the scientific community is relatively small and has relatively efficient processes for sharing information; however, one can imagine that other threats – for instance, identifying the actions of an adversary who is actively attempting to hide – might be far more difficult to assess.

Regardless, this was the easy part: as we have seen, for threats from climate change to COVID-19, widespread awareness of the existence of the threat (and even a rough consensus about its nature) emerged far earlier among scientists than among the population as a whole. Achieving an awareness and consensus among the people who can actually make decisions, however, is vital to implementing a successful threat response. This means not only making leaders aware and convinced of the problem, but also making them want to act on it – which, at least for democratic societies, means that the population must be aware of and convinced of the problem. It is here that the communication problem is most severe.

To see why this is such a problem, consider the issue from the perspective of a random individual in society, maybe someone like yourself. You are being told information – often conflicting, often lacking in details – about a threat which you yourself cannot perceive directly and do not have the training or data to analyse. What and who do you believe? Can you trust that they're correct? Are they lying to you?

We devote substantial cognitive resources to figuring out if somebody is lying to us, but most of our skills have been adapted for in-person communication and are poor even there (DePaulo et al., 2003). The best way to figure out if someone is lying to you is check what they say against your experience: this is why nobody lies about things like the colour of grass. Unfortunately, we can't do that kind of checking with most societal threats, since we primarily know about them from other people. Did Trump engineer a coup? Are vaccines safe? Is climate change real? For all of these things a single person has only a tiny bit of direct evidence at most (e.g., having been vaccinated and not died). Moreover, few of us would have the expertise or time to evaluate indirect data even if we had it.

We handle this very difficult situation by relying on numerous cognitive short-cuts and heuristics in order to figure out what is true and who is trustworthy. One such heuristic is *fluency*: the easier information is to process, the more confident we feel that we understand it, the less we scrutinise it and the more we believe that it is correct (Song and Schwarz, 2008). Not only that, we tend to like it more

(Petty et al., 2007), feel safer (Song and Schwarz, 2009), and feel more confident (Koriat, 2008) and interested in it (Shulman and Sweitzer, 2018). As a result of this preference for simplicity and fluency, we find straightforward, easy-to-understand narratives most believable (Bullock et al., 2021). Unfortunately, because today's world is so complex, communicating the truth often means communicating the complexity. Basing believability on fluency and simplicity can therefore lead us seriously astray.

Another common heuristic is *frequency*, or *familiarity*: the more frequently we encounter some information, the more we believe it (Pennycook et al., 2018; Pillai and Fazio, 2021; Swire et al., 2017; Ransom et al., 2021). This phenomenon is very strong and often unconscious (Hasher et al., 1977). It is why advertising works, why spending time in a prejudiced environment makes it easy to adopt those prejudices, and why repeated positive affirmations make us feel better about ourselves. Relying on frequency is sensible if you're in an environment where false things are penalised, but we are not in one. Instead, today's media environment has disentangled frequency from truth: thanks to algorithms, the structure of our social networks, and our cognitive biases, seeing something multiple times means little about whether it is true (Vosoughi et al., 2018; Mosleh and Rand, 2021; Hindman et al., 2022). Indeed, even without these factors, more extreme views tend to get amplified as they spread (Navarro et al., 2018).

We also rely on heuristics that have nothing to do with the information itself. For instance, even as young children, we are much more likely to believe information when it comes from a trusted source (Heesacker et al., 1983; Koenig et al., 2004; Shafto et al., 2012). This is not actually irrational if you have no way of verifying something for yourself. In a world where you interact with few people over an extended period of time, trust is a good metric: you can be somewhat confident that the people who were reliable and truthful before will continue to be so. The problem is that this metric does not work so well today. Very often we must evaluate media sources that we don't have any personal direct experience with at all. That means we have to do research on their reputation itself, which is prone to the same errors that any other sort of information is. Information about what somebody is *really* like may not be out there, and if it is, it can be hard to discern truth from falsehood. This is especially the case when there might be substantial incentive to spread rumours or ruin a reputation in order to sow distrust.

Most people approach this problem by, again, relying on cognitive heuristics, this time for deciding who to trust: that is, who is both competent and moral (Wojciszke, 1994; Fiske et al., 2007). For this we often try to rely on nonverbal cues like body language or confidence, but such cues are extremely unreliable and we consistently overestimate our ability to use them (Vrij et al., 2019). They are even more useless when filtered through the selective editing of the media. Indeed, information providers sometimes distort the truth in order to *increase* trust if they feel the truth is unbelievable (Perfors et al., 2018). Perhaps because of these factors, another heuristic we use for trust is identity: we consistently assume that people from our own group are more trustworthy (Tajfel and Turner, 2004; Cohen, 2003; Montrey and Shultz, 2022). This is often not without reason: they are more likely to share our values and experiences. But it can lead us astray, since somebody from our group might not have our best interests in mind. Even if they do, they might be incorrect themselves and sharing misinformation as a result.

This latter danger is especially dangerous because our cognitive biases for whether we should *believe* something interact very badly with our biases about whether we should *share* something. Although most of us prefer to share things that we think are true (Pennycook et al., 2021), truth is not the only thing we care about. We also prefer to share things that elicit strong and active emotion, especially negative emotion like self-righteous anger or fear (Peters et al., 2009; Brady et al., 2017; Bebbington et al., 2017; Fay et al., 2021). We also want to share information that achieves social aims, like connecting with others (Baumeister and Leary, 1994) or advertising who we are (Berger and Heath, 2008; Coppini et al., 2017). Given these tendencies, it is easy to see how false news that makes us angry or scared – or that appeals to our social group – might spread much faster and much further than the boring truth (Vosoughi et al., 2018). This makes us see it more often, which due to the increased frequency makes us believe it even more.

One might ask where the role of reason and rationality is in all of this. Surely we are capable of evaluating the truth of a thing based on whether it makes sense and whether it is consistent with other things that we know? Indeed we are, but again it is difficult when we don't have direct access to the actual ground truth. The best we can do is analyse and interpret information provided by the people who do have that data. This involves evaluating whether it coheres with your existing beliefs and knowledge. Very little knowledge occurs in isolation; our beliefs form a thick web of interconnected theories and ideas and premises (Carey, 2009). When that web is broadly correct, then analysing new information by comparing it against what we already know is a very sensible thing to do.

However, there are several problems with this strategy. On an individual level, it is possible to get sucked so far into a web of misinformation that the entire web is incorrect – or, at best, there are a core of truths but they are so entangled and intertwined with falsehood that analysis itself goes astray. When your premises are false, only further falsehood makes sense. The end result is that people who tend to believe in one conspiracy also believe in others (Goertzel, 1994; Dagnall et al., 2015; Williams et al., 2022). People who believe in conspiracy theories do so for all sorts of reasons. These include a lack of trust in government (Moore, 2018); a way to deal with feelings of anxiety (Liekfett et al., 2021) or powerlessness (Abalakina-Papp et al., 1999); an urge to strengthen group belonging, especially if it is perceived as under threat (Uscinski and Parent, 2014); and a desire to understand global and national events (Leman and Cinnirella, 2013). Importantly, however, these people have the same basic cognitive and emotional apparatus that we all do (Oliver and Wood, 2014). The difference is that slight initial variations in information they were exposed to, the nature of the group they were born into, or their emotional needs got magnified and corrupted through this process, resulting in an epistemically closed belief structure built on a web of falsehoods (Douglas et al., 2019).

The other problem with the strategy of analysing information based primarily on whether it is consistent with existing beliefs is that if most people do this it leads to polarisation. Simulation studies demonstrate mathematically why this occurs: small initial differences in opinion are magnified, leading to further division, which continues until the end result is a population composed of subgroups that no longer engage with each other at all (Axelrod, 1997; Baldassarri and Bearman, 2007; Kashima et al., 2021b). The mechanism, which has been observed repeatedly in people, is that if someone provides information that is inconsistent with our prior beliefs, we trust the person less and sometimes believe that information less as well (Nyhan and Reifler, 2010; Bail et al., 2018; Collins et al., 2018). In the absence of an external mechanism that can provide a selection advantage for true statements, a heterogeneous population will thus polarise into smaller and more homogeneous groups (Perfors and Navarro, 2019).

Polarisation and conspiracy theories are not the only ways our reason can lead us astray: once we believe something incorrect, it is extremely hard to eliminate that belief. This is not just because of a potential lack of trust in the source of the new information; even if we fully trust it, the initial incorrect impression lingers and returns (Johnson and Seifert, 1994; Ecker et al., 2011). The reason for this phenomenon, known as the continued influence effect, is not fully clear, but it probably has something to do with the way we remember and construct our theories of the world (Guillory and Geraci, 2010; Ecker et al., 2015; Gordon et al., 2017). When we learn an untruth, we make sense of it by integrating it with the rest of our knowledge; thus, even if we are told it is incorrect, it is still more salient and the links to everything else are difficult to erase. For this reason, presenting detailed alternate explanations about why the original misinformation was incorrect is more likely to lead to successful correction, but even these strategies are fragile (Lewandowsky et al., 2012; Chan et al., 2017). Initial beliefs are the seeds from which future beliefs grow; once they have taken root, it is very difficult to uproot them completely.

All of these problems plague the communication of *any* kind of information, but information about societal threats is especially vulnerable to complication. This occurs for a few reasons.

First, information about a threat is scary, and fear and uncertainty makes people less likely to think out of the box or consider creative solutions (Mueller et al., 2012). For societal threats especially, the solution is very often not obvious and may require fundamental changes in aspects of how we operate. We are less likely to consider those or evaluate them fairly in a climate of fear.

Moreover, fear is deeply unpleasant, as is anxiety, which is what fear turns into when maintained over the long term in a dynamic and uncertain environment (LeDoux, 1996; Gilbert, 2001). Because these emotions are so aversive, we are highly motivated to want to stop feeling them. One good way to remove them is to deny that the problem exists: this is difficult to do for a problem that is direct and obvious, but extremely easy to do if it is complicated, the effects are not being felt yet, and it is being communicated by people who you don't already trust highly. Another way to remove fear and anxiety is to convince yourself that somebody else will effectively handle it. People turn to a certain kind of leader when they are afraid and no longer trust in the existing power structures: a "strong man" who speaks simply and confidently and who provides an easy explanation and visible scapegoat for the problem (Oesterreich, 2005; Sprong et al., 2019). In other words, fear, anxiety, and anomie are recipes for trusting in authoritarian leaders, as well as tightening norms and being less willing to tolerate ambiguity or dissent (Gelfand, 2018). This sort of polarisation, rigidity, communication breakdown, and lack of society-wide trust is exactly the opposite of what is needed to effectively respond to most societal threats (Kashima et al., 2021a).

This brings us to another aspect of societal threats, especially the ones we face nowadays: they are *heterogeneous*. In other words, the effects and knowledge about them are distributed very differently among different parts of the population. Consider the case of the COVID-19 pandemic. The sort of people who have the relevant expertise are part of a group (e.g., scientists and “elites”) that is often distrusted by certain segments of the population. And the effects are felt and seen much more by some than others, such as people with underlying health conditions, minorities, and the poor (Pan et al., 2020; Upshaw et al., 2021). This means that those who fail to perceive the effects might not only fail to pay any immediate price for getting things “wrong”, but might not even notice any problem at all. For those people, any responses to the threat could pose far higher immediate costs than the threat itself. This situation holds not just for our recent pandemic, but for any societal threat where the knowledge are unevenly distributed: there are strong incentives for the people who are unaffected to discount the reality of the threat.

So far the examples we have considered have been environmental threats like COVID-19 or climate change. However, a savvy malign agent – a foreign enemy or domestic terrorist – could easily exacerbate and cultivate the same factors that make these societal threats so hard to assess. Successful aggression in the grey zone (e.g., interfering in an election through propaganda, media manipulation, and so forth) would thus have the same character as these other threats. Evidence would be distributed throughout the population, so most people see nothing and nobody sees everything. It would subvert and destroy the very tool we need to have in order to respond to a societal threat – our trust in each other and our ability to communicate as an undivided and unpolarised society. Lacking that tool, we would become crippled in the next step: our ability to make good decisions and take appropriate actions to overcome the threat.

HOW DO PEOPLE MAKE DECISIONS ABOUT THREATS?

Making decisions on a societal level requires effective communication for the same reasons that identifying threats on a societal level does, and as a result the same issues arise. These are magnified yet again for decision-making. Even as an individual, making decisions in uncertain environments is hard, but societal-level threats have additional characteristics that make them especially hard. To understand why, let’s first discuss what we know about how individuals make decisions normally.

Short-term decision making in individuals

How people make decisions (and how they *should* make decisions) is a very complex problem that has been studied for many decades, so the treatment here is of necessity superficial. But, broadly speaking, three things matter when making a decision about what action would be best to take: utility (how good or bad the outcomes of the various actions are), risk (the probability of those outcomes), and uncertainty (how sure we are about the utilities, probabilities, and outcomes). Let’s discuss each in turn.

At the simplest level, we can think of the problem of decision making as the problem of deciding which of some set of actions to take. Each of the actions is associated with an outcome (or set of outcomes) and each of the outcomes is associated with a utility: whether it is good or bad, and how good or bad it is. Most models of decision-making try to quantify utilities by putting some kind of numeric value on them so that good ones are positive numbers and bad are negative (see, e.g., Kahneman, 2013). This means that while in principle any outcome can be assigned a utility, it is a lot easier to study outcomes that are easily quantified, such as those which we can assign a monetary value to. In more complicated analyses, the costs of an outcome are often included in its utility as well, so that an action (e.g., studying hard for an exam) might have outcomes with high utility (e.g., getting a good grade) but also some costs, denoted by low utility (e.g., having to skip a party). Outcomes also are associated with probabilities: the chance of actually achieving that outcome after taking that action. These probabilities matter because there is a big difference between an action that will result in some reward only 10% of the time vs one that will result in it 100% of the time, and people are sensitive to this.

Within this framework we can calculate normatively what action one *should* take: which one is associated with the highest utility, taking its probability into account. Of course, people don’t do all these calculations, particularly when there are multiple possible actions with multiple outcomes. Instead, we use heuristics, and as a result don’t treat all options as their utilities might lead us to expect.

One of the most robust characteristics of our decision-making is that losses loom higher than gains: that is, losing \$5 feels worse than gaining \$5 feels good (Kahneman and Tversky, 1979). As a result, people are loss averse: like the old saw says, “a bird in the hand is worth two in the bush.” People also tend to be risk averse, preferring outcomes that are certain to happen over outcomes that aren’t certain, even

if they have much higher utility (Pratt, 1964; Kahneman and Tversky, 1979). Because of these factors, we are susceptible to things like framing effects: a medical intervention with 80% chance of survival sounds much better to us than one with 20% chance of death even though they are exactly the same (Levin et al., 1998). We also use other heuristics that result in flawed reasoning. For instance, because we can't necessarily keep all possible actions, outcomes, and their probabilities in mind, the ones that are most vivid or salient gain more importance than they might deserve (Tversky and Kahneman, 1973). This is why people tend to be more frightened of things like terrorist attacks than they are of car accidents, even though the former are far more rare than the latter.

Real-world decisions have an additional factor that makes everything even harder: the presence of uncertainty. Uncertainty does not just mean low-probability: it occurs when a reasoner doesn't know what values to give something. We can be uncertain *about* the probability of an action or outcome: we might know that working hard to study will improve the chance of a good grade, but we don't necessarily know what the exact probability will be or by how much it will improve. We can also be uncertain about utility: would having a baby be really good, really bad, or something in between? Because utilities are subjective and often differ wildly between people, it can be very hard for an individual to estimate them for outcomes they have not previously experienced. Indeed, we often even have uncertainty about what the possible space of actions and outcomes is; there may be many we just don't think of.

Most real-world decisions are marked by a lot of uncertainty about all of these things, and people find uncertainty itself highly unpleasant; indeed, we often go to extra effort to seek out information even when we know the information will be useless for decision-making (Lanzetta and Driscoll, 1966; Grant et al., 1993; Bennett et al., 2016). We also prefer to avoid outcomes or situations that are uncertain or ambiguous, compared to more certain or less ambiguous ones (Ellsberg, 1961; Camerer and Weber, 1992). Probably in part because of this dislike of uncertainty, we tend to discount the future relative to the present: future losses or gains matter less to us than present ones, and the farther in the future they are, the less they matter (Samuelson, 1937; Frederick et al., 2002). This isn't necessarily irrational, since uncertainties multiply, and the further in the future something is, the less likely you were to have correctly estimated the utilities and probabilities associated with it.

All of the factors discussed so far can be quantified, at least in theory, even if it's not easy to do so in practice. However, not everything that matters to human decision-making can be given a numerical value. For instance, when assigning moral weight to choices, we tend to think that intent is important (Malle and Knobe, 1997; Mikhail, 2007) and that active choices are more morally relevant than passive ones (Spranca et al., 1991; Cushman et al., 2006). Thus, somebody who actively and intentionally murders is much more blameworthy than somebody who allows someone to die through accident or inaction, even if the utilities and probabilities are the same. Conversely, somebody who actively saves a life is more heroic than somebody who inadvertently does so.

This overview describes individual decision-making, but since society is made up of many individuals, societal decision-making involves the same sort of reasoning processes and falls prey to the same kinds of mistakes. However, it is even harder, especially when it comes to evaluating societal threats in particular.

Societal-level decisions in response to threat

When it comes to making decisions about how to respond to societal threats, the same decision-making processes and heuristics occur but the task is even more challenging. For one thing, all of the issues caused by the need to communicate with one another are still there. In this context, failure to communicate accurately about the situation means that the communication can distort the information we need to make decisions about, so whoever is making the decision – whether it be a group of people in power or everyone in an election – is not doing so based on accurate information. The utilities, costs, and probabilities may be wrong, and especially wrong in certain directions. Since things that are simpler, more emotional, and more negative are passed on more, that means that we should expect losses and negative utilities to be over-emphasised and made more salient, and nuance about probabilistic outcomes lost.

Another thing that makes societal-level decisions about threats particularly difficult is that very often these are decisions with high uncertainty (Lewandowsky et al., 2015). Climate change, COVID-19, terror events – all of these involve very complex situations with a lot of unknowns, particularly at first. In uncertain situations, our tendency is to make optimistic assumptions about how they will turn out (Lench et al., 2014) and want to do nothing while we gain more information (Freudenburg and Muselli, 2013). This is sometimes a good strategy, but for many kinds of societal threats, early action is far superior to

late action. When those threats build on each other over time – as with the exponential growth seen in COVID-19, or complex ecological systems like the climate – then by the time the evidence is visible and unarguable, it may be too late for simple solutions. Moreover, uncertainty makes people uneasy and thus tends to make cooperation less likely (Raihani and Aitken, 2011). For all of these reasons, uncertainty itself can lead to poor adaptation and increased vulnerability to threats (Kennett and Marwan, 2015).

Another factor that makes societal threats particularly pernicious from a decision making point of view is that they are often associated with very low probability but very bad outcomes. War, pandemic, climate change: for all of these the probability of the collapse of the country or even civilisation is tiny, but not zero. This not only makes them terrifying, but these sort of situations are also extremely hard to know how to evaluate (Hertwig et al., 2004; Madan et al., 2014). What kind of negative utility should one put on this sort of outcome? It feels like a qualitatively different *kind* of thing than just, say, a recession, and thus the temptation is to say that avoiding that outcome is priceless; but if that is the case, it means that literally *anything* one does to avoid that, no matter the costs, even if the odds of such a thing are miniscule. Conversely, though, these tiny probabilities are not only hard to estimate, but they are easy to overlook entirely. It is not even clear, normatively, how one *should* handle these situations (Buchholz and Schymura, 2012). People tend to err towards ignoring tiny probabilities, but there are clear individual differences (Croft et al., 2021) and susceptibility to framing effects (Perfors and van Dam, 2018), both of which make societal consensus even harder to achieve.

In all of this, we should not lose sight of a final major factor that makes societal-level decision making much more difficult than individual-level decision making: the costs and benefits are so unevenly distributed. A single person is at least one coherent entity, which means that whatever costs and benefits result from the decision affect the decision-maker. Societal-level threats are very different. Society is made up of many different – and highly heterogeneous – subgroups: defined by race, by class, by gender, by education, by language, by geography, and so on. This heterogeneity leads not only to the communication and assessment issues described above, but also means that any threat *itself* will be uneven. Some groups – usually the poor and otherwise disenfranchised, but not always – will bear more of the negative outcomes (Pan et al., 2020; Upshaw et al., 2021). Moreover, decision-making capacity is also distributed: while the ideal of a perfect democracy is that everybody gets an equal say, that ideal is far from the reality. Most decisions are made in real time by the people with the power, who may have little idea or concern about the impacts of their decision on people who are unlike them (Nielsen et al., 2021; Zelner et al., 2022). And even if they do care, the difficulty in quickly and accurately communicating information from all of these heterogeneous subgroups means they often do not have the information they need.

Finally, making decisions about societal threats is difficult because societies are complex systems (Urry, 2005). What this means is that it is impossible to fully predict the outcome of any decision. Most of this section – and most psychological research on decision-making – has presumed that the set of outcomes and probabilities can be enumerated with reasonable accuracy; however, that presumption is almost certainly false, especially as the complexity of the problem grows and the future horizon expands. In such a situation, it is especially important that decision-making be responsive to new information, something that the communication and aggregation issues discussed earlier makes difficult.

In sum, because of the way individuals and societies process information, responding to societal threats is a wicked problem. The vast majority of this chapter has focused on the factors that make it so hard – not because we need to despair, but because if we are to have any hope of solving it, we need to have a frank and honest look at what makes it difficult to solve.

REFERENCES

- Abalakina-Papp, M., Stephan, W., Craig, T., and Gregory, W. (1999). Beliefs in conspiracies. *Political Psychology*, 20(3):637–647.
- Arrhenius, S. (1896). On the influence of carbonic acid in the air upon the temperature of the ground. *Philosophical Magazine and Journal of Science*, 41(251):237–276.
- Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *Journal of Conflict Resolution*, 41(2):203–226.
- Bail, C., Argyle, L., Brown, T., Bumpus, J., Chen, H., Hunzaker, F., Lee, J., Mann, M., Merhout, F., and Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221.
- Baldassarri, D. and Bearman, P. (2007). Dynamics of political polarization. *American Sociological Review*, 72(5):784–811.
- Baumeister, R. and Leary, M. (1994). The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin*, 117(3):497–529.
- Bebbington, K., MacLeod, C., Ellison, T., and Fay, N. (2017). The sky is falling: Evidence of a negativity bias in the social transmission of information. *Evolution and Human Behavior*, 38(1):92–101.
- Bennett, D., Bode, S., Brydevall, M., Warren, H., and Murawski, C. (2016). Intrinsic valuation of information in decision making under uncertainty. *PLoS Computational Biology*, 12(7):e1005020.
- Berger, J. and Heath, C. (2008). Who drives divergence? Identity signaling, outgroup dissimilarity, and the abandonment of cultural tastes. *Journal of Personality and Social Psychology*, 95(3):593–607.
- Brady, W., Wills, J., Jost, J., Tucker, J., and van Bavel, J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28):7313–7318.
- Bryant, R. and Harvey, A. (1997). Attentional bias in posttraumatic stress disorder. *Journal of Traumatic Stress*, 10(4):635–544.
- Buchholz, W. and Schymura, M. (2012). Expected utility theory and the tyranny of catastrophic risks. *Ecological Economics*, 77:234–239.
- Bullock, O., Shulman, H., and Huskey, R. (2021). Narratives are persuasive because they are easier to understand: Examining processing fluency as a mechanism of narrative persuasion. *Frontiers in Communication*, 6:719615.
- Camerer, C. and Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5(4):325–370.
- Carey, S. (2009). *The origin of concepts*. Oxford University Press.
- Chan, M., Jones, C., Hall Jamieson, K., and Albarracín, D. (2017). Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation. *Psychological Science*, 28(11):1531–1546.
- Cohen, G. (2003). Party over policy: The dominating impact of group influence on political beliefs. *Journal of Personality and Social Psychology*, 85:808.
- Collins, P., Hahn, U., von Gerber, Y., and Olsson, E. (2018). The bi-directional relationship between source characteristics and message content. *Frontiers in Psychology*, 9:18.
- Coppini, D., Duncan, M., McLeod, D., Wise, D., Bialik, K., and Wu, Y. (2017). When the whole world is watching: A motivations-based account of selective expression and exposure. *Computers in Human Behavior*, 75:766–774.
- Croft, J., Grisham, J., Perfors, A., and Hayes, B. (2021). Risking everything in obsessive-compulsive disorder: An analogue decision-making study. *Journal of Psychopathology and Behavioral Assessment*, 22:1–12.
- Cushman, F., Young, L., and Hauser, M. (2006). The role of conscious reasoning and intuitions in moral judgment: Testing three principles of harm. *Psychological Science*, 17(12):1082–1089.
- Dagnall, N., Drinkwater, K., Parker, A., Denovan, A., and Parton, M. (2015). Conspiracy theory and cognitive style: A worldview. *Frontiers in Psychology*, 6:206.
- DePaulo, B., Lindsay, J., Malone, B., Muhlenbruck, L., Charlton, K., and Cooper, H. (2003). Cues to deception. *Psychological Bulletin*, 129(1):74–118.
- Douglas, K., Uscinski, J., Sutton, R., Cichocka, A., Nefes, T., Ang, C., and Deravi, F. (2019). Understanding conspiracy theories. *Advances in Political Psychology*, 40(1):3–35.
- Ecker, U., Lewandowsky, S., Cheung, C., and Maybery, M. (2015). He did it! she did it! no, she did not! multiple causal explanations and the continued influence of misinformation. *Journal of Memory and*

- Language*, 85:101–115.
- Ecker, U., Lewandowsky, S., Swire, B., and Chang, D. (2011). Correcting false information in memory: Manipulating the strength of misinformation encoding and its retraction. *Psychonomic Bulletin and Review*, 18:570–578.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669.
- Fay, N., Walker, B., Kashima, Y., and Perfors, A. (2021). Socially situated transmission: The bias to transmit negative information is moderated by the social context. *Cognitive Science*, 45(9):e13033.
- Fiske, S., Cuddy, A., and Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, 11:77–83.
- Flannelly, K., Koenig, H., Galek, K., and Ellison, C. (2007). Beliefs, mental health, and evolutionary threat assessment systems in the brain. *The Journal of Nervous and Mental Disease*, 195(12):996–1003.
- Frederick, S., Loewenstein, G., and O’Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2):351–401.
- Freudenburg, W. and Muselli, V. (2013). Reexamining climate change debates: Scientific disagreement or scientific certainty argumentation methods. *American Behavioral Science*, 57:777–795.
- Gelfand, M. (2018). *Rule Makers, Rule Breakers: How Tight and Loose Cultures Wire Our World*. Scribner.
- Gilbert, P. (1993). Defence and safety: Their function in social behaviour and psychopathology. *British Journal of Clinical Psychology*, 32(2):131–153.
- Gilbert, P. (2001). Evolutionary approaches to psychopathology: The role of natural defences. *Australian and New Zealand Journal of Psychiatry*, 35(1):17–27.
- Goertzel, T. (1994). Belief in conspiracy theories. *Political Psychology*, 15(4):731–742.
- Gordon, A., Brooks, J., Quadflieg, S., Ecker, U., and Lewandowsky, S. (2017). Exploring the neural substrates of misinformation processing. *Neuropsychologia*, 106:216–224.
- Grant, S., Kajii, A., and Polak, B. (1993). Intrinsic preference for information. *Journal of Economic Theory*, 83(2):233–259.
- Guillory, J. and Geraci, L. (2010). The persistence of inferences in memory for younger and older adults: Remembering facts and believing inferences. *Psychonomic Bulletin and Review*, 17:73–81.
- Hasher, L., Goldstein, D., and Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning and Verbal Behavior*, 16(1):107–112.
- Heesacker, M., Petty, R., and Cacioppo, J. (1983). Field dependence and attitude change: Source credibility can alter persuasion by affecting message-relevant thinking. *Journal of Personality*, 51:653–666.
- Hertwig, R., Barron, G., Weber, E., and Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8):534–539.
- Hindman, M., Lubin, N., and Davis, T. (2022). Facebook has a superuser-supremacy problem. *The Atlantic*.
- Johnson, H. and Seifert, C. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20:1420–1436.
- Kahneman, D. (2013). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(3):263–291.
- Kashima, Y., Dennis, S., Perfors, A., and Laham, S. (2021a). Culture and global societal threats: Covid-19 as a pathogen threat to humanity. *Group Processes and Intergroup Relations*, 24(2):223–230.
- Kashima, Y., Perfors, A., Ferdinand, V., and Pattenden, E. (2021b). Ideology, communication, and polarisation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 376(1822):20200133.
- Keeling, C. (1960). The concentration and isotopic abundances of carbon dioxide in the atmosphere. *Tellus*, 12(2):200–203.
- Kennett, D. and Marwan, N. (2015). Climatic volatility, agricultural uncertainty, and the formation, consolidation and breakdown of preindustrial agrarian states. *Philosophical Transactions of the Royal Society A*, 373:20140458.
- Koenig, M., Clement, F., and Harris, P. (2004). Trust in testimony: Children’s use of true and false statements. *Psychological Science*, 15(10):694–698.
- Koriat, A. (2008). Metacognition and consciousness. In *Cambridge Handbook of Consciousness*, pages

- 289–325. Cambridge University Press.
- Lanzetta, J. and Driscoll, J. (1966). Preference for information about an uncertain but unavoidable outcome. *Journal of Personality and Social Psychology*, 3(1):96–102.
- LeDoux, J. (1996). *The emotional brain*. Simon and Schuster.
- LeDoux, J. (2000). Emotion circuits in the brain. *Annual Reviews in Neuroscience*, 23:155–184.
- Leman, P. and Cinnirella, M. (2013). Beliefs in conspiracy theories and the need for cognitive closure. *Frontiers in Psychology*, 4(378).
- Lench, H., Smallman, R., Darbor, K., and Bench, S. (2014). Motivated perception of probabilistic information. *Cognition*, 133:429–442.
- Levin, I., Schneider, S., and Gaeth, G. (1998). All frames are not created equal: A typology and critical analysis of framing effects. *Organizational behavior and human decision processes*, 76(2):149–188.
- Lewandowsky, S., Ballard, T., and Pancost, R. (2015). Uncertainty as knowledge. *Philosophical Transactions of the Royal Society A*, 373:20140462.
- Lewandowsky, S., Ecker, U., Seifert, C., Schwarz, N., and Cook, J. (2012). Beliefs in conspiracy theories and the need for cognitive closure. *Psychological Science in the Public Interest*, 13:106–131.
- Liekefett, L., Christ, O., and Becker, J. (2021). Can conspiracy beliefs be beneficial? Longitudinal linkages between conspiracy beliefs, anxiety, uncertainty aversion, and existential threat. *Personality and Social Psychology Bulletin*.
- Loewenstein, G., Weber, E., Hsee, C., and Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, 127(2):267–286.
- MacLean, P. (1990). *The Triune Brain in Evolution*. Plenum.
- Madan, C., Ludvig, E., and Spetch, M. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin and Review*, 21(3):629–636.
- Malle, B. and Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology*, 33:101–121.
- Manabe, S. and Wetherald, R. (1967). Thermal equilibrium of the atmosphere with a given distribution of relative humidity. *Journal of the Atmospheric Sciences*, 24(3):241–259.
- Mikhail, J. (2007). Universal moral grammar: Theory, evidence and the future. *Trends in Cognitive Sciences*, 11:143–152.
- Montrey, M. and Shultz, T. (2022). Copy the in-group: Group membership trumps perceived reliability, warmth, and competence in a social-learning task. *Psychological Science*, 33(1):165–174.
- Moore, A. (2018). On the democratic problem of conspiracy theory politics. In Uscinski, J., editor, *Conspiracy theories and the people who believe them*, pages 111–121. Oxford University Press.
- Mosleh, M. and Rand, D. (2021). Falsehood in, falsehood out: Measuring exposure to elite misinformation on Twitter. *PsyArxiv*.
- Mueller, J., Melwani, S., and Goncalo, J. (2012). The bias against creativity: why people desire but reject creative ideas. *Psychological Science*, 23(13–17).
- Navarro, D., Perfors, A., Kary, A., Brown, S., and Donkin, C. (2018). When extremists win: Cultural transmission via iterated learning when populations are heterogeneous. *Cognitive Science*, 42:2108–2149.
- Nielsen, K., Nicholas, K., Creutzig, F., Dietz, T., and Stern, P. (2021). The role of high-socioeconomic-status people in locking in or rapidly reducing energy-driven greenhouse gas emissions. *Nature Energy*, 6:1011–1016.
- Nyhan, B. and Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior*, 32:303–330.
- Oesterreich, D. (2005). Flight into security: a new approach and measure of the authoritarian personality. *Political Psychology*, 26:275–297.
- Oliver, J. and Wood, T. (2014). Conspiracy theories and the paranoid style(s) of mass opinion. *American Journal of Political Science*, 58(4):952–966.
- Pan, D., Sze, S., Minhas, J., Bangash, M., Pareek, N., Divall, P., Williams, C., Oggioni, M., Squire, I., Nellums, L., Hanif, W., Khunti, K., and Pareek, M. (2020). The impact of ethnicity on clinical outcomes in covid-19: A systematic review. *EClinicalMedicine*, 23:100404.
- Panksepp, J. (1998). *Affective neuroscience: The foundations of human and animal emotions*. Oxford University Press.
- Pennycook, G., Cannon, T., and Rand, D. (2018). Prior exposure increases perceived accuracy of fake

- news. *Journal of Experimental Psychology: General*, 147(12):1865–1880.
- Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A., Eckles, D., and Rand, D. (2021). Shifting attention to accuracy can reduce misinformation online. *Nature*, 592:590–595.
- Perfors, A. and Navarro, D. (2019). Why do echo chambers form? the role of trust, population heterogeneity, and objective truth. In Goel, A., Seifert, C., and Freksa, C., editors, *Proceedings of the 41st Annual Conference of the Cognitive Science Society*, pages 918–923.
- Perfors, A., Navarro, D., and Shafto, P. (2018). Stronger evidence isn't always better: A role for social inference in evidence selection and interpretation. In Kalish, C., Rau, M., Zhu, J., and Rogers, T., editors, *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, pages 864–869.
- Perfors, A. and van Dam, N. (2018). Human decision making in black swan situations. In Kalish, C., Rau, M., Zhu, J., and Rogers, T., editors, *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, pages 870–875.
- Peters, K., Kashima, Y., and Clark, A. (2009). Talking about others: Emotionality and the dissemination of social information. *European Journal of Social Psychology*, 39(2):207–222.
- Petty, R., Brinol, P., Tormala, Z., and Wegener, D. (2007). The role of metacognition in social judgments. In *Social Psychology: Handbook of Basic Principles*, pages 254–284. Guilford Press.
- Pillai, R. and Fazio, L. (2021). The effects of repeating false and misleading information on belief. *Wiley Interdisciplinary Reviews in Cognitive Science*, 12(6):e1573.
- Pratt, J. (1964). Risk aversion in the small and in the large. *Econometrica*, 32(1/2):122–136.
- Raihani, N. and Aitken, D. (2011). Uncertainty, rationality and cooperation in the context of climate change. *Climate Change*, 108:47–55.
- Ransom, K., Perfors, A., and Stephens, R. (2021). Social meta-inference and the evidentiary value of consensus. In Fitch, T., Lamm, C., Leder, H., and Teßmar-Raible, K., editors, *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*, pages 833–839.
- Revelle, R. and Suess, H. (1957). Carbon dioxide exchange between atmosphere and ocean and the question of an increase of atmospheric CO₂ during the past decades. *Tellus*, 9:18–27.
- Samuelson, P. (1937). A note on measurement of utility. *The Review of Economic Studies*, 4:155–161.
- Shafto, P., Eaves, B., Navarro, D., and Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*, 15:436–447.
- Shulman, H. and Sweitzer, M. (2018). Advancing framing theory: Designing an equivalency frame to improve political information processing. *Human Communication Research*, 44(2):155–175.
- Song, H. and Schwarz, N. (2008). Fluency and the detection of misleading questions: Low processing fluency attenuates the moose illusion. *Social Cognition*, 26(6):791–799.
- Song, H. and Schwarz, N. (2009). If it's difficult to pronounce, it must be risky. *Psychological Science*, 20(2):135–138.
- Spranca, M., Minsk, E., and Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, 27(1):76–105.
- Sprong, S., Jetten, J., Wang, Z., Peters, K., Mols, F., Verkuyten, M., Bastian, B., Ariyanto, A., Autin, F., Ayub, N., Badea, C., Besta, T., Butera, F., Costa-Lopes, R., Cui, L., Fantini, C., Finchilescu, G., Gaertner, L., Gollwitzer, M., Gomez, A., Gonzalez, R., Hong, Y.-Y., Jensen, D., Jasinskaja-Lahti, I., Karasawa, M., Kessler, T., Klein, O., Lima, M., Megevand, L., Morton, T., Paladino, P., Polya, T., Renvik, T., Ruza, A., Shahrazad, W., Shama, S., Smith, H., Torres, A., van der Bles, A., and Wohl, M. (2019). “Our country needs a strong leader right now”: Economic inequality enhances the wish for a strong leader. *Psychological Science*, 30(11):1625–1637.
- Steimer, T. (2002). The biology of fear- and anxiety-related behaviors. *Dialogues in clinical neuroscience*, 4(3):231–249.
- Swire, B., Ecker, U., and Lewandowsky, S. (2017). The role of familiarity in correcting inaccurate information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(12):1948–1961.
- Tajfel, H. and Turner, J. (2004). The social identity theory of intergroup behavior. In Jost, J. and Sidanius, J., editors, *Political psychology: Key readings*, pages 276–293. Psychology Press.
- Tversky, A. and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2):207–232.
- Upshaw, T., Brown, C., Smith, R., Perri, M., Ziegler, C., Pinto, A., and Camacho-Rivera, M. (2021). Social determinants of covid-19 incidence and outcomes: A rapid review. *PLoS One*, page e0248336.

- Urry, J. (2005). The complexity turn. *Theory, Culture and Society*, 22(5):1–14.
- Uscinski, J. and Parent, J. (2014). *American conspiracy theories*. Oxford University Press.
- Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Vrij, A., Hartwig, M., and Granhag, P. (2019). Reading lies: Nonverbal communication and deception. *Annual Review of Psychology*, 70(1):295–317.
- Williams, M., Marques, M., Hill, S., Kerr, J., and Ling, M. (2022). Why are beliefs in different conspiracy theories positively correlated across individuals? Testing monological network versus unidimensional factor model explanations. *British Journal of Social Psychology*.
- Wojciszke, B. (1994). Multiple meanings of behavior: construing actions in terms of competence or morality. *Journal of Personality and Social Psychology*, 67:222–232.
- Zelner, J., Masters, N., Naraharsetti, R., Mojola, S., Chowkwanyun, M., and Malosh, R. (2022). There are no equal opportunity infectors: Epidemiological modelers must rethink our approach to inequality in infection risk. *PLoS Computational Biology*, 18(2):e1009795.