// THE MEANS TO REACH FURTHER

Foresight biases and the problem of misfuturing

by

Jayar La Fontaine

Submitted to OCAD University
in partial fulfillment of the requirements
for the degree of

Master of Design
in
Strategic Foresight and Innovation

Toronto, Ontario, Canada
April 2014

/ AUTHOR'S DECLARATION

I hereby declare that I am the sole author of this Master's Research Project. This is a true copy of the MRP, including any required final revisions, as accepted by my examiners.

I authorize OCAD University to lend this MRP to other institutions or individuals for the purpose of scholarly research.

I understand that my MRP may be made electronically available to the public.

I further authorize OCAD University to reproduce this MRP by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Signature _____

ABSTRACT

In this paper, I outlined a framework for understanding how and why future-related biases took root in the human mind, and made some suggestions about possible ways we might be able to mitigate their effect in the context of the strategic foresight practice. My primary research methods consisted of a broad literature review covering pertinent areas of cognitive science, evolutionary psychology, management science, and future studies, as well as expert interviews with experts representing these fields. I concluded the paper with suggestions about how to integrate the findings from fields that scientifically investigate human foresight and decision-making under uncertainty with the practice of strategic foresight.

ACKNOWLEDGEMENTS

Special thanks to my primary advisor Suzanne Stein, for encouraging the line of inquiry in this project, and also for the opportunity to work on several fascinating, exciting projects that shaped my perspective on the power and promise of future thinking. Thanks also to my second reader, Robert K. Logan, for his illuminating comments and directions, and especially for his encouragement to seek allies for my position in the field of media ecology.

Additional thanks the foresight practitioners, scientists, and scholars with whom I had the opportunity to correspond or interview for this project, including Madeline Ashby, Peter Bishop, Stuart Candy, Jim Dator, Mark P. Healey, Alexander Rose, Douglas Rushkoff, Karl Schroeder, Wendy Schultz, Keith Stanovich, Thomas Suddendorf, and Maggie Toplak.

My time at OCAD was immeasurably enhanced by my friendships and working relationships with other students – both in my cohort and beyond. There are too many here to enumerate, so I'll simply say: Thanks. Good luck. And see you in the field.

To **my parents**, whose unwavering help and encouragement allowed me once more to stray from the path in my sights and reimagine my own future.

To **Rachel**, whose words and actions have shown me that the task of exploring the future is best done together.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1. THE CHALLENGE OF FORESIGHT

Strategic foresight has a clear *normative* component. The key justification for conducting a strategic foresight exercise is that certain ways of imagining and thinking about the future are *better* than others for the purpose of guiding decision-making, and that we can develop better future thinking through application of the strategic foresight methodology. The first clause above is reasonably uncontroversial. As we will see, large swaths of evolutionary history were driven by difference in survival rates of creatures adept at anticipating and acting on the future and those less able; it's clear that *good futuring* makes a real difference in decision making and, therefore, survivability. But how strong is the second claim: that strategic foresight *itself* leads to better future thinking? In order to assess this claim, we need to first understand a little bit more about the way strategic foresight works.

Foresight strategist Richard Slaughter defines his craft as "the ability to create and maintain a high-quality, coherent and functional forward view, and to use the insights arising in useful organizational ways" (Slaughter 2003, p, 104). As an overview, Slaughter's rendering is fairly comprehensive, but it necessarily only *hints* at the two major complexities inherit in the strategic foresight methodology.

First, the "forward view" that strategic foresight enables us to

create and maintain is, in most cases, actually *several* forward views, or *scenarios*, exploring a space of possibilities. Practitioners of strategic foresight often stress the need to *pluralize* the future (Van Alstyne 2010, 70). At first blush, this sentiment can either confuse or exasperate, depending on the hearer's temperament and tolerance for ambiguity. What could it possibly *mean*?

The call to pluralize the future becomes especially confusing when, as often happens, people elide over the difference between and *foresight strategists and a certain brand of futurists that believe its possible to make meaningful point predictions about the future*.[1] Methodologically, and ideologically, there are many important differences between these two stances, a thorough discussion of which is beyond the purview of this project.

A common touch point, however, for both foresight strategists and these predictive futurists is their mutual admiration for science fiction narratives, and their eagerness to leverage them in their respective practices. The difference in the way they go about this reveals something of what foresight strategists typically mean by "pluralizing the future."

Predictive futurists, naturally, stress the predictive aspect of science fiction in order to make a case for the knowability of the future through keen insight, sound research, and proper inference. Foresight strategists,

[1] I'm taking care here not to paint all brands of futurism with the same brush; as futurist Wendy Schultz pointed out to me, there are many practitioners self identifying as futurists who are "interested in images of our futures broadly extent [sic] in society, as well as cultural biases, and depth analysis/critique of

on the other hand, celebrate the ability for science fiction narratives to emotionally and mentally draw us into unfamiliar worlds and imagine others or ourselves therein as agents (Schroeder, 2012). This is illustrative of the contrasting professional aims of predictive futurists and foresight strategists: predictive futurists try to corral the future *itself* by generating arguments which point to some probable future state or range of future states, while foresight strategists seek instead to multiply and vivify what psychologists call our *episodic thinking about the future,* and use these episodes as frames for assessing strategic decisions.

The call from foresight strategists to pluralize the future, therefore, is not a prelude to a lecture about the importance of considering the function of un-collapsed probability waves, or the virtues of the "many worlds" hypothesis, or the importance of Lewisian modal realism in the context of strategic work.[2] It's not about *the future,* as a dimension of time, *at all*. Rather, it is meant to encourage in the hearer a mindset conducive to the creation and curation of multiple *mental images of* possible future worlds. The cultivation of these mental images, according to Kees Van Der Heijden, accomplishes the first objective of foresight, which is to provide the "requisite variety in mental models necessary to see and perceive the outside world beyond the traditional business models," or governance models, or community building models, etc.

It also serves the other key objective of foresight, which is to use

[2] For better or worse, foresight strategists can only perform foresight exercises in the actual world, and not in merely possible worlds.

these broadly aligned mental models as backdrops against which to pre-perform future strategic plays (Van Der Heijden 2005, 132). The quality, coherence, and functionality of these mental images therefore has less to do with whether one or the other of them *come true* in the long run and much more to do with the way they interact with the cognitive toolkits, beliefs, and broader mental models of people who interact with them.

That foresight generates a multiplicity of views of the future rather than a probabilistic prediction is something that foresight strategists do a good job of communicating to clients as a baseline expectation. However, that the success of the engagement depends fundamentally on the ability for participants to examine, probe, and alter their *personal* mental images of the future, though sometimes broadly articulated, is less often thoroughly explored.

Thinking About the Future, the indispensible manual of strategic foresight edited by Andy Hines and Peter Bishop, does a better job of addressing the relationship between these two goals of the practice. The authors identify mental models as "the deeply ingrained assumptions, generalizations, or images that influence how one makes sense of and responds to the world, and suggest that changing the mental models of decision makers is a key outcome of the foresight activity. They go on to suggest that mental models are "usually biased toward the past, and are often vague or based on faulty assumptions about the future" (Hines and Bishop 2006, 30). These faulty assumptions give rise to an *official* or *default*

future based on an uncritical belief in incremental change "that pretty much preserves the current paradigm or way of doing things" (Ibid 2006, 85). The process of expanding "the range and depth of possibilities for the organization to consider" can challenge these default futures, thereby helping to reduce "the likelihood and magnitude of surprise" in the future (Ibid 2006, 85).

Here, Hines and Bishop articulate a key assumption of foresight: creating and utilizing multiple, robust "forward views" *naturally* loosens the grip that inadequate default futures have on our minds. The assumption has strong *prima facie* appeal. But I think it is mistaken, and that the matter is not nearly so straightforward. Under the right conditions – perhaps even in *most* cases – consideration of several forward views can have the effect suggested by Hines and Bishop. But, as we will see, it may also have *no* appreciable effect on an individual's commitment to the default future, or may even serve to *strengthen* their existing commitment to it in some cases.

Why might this be the case? By way of an explanation, let's return once more to the contrast between foresight and predictive futurism.

Since foresight strategists need not concern themselves with the question of the precise likelihood that any of their scenarios might accurately predict the future, it might seem like their project is comparatively less taxing than the work of futurists. I do not believe this is the case, at least in cases where foresight is robustly performed.

Accurately predicting the future is, of course, gruelingly hard work. Even in a relatively circumscribed domain of inquiry, accurate prediction work improves only through highly original thought, painstaking research and development, piecemeal engineering efforts, and relentless troubleshooting and fine-tuning.

Take forecasting the weather as an example. Just thirty years ago, hurricane prediction systems would routinely miscalculate landfalls by as much as 350 miles. Today, the average miss is about 100 miles: a threefold improvement in the predictive power of hurricane modeling. The amount of effort expended to effect even this change boggles the mind. Wrangling nonlinear, dynamic systems like hurricanes means, for instance, accounting for fluctuating barometric pressures to the fourth decimal place; meteorologists discovered that rounding to the third decimal would lead to confusing results, in which the same predictive weather model "would somehow forecast clear skies over Colorado in one run and a thunderstorm in the next" (Silver 2012a).

Computing the complex interactions of social, technological, environmental, economic, and political factors in order to make real, measurable progress in the rates of accurate prediction of world events is an exponentially harder task, and one that we've barely embarked on. The current state of the art, if one could call it that, is quite dismal: Philip Tetlock's long term study of political forecasts found, for instance, that when political experts described an event as being *absolutely certain*, it

failed to happen *one fourth* of the time. What's more, their performance was only marginally better than that of dilettantes in the subject area, and marginally *worse* than algorithms enacting barebones models of change (Tetlock 2005, 55). If weighty political experts were the example *par excellence* of our predictive capabilities in the realm of political events, this would be as if our most sophisticated hurricane modeling systems barely outperformed someone sticking their finger in the air in order to judge the direction of the wind.

Forecasters have it rough. But delivering robust, representative, and actionable strategic foresight projects requires grappling with a system whose nonlinearities and dynamism make weather systems look simple by comparison. The system is not in the external world, but in the *minds* of individuals engaged in the foresight process. Foresight strategists don't necessarily have to concern themselves with the probability of external events taking place; when asked, it's a professional convention to demure from assigning probabilities to scenario work. But they *ought* to concern themselves, for instance, with the *subjective weight* participants assign to the possibilities discussed in foresight activities, and specifically the automatic and unexamined psychological and cognitive factors that go into determining these weightings; when left unaccounted, these factors can easily "poison the well" for a foresight exercise, causing it to become an exercise in the amplification and confirmation of pre-conceived notions about the future.

Put simply: Futurists' greatest occupational hazard is the brute probabilistic nature of the world. For foresight strategists, it's the raft of *biases* and *heuristics* deeply embedded in human thinking and decision making, the effects of which can steer us into endorsing and acting on images of the futures that do not serve our goals, feed on our prejudices, play to our fears, and satisfy our egos rather than prepare us to clear-headedly face uncertainty.[3]

How big is the problem? The assumptions, generalizations, images, and biases about the future that Hines and Bishop refer to are deeply ingrained *indeed*. So deeply ingrained, in fact, that they are the way much of our mind naturally *works*. It is not only the received future within an organization that we are working against. If that were the case, the act of presenting alternative visions of the future would itself likely be sufficient to effect the change of mental models. The cause of the uncertain relationship between scenarios and mental models is the susceptibility of our minds to biases and heuristics that *distort* our thinking and lead us to misperceive the world in various ways. When not engaged in slow, deliberative thinking, our minds instead default to processes such as "the automatic firing of over-learned associations, behavioral regulation by

[3] Briefly, a heuristic is an information-processing rule that allows us to take a "mental shortcut" to reaching a conclusion. A bias, on the other hand, arises as a result of the misapplication of heuristic rules. In other words, heuristic reasoning doesn't necessarily lead us to biased decision making, but it opens the door to the effects of bias. As we will see, the use of heuristics when reasoning about the future appears to be particularly troublesome, leading to "severe and systemic errors" (Kahneman & Tversky 1973, 241).

emotions and processes of implicit learning", and a raft of other unconscious processes that impact our thinking without our knowing (Stanovich et al 2010, 16). It's my contention that the shortsightedness of organizational thinking is ultimately rooted in this tendency for individual minds to engage in biased future thinking. If we foresight practitioners concentrate our efforts on disrupting an organization's image of the future without at the same time furnishing them with the tools to mitigate the inevitable slide back into flawed future thinking, then we fall short of the transformative effects of a strategic foresight engagement.

But perhaps strategic foresight is not yet up to this more fundamental task. That the discipline of strategic foresight has historically lacked a serious engagement with the science of human foresight and decision making under uncertainty raises questions about how well it addresses the root of poor organizational future thinking. Considering strategic foresight in light of these research programs raises a host of fundamental questions about the received methodology: Does application of the Delphi method mitigate the individual biases of expert groups? Does the way scenarios are built adequately address the errors human beings naturally make when engaged in episodic future thinking? Do windtunnelling exercises encourage cognitive miserliness or focal biases instead of robust strategic thinking? Is foresight truly *anyone's* game, or does effective foresight depend as much on the participation of individuals with certain types of knowledge and predispositions as it does

on the participants with adequate clout and decision-making power? In order to start uncovering answers to some of these more tactical questions regarding the strategic foresight methodology, this paper will focus on two fundamental questions in this space:

What can we learn about the efficacy of strategic foresight by examining its methods in light of what scientists and researchers are learning about the function and limitations of human foresight?

How might we use what we learn from these disciplines to transform strategic foresight into an evidence-based practice?

For this paper, I used two research methods in order to explore these questions. I conducted a broad literature review surveying investigations into human foresight from several key scientific fields – including evolutionary anthropology, clinical psychology, and cognitive science – in order to discover how the strategic foresight methodology both succeeded and failed in accounting for the way that humans in fact think about the future. Second, I spoke with experts in each of these fields and in strategic foresight in order to get a better sense of how (or whether) these ideas were being synthesized and to test-drive the insights I'd uncovered in the course of my research. Conversation topics ranged from the state of the strategic foresight practice, to the natural limits of human foresight, to the systematic cognitive biases that impact foresight activities. I also spoke with two individuals representing different poles of a debate around how far we can and should attempt to extend human foresight to take account of larger periods of future time. The complete list of expert interviewees I spoke to in the course of this project is as follows:

- **Peter Bishop** (Assistant Professor, University of Houston, Future Studies)
- **Stuart Candy** (Assistant Professor, OCAD University, Strategic Foresight and Innovation)
- **Jim Dator** (Professor, University of Hawaii, Future Studies)
- **Mark P. Healey** (Lecturer, Manchester Business School, Strategic Management)
- **Alexander Rose** (Director, Long Now Foundation)
- **Douglas Rushkoff** (Media Theorist)
- **Karl Schroeder** (Senior Foresight Strategist, Idea Couture)
- **Keith Stanovich** (Emeritus Professor, University of Toronto, Applied Psychology)
- **Thomas Suddendorf** (Professor, University of Queensland, Psychology)
- **Maggie Toplak** (Associate Professor, York University, Psychology)

My claim is that at this point we *simply don't know* the answers to the types of questions posed above because we haven't yet done the work to connect the strategic foresight methodology to our growing scientific knowledge of the evolutionary history of human foresight and the research scientists have conducted around mental biases and use of heuristics in human future thinking. There's much to be gained from working toward this synthesis. I believe it is the next logical step in transforming strategic foresight (not to mention many other sensemaking and decision making methodologies) into a discipline that is consonant with what we're uncovering about the operations of the human mind.

In the rest of this paper, I hope to provide a preliminary framework for conducting that synthesis, and point to concrete examples where foresight methods subvert their own aims through insensitivity to the ways humans naturally think about the future.

2. THE 'TWO SYSTEMS' VIEW

If we are going to recast foresight as a methodology for rooting out and addressing the biases and heuristics that impair the way we collectively conceive of and respond to the future, we must first establish a theoretical basis for understanding the multifarious ways our mind constructs and interprets the future. In this work, I will adopt the perspective developed in seminal works of cognitive psychology – most notably, the work of theorists and experimentalists like Daniel Kahneman, Amos Tversky, and Keith Stanovich– which hypothesizes that the human mind is broadly divided into two separate but interacting systems. Daniel Kahneman provides a succinct overview of the characteristics of these two systems in the opening pages of his book *Thinking, Fast and Slow*:

System 1 (hereafter *S1*) operates automatically and quickly, with little or no effort and no sense of voluntary control. Processes of S1 are sub-personal in the sense that they do not depend on input from high-level control systems (Stanovich et al 2010, 16).

System 2 (hereafter *S2*) allocates attention to the effortful mental activities that demand it, including complex computations. The operations of S2 are often associated with the subjective experience of agency, choice and concentration (Kahneman 2011, 22).

S1 enables us, for instance, to automatically read the facial expressions of the people around us and parse their emotional states, orient ourselves to the source of a sudden sound, keep our car in the right lane while driving, compute simple arithmetic like 2 + 2, and step over uneven ground when hiking through the woods. Kahneman has described thoughts arising from S1 as *automatic,* and also as being *unauthored.*

S2, in contrast, is associated with the subjective experience of labored thinking most people experience when they try to book a flight and hotel, multiply a two-digit number in their head, or recall the details of what they did last Tuesday. S2 is coextensive with *mental work*: its deliberate, effortful, and orderly. Thoughts arising from S2 come with the subjective feeling of having been *authored* by the subject. Because it is hard to *author* two things at once, S2 operations also have the hallmark of interfering with one another when we try to attend to them simultaneously. Our conscious mental working space quickly becomes overloaded when, for instance, we try to overtake a transport trailer on a narrow highway while having a conversation with our passenger about German Neo-Kantianism. This is why most passengers in that situation will naturally allow for a break in conversation as the driver executes a passing maneuver. "They know that distracting the driver is not a good idea," writes Kahneman, "and they also suspect that he is temporarily deaf and will not hear what they say" (Kahneman 2011, 25).

There are a few important *provisos* to the "two systems" view of thinking. First, conceiving of the cognitive mind as comprising two distinct but interacting systems does not commit us to believing that they are separate physiological structures in the human mind. Not only is there is no evidence to support this claim, it would be trivial if true for a very important reason: with enough exposure and rehearsal, thinking that was once experienced as *authored* can become "the automatic firing of over-learned associations" (Stanovich et al 2010, 18). Broadly speaking, this is what we mean by expert competency or knowledge. For instance, a car park attendant who backs vehicles into narrow spaces all day long will, over time, no longer experience significant mental effort when performing the task.[4] The same can be said of individuals who have deep experience of the structure and operations of their organization and industry: they may no longer exert significant mental effort in either understanding or operating in the systems in which they're embedded.

Moreover, there is a strong tendency among individuals to defer to S1 whenever possible. As Keith Stanovich points out, this is simply a computational bias: S1 operations are cognitively *easier* to compute, so there is a strong prima facie case for defaulting to S1. This "cognitive miserliness" helps explain the tenacity of "my-side bias", which is the

[4] However, there are physiological hallmarks of purposeful, active thinking. Individuals engaged in mental effort will tense up. Their heart rates increase, and their pupils dilate. This effect is incredibly reliable. By observing a close-up of a subject's pupil, Kahneman and Tversky were able to accurately predict when they were was engaged in effortful mental tasks (Kahneman 2011, 34).

tendency to implement knowledge that props up existing mental models and disregards disconfirming evidence: quite simply, "the easiest models to represent are those closest to what a person already believes and has modeled previously" (Ibid, 19).

This brings us to an important, related note, which is that influence can, and often does, run in the opposite direction: thoughts that *feel* authored can actually be mere recapitulations of judgments rendered by S1. In this situation people are likely to confabulate purposeful reasons for their judgments; they hide the true origin of their judgments *without knowing* they're doing so. S2 – the *authoring* system that feels like *you* – may endorse a judgment rendered on the value of its cognitive affordability by S1 for an entirely *ad hoc* reason. In the domain of *futuring*, the pull of the default future might be as much about its *computational affordability* as it is about its assumed plausibility.

All is not lost, however, because the processes of S2 can also be martialed to override those of S1. The innocuous system level description of this ability disguises the enormity of its import; using S2 to reign in S1 is a drab operational description for no less than the ability for humans to exercise "rational self-determination" (Stanovich 2004, 275).

We have already seen that S1 can pass judgments or make decisions that are anathema to our higher-level desires. This commonly occurs in instances where we individuals are said to be *lacking willpower*. For instance, the smell of Belgium waffles might trigger a desire in my S1 to

gorge on them. All things being equal, unless S2 intervenes with some higher level desire – such as "I want to stick to my diet", or "I don't want to have a sugar crash later" or "I want to live to see my seventieth birthday" – and a plan of action for circumventing S1 – such as "I'll plug my nose till I'm around the block" or "I'll reward myself with a handful of almonds later" or "I'll just grit my teeth and keep walking" – then, all other things being equal, S1 is likely to get what it wants.

Note just how many forces S1 can martial against slow and deliberative S2 in this situation. S1, as will see, is evolutionarily ancient. It's rules for governing an organism's behavior have been shaped and cemented over hundreds of millions of years of evolutionary history for the sole purpose of increasing reproductive success; S1's that enacted underperforming or flimsy rules for guiding behavior in the evolutionary landscape simply didn't make it into the next generation. It's no wonder, then, that humans often find themselves acting without thinking, that strong emotional responses like fear and greed so strongly color their attempts to construct rational judgments, and that they have a hard time resisting impulses.

Now, S1 really, *really* wants us to eat those Belgium waffles. Why? In the evolutionary landscape, humans needed as much fat as they could get in order to survive because it was a relatively scarce commodity (Pinel et al. 2000, 1109). Fat was mostly bound up in cognitively sophisticated, maneuverable, and rightfully suspicious quadrupeds that were not eager

to give up their fat stores to our human ancestors. So when humans *did* hit the fat jackpot it made good sense from the perspective of survival and reproduction to consume as much of it as possible: more fat was directly correlated with greater longevity and therefore, more opportunities for genetic reproduction.

In addition, S1's case for eating waffles is also strengthened by the aforementioned bias toward computational conservatism; deciding to eat the waffles is far less cognitively expensive than overriding this desire and then coming up with both reasons and strategies for abstaining. It is, quite literally, *easier* to give in to the urges emanating from S1's desires.

Perversely, S1 also gets a boost to its influence when S2 is engaged is some other deliberative task. Studies have shown that individuals who are given a psychologically demanding task and then tempted with an array of snacks will tend to choose fat and sugar laden items like candy and chocolate over items like carrot sticks and nuts (Kahneman 2011, 43). S1 is opportunistic; it waits for moments when S2 is occupied, and drives behavior toward its preferences.

Lastly, the effects of words and messaging surrounding the Belgium waffles display can also bolster the influence of S1. If the word "Irresistible!" appears next to a photo of the waffles, this will subtly lessen the effectiveness of S2 in asserting its control over S1. This is why so much messaging and branding in advertising is aimed at S1: it's far less

discerning a consumer than a properly functioning S2 enacting a project of rational self-determination.

I use the example of waffles here because it is illustrative of how the cards are stacked against rational self-determination in a very familiar scenario. But this is not an isolated phenomenon; it's cognitively and behaviorally global. Once we start looking it's easy to find the influence of S1 and its attendant biases in all of our deliberations, from choosing a neighborhood to live in, to purchasing financial services, to building visions of the future and imagining how we might respond to them.

The struggle within strategic foresight activities between deep-seated biases and heuristics, on the one hand, and a program of rational self-determination, on the other, is therefore a very particular example of a general set of issues in human decision-making. That being said, I think there are two compelling reasons to re-examine foresight in light of the science of biases and heuristics. First, because foresight is a fairly well encapsulated methodology, we stand a reasonably good chance of building parameters, protocols, and checks into it that effectively reduce the irrational influence of S1 over the futures we construct, thereby improving the demonstrable utility of foresight engagements. Second, the stakes are high: powerful organizations whose actions have non-negligible impacts on global events are increasingly turning to foresight as a means to develop longer-term thinking and forward-facing strategies. There is a real danger for the both the viability of these organizations and the long-

term outlook for humans generally if strategic foresight serves to amplify the judgments of S1 and furnish it with effective strategies for achieving its goals. S1 goals are not co-extensive with human goals, though since they have common evolutionary origins there are of course many places they overlap.[5] If foresight is to be the preferred lens through which to view the future's possibilities, we should do our best to ensure that lens isn't seriously flawed.

It's doable, but it certainly won't be easy. The literature exploring biases and heuristics from a "two system" approach is vast. Even selecting a subset of biases that are explicitly *about* future thinking does little to narrow the field. Here is a small sampling of a few such biases and heuristics:

> **Availability/unavailability heuristic**: judging the probability of an event based on the ease or difficulty with which examples of it come to mind

[5] Some futurists are very worried about the possibility that future AGI (artificial general intelligence) will lack goals that are properly scaled to human goals. The potential design space for minds, they point out, is likely unimaginably vast. Human minds (and their attendant goals) occupy a vanishingly small point in that space. The likelihood that any recursively self-improving intelligent agent will develop goals that map to human goals is very poor, and we should therefore expect that an AGI will enact many of its goals at the *expense* of human flourishing, in much the same way that Joe's goal to get chips at the corner store is achieved at the expense of the flourishing of ant colonies living next to the sidewalk. I do not think their concerns are unfounded, but I would simply point out that we are already grappling with this problem, except that instead of contending with AGIs we are up against genetic and memetic interests nested in our own bodies and minds. See Stanovich (2004) for a thorough exposition of this position.

Hot hand fallacy: misperceiving stretches of uniformity consistent with randomness in a sequence as evidence for a "hot streak", also known as the *gambler's fallacy*

Impact bias: the tendency to overestimate the length or intensity of future feeling states

Planning fallacy: a tendency to underestimate the length of time required to complete a task

Sunk cost fallacy: justifying the assigning of additional resources to a project based primarily on the amount of resources that it has already been allocated

Status quo bias: adopting the current baseline as a reference point, and perceiving any change from that baseline in either direction as a loss

Overrepresentation of these biases and heuristics will cause us to commit errors in generating useful mental images of the future and planning around them accordingly.

We are all subject to these biases. Any, or all, of them can influence the formulation and use of foresight scenarios and strategies in ways that reduce their thoroughness, credibility, and usefulness, or worse, that lead to strategic decisions which are antithetical to the organization's goals, or even to *human* goals broadly conceived. In this sense, and in the context of foresight, *thinking clearly about the future and the way we ought to respond to it* is therefore synonymous with *developing clearly authored, reasonably de-*

biased visions of the future and formulating explicitly rational human-regarding strategies for flourishing within whatever future we happen to get.

This is, of course, an ideal state of affairs that will require considerable rejigging of the methodology and its guiding framework. In the interim, it is essential that we begin the conversations that will bridge this gap in understanding between what we're learning about our native foresight abilities and the foresight methodologies that purport to make us better at thinking about the future.

We have nowhere to go but up: as Ronald Bradfield points out in his review of cognitive barriers to scenario development, as it stands "[foresight] practitioners at large … appear generally unaware of, or at best only vaguely aware of, these phenomena and their ramifications on the scenario development process" (Bradfield 2008, 14), and the influential evolutionary psychologists and cognitive scientists I've approached for interviews in this project have likewise professed an ignorance of the strategic foresight methodology. It's time to change this state of affairs. We might begin with a survey of exactly what evolutionary psychologists are learning about the structure of human future thinking, which we'll turn to now.

3. A NATURAL HISTORY OF PROSPECTION

The cornerstone of strategic foresight is an exploration of the human ability to perform feats of future-oriented *mental time travel*. So far as we can tell, humans might well be the only creatures on earth that can "pre-experience the future by simulating it in our minds" (Gilbert 2007, p. 1352).[6] The practice also appears to be something we engage in more often than even we might realize. Humans obsessively think about the future; mental simulation of the future, or prospection, consumes "nearly a third of our spontaneous cognition at rest and supports a range of adaptive behaviors, from planning to problem solving" (Race et al 2013, 1), which is an important indicator of its centrality to our success as a species.

But even though we are overachieving future thinkers, we shouldn't lose sight of the fact that the human ability to imagine potential future scenarios that are decoupled from their representation of the present doesn't come from nowhere. Future-oriented mental travel is an ability constructed on the scaffolding of much older prospective abilities that originate deep in our evolutionary history and support various degrees of planning and problem solving capacities. It is therefore worthwhile, before discussing how the strategic foresight methodology

[6] There is a healthy debate about whether other animals have the ability for mental time travel. Scrub jays, for instance, exhibit behaviors that suggest they have some sense of themselves continuing into the future, and can act in ways that take into account their future (as opposed to current) drive states. For a survey of the evidence, see Suddendorf and Corballis (2007).

both supports and inhibits effective and de-biased mental time travel, to take a look into the evolutionary history of prospection. This will help us understand both its structure and historical function, and locate potential strengths and weakness in the scaffolding of which we should be mindful when considering how we might augment and extend the foresight methodology. Ultimately, our goal should be to shape foresight to drive the conversion of automatic, rash, and sub-personal System 1 processes into careful, deliberative, and self-aware System 2 processes.

First, we'll investigate some the of peculiarities of the machinery of prospection, and consider evidence for the theory that memory and prospection share much of the same machinery and so cannot easily be considered in abstraction from each other; in fact, there is good reason to postulate a mirrored prospection system for each of the different types of human memory. This is called the Janus Hypothesis, and it informs the structure given to human prospection in this chapter. As we'll see, absent serious injuries or diseases most people are able to deploy *procedural*, *semantic*, and *episodic* foresight when conceiving of and responding to the future. However, the picture becomes more complicated when we consider the shared machinery of memory and prospection. The ability to think about the future is not well compartmentalized from memory in the human mind; an important fact that suggests we look closely at the use and misuse of notions of time in foresight exercises.

After a discussion of the Janus Hypothesis, we'll survey the current

thinking about the three different broad foresight systems, corresponding to established and well-researched memory systems: *procedural* foresight, which encompasses the set of automatic and stereotyped responses to imminent future events; *semantic* foresight, which enables script-based, depersonalized investigation and planning for the future based on both acquired and implicit knowledge; and *episodic* foresight, which enables the development of ego-centered – or *autonoetic* – scenarios, effectively placing future scenarios on an individual's future timeline.

Finally, I will discuss how individual differences in temporal orientation is an added important wrinkle to the complexity of foresight. There are considerable differences to consider in the way individuals are oriented to time, which can impact their contributions to foresight activities and amplify (or, with careful planning, dampen) bias in group ideation and strategizing.

3.1 The Janus Hypothesis

As Thomas Suddendorf and Michael Corbalis point out in their seminal work on the evolution of foresight, while episodic memory – the reliving of past events – has been the topic of intense research efforts, "the mental construction of potential future episodes has only very recently begun to draw attention" (Suddendorf & Corbalis 2007, 299).

Granted, there is strong evidence to suggest that both abilities are enabled by many of the same cognitive resources. To *us*, time might seem as though it stretches out in two different directions, both into the past

and future. But it seems that our brain itself is less sensitive to the difference, since it mobilizes much of the same machinery when dealing with one or the other. Thomas Suddendorf, a psychologist who has drawn attention to this surprising feature of the brain, calls this 'The Janus Hypothesis' after the Roman god with two faces on opposite sides of his head, one staring into the past and the other into the future. The human brain, like Janus, seems to recycle at least some aspects of its temporal perspective on events and facts, whether what it is 'looking at' is 'behind' it or 'ahead' of it in time.

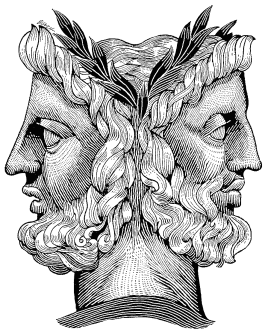

Figure 1. The Roman God Janus

Some compelling evidence for the hypothesis can be found by studying individuals who've had catastrophic injuries that have left with unable to form new memories. Molaison – or H.M., as he is known in the literature of psychology – is the most well known subject in the history of the study of memory. As a teenager, H.M. underwent brain surgery to treat his constant, debilitating epileptic seizures. The surgery was successful in treating his condition, but left him with serious memory deficits. Most conspicuously, he became virtually unable to form new

episodic memories: those memories of events in our lives that we draw on when, for instance, relating what we did on our vacation last week. H.M.'s episodic memory was so impaired that such a recollection was beyond him. By the early afternoon of any given day he was completely clueless as to what he may have done with his morning. New semantic memory – the explicit knowledge of new facts – also seemed to elude him; due to significant *retrograde* amnesia, his knowledge of world events dried up during the period several years *before* his surgery. (He could, however, acquire new procedural memories – the body memory that motor skills like piano playing and painting are built upon – although he'd have no recollection of how he might've learned them).

As for his view of the future: when asked what he thought he might do tomorrow, H.M. evaded the question with the odd nondescript reply, "Whatever's beneficial." In much the same way that H.M. was unable to conjure specific scenarios from his past, he could neither construct potential scenarios of his future, even those that were relatively immediate.

Edwin Tulving, an influential pioneer the study of memory in cognitive psychology, was the first researcher to notice the deep link between memory and foresight. His insight also came from his study of a brain-injured patient. N.N., as the man was known, still had memory for basic facts about the world (unlike H.M.), and even retained the ability to learn new skills. But N.N.'s memory of the episodes of his own life was

completely lost to him. Curiously, Tulving noticed that his ability to imagine potential future episodes in his life – even innocuous ones, like his plans for the following day – was similarly extinguished. Science writer Carl Zimmer reports on an exchange between N.N. and Tulving that demonstrates the extent of the former's impairment (Zimmer 2011).

Tulving: What will you be doing tomorrow?

N.N.: (long pause) I don't know.

Tulving: Do you remember the question?

N.N.: About what I'll be doing tomorrow?

Tulving: Yes. How would you describe your state of mind when you try to think about it?

N.N.: (long pause) Blank, I guess. It's like being in a room with nothing there and having a guy tell you to go find a chair.

It seems that, as the memory of the past is lost, so too goes anticipation for the future. Both H.M. and N.N. were stranded in the present, with no stories to tell or plans for the future.

At first glance, the structural features hinted at in the Janus Hypothesis seem like a poor way for our brains to deal with two very different domains of inquiry. The past is the domain of certainties, and the future that of uncertainties. Why, then, do our brains seem to mix them together?

The messy way we conceive of the past and future only begins to make sense when we understand that our brains weren't built to grasp at

objective reality, but rather to get us through the next day and, more importantly, get our genetic material into the next generation. Objectivity is a scientific concept, and scientists didn't build our brains. The shared machinery of past and future in our minds is instead the result of a persistent *tinkerer*: biological evolution. As psychologist Gary Marcus might say, our natural foresight ability is a hacked-together evolutionary *kluge*: a quick (at least on evolutionary time scales) and cheap solution to the pressing problem of navigating a complex social and physical environment, and a clumsy and inelegant solution to the problem of anticipating the future (Marcus 2008).

The evolutionary understanding of future thinking also suggests another remarkable hypothesis: when you consider the boost to evolutionary fitness – an animal's likelihood of survival and reproduction – foresight confers, it seems more likely that the ability to construct memories of the past exists primarily to serve as *input* into the foresight process, a perspective Suddendorf himself echoed in conversation: "The past is totally irrelevant unless it impacts survival and fitness," he said, "and from a fitness perspective what matters is foresight (Suddendorf 2014). In other words, we might only have memory of the *past* for the sake of helping us better understand and plan for the *future*: a stunning inversion, considering the amount of psychological research that has been conducted on human memory, and the relative scarcity of studies on human foresight.

There is some intriguing evidence for this hypothesis. As vision scientist Stanley Klein and his co-researches suggest, if the primary function of memory is to assist in planning for an uncertain but potentially controllable the future, than we should expect it to be more efficient when its employed for this purpose relative to other tasks (Klein et al 2010, 14). In order to test this hypothesis, they randomized subjects to four different encoding conditions – past, atemporal, planning, and surviving - which were variations on a prompt of how one would go about deciding which items from a prepopulated list to bring on a camping trip: *past encoding* prompted participants to remember a specific time in their past that they had camped, *atemporal encoding* prompted them to use their imagination to form a picture of a campsite in a forest, *planning encoding* prompted them to imagine that they were making plans to take a camping trip, and *survival encoding* prompted them to imagine that they are stranded in a forest without any basic survival material. After completing the portion and a short distraction exercise involving completing anagrams, subjects were prompted to recall the words they had been asked to prioritize at the beginning of the experiment. A comparison of the mean recall achieved on each encoding task found that subjects performed best when primed with planning encoding, followed by survival encoding, atemporal encoding, and lastly past encoding. That the biggest recall gap uncovered in the experiment was between past encoding and future encoding bafflingly suggests that we are *least* good at using memory when thinking about *the*

past itself (Ibid, 17).[7] This is the opposite result we should expect if memory's function was to record the past in a purely reproductive manner, like a video camera. But if we understand memory is primarily designed for future planning rather than reminiscence, than its distortions and shortcomings make sense: "natural selection can only work on what memory can offer for present and future fitness rather than on the accuracy of past record *per se"* (Suddendorf et al. 2009, 1317).

So while memory and prospection are deeply physiological linked, there are also interesting asymmetries in their functioning. It seems reasonable to assume that the relevant pressure that gave rise to this asymmetry is the directionality of time perception itself; our ability to process and plan around temporality was no doubt shaped by relationship with time, in which, as philosopher Hans Reichenbach observed: "There are no past impossibilities and there are no future facts" (Reichenbach 1951, 241).

Simply put, the brain treats past and present similarly because it's convenient for it to do so, and the mistakes that accrue as a result of the overlap weren't enough to get us killed in the ancestral world before we could produce offspring. It seems clear, however, that our myopic

[7] The atemporal encoding used in this experiment closely resembles the "scripting" style of planning that relies primarily on semantic memory. I suggest later in this paper that foresight based on scripting is less memorable and more prone to a specific raft of biases that leads us to *underimagine* the future. But it is intriguing to note here that scripting makes even memories seem less *memorable,* which is further evidence for the Janus Hypothesis.

foresight abilities can get us into quite a lot of trouble in the modern world, as we'll see in the final chapter.

The current dominant paradigm of memory studies postulates three separate, interacting memory systems in human beings: procedural memory, semantic memory, and episodic memory. Proceeding on the assumption that the Janus Hypothesis is true, I will explore some of what we know about the *prospective* counterparts of each of these memory systems, in order to give us a better sense of the origins of the many biases we are subject to when thinking about the future.

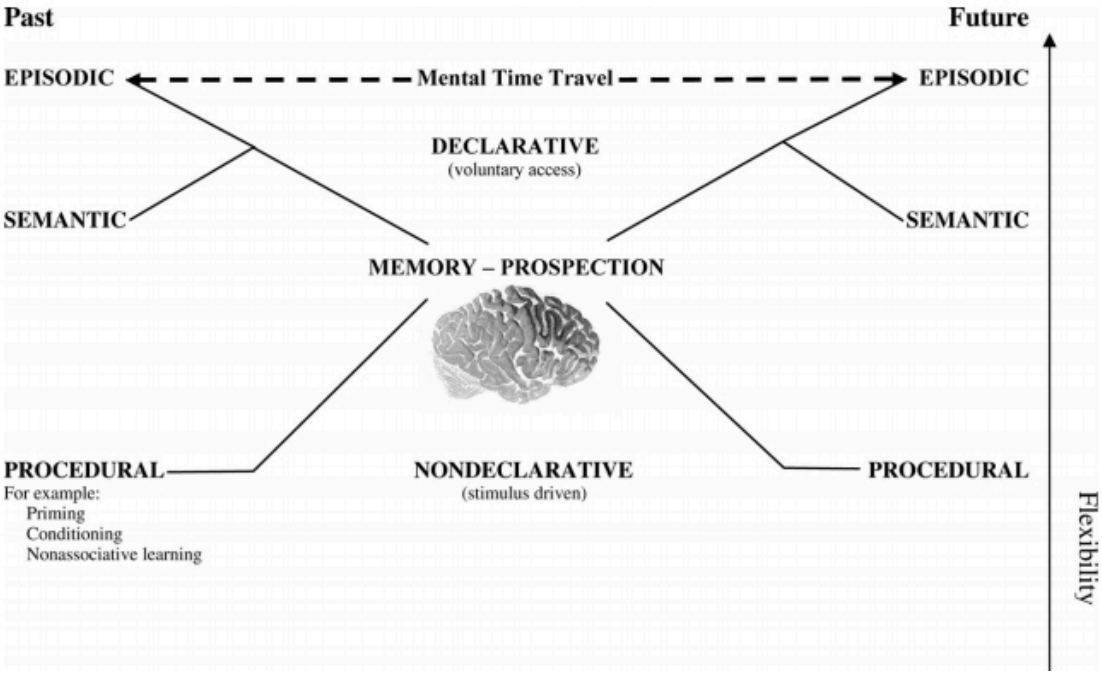

Figure 2. The memory and prospection systems. The common taxonomy of memory systems (left), after Squire (1992), and its proposed prospective counterpart (right). The figure illustrates the Janus hypothesis, which is that procedural memory, semantic memory, and episodic memory each has mirrored counterparts in our foresight abilities. (Suddendorf & Corballis 2007, 301)

3.2. Procedural Foresight

The evolutionary development of prospection is a long and complex tale with one simple driving force: selection pressure. We begin the story (which we can only retell in broad overview here) with an account of procedural memory.

Certain types of futures – namely, those that were imminent and portended highly regular or very probable consequences – were so influential on the evolutionary path taken by our distant ancestors that we have inherited their clever strategies for coping with those futures in the form of *hard-wired behaviors*.

When we think of the structures that enable the guidance of behavior in animals, we naturally think of brains. Brains, as philosopher Daniel Dennett succinctly put it, are *anticipation* machines: their primary purpose is to "produce future." (Dennett 1996, p. 177) In the evolution of life on planet Earth, the development of brains is a particularly elegant solution to the problem of navigating and surviving a complex environment. But brains are far from the only solution, and perhaps aren't even the most common one deployed in the whole of evolutionary history. Some organisms, like coconut trees and clams, employ an "armor and wait" survival strategy. Others – cnidarians like coral and jellyfish – evolved diffuse, highly interconnected neural nets, which register and respond to threats and opportunities through no other channel but physical touch (Ginsburg 2008, p. 223). A jellyfish won't withdraw its

tentacles in response to a *potential* impact, only an *actual* one. These are not future-facing flora and fauna. Plants, protists, fungi, eubacteria, and archaebacteria: the brainless kingdoms in the domain of eukaryotic life are completely incurious about the future. They live instead in an immediate, unextended now.

Even the seemingly complex behavior of brained and mildly future-regarding animals often disguises a series of inflexible behavioral triggers that respond to immediate, temporally unextended stimuli. Ticks, for example, will find a high perch and remain there indefinitely, or until the smell of butyric acid, secreted by the skin of all mammals, triggers them to loosen their grip and drop down from their perch. Landing on some solid surfaces extinguishes the first behavior and serves as a cue to initiate a search for heat. When the tick locates a warm spot, this serves as a third environmental trigger, and the tick begins burrowing down toward the source of the heat. Each response is law-governed and doesn't allow for cognitive flexibility (Carruthers 2000, 124). Jacob von Uexkull, who wrote imaginatively about the perceptual worlds (*Umwelt*) of many creatures – evocatively describes the tick's world: "The whole rich world around the tick shrinks and changes into a scanty framework consisting, in essence, of three receptor cues and three effector cues – her Umwelt. But the very poverty of this world guarantees the unfailing certainty of her actions, and security is more important than wealth" (Von Uexkull 1934, 12).

Even the most rudimentary forms of procedural foresight inch the life of the mind forward in time; whereas animals like jellyfish and coral need to wait for *physical contact* in order to initiate behaviors and build sensitivities and habituations, animals that anticipate imminent future states behave with respect to *possibilities*. The space opened up between *now* and the *imminent* future is where planning – or perhaps more accurately, protoplanning – begins. Protoplanning is interesting from the perspective of foresight because it marks the moment that *images of the future*, the stated quarry of strategic foresight, enter onto the evolutionary scene.

It is hard to understate how seemingly miraculous is the decoupling of behavior from immediacy. "When a mouse hides before a cat enters the room," writes psychologist Dan Gilbert, "it is responding to an event that has not happened yet, and its ability to do so is one of evolution's most remarkable achievements." (Gilbert & Wilson 2007, p. 1351) The mouse can't of course, *see* the future in which it is eaten by the cat. However, it can take action *now* in such a way as to reduce the possibility of being caught flat-footed in such a future.

But we shouldn't confuse all that appears to be procedural foresight with *planning*. All that is necessary for procedural foresight is that an organism's behavior be "modulated by experience such that the organism gains a future advantage". Procedural foresight therefore only enables learning with reference to the "current indicators of upcoming events"; it

is "stimulus-bound, or better, bound to the perceptual tracking of stimuli" (Suddendorf et al 2007, 300). Stimulus-bound responses are often referred to as *instincts* or fixed action patterns. We should not lose sight of the fact that future advantage can be – and often is – achieved through mindless behaviour.

Take the complex choreography of the female digger wasp (also known as *Sphex ichneumoneus*) as she prepares to lay and hatch her eggs. Keith Stanovich relates the complex sequence of actions:

> First she digs a burrow. Then she flies off looking for a cricket. When she finds a suitable one she stings it in a way that paralyzes it but does not kill it. She brings the cricket back to the burrow and sets it just outside at the threshold of the burrow. Then she goes inside to make sure things are safe and in proper order inside the burrow. If they are, she then goes back outside and drags in the paralyzed cricket. She then lays her eggs inside the burrow, seals it up, and flies away. When the eggs hatch, the wasp grubs feed off the paralyzed cricket which has not decayed because it was paralyzed rather than killed (Stanovich 2004, 74-75)

The complexity of the digger wasp's behaviour is no doubt impressive. But something strange happens when experimenters add barriers or interruptions to the wasp's dance. For instance, if when the wasp descends into her burrow to make sure it is safe an experimenter moves the paralyzed grasshopper an inch or two away from the burrow's edge, the wasp will come out of her burrow, drag the grasshopper again to the edge, and then go once more into the hole to check its safety. If the grasshopper is moved again, the wasp will again repeat the same behaviour. In one experiment, "the wasp checked the burrow forty times and would still not drag the cricket straight in" (Ibid, 75). Although it may

seem to us as though the wasp is engaged in complex planning to secure the future viability of its offspring, in reality it is going through a set of actions "choreographed by rigid and inflexible preprogrammed responses to specific stimuli in the Sphex environment" (Ibid, 75).

Reflecting on the complex, yet mechanized, behaviour of digger wasps, philosopher Daniel Dennett refers to "that spooky sense one often gets when observing or learning about insects and other lower animals: all that bustling activity but there's nobody home!" (Dennett 1984, 13). Dennett also suggests that the proto-planning behaviour of the wasp also holds up a mirror to our own behaviour: "What makes you sure you're not sphexish – at least a little bit" (Ibid, 11)?

Dennett's is right to ask; we are sphexish, and often more than just a little bit. These types of sphexish procedural foresight strategies, of course, persist in humans, and are still highly persuasive in the situational domains in which they were shaped. The ducking reflex is one such strategy. Ducking, as it turns out, is a particularly robust strategic response to a possible future where one is struck in the head by, say, a hurled rock. Those individuals who failed to execute ducking strategies in a timely matter – and there must have been many of them – left fewer offspring with poor ducking skills.

Another inherited strategy based on prospecting imminent threats is the phenomenon of hypnic jerk. Many people, when on the verge of sleep, experience a sudden falling sensation that triggers an involuntary

startle reflex. One hypothesis for the tight causal coupling of the falling sensation with a spastic flailing of the limbs is that those of our tree dwelling ancestors who quickly roused themselves from sleep and flailed their limbs were more likely to catch a branch and avoid being maimed or killed by impact with the ground. As a consequence of our ancestors' slowly acquired ability for effective prospection and anticipation, we moderns are naturally good at not being hit in the head by foul balls or falling out of our chairs when we doze off.

But even these deeply rooted, imminent future-oriented strategies can be misapplied. Note that while the ducking reflex is most often *good futuring*, the hypnic jerk is most often an example of *misfuturing*; generally, when we reflexively duck, it's because we were really in some danger of being knocked on the head (or at least a near collision). But while relaxing in bed, we likely aren't in any real danger of a fatal ground impact.[8]

Again, the examples of ducking and hypnic jerk here are meant to be evocative illustrations of a general principle: procedural foresight strategies can cause us to underrepresent the future and blind us to possibilities. An example that more directly impacts the practice of foresight is the implicit set of procedural strategies encoded in emotional content.

[8] It's easy to imagine scenarios where the reverse is true. Moviegoers who duck when three-dimensional objects rush at them are misfuturing. And if you experience hypnic jerk because you and your bed have just sailed out of a 10th story window, well, that's *good* futuring.

Emotions are not action agnostic; they carry specific hardwired, evolutionary adapted "action tendencies" or implicit goals, the purpose of which is to "save cognitive processing by triggering time-tested responses to universal experiences (such as loss, injustice, and threat)" (Lowenstein and Lerner 2003, 628). The limit case of extreme emotional agitation or stress is instructive. As emotions intensify, they exert an ever-stronger tendency to over-ride rational deliberation and consideration. Reflecting on the influence of intense emotions on their actions, people often report that they have felt "out of control", or even as if they were "acting against their own self-interest" (Baumeister et al 1994, 183).[9]

In a striking experiment, psychologist Dan Ariely asked heterosexual undergraduate males a series of questions about whether they'd engage in immoral and risky behaviors for the purposes of increasing their chances of having a sexual encounter. The men completed the survey twice, once in a state of self-reported sexual arousal, and again in an un-aroused state. The experiment found that, while in a state of arousal, the men surveyed were more likely to report a willingness to lie to dates, to encourage them to consume more alcohol, or even to drug them if it would increase their chances of having sex. Incredibly, the men surveyed consistently *under-predicted* the effect that arousal would have on their judgments. "Across the board," Ariel writes, "they revealed in their

[9] Opponents of two systems theories of cognition would do well to ask themselves in light of this example: are these individuals simply mistaken in their belief that they weren't acting in their own best interest in a heated moment? If not, then in the interest of *who* or *what* were they acting?

unaroused state that they themselves *did not know what they were like* once aroused" (Ariely 2008, 97).

In more moderate levels of intensity, emotions and stress appear to play an *advisory*, rather than a *dictatorial*, role. In these cases, emotions act as information input into decision-making process. In many cases, the advisory role of emotion is legitimate. For example, assessing how you feel now is a relevant input into your decision regarding whether to see a movie this evening. Still, if the contribution of an emotion or stress state to judgment making is not explicitly recognized, its influence can become tacitly endorsed in an individual's supposedly authored judgments without their knowing.

A study of experienced parole judges in Israel found that the rate at which they granted parole to prisoners slowly declined from a rate of 65% to almost zero throughout morning sessions. After a food break, the rate at which they granted paroles again rose to around 65% (Danziger et al. 2011, 6890). The authors conclude that "making repeated rulings can increase the likelihood of judges to simplify their decisions", and that "they will be more likely to accept the default, status quo outcome: deny a prisoner's request" (Ibid, 6889). While they might not be dictatorial, unacknowledged mildly and moderately intense emotional states, it seems, can still exert a powerful *Machiavellian* influence on our decision-making.

The act of building multiple representations of the future, as we have seen, is computationally expensive, and it is often a bad strategy in time intensive situations. Stereotyped, low cost, low risk strategies based on shallow, automatic analysis of events were often the better strategic option in these types of situations. There is good reason to believe therefore, that decision-making under time-sensitive, stressful, or emotional conditions will create biases toward drawing on unauthored procedural foresight strategies. As Daniel Gilbert writes: "People use their immediate hedonic reactions to simulations as predictors of the hedonic reactions they are likely to have when the events they are simulating actually come about" (Gilbert et al 2007, 1352). And even if individuals subsequently adjust their outlook to take account of the time that will elapse between now and the event being simulated, they are still apt to use their immediate hedonic reaction as an anchor for judgment. As a result, "when we attempt to predict our future feelings, we expect our future to feel a bit more like our present than it actually will. (Gilbert 2006, 151)"

The tradeoff between the security and reliability of automatic decision-making and the wealth of possible representations of future states is a recurring theme in evolution, and indeed in the minds of humans. These emotions and reactions are common enough in the process of group decision making that it is worth investigating the extent of their impact and developing processes to mitigate their effect.

3.3 Semantic Foresight

The roots of procedural foresight, as we've discussed, are evolutionarily ancient, and so we naturally share many of its mechanisms with distant cousins like tree shrews and toucans that arguably lack the flexible communication skills and robust self-awareness of humans. We have seen that it *is* possible to get by without even procedural foresight in situations where organisms only need to employ very regular and stereotyped responses to a static environment. However, procedural foresight becomes indispensible in environments that impose large penalties – such as starvation, dismemberment, and death – when organisms misread the imminent future.

Procedural foresight encompasses all those aspects of our future-oriented cognition that can't be explicitly verbalized in language; in the parlance of cognitive science, it is a non-declarative set of skills. Declarative foresight – a category which covers both *semantic* foresight and *episodic* foresight – is more flexible than procedural foresight, because it can "be triggered top-down from the frontal lobes, rather than bottom up through perception" (Miyashita 2004, 435). In other words, whereas procedural foresight is always attached to the current context, organisms with semantic and episodic foresight capacities can employ them more or less *at will* (hence the "declarative" moniker) in order to simulate the future. Next, we will consider both of these declarative foresight abilities in turn.

Semantic foresight is the ability to use our knowledge of the world in order to construct depersonalized scripts of how future events might unfold; if (as we'll see) episodic foresight is "projecting into" the future, semantic foresight is "knowing about" the future. In Klein's experiment described above, for instance, the participants' abstract knowledge about events that are likely to happen during a camping trip, and things it'd be handy to have in those situations, are both examples of semantic knowledge. With respect to the future, it is possible that our conception of how it unfolds "is predominately derived from what could be termed "script-based" knowledge – our knowledge of familiar, routine events" (Atance et al. 2005, 127). This may go a long way to explaining the seeming ubiquity of the planning fallacy: because the future is only rendered in script, it's easy to underestimate the time and resources necessary to complete tasks in a predetermined schedule and to imagine the task taking place "in a vacuum" without unexpected setbacks.

Semantic foresight is much more rare than procedural foresight, in part because it appears to depend on the presence of brain structures that are underdeveloped or absent in much of the animal kingdom, especially the frontal and pre-frontal cortex. There are, of course, no straight lines in nature. Deciding whether some instance of foresight in animals is merely procedural or amounts to semantic or episodic foresight is difficult to do.

The study of prospection in rats provides a good example. Rats move around a maze or meadow they encode a map of the territory in

their hippocampus. Neurons there become active as the rats pass through touch points in their journey, and the same neurons will fire again in the same sequence if they pass through the same territory again. Remarkably, when the rats stop for rest, readings from their hippocampus suggest that they are both rapidly replaying their route and imagining themselves running in different directions, "projecting themselves into different futures to help them decide where to go next" (Zimmer 2011; Redish et al. 2007).

But what kind of representation of the future are these rats using? It's hard to say. While there is no evidence to suggest that rats possess autonoetic consciousness – a sense of *themselves* as situated in a personal timeline stretching from past to future – there are experimental results suggesting rats may have the ability to reflect on their own *semantic* knowledge about the world when making decisions (Foote 2007, 551; Buckner et al 2006, 54-55).

The opacity of mental life is a key difficulty to assessing the results of both laboratory behaviour experiments and strategic foresight activities, among other things.

If there is a general rule of thumb regarding semantic foresight, it's this: *if* a problem can't be solved through procedural foresight, and *if* an individual can develop a course of action without having to imagine their personal future in detail, then they will do so. This is in keeping with the principle of *cognitive miserliness* identified by Stanovich et al.: "We often

make the easiest (incorrect) inference from the information given and do not proceed with the more difficult (but correct) inference that would follow from considering all of the possibilities" (Stanovich 2010, 18). When it comes to foresight, concerns of computational conservatism will often trump those of thoroughness.

There is a cluster of heuristics surrounding semantic foresight that will tend to make people more reliant on script-based futures. First, scripts are usually constructed in domains with which we have a great deal of familiarity. Because things that are familiar to individuals naturally come to mind more easily than things with which they aren't familiar, we should expect them to employ an availability heuristic or 'my-side bias' when judging the course of the future. Simply put, futures that include a place for the knowledge that they've scripted through experience and exposure will come to mind more easily than those that do not rely on this knowledge, and we know from experimental findings that individuals mistakenly conflate *ease of recall* with *greater plausibility* (Ibid, 19).

The affect heuristic is also at work in semantic foresight. Optimistic scenarios will be disproportionately based on the results of semantic foresight because a more detailed analysis in most cases would render any given scenario less optimistic, since the increased contextual information might introduce blockers. Since thinking about optimistic scenarios *feels good*, we should also expect a strong emotional bias toward endorsing the results of semantic foresight (Mesoudi 2007).

Understanding the creation and use of script-based, semantic foresight, and the biases and heuristics attached to this type of prospection, can help us to formulate the ontology of default futures in organizations and learn how to better combat them.

3.4 Episodic Foresight

Investigating the evolution of episodic foresight – or episodic future thinking – has become a cottage industry since the publication of Suddendorf and Corballis' 2007 landmark call to arms for research in the area (Suddendorf & Corballis 2007). As the authors point out in that survey, while episodic memory has been the focus of intense research interest since the pioneering work of Edward Tulving, "mental construction of potential future episodes has only very recently begun to draw attention. (Ibid, 299)" Due to the relative youth of the field, much of what it has uncovered has yet to make it into the practice of strategic foresight despite the potential it has to transform the practice.

As we might expect, in light of the Janus Hypothesis our abilities for episodic foresight bears commonalities with our abilities for episodic memory. As in the case of episodic memory, the ability for episodic foresight appears to be contingent on possessing what Tulving calls *autonoetic consciousness*, which is "the kind of consciousness that mediates an individual's awareness of his or her existence and identity in subjective time extending from the personal past through the present to the personal

future", creating the "characteristic phenomenal flavor of the experience" (Tulving 1985, 1). But what does autonoetic consciousness substantively add to our ability to cope with the future? In order to make its contribution clear, let's quickly recapitulate the contribution of each type of foresight to our ability to grasp the future.

Procedural foresight is stimulus bound, and so is largely inflexible to contextual information; place an audio device with continuous playback of the sound of running water in a beaver's territory, and the animal will cover it with sticks and mud in an attempt to dam it up (Richard 1983, 107). If the beaver could articulate semantic knowledge about the world – for instance, about what running water *looks* like apart from what it *sounds* like – it may be able to render a judgment that building a dam in this situation is inappropriate and so not engage in dam building behaviour.

Semantic foresight allows for more flexibility in planning, but due to the low resolution afforded by mere knowledge without experiential context, it can also fall short in certain situations. The Bischoff-Köhler hypothesis states that animals other than humans are unable to decouple their representations of the future from their current goal states, such as securing food or sex. The inability to imagine themselves in goal states other than the one they are pursuing currently – in other words, to engage in episodic future thinking – means that their plans can only support their current activities (Raby & Clayton 2008, 318). We have already seen a

similar planning deficit in humans with brain injuries who have lost the ability for episodic foresight but not semantic foresight. Those individuals might be able to provide a script for activities like going to a restaurant, but the task of marrying that script to a personal plan for the future seems insurmountable to them.

Episodic foresight adds another level of resolution to planning. With it, we can simulate ourselves in the future, in different *contexts*, facing different *challenges*, and pursuing different *goal states*. This affords the ability to finely tune our behaviour in anticipation of responding to multiple possible futures in which we may find ourselves; with this capability, we can "make specific plans, and compare different scenarios" in order to develop contingencies and rehearse responses (Suddendorf et al 2009, 1321). Similar to how episodic memory allows us to re-experience our past, "episodic future thinking allows us to pre-experience our future" (Atance & O'Neill 2001). The best way to understand episodic foresight is as a further improvement on both procedural and semantic foresight that allows for the further fine-tuning of behaviour, a third "dial" that allows us to modulate our response to the future by shifting the temporal and situational context of our actions. Unlike other animals, we can decouple both our knowledge of the worlds and *ourselves* from the present moment.

Episodic future thinking is an incredibly powerful tool. So powerful that as "perhaps the only species with such foresight, humans alone may be driven to consciously guide the planet into the future and

thus be burdened with the responsibility of getting it right" (Suddendorf & Corballis 2007, 312; Dawkins 2000). At the same time, it is also an incredibly risky tool to employ. Consider, as Keith Stanovich related to me, that hypothetical thinking demands that we "decouple from the primary perception of the world and run a so-called mental simulation of an alternative world" (Stanovich 2014). This action leaves us open to incursion from the un-simulated world. Simply put: "Animals that took attention away from primary perception tended to get eaten" (Ibid 2014). This, Stanovich suggests, is why it is so difficult to exercise foresight, and why, for instance, building and exploring mental simulations causes us to close our eyes, or look up at the ceiling in order to ease the effort (Ibid 2014). Our minds recognize that foresight is a risky and costly operation, so they do their best to create a sense of aversion to engaging in protracted and deep speculation.[10]

Thomas Suddendorf also points out that our minds have a series of mechanisms for getting us out of the mode of mental simulation as quickly as possible. The sufficing criterion is one such mechanism. "People tend to go with the first solution they hit upon that suffices for dealing with their problem, rather than continue on a more exhaustive

[10] There is a strange push and pull in our ability to think about the future. On the one hand, our minds *want* us to engage in future thinking, and so we sometimes find it easy and pleasurable. As I mentioned earlier, we tend to engage in foresight whenever we have free thinking cycles; we call this daydreaming. And people also find that planning for the future greases the wheels of other mental faculties, such as our memory. That being said, our minds clearly do not want us to decouple so completely from the environment that we forget it entirely for lengths of time, lest we be gobbled up.

search to find an even better solution," says Suddendorf. That we would naturally put a stop to mental simulation based on the sufficing criteria stops us from endlessly entertaining possible scenarios. "When you can conceive of various scenarios," he says, "you have to have a way to stop yourself from doing that. As soon as something is sufficient, it makes sense to pull the plug and think about the next problem" (Suddendorf 2014).

The sheer riskiness of engaging in foresight should give us pause; given the immediate danger to which it exposes us, foresight must have been an incredibly important feature of the human mind in the evolutionary environment in order to persist in the human mind. How do humans alone, as it appears, possess this astonishing skill?[11]

In order to answer this question, we have to set aside our assumptions about human exceptionalism. Humans, like all other creatures, have an evolutionary story, and so the answer must be sought in the prosaic and incremental story of our descent. Mathias Osvath and Peter Gädenfors suggest "the cultural niche that was created by the use of Oldowan tools, including the transport of tools and carcasses, lead to a selection for anticipatory cognition, and in particular anticipatory

[11] Suddendorf and Corballis raise the chilling and very plausible scenario in which our ancestors, after establishing a level of dominance over the environment, began to compete primarily with one another in inter-band arms races. The principally hostile forces in nature our ancestors faced in this retelling of our cognitive evolution were *other humans*. This might help explain the human tendency to violently clash with and displace other humans. In short, the characteristic most central to our human identity – the ability to inhabit the future – may have been won at the cost of perpetual inter-species warfare.

planning" (Osvath and Gädenfors 2005, 4; Plummer 2004, 135).

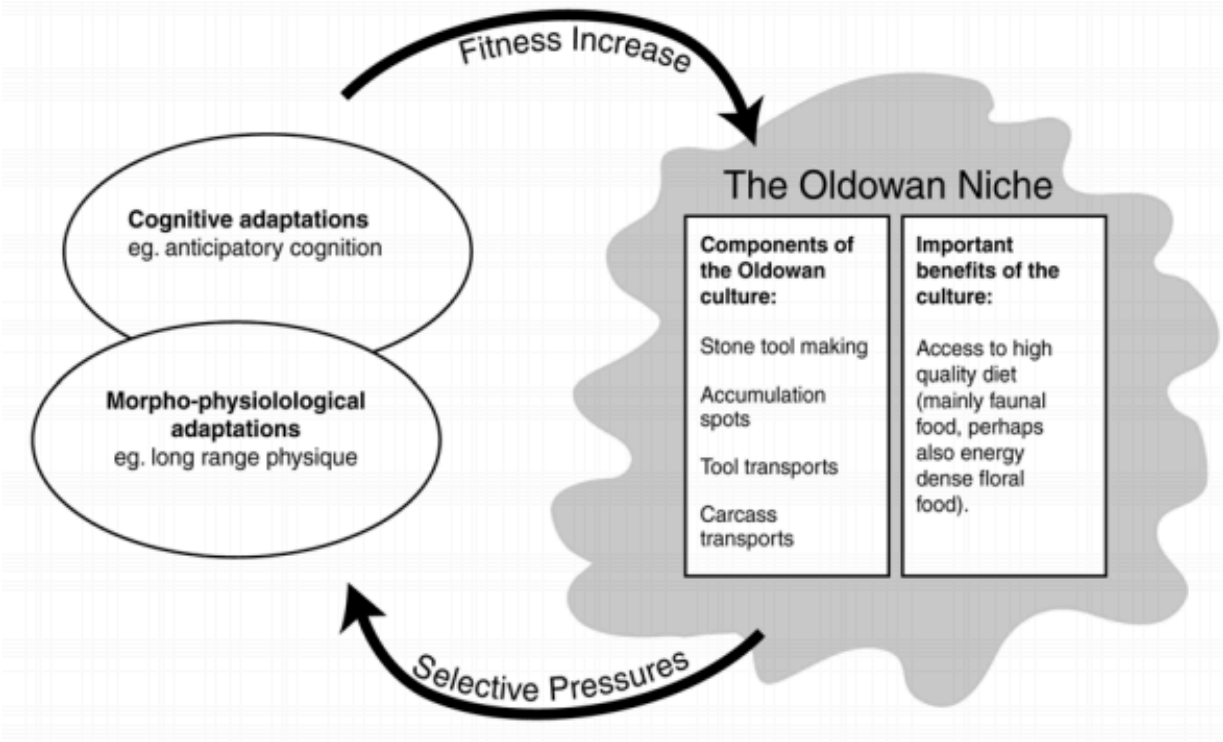

Figure 3 The Oldowan niche had its own selective pressures. Adaptations to these pressures increased the fitness within the niche, and the environment in the niche changed as a result of the adaptations. This diagram shows the closed causal loop between the selective pressures of the Oldowan niche that drove our evolution, and ensuing elaboration of that very niche with our newly equipped bodies and brains. (Osvath and Gädenfors 2005, 7).

The Oldowan culture represents a watershed moment in the history of evolution. The niche first appeared as a result of several key changes in the environment of early humans: deforestation and expanded savannahs in Africa drastically reduced access to floral food while increasing access to large grazing herbivores. These environmental changes "resulted in selective pressures on the hominids that lead them to change their diet from predominately vegetarian to more protein and fat based" (Osvath and Gädenfors 2005, 6). Adapting to such a radically different niche must've required a genetic advantage conferring mental flexibility, even in

this very early stage of the evolution of anticipatory cognition. As philosopher Robert Arp points out, a reliance on mental associations or trial and error wouldn't have sufficed, since "the environments in which these hominins found themselves were *wholly new*, and there would have been no precedent by or through which one could form mental associations utilizing past information" (Arp 2008, 146). It should come as no surprise that the fossil record dating to this period suggests that thousands of mammalian species went extinct, having failed to make the transition from an arboreal to a savannah-based lifestyle (Novacek 2002). Tracking, killing, and transporting animals requires a vastly different skill set than foraging for edible plants in dense forests. It is difficult to imagine how any creature – using only those physical skills acquired in forest life – could make the jump over to savannah life without finding a way to creatively modify their existing capacities in order to survive. Arp suggests the ability that saw our ancestors through this tumultuous period was a capacity for non-routine, creative problem solving, driven by the ability to visualize scenarios through "intermixing of visual information from mental modules" and "the active selection and integration of that information" for solving problems (Arp 2008, 146). It's no surprise then that it's in the Oldowan niche evidence of the manufacture and use of stone tools, the transportation of artefacts and stone tools, the transport of carcasses, and the use of accumulation spots (i.e. tool caches) appears in the pre-historical record for the first time. The abilities to combine visual

data from the past and present and imagine future outcomes, even if only in relation to very brief windows of time, must have been considerable advantages in adapting to a new niche.

The new niche changed our ancestors in fundamental ways. As Osvath and Gädenfors suggest, the effects of the Oldowan niche on human morphology was likely considerable. In an era when transport and mobility bestowed a considerable edge in survivability and reproduction, we became taller, and shorter and broader through the middle; adaptations suggestive of both long range travel and load bearing. The increase in brain volume and shrinking of our lower jaw suggest a high quality, easily digestible diet comprised of animal meat and foraged plant food (Ibid 8).

More importantly for our story, it also had a profound effect on our cognitive toolkit. The Oldowan culture was marked by an extension of the considerations of both time and space in relation to human behaviour. Evolutionary pressures toward developing expanded views of both time and space might've included "long delays between the acquisition and the use of the tool, as well as considerable geographical distances between the sources of tool raw material and killing sites" (Ibid 8). Because early humans could only carry so much, these realities made it necessary to economize and curate both tools on-hand and caches of tools in strategic (that is, mentally anticipated) locations (Plummer 2004, 133). Optimizing strategic carrying and caching meant keeping track mentally of the

resource available in a given accumulation spot so that hominids weren't either trying to access depleted caches or restock caches that were already stocked (Osvath and Gädenfors 2005, 9). It also meant suppressing current goal states, such as indulging in lone calorie rich meals, in favor of projected goal states, such as bringing a kill home to one's family – an action that has clear ramifications for reproductive success (Ibid, 10).

Osvath and Gädenfors also speculate that the need to collaborate on distal goal states – both temporally and spatially speaking – may have driven the development of symbolic communication. Language, as they point out, is based on representations as stand-ins for entities, real or imagined, that can override the need to indicate cues in the immediate environment for communication purposes (Ibid, 10). Animals that aren't able to represent detached goal states can of course collaborate, as prides of lions do when strategically cutting off the escape of prey animals. But if the goal is remote in time and space "then a *common* representation of it must be produced before co-operative action can be taken" (Ibid, 10). As far we can tell, only humans can achieve this feat.

Like all behaviors that have a considerable net positive impact on reproductive success, conversing about the future as a group can *feel good*.[12] In conversation, science fiction author and foresight practitioner Madeline Ashby has pointed out individuals in a foresight exercise "come

[12] Of course, there are also potential pathologies of future thinking. Some theorists suggest that obsessive-compulsive disorder and chronic stress are positively correlated to excessive future thinking (Zimbardo 1999, 1285)

alive" when prompted to think strategically about the future (personal communication, October 16). This natural exuberance for group future thinking might be as evolved a response as the tendency for modern humans working with stone samples – for example geologists – to unconsciously pick up and heft stone samples of optimal throwing size (Ibid, 5; Bingham et al, 507). But not everything that *feels good* is, of course, good for you. Just as the optimal function of our evolved feeding mechanisms have been disrupted by the abundance and ease of access of calories in our modern environments, so too can our evolved capacity for visualizing future scenarios fall short when faced with modern uncertainties.

On an individual level, our procedural foresight might be overactive. We may rely too much on semantic foresight, leading us to envision *thin futures*. And even though it is the most sophisticated of our evolved foresight capacities, the machinations of episodic foresight can distort our view of the future, leading us to make poor predictions and poor decisions.

Take the relatively simple example of predicting hedonic experiences. Episodic foresight allows us to form an emotional impression of how we might feel in some future scenario. For instance, if I am trying to decided whether to go to Disneyland tomorrow, I can imagine myself being there and "consult" the ensuing feeling: am I having a good time in this future scenario, or not? Of course, the "prefeeling" I am experiencing

in this visualization is not the *actual* feeling I will have when I am at Disneyland. It might be similar enough in order to help guide my decision making, or it might be somewhat, or even wildly, inaccurate. And as it turns out, there are many ways in which this strategy of *prefeeling* can lead us astray. As psychologists Dan Gilbert points out, the hedonic feeling I'm experiencing in the scenario (H_1) will only reflect my actual feelings (H_2) in the future if two conditions are met: first, that the contextual factors surrounding the time of scenario visualization ($\bar{e}_1$) – such as the weather, my level of stress, my mood, etc. – match with the contextual factors of the actual event ($\bar{e}_2$); second, that my *simulation* of the event (e_1) matches with my *perception* of the event itself (e_2).

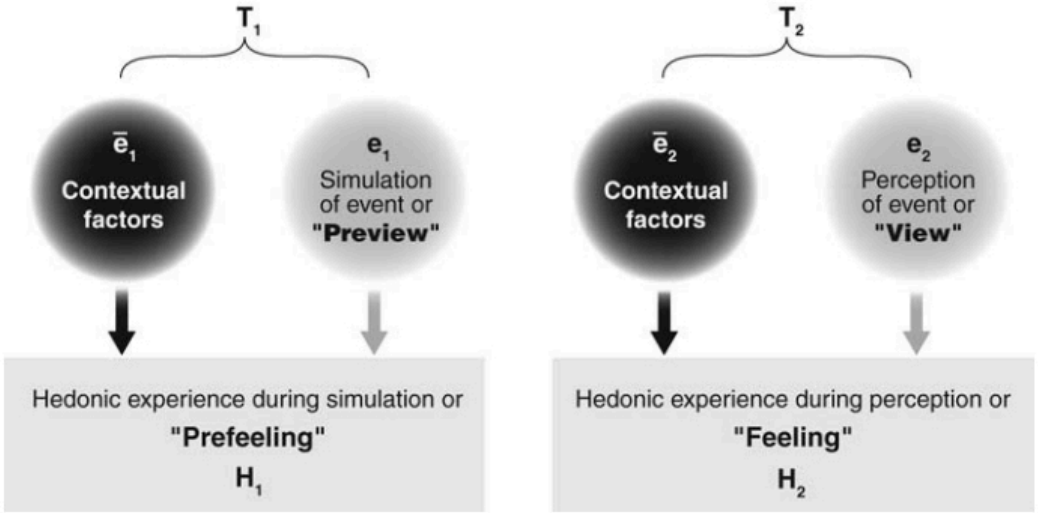

Figure 4 Gilbert's diagram showing how hedonic experience is influenced by both mental representations (simulations and perceptions) and contextual factors. The elements that go into our prefeeling of an event – namely the current contextual factors of our imagining and our simulated preview of that event – will only match the way we *actually* feel once the event arrives if both the contextual factors and our perception of the event are reasonably similar to the conditions in which we imagined it. Otherwise, our prefeeling of the event will not be a good analog for how we actually feel – a state of affairs that will adhere more often than not. (Gilbert et al. 2007, 1352).

Errors in prospection, says Gilbert, "arise from the fact that people use their prefeelings to make hedonic predictions even when one or both of these conditions is not met" (Ibid, 1352). Gilbert isolates four different ways that we fail to meet these conditions.

First, we tend to use *unrepresentative* material from our past experiences when building scenarios of the future. The past material that is most "ready-to-hand" for building scenarios of the future are generally those memories that have some kind of salience; they stand out from the crowd of memories, and when we select them we are submitting to the availability heuristic ("This memory is the first that came to mind, so it's the one I'll use."). As Gilbert observes: "It seems that everyone remembers their best day, their worst day, and their yesterday. Because unusual events and recent events are so memorable, people tend to use them when construction simulations of future events" (Ibid, 1353).

Second, since it takes a great amount of cognitive effort to imagine scenarios in detail, we tend instead to create *essentialized* scenarios based on what we believe to be it's key features. Omitting all these details – which can influence hedonic experience as much as essential features – causes us to "predict that good events will be better and bad events will be worse than they actually turn out to be" (Gilbert et al. 2007, 1353). The effect of essentializing might also be amplified by temporal delays due to our tendency to think of far futures more abstractly than near futures: assessing the hedonic experience of taking a vacation next week might

include the prefeeling of airport security, while imagining a vacation five years from now might not (Liberman et al. 2002, 532). Hyper-essentialism regarding distant future scenarios might explain why people tend to orient to either highly optimistic or highly pessimistic beliefs about the far future. More prosaically, it also neatly explains why "people so often make future commitments that they regret when the time to fulfill them arrives" (Gilbert 2007, 1353).

Third, our simulation of future events is naturally abbreviated; if they weren't, it would take as long to *simulate* the event as it does to actually *experience* it. In addition, the moments that people select when simulating the future tend to be crowded toward the beginning of scenarios. For instance, when asked to imagine losing mobility, able-bodied people predict that they will be much more unhappy than individuals who've lost mobility actually report feeling. This seems to be because the able-bodied individuals considering scenarios of mobility loss construct their hedonic response based on "the initial – and typically the worst – moments of these events" (Ibid. 2007, 1353). This causes us to construct very pessimistic outlooks for these scenarios. As psychologist Maggie Toplak points out: "People just can't imagine that if they were met with some tragic and unhappy event in the future that they would ever be able to cope" (personal communication, March 6).

The inability to predict adaptive responses to situations is among the most commonly observed errors in research on hedonic predictions. It

seems our brains are not good at writing endings, but they love a good opening. In other words, when considering the affective components of a future scenario, we are much more likely to focus on events that come *early* in the narrative, and far less likely to envision the mid-point or ending. We easily imagine ourselves, full of energy, bounding through the gates of Disneyland, but are not as likely to imaginatively place ourselves in second hour of our wait in line for Splash Mountain.

Finally, our scenarios of the future tend to be decontextualized, floating free of the broader factors that color our perception of present events. Because scenarios of the future are unrepresentative, essentialized, and abbreviated, they also tend to lack the stable "background" against which we can assess our hedonic response. This leads people to overweigh current contextual factors when considering hedonic responses to the future. For instance, low blood sugar in the present might cause us to feel blasé about the prospects of going to Disneyland tomorrow, after we will have eaten a hearty lunch. Alternately, it can also cause us to underweigh the impact of contextual factors that will surround future events. People overestimate how unhappy they will be if their favorite team loses a football game, because they do not consider that their hedonic experience at the time will be influenced by more than what is simply on the scoreboard (Wilson et al. 2000, 825).

Our poor showing when it comes to accurate affective forecasting, as Maggie Toplak points out, is a particular manifestation of a much

broader set of issues in human cognition: thinking is hard and unpleasant work, so humans use any excuse to behave like "cognitive misers" in order to get out of it. Instead, they'll engage their availability bias, and take the high fluency they experience in coming up with answers as an indication of their veracity (personal communication, March 6). A further complication is that indulging our heuristic thinking often *feels good*: cognitive scientist Valerie Thompson calls this affect "the feeling of being right", which is strongest when we quickly and fluently read problems and generate responses (Thompson et al 2012, 237).

Serious, deliberative thinking does not have these virtues. The problem is compounded when we are being asked in addition to think about unpleasant events; in this case, we are apt to experience negative affect issuing both from the process of thinking and the object of thought itself. "Thinking about happy things that are going to happen in your future is affectively pleasant and may not be very taxing on you," says Toplak, "and that's a different activity than the hypothetical, cognitive decoupling activities that are actually very hard and create negative affect" (personal communication, March 6).

These tendencies to misfuture might have minimal impacts in cultural and technological environments where not much changes over an individual's lifetime. This appears to have been the case during the period in which the cognitive ability for episodic future thinking primarily developed. In the everyday Paleolithic, today's hunt probably looked

quite a lot like the hunt of yesterday, and highly salient events – such as failed raids and near tramplings – were instructive raw materials for planning and constructing future scenarios. But in cultures where the pace of change in a single individual's lifetime completely remodels the technologies they can access and the social milieu in which they find themselves every few years (or even more often), the capacity for episodic future thinking we inherited from our Paleolithic ancestors will be far less reliable. The native futuring abilities of our minds were not "designed" to account for the rate of change and sudden discontinuities to which we are exposed in the modern world. To illustrate, consider that it took over 1 million years for hominids to move from blunt, clumsy Oldowan technologies to slightly more refined Acheulean stone tools. Over this same time period, the distances observed between the location of raw materials for stone tools and their geological sources increased from 10 km to approximately 20 km, doubling the range of typical hominid bands. In contrast, a "mere 12,000 years separate the first bow and arrow from the International Space Station" (Ambrose 2001, 1752), and the ability for a single human to range over the whole planet, and beyond.

3.5 Time Perspectives

The common shortcomings and dysfunctions of prospection, however, are only part of the story. It seems there are also crucial

individual differences in the way each of us perceives time that can lead us to exhibit individual quirks in the way we think – or fail to think – about the future.

As it turns out, life staging turns out to be an important factor in the way we think about both the past and the future. Psychologist Donna Addis and her co-researchers tested college age students (average age: 25) and older adults (average age: 72) on their ability to generate both semantic and episodic memories of the past and prospection for the future. When cued with a word and asked to use it both to remember the past and project into the future, they found that young adults were more likely to employ episodic memory and episodic future thinking in their responses. Older adults, on the other hand, showed a reduction in the episodic specificity of past events that extended to the way they imagined future events. They were more likely to both recall the past and imagine the future with heavier reliance on their semantic knowledge (Addis et al. 2007, 40).

Age group	Event type	Cue word	Event description	
			Non-episodic information	Episodic information
Young	Past	Tree	…because I love cheese."	"I went hiking in Muir Woods in California … with my boyfriend then and his room-mates … we went through all these different ecosystems … and you would see different kinds of plants, so like we would see orchids … and I said, "Wow it's so beautiful and it's like wild" … that part was like a jungle: wet, very lush and green. At the end of that trail, was like the beach … so we had bread and cheese and it was very fun and good…
Young	Future	Oven	None	"I'm going to bake my first loaf of bread. It's going to be probably Friday afternoon before Yom Kippur … I'll get the recipe from an old cookbook… The room's going to be hot even though we have the windows open, because we're going to turn on the oven in the middle of the summer. The light, bright kitchen light will be off, and instead, we're going to light a candle. BBC will be on in the living room, it'll be kind of static…"
Older	Past	Toy	" This reminds me of those toys that our grandchildren have. I think they're spoiled… Our son, the teacher, doesn't have a lot of money, and I think his wife is just can't say no to the kids. Every time she goes out, there's a toy coming home… I generally will give my son money for specific things…	…Like he had a problem with his knee and I, so, to help him with his doctor's bill, I gave him some money, and on the check I wrote, 'Don't spend on toys with this check'. "
Older	Future	Engine	" In the next few years I hope we have an engine that doesn't have to use gas to run. I hope we come up with an alternate source of energy to run vehicles. Because they're a polluter, and its getting to be very expensive to drive, and there's a lot of driver irritability over stop and go driving, having the carbon poisoning happening…	…The scene is I'm just driving along, in the Saab, and… not worrying about high energy costs…"

Table 1. A comparison of semantic and episodic recall and prospection in both college age and older adults. In these examples, individuals are prompted with a cue word and a request to describe some element of either their future or past relating to that word. From these samples, we can see that young participants generated substantial episodic information in both past and future events, while older participants tended to provide non-episodic, semantic descriptions related to cuing words. From Addis et al. 2007, reported in Schacter et al. 2007.

If our life stage affects our use of either semantic or episodic elements in constructing future scenarios, it also impacts the way in which we evaluate how we might act in those futures.

Psychologist Laura Carstensen's research further suggests that individuals are influenced in their decision making by their perception of time as either open ended and expansive, or closed and limited. Carstensen's socioemotional selectivity theory predicts "perception of gains and losses is influenced importantly by individuals' temporal frameworks" (Carstensen 1999, 177). The perception of time as either expansive or limited is, of course, tightly correlated with life staging: young people are more likely to perceive time as only abstractly bounded,

while older people are more likely to grasp the finitude of life. As we might expect, this leads to differences in decision-making. Individuals who perceive the future as open ended and expansive are more likely to pursue avenues that maximize their future possibilities, regardless of whether they have an impact on present circumstances. Those who perceive time to be limited will tend to more heavily weigh present gains over future possibilities, and therefore may choose, for instance, to invest more time and energy in cultivating existing relationships rather than seeking out new social connections. We shouldn't be fatalistic about these frames, but changing them does require willful thought: "age related patterns do emerge, but even these age patterns can be altered when individuals adopt a time perspective different from what is predicted by their place in the life cycle" (Carstensen 1999, 166).

Finally, there are also individual temperamental dimensions to time perception, apart from life staging, that bear consideration. Psychologist Philip Zimbardo – most famous for his controversial Stanford Prison Experiment – developed an inventory for assessing individual differences in the perception of time. The aim of the Zimbardo Time Perception Inventory (ZPTI) is to measure "the often nonconscious process whereby continual flows of personal and social experiences are assigned to temporal categories, or time frames, that help to give order, coherence, and meaning to those events" (Zimbardo & Boyd 1999, 1271). These frames, in turn, are used "in encoding storing, and recalling experienced

events, as well as in forming expectations, goals, contingencies, and imaginative scenarios" (Ibid, 1271-1272). The Inventory isolates three dominant "frames" through which individuals tend to view experiences and construct responses to them. Individuals employing a "past frame" will tend to recall analogous prior situations when forming judgments, paying special attention to the costs and benefits connected with decisions made at the time. Those with a "future frame" focus instead on "anticipations and expectations constructed to embody an extension of the present into the future when the calculated costs of this current action will be paid or reward will be reaped" (Ibid, 1272). Individuals employing future frames are more likely to draw on considerations of alternative goal-states, means-ends relationships, and probabilistic assessments – in other words, to consider possible and probable future scenarios. Finally, individuals relying on a "present frame" will tend to focus instead on their current mindset and affects, de-emphasizing the influence of both past and prospective events. Where past and future oriented individuals will be adept at delaying gratification and disconnecting themselves from current goal states, present-focused individuals may be more highly influenced by "sensory, biological, and social qualities associated with the salient elements of the present environment" (Ibid, 1272).

People, of course, can flexibly deploy different temporal frames for different situations. This can be a prudent strategy: you might gleefully adapt a present frame that drives you to indulge when confronted with a

Klondike Bar while on family vacation, but employ a past frame – hearkening back to that time you ate a Klondike Bar and laid prone on the couch groaning for an hour afterward – or a future frame – thinking about your goal to run a full marathon in three month's time – when in the same situation at the corner store down the street from your house. However, in certain cases, the use of specific frames can become chronic, leading to "a dispositional style, or individual-difference variable, that is characteristic and predictive of how an individual will respond across a host of daily life choices" (Ibid, 1272).

Zimbardo's inventory also suggests several different valences to time perception: past-negative frames are associated with depression, anxiety, unhappiness, and low self esteem, whereas past-positive frames are characterized by "glowing, nostalgic, positive construction of the past" but tended to not be positive correlated with future present frame traits like novelty seeking, sensation seeking, or preference for consistency" (Ibid, 1278); present-hedonistic frames orient us toward enjoyment, pleasure, and excitement, with a lower emphasis on future considerations and a low preference for consistency, while present-fatalistic frames cause us to be both muted in our future outlook and passive in the present, subject to the whims of a perceived predetermination.

In this chapter, we've reviewed many of the common features of future thinking shared by most people with typical cognitive function. The Janus Hypothesis suggests that our ability for prospection has deep

functional and physiological ties with our memory; the better we understand this connection, the more surprising commonalities we will no doubt discover between past and future oriented cognition. Using the Janus Hypothesis as a basis for understanding, we unpacked the theoretical layers of future cognition, examining procedural, semantic, and episodic prospection in turn and suggesting ways in which each of these abilities are subject to biases and heuristics which can cause us to misfuture in various ways. We then examined some of the factors underlying individual differences in time perception, and goal and motivation structures related to temporality.

Our takeaways should be that treating "future thinking" as a single, monolithic skill is far too simplistic, and assuming that humans *qua* humans share a generic conception of the future ignores important distinctions that may impact the way strategic foresight is conducted. As it stands, strategic foresight is a blunt tool for addressing these intricacies; in important ways, the methodology reflects pre-critical assumptions about future thinking which might make it ineffective in addressing certain biases of future think, or worse, exacerbating the very biases it purports to address. In the next chapter, we'll examine some of these shortcoming in the methodology, and explore how developing targeted methods and protocols in order to mitigate the effect of systemic biases in future thinking is essential in order for strategic foresight to have a reasonable claim to helping us "create and sustain a variety of high-quality forward

views and to apply the emerging insights in useful ways" (Hines and Bishop 2006, v)

4. THE VIRTUES AND ERRORS OF STRATEGIC FORESIGHT

Many of the methods that make up strategic foresight to this day predate widespread dissemination and understanding of the key findings of the research we surveyed in the last chapter. While Herman Kahn was pioneering the institutional use of scenario planning at the RAND Corporation, Amos Tversky and Daniel Kahneman were still more than a decade and a half away from embarking on their era-making studies of human decision making (Kahneman 2011, 8; Millet 2009, 62). The University of Toronto conference at which Endel Tulving would first propose semantic memory as a distinct system from both acquired skill and autobiographical memory, setting in motion research into the tripartite division of human (and animal) memory and foresight, was likewise as far in the future. Philip Zimbardo's research on time perception was even more remote; in the early seventies, he was immersed in research on the psychology of prison life that would lead to the infamous Stanford Prison Experiment, for which he is still best known. These research projects, in conjunction with countless other new theoretical models and pieces of experimental evidence, contributed to the growing "cognitive revolution" in psychology that nowadays shapes

everything from the study of group thinking and radicalization to the creation of public policy and store layouts.[13]

Since the cognitive revolution was, at its core, a counter-revolution, we should quickly take stock of the paradigm against which it was a reaction. The intellectual climate in which scenario planning – which remains at the core of strategic foresight – was developed was an era of scientific management and behaviour-based approaches to the study of human psychology. With the benefit of hindsight, we can now see how the technological and scientific developments made in the early decades of the 20th century formed a powerful juncture in which humans could attempt the scientific management of human affairs. Strategy and planning efforts of the early 20th century were influenced by an intellectual milieu that included huge breakthroughs in physics and engineering sciences as well as the rise of reductionist psychological research programs such as psychoanalysis and behaviorism. Driven by the success of their brethren in the physical sciences, psychologists especially were seeking a new paradigm that could transform their discipline into a bona fide, quantitative science. In his personal reflections on the behavioral revolution, psychologist George A. Miller (an early entrant in the cognitive revolution) writes that a group of experimental psychologists, "influenced by Pavlov and other physiologists, proposed to redefine

[13] I include as an extension of the "cognitive revolution" the erstwhile "affective revolution" which is returning the study of emotion and arousal to psychology in an experimentally rigorous way. See, for instance, the work of Dan Ariely referenced in Chapter 3.

psychology as the science of behaviour", arguing that "mental events are not publically observable" and so "the only objective evidence available is, and must be, behavioral" (Miller 2003, 141).[14]

As antecedents, these developments made a strong prima facie case that the application of the right formal methodology could render the world knowable - or, in the case of foresight, could make *known* those things that were *unknown* - and that human behaviour and decision-making could be effectively organized and operationalized through these methodologies analogously to the way in which the flow of electrons was managed in a circuit. The core foresight methodology, with scenario planning at its heart, remains a product of the intellectual era in which scientific management reigned; for reasons that are hard to understand – perhaps due to its relative simplicity when stacked against other tools for organizational change, perhaps because of the convenient time lag between foresight exercises and the time horizons they investigate - it has been surprisingly resistant to change.

It's also important to understand that the aims of scenario planning in its initial inception were *not* to root out and dampen the pernicious effects of biases and heuristics in the way we think about the future, but rather to *think the unthinkable*, to paraphrase the title of Kahn's most well

[14] The fact that Gaston Berger was concurrently developing a similar set of tools in France he called "La Prospective" is a wrinkle in this story, since by all accounts European experimental psychologists were not held in sway by behaviorism – but like the Americans, they were neither privy to any clear, experimental findings on the pervasive and persistent nature of cognitive biases.

known book. The aim of thinking the unthinkable is still very much in line conceptually with the most commonly and plainly stated objective of strategic foresight: to *reduce all the uncertainty that you can, but not more than that*.[15] Intuitively, it's easy to make the case that imaginatively exploring the most critical uncertainties in a space in order to form images of the possible can help us to manage uncertainty. But without proper checks and balances in place against the excesses and shortcomings of human foresight in the strategic methodology, there's really no telling in the end if we've indeed managed it, or made it worse. Human minds, it turns out, are more unruly things than electrons.

No doubt, it's a testament to Kahn's genius that many facets of his methodology for building mental models of the future seemed to naturally address the excesses of both overly rigid methodologies and the fixation-prone human mind. That Kahn incorporated elements of Hollywood screenwriting into his exploration and strategic planning of the consequences of nuclear war – a topic many people thought warranted only sober-minded discussion – demonstrates his impulse to break from the protocols of the scientific management paradigm. At the same time, Kahn railed against the sentiment that the prospect of nuclear war was too horrible to be discussed using normal, neutral, professional, everyday language. Kahn sought a solution somewhere between the poles of creativity and analysis: "Awe is fine for those who come to worship or

[15] Thanks to Peter Bishop for suggesting this useful formulation as an improvement over the simple and misleading injunction "to reduce uncertainty".

admire," he wrote, "but for those who come to analyze, to tamper, to change, to criticize, a factual and dispassionate, and sometimes even colorful, approach is to be preferred" (Kahn 1962).

At the same time, Kahn was also sensitive to the ways in which over-reliance on available information or the rigid use of thinking tools could create barriers for thinking deliberately and rationally about a subject. For instance, Kahn's insistence that probabilities not be attached to scenarios can be read as an attempt to stop participants from becoming overly mentally invested in certain visions of the future (Kahn 1962, 150). This safeguard that suggests an intuitive understanding of the tendency for people to become *anchored* in even meaningless numbers and use them as an uncritical benchmark with which to measure difference and deviation, a phenomenon that Kahneman and Tversky would be go on to thoroughly experimentally explore (Kahneman 2011, 119).[16] "If there's one message we can take from [Kahn's book On Thermonuclear War]," strategic management researcher Mark P. Healey told me, "it's that people are sensitive to *possibilities*, not *probabilities*" (personal communication, November 20 2013).

Kahn also recognized the danger of becoming fixated on a subset of environmental information to the point of missing larger, looming threats. In her book *The Worlds of Herman Kahn: The Intuitive Science of Thermonuclear War*, Cold War scholar Sharon Ghamari-Tabrizi relates an

[16] However, as we'll see, avoiding discussion and use of probabilities entirely can lead to its own brand of myopia.

anecdote of Kahn in his natural environment, delivering a slideshow lecture on the pitfalls of uncertainty and unknowability to a group of air force officials sometime in the 1950s:

> The next cartoon shows a man steering a roadster off a crook in a mountain road, distracted by a buxom woman gazing at the view. [Kahn] jabs a finger at the drawing and tries to suppress a giggle. "Another mistake which is very very important is over-concentration. This is the kind of thing that, for example, you see: he's just concentrating not on the wrong thing – it's worth looking at, but not exclusively. We don't object to you looking at the blonde. We'd look at her ourselves [but] you should look at something else. There's a cliff over here. And the point is look around, look for loopholes, see what's happening. (Ghamari-Tabrizi 2005, 13)

Kahn's use of language in this example is perhaps telling, since he references the behavior of *looking* when discussing what is really myopia caused by excessive mental attention and fixation on certain information. Despite Kahn's ranging intellect and willingness to incorporate diverse bodies of knowledge, the foresight methods we've inherited from his work, and the work of others at the time attempting to develop intuitive ways to reduce uncertainty, remain situated in a place in intellectual history marked by the tail end of scientific management paradigms and the behavioral revolution; the most we can say is that some of the methods are suggestive of the ur-science of human judgment and decision-making that would only come to fruition years (and in some cases, decades) later.

Historically, these shortcomings have largely escaped the notice of commentators and practitioners of strategic foresight. In his book on organizational excellence, Competitive Advantage, Michael Porter calls strategic foresight "a powerful device for taking account of uncertainty in

making strategic choices" (Porter 1985, 447). Kees van der Heijden has stressed the value of scenario planning as a means of creating cultures of learning and building organizational foresight practices, even praising the intuitively *causal* elements of scenarios for these purposes (van der Heijden 2005, 139). Though van der Heijden himself refers to scenarios as a "cognitive device", a critical engagement with the psychological literature on human biases and heuristics is almost entirely missing from the foresight oeuvre. The overall impression of strategic foresight as a "panacea for strategic decision making under uncertainty" ignores that "using scenarios to inform organizational decisions is a complex matter and can yield mixed psychological effects, some of which might actually impair judgment and decision making" (Healy and Hodgkinson 2007, 556). As we will see, it's not only scenarios that bear more scrutiny; every step of the strategic foresight methodology is rife with opportunities for cementing foresight biases. What follows is a preliminary examination of some of the ways that strategic foresight can run afoul of the way humans spontaneously and naturally think about the future.

4.1 Delphi Polling, Desirability Bias, and Group Polarization

Originally developed by the RAND Corporation in the 1950's to help forecast the effects of technology on the future of warfare, The Delphi Method of expert polling has ever since been used as a means of building

consensus views about the future among experts on everything from foreign policy to the future of food. The method's rationale is simple: anonymous polling and reflection avoids the influence exerted by biasing factors in open discussions, such as deference to authority or dominant personalities, on the one hand, and the tendency for groups to get mired in their differences and fail to reach consensus on shared views, on the other. If we can discover those assertions about the future shared among a broad range of experts, so the thinking goes, then we will have articulated a "core vision" of highly plausible expert judgments that we can then use to inform the rest of our explorations: "… through the exchange of expert knowledge, iteration in the survey process, provision of controlled feedback, and convergence of probability assessments, the adverse effects of cognitive limitations on probability assessments such as overconfidence, can be reduced" (Ecken et al 2010, 1654).

However, this simple picture of building a plausible picture of the future based on the consensus of experts overlooks several important confounding factors. First, what assurance do we have that the consensus view is actually based on anything like the most *plausible* of the aggregated expert opinions? The assumption seems to be that opinions shared among multiple experts must be in some way uncontroversial, but a recent study of Delphi raises serious questions about whether this is the case.

Researcher Phillip Ecken and his colleagues examined the data from six Delphi exercises in order to determine the effects of desirability bias on expert judgment. Desirability bias, a simple and pervasive quirk of human reasoning, is the tendency to overestimate the probability of favorable outcomes and underestimate the probability of unfavorable ones.[17] In the six Delphi exercises, the researchers asked "a total of 200 qualified experts from business, academia, and government or public authorities" to assess the probability of a given development on a percentage scale from 0% to 100%, and also rate its desirability on a five-point scale. In total, the researchers gathered data on 8300 paired desirability-probability assessments throughout the testing.

What they found was a considerable positive correlation between how favorable a development was and the perceived probability of its taking place some time in the future. What this means is that, absent the right controls, we should expect that the consensus view of experts in a Delphi exercise should be highly influenced by the *desirability* of the agreed-upon predictions, and this effect will be amplified in groups that are more homogenous (i.e. groups in which all the experts represent

[17] A quick and easy way to see desirability bias in effect is to ask a friend the probability of their sports team making it to playoffs, and then contrast this judgment with the judgment of professional odds-makers; chances are very good that your friend's probability assessment will be significantly higher than that of individuals who stand to make money on the accuracy of their assessment. One easy way to make money is to ask your friends to lay down bets based on their inflated assessments. This is not, however, a good way to make friends.

similar viewpoints in the same industry). Ecken's group provides an illustrative example:

> Consider the following Delphi projection: "In 2030, Chinese car manufactures dominate the automobile industry". Experts X, Y, and Z all estimate the probability of this projection. Expert X is personally involved in the Chinese car industry and is likely to find this projection desirable and thus assesses the probability higher than expert Y, whose desirability about this projection is neutral. Vice versa, expert Z who is personally involved in the US car industry might find this projection undesirable and gives a lower probability of occurrence to that event than Y. The point is that even though X and Z are led by desirability bias, Delphi reports the average of X, Y, and Z and thus the subjective bias of desirability bias that influenced X and Z could be offset. Thus, in this case Delphi's averaging characteristic eliminates, or at least reduces, the desirability bias. However, if we change our example slightly, the following occurs: X, Y and Z are all personally involved in the Chinese car industry. They all share a high desirability for the projection and would all overestimate the probability. In this case, Delphi's averaging would not reduce or eliminate the bias and desirability bias becomes a "dysfunctional shared representation" (Ibid, 1664).

The researchers also point out that, interestingly, the effects of desirability bias in Delphi polling seem to be stronger the longer the timeframe is under consideration; it seems counterintuitive, but in the face of increased uncertainty and lacking strong evidence to justify their predictions, polled experts will rely *more* heavily on their desirability bias, leading to ever more heavily biased projections and probability assessments based on affective sorting.

For these reason, we should not assume that use of the Delphi Method always produces consensus around relatively uncontroversial predictions. Especially in homogenous Delphi groups (i.e. in which all experts are working in or around a given industry), the exercise may well

produce results that reflect the tendency for homogenous groups to build consensus based on increasingly polarized views. In his book Going to Extremes: How Like Minds Unite and Divide, US legal scholar Cass Sunstein describes the phenomenon of group polarizations as follows: "When people find themselves in groups of like-minded types, they are especially likely to move toward extremes" (Sunstein 2009, 2). It is easy to see how this phenomenon might take hold in open group discussions, but group polarization effects don't *require* face-to-face interactions. This effect can even be translated through second-hand exposure to the opinions of others belonging to one's group (as is often the case in Delphi exercises), and even across distinct sets of individuals. Sunstein provides the following synopsis of a study of group polarization in a church setting:

> About a hundred church members were given a survey of sixteen church-related opinion statements, such as "ministers should feel free to take a stand from the pulpit on a political issue." Three weeks later, 169 other church members were given either the average of the hundred responses or a frequency distribution of the hundred responses, and then asked to make their own responses. They showed significantly more extreme attitudes than the original hundred (Ibid, 164).

Because Delphi exercises tend to be deployed in the early stages of insight gathering and so feed into later exercises in a foresight engagement – such as scenario planning, windtunnelling, and tangible futures – and because the opinions of experts may carry inordinate weight both for foresight practitioners and individuals engaged in a foresight activity, the effects of desirability bias in an uncontrolled Delphi exercise stand a high chance of surviving through the course of a foresight exercise, heavily

coloring both the mental models of participants and the outputs of the activity. Systemic desirability bias in the context of foresight can lead us to the conclusions that the futures we most *want* are the futures that we are most *likely* to get, a situation that might lead to short tem psychological benefits and happy clients without getting us any closer to a future-proofed strategic vision.

Unless the aim of the foresight activity is to unearth the *preferred* future of an industry or organization, foresight practitioners should take care to build additional controls into Delphi exercises in order to protect against the effects of group polarization and desirability bias. In a chapter on preventing extremism in groups, Sunstein suggests that groups also discuss the consequences, both negative and positive, that might follow if their predictions were to come true (Ibid, 134).

Another strategy for reducing these effects would be to have Delphi participants declare the desirability of the states of affairs represented in their predictions alongside their confidence ratings, as suggested by Ecken and his co-researchers. Simply highlighting the common effects of the desirability bias on predictions prior to the Delphi exercise may itself make experts subject to its effects. In addition, foresight practitioners can also perform a post-hoc analysis in order to determine the correlation between the assigned probabilities and desirability of each prediction,

determine the effects of desirability bias, and adjust biased predictions for those effects, as described by Ecken and his co-researchers (Ibid, 1664).[18]

Finally, paying particular attention to ensuring diversity of opinion within the expert group may be the simplest means of reducing desirability bias and polarization effects. As Sunstein points out, curating diversity of opinion is common practice among groups where the stakes of being led astray by biases are high:

> Well-functioning groups attempt to ensure a diversity of views, if only to protect themselves against blunders and confusion. If teams of doctors want to make accurate diagnoses, they will promote a norm of skepticism, even among younger and less experienced members. If corporations want to avoid disaster, they do best to create diverse boards that do not defer to the CEO (Ibid, 147).

Just as group diversity reduced the aggregated effect of the desirability bias in the above anecdote on the future of Chinese car manufacturing, diversity of opinion can also dispel the tendency for group polarization effects to push groups to more extreme predictions.

4.2 Trend Extrapolation and Non-Regression

One of the hallmarks of human reasoning is that we will behave like cognitive misers if given the slightest opportunity to do so. In order to preserve our precious cognitive resources, we'll substitute a hard problem

[18] The method itself presumes familiarity with probability theory, but foresight practitioners shouldn't be deterred; throughout this paper, I have suggested that the future of foresight as a proven methodology depends crucially on the ability of its practitioners to improve both the methodology itself and its results through learning, iteration, controlled experimentation, and quantitative methods.

for a comparatively easy one, doing our best to hide the bait-and-switch from everyone (including our own conscious mind) so that we aren't forced to do the harder mental work thorough analysis would involve (Stanovich et al 2010, 19). In the domain of future thinking, one (among many) potential result of this tendency, as Daniel Kahneman points out, is that "intuitive judgments can be made with high confidence even when they are based on a nonregressive assessment of weak evidence" (Kahneman 2011, 185). In other words, we often feel good about our judgments about the future even when we have little evidence to go on and our judgment of future outcomes are exactly identical – that is, nonregressive – in comparison to the judgment we'd make if we were prompted to make an assessment about the *present* rather than the *future*. An experiment conducted by Kahneman and Tversky illustrates and explains the phenomenon (Ibid, 187). They asked participants to judge the descriptions of eight college freshmen, allegedly written by a counselor assessing them for enrolment in a class. Each description consisted of merely five adjectives, such as:

intelligent, self-confident, well-read, hardworking, inquisitive

They then prompted one set of participants to answer the following two questions:

How much does this description impress you with respect to academic ability?

What percentage of descriptions of freshmen do you think would impress you more?

Kahneman suggests that, in the absence of any further information, people tend to answer this question by comparing the descriptions of the students to their internal norm about how counselors describe students. As he points out, the very existence of such a norm is in itself remarkable; it's quite likely that participants engaged in this experiment have never done any serious thinking on the subject, and may be creating a token of this norm in their minds completely out of whole cloth. Despite this, most participants had the sense that these adjectives, while positive, were not the most superlative adjectives they could use to describe students; they could have been brilliant, spectacular, erudite, shockingly original, etc. As a result, most participants gauged that students ascribed the five word description above were likely to be in the top 15%, but likely not in the top 3%. As Kahneman reports, there was impressive consensus around this range of judgments (Ibid, 188).

Kahneman and Tversky asked the second set of participants a slightly different set of questions:

What is your estimate of the grade point average that the student will obtain?

What is the percentage of freshmen that obtain a higher GPA?

The difference in the two sets of statements, Kahneman points out, is that in the first, participants are simply asked to make their evaluation

based on the evidence in the descriptions. The second set of statements treads through a great deal more uncertainty, such as: "What happened during the year since the interview was performed? How accurately can you predict the student's actual achievements in the first year of college from adjectives? Would the counselor herself be perfectly accurate if she predicted GPA from an interview" (Ibid, 188)?

The experiment was designed to assess the percentile gap participants made when merely evaluating evidence, in the first set of questions, and making predictions about the future performance of students based on their descriptions, in the second. Kahneman and Tversky found that the judgments participants made in both cases were identical. Why might this be the case? Cognitive miserliness helps to explain the result: "People are asked for a prediction but they substitute an evaluation of the evidence, without noticing that the question they answer is not the one they were asked" (Ibid, 188). Accounting for the above factors in rejigging our judgments is hard cognitive work, and so we fool ourselves into using our present evaluations as a "close enough" stand in.

Of course, we might condemn this as poor futures thinking, simply on the grounds that it does not account for the stories we might be able to tell about the paths students might take through their education. But that is only half the failing of this type of heuristic thinking; the other half is something that is more likely to be overlooked by foresight practitioners,

since it often left out of an accounting of the *causal* stories about future change that are central to strategic foresight. Critically, using our present evaluation as a stand-in for predictions ignores the brute phenomenon of "regression to the mean", a phenomenon so counterintuitive that it wasn't understood till "two hundred years after the theory of gravitation and differential calculus" (Ibid, 179).

Simply stated, regression to the mean is the probabilistic tendency for most measurements to cluster around an average over time, rather than toward the upper or lower limits of a range. One evocative way to understand regression to the mean is to consider physical attractiveness in offspring. With reference to regression, we should expect the offspring of attractive individuals to be, on average, less attractive than their parents, and offspring of unattractive individuals to be, on average, more attractive. There is no "causal" accounting for this; its simply because attractiveness and unattractiveness are polarized descriptions remote from an average, and that most people are average looking (which, of course, is what *average* means). To relate it back to Kahneman and Tversky's experiment, subjects asked to provide a prediction fail to account for the fact that a student who is above average in their freshman year is more likely to *under-perform* based on the expectations of this assessment than they are to over-perform. Mercifully, the opposite is also true: students who are below average are likely to improve toward the mean rather than become worse.

Since our minds naturally seek causal explanations for events, we want to believe that the adjectives used to describe freshmen students point to personal attributes that will sure-footedly carry them through to the 85th percentile in their classes. Because regression to the mean has a statistical, rather than a causal, explanation, we do not naturally think to account for its effects when rendering our judgments about the future. Just as seeing the image of Jesus on a piece of toast is a result of our over-active tendency to spot agents in the environment, we are fooled by the hardwiring of our brain into over-stating the case for the causal stories we concoct.[19]

There are many, many opportunities in the strategic foresight methodology for our cognitive miserliness to take over, leading us to substitute easy present evaluations for hard projections. The most obvious stage at which this might happen is in the process of projecting emerging trends into our future scenarios. In order to understand why we're vulnerable to substitution in this part of process, we should first note that present evaluations are "easy" only relative to the labor of thinking through the messiness of the future when extrapolating trends. Daniel Kahneman provides a partial catalog of the cognitive work we need to

[19] I'm not suggesting that we should rely more heavily on those images of the future that represent a world that is "like today, only more so." The work of "going out" to extremes in scenario building remains valuable, since it can help acquaint us with the broader possibility space under consideration. Rather, we should take any "raw" imaginative scenario as an un-regressed, highly causal story and interrogate its assumptions in order to bring it back closer to the mean.

perform when making a present evaluation (Ibid, 186). If we're assessing, for instance, whether or not the trend of wearable computing is worthy of consideration, we'll have to perform some or all of the following operations:

- We seek a causal connection between the evidence of the trend's current influence and it's importance for the future context of our area of exploration. Once a link is discovered, our associative memory "quickly and automatically constructs the best possible story from the information available" (Ibid, 186).
- The evidence for the trend is evaluated in relationship to a relevant norm. For instance, we might do a quick mental check to find other examples of information technologies that have had significant impact in the past, and add wearable technologies to the pile.
- Next, we substitute our weighing of the potential impact of wearables with the historical impact of those technologies around which we've constructed the norm. We come to believe that the performance of the technologies we've taken as the norm gives us license to ascribe similar potential performance to wearables.

It shouldn't be surprising, then, that the work we've done evaluating trends tends to be simply "ported over" by our inner cognitive

misers to reflect their appearance in our future scenarios. *After all*, our minds seem to assure us, *how much could the passage of time really matter?* But of course, the passage of time *itself*, quite apart from any driver or blocker you could name, counts for quite a lot.

This is certainly not an unknown problem to foresight practitioners. As it stands, astute foresight practitioners attempt to stop our trend evaluations from becoming overpowered in our images of the future by encourage organizations to think as well about *countertrends* and *blockers* to their chosen trends. Countertrends exist on the same level of analysis as our trends, acting as an opposing or competing force; for instance, religious conservatism might be a countertrend to liberal secularism. Blockers are higher level animating forces on par with drivers that might act as a damper to the development of our trends; for instance, future budgetary constraints might act as a blocker to the expansion of the welfare state. A consideration of countertrends and blockers is no doubt essential; it ensures that we do not grant ourselves license to create a frictionless passage for our trends as they move into the future by throwing opposing causal forces in their path. Countertrends are important for another reason: they make it less likely that our future images will be homogenous, since it insists that our trends live alongside their mirror images, making them more robust and evocative (though this might not always be what we want, as we'll see below).

On the other hand, both countertrends and blockers still fit into a perfectly *causal* story about how events unfold. They might be satisfying to our minds, which naturally seek to grasp onto causal explanations for why things turn out the way they do, but they are far from the whole story. Chiefly, the phenomenon of regression to the mean is left out of this causal story. Regression doesn't occur because countertrends exert more influence than do our trends, or because blockers get in the way of the growth of our trends. In addition to all of this, regression *just occurs* as a probabilistically inevitable consequence of the role that chance, luck, and time play in moving things closer to the mean from either direction: as Kahneman puts it, "regression to the mean has an explanation, but does not have a cause" (Ibid, 178).

As a general rule, we should assume that trends we feel will have low impact will turn out to be more influential than we anticipate, and that trends we intuit will have high impact will be less influential; both naturally regress toward the mean in the course of time. If there is a danger, it is in clinging to a default future of the world, it's also there in the propensity for defaulting to substitution when projecting weak signals into the future; in the former case, relying on default futures can cause our view of the future to be too similar to the present, and in the latter, too dissimilar. When substitution is coupled with the charisma of causal explanations of change – especially the charisma of those we craft ourselves – we have a recipe for a kind of future myopia quite different

than the one foresight practitioners typically warn against. Kahneman puts it succinctly: "Be warned: your intuitions will deliver predictions that are too extreme and you will be inclined to put far too much faith in them" (Ibid, 194).

4.3 Default Futures and Alternate Scenarios

One of the key claims of strategic foresight is that the act of taking into consideration several possible futures naturally loosens the hold of an individual's default mental model of the future, making room for more flexibility in the way they conceive of and respond to the future. Just as we can reduce the effects of hindsight bias – the feeling that, in regards to *how things actually turned out*, "we knew it all along" – through exploring carefully crafted counterfactuals (Fischhoff 1976), so too can we reduce attachment to the *default future* within an organization through creating multiple carefully crafted future scenarios.

It's a claim with strong prima facie plausibility, and it's incorrect. Or at least only correct under specific circumstances. It's a critical mistake to assume that investment in the default future *always* decreases as we generate and explore alternative futures; as it turns out, there are several situations in which considerations of alternate scenarios can actually further entrench commitment to the default future.

Psychologists Neal Roese and James Olson uncovered just such a counterintuitive inversion in the relationship between hindsight bias and

counterfactuals. Hindsight bias is the name for the common feeling of "having known it all along"; individuals tend to overrate their confidence in having predicted the outcome of an event after the fact, seeing it in retrospect as a virtual inevitability. The hindsight bias is more than the simple inability to learn from experience. It also represents "an inability to retrieve one's own pre-outcome explanatory perspective" (Rose & Olson 1995, 198). An oft-quoted passage in the cognitive science literature from historian Georges Florovsky keenly sums up the phenomenon: "In retrospect, we seem to perceive the logic of events which unfold themselves in a regular or linear fashion according to a recognizable pattern with an alleged inner necessity. So that we get the impression that it really could not have happened otherwise" (Florovsky 1969, 369).

It seems reasonable to assume that exposing individuals to counterfactual thinking should reduce the effects of hindsight bias. For instance, presenting plausible alternate scenarios for how a baseball game may have gone otherwise should make even die-hard fans of the victors pause to consider ways in which the team might've blown the game. One reason we might expect this is that counterfactuals increase the pool of potentials over which to distribute probabilities. If you present the baseball fan with several alternate visions of how their team might've lost, they'll be compelled to reduce their confidence on the likelihood of the actual outcome. As we'd expect, it has been demonstrated experimentally

that the more alternatives to a stated outcome that subjects consider, the lower the median likelihoods ascribed to these outcomes (Fischhoff 1976).

But it's not so simple as that. Counterfactuals might reduce hindsight bias in certain circumstance, but psychologists Neal Roese and James Olson have experimentally demonstrated that counterfactuals can in some cased actually *heighten* the effect of the hindsight bias. When individuals have established strong causal chains of inference connecting events in the past to actual outcomes, the hindsight bias is stronger relative to instances where more random elements are at play, and the clearer the outcome is wedded to causal antecedents, the stronger the hindsight bias becomes. This effect is strongest when the counterfactuals were targeted at *undoing* the actual outcome.

In order to demonstrate this effect, Roese and Olson had subjects read a scenario depicting a student preparing for an important exam. The student "engages in several preexam undertakings (the target causal antecedents), some facilitative and other inhibitory of success on the exam" (Roese & Olson 1996, 203). One such action, taking pills for combating panic attacks, was alternately portrayed as being either consistently or inconsistently correlated with her success on exams in the past. In this scenario, the student forgets to take the pills. This leaves it open for subjects to contain two counterfactuals: "if she had taken the pill, she may have performed better" (the undoing condition) or "if she had taken her pill, her performance would have been the same" (the no-

undoing condition). As Rose and Olson hypothesized, when subjects considered the *undoing* counterfactual, it *heightened* their hindsight bias regarding the actual outcome. When counterfactuals are highly causative, and when they are portrayed as possibilities that would've undone the actual outcome, they can further entrench the feeling that we knew it all along.

This experimental result points to troubling implications for human foresight, as well. Working on the implication of the Janus Hypothesis, which suggests that the vagaries of memory can also manifest in the way we deal with the future, perhaps foresight bias (which, depending on the individual in question, might be based in status quo bias or optimism bias) will similarly turn out to be *reinforced,* rather than weakened, by certain types of scenarios or scenario generating methods. If hindsight bias is "a projection of new knowledge into the past accompanied by a denial that the outcome information has influenced judgments, (Wasserman, Lempert, & Hastie 1991, 30) then "foresight bias" might be the projection of current knowledge into the future accompanied by an assertion that incoming information won't influence judgments.

Of course, the most important asymmetry between judging the likelihood of things that have already happened and those that haven't happened yet is *outcome knowledge*. Simply put, if you know the outcome of a causal chain, you are more likely to highly rate the likelihood of that outcome against other possible outcomes. But though outcome knowledge

is never the case in foresight, there are very likely other factors that together cause us to more highly rate the likelihood of certain futures over others, such as optimism bias and overconfidence in entrepreneurial individuals, or status quo bias in managerial types (Busenitz & Barney, 1997), polarized organization thinking, groupthink, or any of the raft of biases mentioned in Chapter 3. Couple these biases with an organizational attachment to the results of market and financial forecasts, which provide a causal story about the future, and you have all the components of a robust, default future. Put plainly: creating plausible, strongly causative possible futures engineered to steer organizations away from their default view of the future might actually lead them to perceive the default future as *more*, not *less*, likely. Luckily, there is nothing stopping foresight practitioners from testing this hypothesis using the same experimental design as Rose and Olson, simply flipping the scenario script from a consideration of the past to a scenario of the future. The Janus Hypothesis strongly suggests that we will uncover the same effect, and indeed that is what we are finding (Shnaars and Topol 1987; Kuhn and Sniezek, 1996).

But neither should we exclusively build scenarios that are deliberately *unlikely* in an attempt to offset commitment to the default future. Consideration of *wildcard* scenarios – highly impacting, highly improbable future events, or Black Swans to use Nassim Taleb's coinage – can also prop up the default future, as can scenarios that are improbable and merely run of the mill in terms of their impact. Psychologist Michael

Dougherty and his co-investigators examined how the likelihood that individuals attached to focal causal scenarios about the past was altered through consideration of alternate scenarios. In order to frame their discussion, they discuss speculation surrounding the crash of Flight 800 over the Atlantic Ocean in 1996. The initial investigation surrounding the crash focused on a causal scenario in which a terrorist had planted and detonated a bomb on the plane, though after the most exhaustive and expensive air disaster investigation in U.S. history it was determined that the most probable cause was an explosion of flammable air vapors inside a fuel tank. Dougherty and his co-researchers hypothesized that there were at least two ways that newscasters and investigators may have initially become overconfident in the terrorism-based causal scenarios explaining the crash.

First, "the newscasters and investigators may have completely failed to generate alternative causal scenarios", in which case the cause of overconfidence is a kind of causal myopia (arguably similar to the phenomenon of default futures). Second, the newscasters and investigators might have generated alternative scenarios, but "judged them to be so unlikely that they were discounted or eliminated from serious consideration" (Dougherty et al. 1997, 137). This suggests that there is a sort of perceived plausibility threshold below which scenarios are simply discarded in mental reckoning. This is a phenomenon that should be familiar for many foresight practitioners: one or more scenarios

is left dangling outside the discussion because it has failed to make it into the ring of "leading contenders" for plausible futures, for any number of reasons. The research by Dougherty and his co-investigators strongly suggests that not only do these scenarios have virtually *no* impact on reducing the perceived plausibility of focal causal scenarios; they may actually *increase* their perceived plausibility.

In order to test this hypothesis, the investigators presented subjects with the following script, constructed to be strongly suggestive of a focal causal scenario, with information about other potential causal scenarios omitted:

> It was the smokiest fire that Bill had seen in his eight years as a firefighter. Bill thought he could handle the fire by himself while the others went to get a second hose. He entered through the main entrance on the second floor. It immediately became clear that he would have to make it to the basement in order to extinguish the fire. The smoke from the fire made it especially difficult for Bill to see where he was going. He soon became disoriented and had no idea how long he had been in the building or how far he had traveled into the building. Nevertheless, Bill hosed down the fire while he waited for help. Unfortunately, by the time his co-workers reached him, Bill was dead (Ibid, 141).

The script is meant to be suggestive of a focal casual scenario in which Bill died from smoke inhalation, and was found to illicit this scenario through beta testing.

In addition to this generic version of the script, the experimenters devised two additional scripts with added content. The second script was intended to make alternative explanations for Bill's death more likely by adding the following elements:

It was a relatively hot fire

The building was made primarily of wood beams, as it was a relatively old building (Ibid, 144).

The third script was intended to make alternative explanations for Bill's death *less likely* through adding the following elements:

It was a relatively cool fire.

The building was made primarily of steel beams and concrete, as it was a relatively new building (Ibid, 144)

As we can see from the figure below, subjects rated the focal scenario as slightly *less likely* when additional information was added making alternatives more likely, and slightly more likely when additional information was added making alternatives less likely.

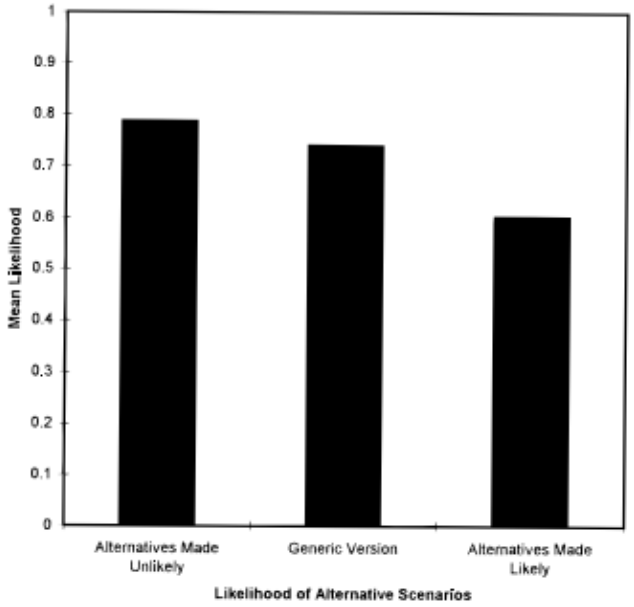

Table 2. Mean likelihood of the focal causal scenarios for the three conditions described in Dougherty et al. 1997. When alternative explanations for a scenario seem less probably than a focal explanation, rating of likelihood of the focal explanation actually *increases* relative to a situation in which no alternative is offered to the focal explanation.

Of course, neither of these findings demonstrate beyond doubt that the same tendencies to react to certain types of alternative scenarios with increased confidence in default scenarios applies in the same way to future thinking. But the possibility should give us foresight practitioners pause, and press us to investigate further. As a start, we should at least cast some skepticism on the common claim that the act of considering multiple possible futures unfailingly loosens the hold that the default future has on us. If these same pitfalls apply to future thinking, it seems more likely that in crafting future scenarios for the purposes uprooting default futures we must carefully and deliberately steer between Charybdis and Scylla.

4.4 Probabilistic Reasoning and Scenario Planning

We can see from the above discussion that another common injunction in foresight – that we *not* attach probabilities to future scenarios and instead consider them as mere possibilities – is asking the impossible. Both Pierre Wack and Herman Kahn wished to avoid attaching probabilities to scenarios for fear of limiting engagement to those scenarios that were deemed most probable; their aim, after all, was to "stimulate senior management thinking about the future, not to provide forecasts of the most likely futures" in line with financial and statistical

methods that had "become the opiate of corporate strategy" (Millett 2009, 62). But whether probability figures are on a workshop whiteboard or implicit in the confidence weighting individuals mentally ascribe to them, scenarios are *always already* being probabilistically assessed. Ignoring these naturally occurring probability weightings only allows them to secretly and freely influence scenario and strategy development, and potentially creates heavy anchors or irrational resistance in an individual's mind, subtly biasing them toward certain scenarios and strategies.

Rather than pretend that these probabilities do not factor into the scenario exercise, we should face them head on and ensure that the influence of scenarios on our thinking is reasonably de-biased; not either disproportionately strong or weak given the mental frame in which we interpret them. One potential solution is to have individuals *declare their Bayesian priors* ahead of scenario development in order to assess their confidence in a selection of propositions that illustrate the default future of their organization. In his well-known work "An Essay Toward Solving a Problem in the Doctrine of Chances", English minister Thomas Bayes formulated an account of how people naturally formulate probabilistic estimates about the way the world works, and provides a method for *rationally* updating these estimates in accordance with new knowledge. *Bayesian priors* are the probabilities we assign to certain events, based on previous knowledge, intuition, and expectations; they are the machinery that powers up our default mental images of the future.

Declaring our priors and thinking through how new knowledge will impact our judgment can help us to connect imaginative future scenarios to a method for rationally grappling with uncertainty. As Nate Silver writes, successful forecasters of any kind "do not think of the future in terms of no-lose bets, unimpeachable theories, and infinitely precise measurements," but instead as "speckles of probability, flickering upward and downward like a stock market ticker to every new jolt of information" (Silver 2012, 238). A Bayesian approach doesn't attach probabilities to scenarios. Rather, it provides us a means of rationally updating *our* assumptions about the world when presented with new knowledge.

After a lengthy period of dormancy, Bayesian models are "on the ascent throughout the cognitive sciences", especially in the psychology of reasoning (Elqayam et al. 2013, 454). This new paradigm represents a break from the traditional deduction-based paradigm, "which took classical, binary logic as an appropriate model for human deductive competence and marginalized the role of prior beliefs and desires (Ibid, 454). The Bayesian paradigm of the psychology of reasoning represents "the shift from truth to belief"; rather than simply assert whether or not you agree with some proposition, it is crucial to report the degree to which you believe or disbelieve it, and provide justification for the *degree of your belief*. Only then can we begin to understand the actual dynamics of interactions between existing mental models and new knowledge.

Let's imagine that an individual from the marketing division in a company that manufacturers biosensors believes there is a 50% chance that in five years time the wearable technology market will be approaching the size of the smartphone market. This is her prior (x). She is presented with a scenario in which Google Glass is launched in early 2014 and is a complete flop. How should she update her priors based on this hypothetical event? First, she needs to estimate the probability of Google Glass flopping if her prediction in the market success of wearables is correct. This possibility makes her feel slightly bearish on wearables, but she still believes there are other avenues to growing the market, so let's say she assigns a 13% likelihood (y) of her initial claim given this event. Finally, she needs to estimate the probability that her initial belief is false given the hypothetical terrible reception of Glass in 2014. She admits that the development would not bode well for the market, so she assigns this 35% likelihood (z). Once we've estimated these values, we can apply Bayes' theorem (the algebraic expression in the final line of the table below) to establish how our imaginary marketer should update her belief based on the hypothetical new knowledge.[20]

[20] This method for elucidating Bayes' theorem is adapted from Nate Silver's account in The Signal and the Noise.

PRIOR PROBABILITY		
Initial estimate of how likely it is that the wearables market will be as large as the smartphone market by 2018.	x	50%
SCENARIO EVENT: GOOGLE GLASS FLOPS		
The probability of wearables matching smartphone market given the failure of Google Glass	y	13%
The probability that wearables will fall short of the smartphone market, given the failure of Google Glass.	z	35%
POSTERIOR PROBABILITY		
Revised estimate of how likely it is that the wearables market will rival smartphones given the failure of Google Glass.	$\frac{xy}{xy + z(1-x)}$	27%

Table 3 Bayesian analysis of Google Glass scenario

Based on an application of Bayesian reasoning, given the flop of Google Glass and her own confidence ratings, our imaginary marketer should be about half as confident about the prospects of the wearables market matching the smartphone market by 2018; a significant change, but not exactly a catastrophic collapse in confidence. Note however that without making her priors clear and thinking deliberately through how her belief would be impacted given this new information, there's a good chance that her confidence levels would be lead by other considerations, such as group norms, or her emotional reaction to elements of the scenario under consideration, causing her either to under adjust or over adjust her estimates. All that Bayes' theorem does is provide an operation for "how

base rates should be integrated with new evidence to produce posterior probabilities" (Ibid, 459).

As foresight practitioner Stephen Millet points out, both scenario generating teams and executives "gravitate toward the scenarios they find to be "most interesting," which typically reflect corporate culture biases and wishful thinking" (Millet 2009, 65). Application of Bayesian reasoning can loosen the hold these scenarios have on us and lead us to carefully think through scenarios that have less impact and appeal. It can also loosen the effects of the anchoring bias, which makes us cling strongly to initial estimates (even implicit ones) and resist movement away from this norm; Bayes' theorem serves as external evidence that we might be overly invested in our initial estimates, especially in light of compelling evidence that we may be mistaken.

The use of Bayesian reasoning also encourages people to give substance to their judgments, which can generate fruitful discussions and "expose hidden assumptions, biases and expectations that too often go unarticulated in the generation of purely intuitive scenarios" (Ibid, 65). Noting where priors are particularly strong or weak can give foresight strategists important clues about the default future in an organization, and how this future is socialized and differently weighted in different areas of the organization. Imagine, for instance, that Bayesian analysis reveals that the marketing team has a far higher prior belief than the engineering team in the pending boom in wearable technology. This finding would no

doubt be an important jumping off point for discussion and broader strategic alignment within the organization.

What we know of the way that people spontaneously update their beliefs suggests we often run afoul of Bayes' theorem. For instance, people frequently neglect base rates when making inferences; if asked whether it is more likely that a man with long black hair wearing a death metal t-shirt is a Satanist or a Christian, people will tend to answer that is he more likely to be a Satanist, despite the fact that Christians number over two billion people and Satanists only a few thousand worldwide. The order in which people are exposed to new information seems to have an impact on their degree of belief (in other words, belief updating is *diachronic)* even though "the temporal order in which information is integrated should make no difference" (Ibid, 459). Lastly, an individual's existing mental models have a powerful inertial quality: "because holding on to the existing model takes up far less cognitive effort than revising or discarding it, beliefs are not always updated even in the face of new information" (Ibid, 460). Bayesian analysis - and indeed all of the methods suggested in this chapter, as well as the foresight methodology itself - presumes that "human rationality is fallible but corrigible" (Ibid, 459). There is evidence to suggest that using Bayes' theorem in a pedagogical way can reduce the effects of phenomenon like base-rate neglect, diachronic belief updating, and mental model inertia. The most obvious way in which this takes place is when subjects explicitly endorse Bayesian

analysis and other methods that impose a normative structure on belief updating. In an introduction to the work of Amos Tversky, psychologist Elda Shafir stresses this point: "The research showed that people's judgments often violate basic normative principles. At the same time, it showed that they exhibit sensitivity to these principles' normative appeal" (Shafir 2003, x).

If foresight is to be a valid process for examining and improving our mental models of the future, then an examination of priors should be an integral part of the process. We shouldn't mistake the acts of scenario creation and socialization as adequate for the purposes of shaking up thinking and changing mental models. If not accompanied with the proper frame for understanding scenarios, foresight activities can actually introduce *more* uncertainty into an organization, not less, by driving individuals and groups deeper into their unconscious biases. Activities for examining and updating priors should go hand in hand with periodically renewed trend analysis and foresight activities. Bayesian reasoning exercises might also be especially powerful when used in conjunction with the results of backcasting techniques; participants could then wind tunnel strategies and assess their beliefs at every temporal point in the backcast to get a sense of how their probabilistic assessments change with each new piece of information.

Millet's ends his examination of the use of probabilities in scenarios by suggesting that "if the use of probabilities with scenarios fits the

corporate culture and stimulates creative thinking about alternative strategies under alternative conditions, then they should be used." He even goes so far as to suggest that probabilities should only enter into the conversation if the scenario team "is familiar and comfortable with the concept of Bayesian probabilities" (Ibid, 66). No doubt, the use of algebraic equations and probability theory would probably frighten foresight teams that are not accustomed to dealing in the quantitative aspects of their organizations. But although the algebraic expression used to update priors is daunting and somewhat arcane for laypeople, the heart of Bayesian reasoning is rather simple. An organization called the Center for Applied Rationality (CFAR) has had success recently in instructing teams and entrepreneurs in the art of applied rationality, including simple and intuitive exercises that introduce them to Bayesian reasoning and demonstrate its utility in everyday situations.[21] Foresight practitioners could learn a lot from their friendly and intuitive approach to probing mental models. For instance, CFAR has devised an exercise that "gamifies" the creation of Bayesian priors and the process of updating them based on new events and evidence by structuring these activities around an unfolding, engaging murder mystery. Similar techniques, using elements of the narrative elements of created scenarios, could serve as a powerful tool for participants to explore how unfolding events would impact their judgments and decision-making processes. For instance,

[21] http://rationality.org/

participants could be sequentially led through the results of backcasting exercises in order to better understand the impact of unfolding events on the strength of their prior beliefs.

4.5 Motivated Reasoning and Tangible Futures

As foresight practitioner Karl Schroeder has pointed out, scenarios are "highly charismatic artifacts" that may draw attention away from the serious and essential work yet to be done in a foresight exercise: using scenarios as a backdrop in front of which to perform strategic assessments (personal communication, November 18 2013). With scenarios, more vivid is therefore not always better. The aesthetics of scenarios might well distract from, rather than enhance, strategic thinking.

But it's not only that scenarios are potential *diversions*. Scenarios also often have a strong *affective* quality; in order to increase their perceived plausibility, they are designed to engage individuals on an emotional level through use of techniques like narrative, science fictioning, and visualization. The affective components of scenarios and other aesthetically striking elements of the foresight methodology are usually understood as positive attributes because they can make provocative ideas more *plausible* and *relatable*; the combination of plausibility and provocation is, after all, the hallmark of good scenario design. Plausible, provocative scenarios are those that are close enough to

reality to achieve buy-in from foresight consumers, but simultaneously jarring or unexpected in some fundamental way that motivates individuals to treat them as puzzles to solve in the process of strategic work.

However, provocation in this context is a little ambiguous. Are we *provoking* careful, considered engagement with our scenario worlds, or are we *provoking* visceral, emotional responses that lead to irrational gut reactions? In our scenarios, are we speaking to System 1, System 2, or both? And which is our primary audience? These affective elements create opportunities for *motivated reasoning* in which individuals might be either unconsciously enticed or repulsed by the emotive components of representations of the future, leading them to biased strategic thinking. "The affective forecasting bias leads us to exaggerate the emotional satisfaction of a future success and the devastation of a failure," notes Thomas Suddendorf. The reality, he notes, is that "we typically don't get as excited as we thought we might at success, and we handle failure much better than we imagine." Strangely, the bias persists even in the face of persistent evidence that we are poor forecasters of our future emotional states: "We experience the reality, but we don't update our minds," Suddendorf notes, "we keep the exaggeration going" (personal communication, January 29, 2014).

We should note that the question of motivated reasoning is not only a problem of scenario exercises. Indeed, as access to multimedia

technologies and their ease of use increase, opportunities arise to inject emotion-generating components into various steps of the foresight methodology, creating new opportunities for foresight consumers to engage in motivated reasoning. The tangible futures method and its sister methods (future artifacts, experiential futures, etc.) are increasingly used in foresight activities and are perhaps the most obvious place to begin asking questions about the impact of the affective qualities of foresight on motivated reasoning.

In an interview with futurist Stuart Candy on the topic, design strategist Vince Lombardi, who pioneered the use of tangible futures, provides the following definition: "Tangible Futures are the output of applying design-fueled disciplines like visualization, drama, and film to represent futures and strategies" (Lombardi 2008).[22] Describing the development of tangible futures, Lombardi provides the framing question that drove him and his co-worker Christina Wodtke: "How can we help managers *experience* futures and strategy so that it can be more substantially understood, shared, and acted on" (Ibid.)? In Bringing the

[22] In the same interview, Lombardi claims that tangible futures are part of the solution to an over-reliance of what he calls the "cognitive" elements of strategy, by which he means that strategy "rarely exists outside of our minds." I do not think that he intends the same thing by "cognitive" as I do in this essay; he seems to rather be using it as a way to describe the space occupied by "proverbial binders of reports" which do not significantly engage managers. All the same, we shouldn't fool ourselves into thinking that artifacting on its own is enough to ensure that individuals think more deeply about the futures they're being presented. It may sometimes do the opposite by appealing directly to System 1, which is my primary claim in this section.

Future to Life, digital futurist Trevor Haldenby articulates another facet of the rationale for using tangible and experiential futures:

> Through the creation of physical artifacts, interactive environments, cinematic narratives, and social communities, futurists with varying approaches and goals are telling stories that make speculative scenarios more engaging to non-specialist audiences (Haldenby 2013, 58).

No doubt, the use of multimedia is effective both for the purposes of driving engagement, as well as democratizing and socializing the findings of foresight activities. For these reasons, as well as the increasing availability and ease of use of the technologies enabling these methods, we should expect that foresight practitioners will incorporate more multimedia elements into the methodology moving forward.

But in this process, we should not lose sight of the fact that foresight generates material which makes participants uncomfortable in that it emphasizes the "uncertainty, instability, and precariousness" of the future over images that "underline surety or make accurate predictions" (Healey and Hodgkinson 2007, 578). And if participants experience this discomfort even when scenarios are merely text on a page, we should expect that the negative affective dimensions of scenarios would only be *amplified* when they're built into vibrant multi-media worlds.

Why is this a bad thing? Of course, scenarios *should* be provocative; Kees van der Heijden and his colleagues suggest that scenarios that exert the strongest influence on decision making are those that "elicit feelings of fear, hope, security, and threat" because these emotional states "create the

jolt needed for action" (van der Heijden et al. 2002, 263). But much of what we've learned about the effects of subjecting people to uncertainty suggests that exposure to provocative material might make it *less* likely that they do the hard work of thinking through its meanings slowly and carefully. Instead, decision makers may instead consider scenarios with reference to "an affective reaction to a salient image, and this feeing (not explicit consideration of the scenario's probability) may guide behavior" (Hsee and Rottenstreich 2004, 28). Healy and Hodgkinson provide a small litany of the potential negative effects of uncertainty on the overreliance on affective reasoning: "… perceived uncertainty over anticipated events, decisions, and their outcomes has been linked with rigidity and slower decision making, escalation of commitment to a failing course of action, and increased interpersonal conflict and reduced performance among decision making groups" (Ibid, 578).

An anecdote related by Healey and Hodgkinson illustrates this potential pitfall. A team of foresight strategists attempted to use scenario planning techniques to help a publishing firm explore their business context so that they could develop an adaptive strategy that addressed critical uncertainties in the volatile publishing industry head on. The result was a set of scenarios that drove the publishing company's management team into a reactive, defensive posture: "In constructing the scenarios, the management team focused on envisioning quite vividly a threatening future in which technological changes would replace their

main offering, to the extent that this triggered defensive avoidance and threat rigidity effects" (Ibid 579).

Defensive avoidance and rigidity are both strategies invoked by highly negative emotional stimuli. In the former, decision makers will default to "choosing strategies that deflect the responsibility on to other individuals or on to factors outside the chooser's control (Foskett and Hemsley-Brown 2001, 41), while in the latter they "reduce their flexibility under a stress situation, sealing off new information and controlling deviant responses" (Janis 1972; Staw et al 1981, 502). Healey and Hodgkinson continue their analysis of the event:

> Consequently, they were unable to reach a consensus on an alternative to the current failing strategy. This proved to be anything but an anhedonic response. The scenario intervention "raised the levels of decisional stress and conflict within the group to unacceptably high levels" (Hodgkinson and Wright 2002, 964). The stress created by attempting to face an uncertain future with a disparate team led the decision makers to adopt a variety of dysfunctional coping strategies, including bolstering commitment to the current failing strategy, procrastinating, and shifting responsibility for maintaining the inert status quo to other stakeholders within the firm (Healey and Hodgkinson 2008, 579-581).

It is hard here not to be reminded of the biological concept of a supernormal stimulus. In her book *Supernormal Stimuli: How Primal Urges Overran Their Evolutionary Purpos*e, psychologist Deidre Barrett provides a quintessential illustration:

> Nobel laureate Niko Tinbergen coined this term after his animal research revealed that experimenters could create phony targets that appealed to instincts more than the original objects for which they'd evolved. He studied birds that lay small, pale blue eggs specked with gray and found they preferred to sit on giant, bright

> blue ones with black polka dots. The essence of the supernormal stimulus is that the exaggerated imitation can exert a stronger pull than the real thing (Barrett 2010, 3).

In effect, tangible and experiential futures are a supernormal experience when contrasted with the native foresight abilities of human beings. If they are doing their job properly by provoking us to confront tangible and experiential futures – an immediacy that has no match in the evolutionary landscape - we should *expect* supernormal responses to these methods.

But those supernormal responses are not different in *kind* to the normal responses we have to prospection. As we've seen, projecting ourselves into the future can help us formulate plans with greater fidelity and prepare for potential uncertainties. But it can lead us to overcommit to futures that we haven't thought through sufficiently, or even to become anxious and paralyzed in the face of future possibilities. Tangible futures can amplify our natural, unreflective optimism about the future; they can also raise the plethora of fears and anxieties we experience in the face of pressing risks and uncertainties. As a method for provoking thinking about the future, creating supernormal stimuli from imagined scenarios is not a poor strategy per se, but if executed improperly it can lead us ever further down the path of poor future thinking and strategic decision making.

There may also be good evolutionary reasons that we are attracted to strong narratives about the future: they galvanize action around a

shared vision. "It's adaptive to exaggerate these emotions because they allow us to coax other people into joining us," Suddendorf says. As a social mechanism, he notes a parallel between affective forecasting and the practice of self deception as described by biologist Robert Trivers. Both present the same puzzle: the practice of creating overly optimistic (or pessimistic) visions of the future, or hiding our true motivations for our actions via self deception, seems like it should be maladaptive, since it causes us to hide key affective and motivational information from even ourselves. Why, then, did these habits persist in our evolutionary history? Trivers hypothesis was that we self deceive in order to be better a deceiving others. Suddendorf suggests that, similarly, "we exaggerate how wonderful it would be if we could achieve X in order to coax people into cooperating and join in our vision." Moreover, "if you believe your own exaggerated emotions, than if people found out you were leading them astray you aren't punished as much" (Suddendorf). This explanation should give us pause. If it's true, then our tendency to engage in unconscious affective forecasting is ultimately about protecting our own hides from group censure and the cheater-punishing tendency inherent in our social intelligence.[23] Foresight strategists steeped in the evolutionary

[23] The human urge to punish cheaters is uncannily strong. Consider the two variations of the Wason selection task. In the basic test, participants are shown a set of four cards place on a table, each of which has a number on one side and a colored patch on the other. The visible sides of the cards show the numbers 3 and 8, and the colors red and brown. Experimenters ask participants to test the truth of the following proposition "if a card shows an even number on one face, then it's opposite face is red" by turning over only those cards which will decide the truth of the proposition. Following the rules of classical logic, the only cards that

rationale for affective forecasting take note: allowing participants to generate unrealistically optimistic futures without putting a check on the process might be more about self-preservation – on both their part and your own – than about creating compelling visions of the future.

None of this should be taken to mean that we should never employ rich, imaginative, multi-media content within the context of strategic foresight; instead, there may be situations in which it is appropriate, and situations in which it is not. Or we may need to embed additional information or controls into the process to ensure that we are not unintentionally creating opportunities for highly motivated reasoning in participants – unless, of course, that is what we want, as may be the case when we are trying to elicit *preferred futures*.

At this point, we simply don't know enough to say with any degree of certainty which of these situations we're in during a foresight activity. Mark P. Healey put his finger on the difficulty of assessing the efficacy of these methods during our interview:

---

can disprove the proposition are the 8 and brown card: if the 8 card is brown on the other side, or if the brown card has an even number on the other side, then the rule is violated. Only 10% of participants correctly choose these two cards. However, if the colors and numbers are replaced "16 years old", "drinking beer", "25 years old", and "drinking water" and the proposition to be tested is changed to "if you are drinking alcohol then you must be over 18", participants fare *much* better. It seems perfectly clear that you should check what the 16 year old is drinking as well as the age of the individual drinking the beer; checking what the 25 year old is drinking or the age of the water drinker can do nothing to detect whether a rule violation is taking place. According to Leda Cosmides and John Tooby, the fact that participants fare much better on the "social rule" version of the Wason selection tests strongly suggests that human reasoning is governed by context-sensitive mechanisms that have evolved to solve specific problems of social interaction, rather than context-free, general-purpose mechanisms.

> We don't know enough about the psychology of how people think about and deal with uncertainty in a more general sense. We know a lot about making probability judgments under uncertainty from the Kahneman and Tversky tradition, but we don't know a lot about the trade offs between the motivational benefits and their essential cognitive downsides. And that makes it difficult to figure out whether these types of techniques are useful. And we really don't know enough about the techniques themselves (personal communication, November 20, 2013).

Whatever we discover about the extent to which cognitive and affective reasoning are impacted through the use of multimedia in illustrating scenarios, we should consider using either facilitation practices of additional methods that allow for careful handling both of the unwarranted optimism and anchoring effects of positive scenarios and "of the anxiety and decisional stress that can arise when users imagine and simulate future threats with scenarios" (Healey and Hodgkinson 2007, 581).

Futurist Jim Dator has argued that all images of the future tend to fall into one of four broad categories: scenarios of continued growth, driven by a faith in new technologies and market forces; scenarios of collapse, rooted in "concerns about overpopulation, energy and other resource exhaustion, and environmental pollution"; scenarios of discipline, based in the belief that "continued economic growth is either undesirable or unsustainable," and so we should revisit "fundamental values – natural, spiritual, religious, political, or cultural" in order to bring growth to heel in the name of "survival and fair distribution"; or scenarios of transformation, which posit that radical technological transformations

such as "robotics and artificial intelligence, genetic engineering, nanotechnology, teleportation, space settlement" will lead to the "emergence of a 'dream society' as the successor to the 'information society'" (Dator 2009, 10).

It is worthwhile noting that each of these scenario types carries *motivational* weight. As Dator himself points out, scenarios of continued growth are especially attractive to those from the worlds of modern government, educational systems, and organizations, while collapse scenarios motivate the efforts of those involved in social or environmental causes. Scenarios of discipline are highly motivating for religious and moral authorities, while scenarios of transformation appeal to those who put their faith in the disruptive power of technological innovation. Each of these groups sees in their preferred future a strongly uni-directional and causal story that can be read as a manifestation of unchecked motivated forecasting. Dator also notes the connection between an individual or group's preferred future and their tendency to be swayed by biased thinking, and agrees that the origin of these generic images of the future might "represent fundamental human biases":

> It is very important to understand that individuals strongly assume that whatever is happening now (in their opinion) will continue: if times are good, the default assumption is they will continue to get better. If times are bad, they will get worse. This presumption of the future as the present continued made total sense for tens of thousands of years and is deeply ingrained in all of us, biologically and psychologically – and culturally (personal communication, February 21, 2014).

Given that overshooting the mean might be a hardwired component of human foresight, we should be looking for ways to mitigate this effect in strategic foresight engagements. Dator's own strategy for fighting bias in a client's vision of the future is to ensure that "preferred futures visioning take place late in the overall futures project, after many other things have been experienced." Most importantly, clients should have "experience in at least two of the generic alternative futures," as described above (Ibid 2014).

Another potential strategy for coping with the affective components of uncertainty that drive us away from the mean is to situate the strategic decisions under consideration within a reference class of similar past initiatives. Reference class forecasting, a method developed out of the theories of Kahneman and Tversky a project's potential outcomes with those of similar, past projects to produce more accurate predictions" (Lovallo and Kahneman 2003, 1). Rather than forecasting trends, reference class forecasting asks us to reassess our projections for a project's success by locating its reference class, assessing the distribution of outcomes in that class, predicting our project's position in that distribution, and then adjusting our intuitive prediction based on how we feel our predictions have performed in the past (Ibid, 1). Lovallo and Kahneman provide a sample reference class forecasting description, which here appears in an abridge form:

1. **Select a reference class:** Identifying the right reference class involves both art and science. The key is to choose a class that is broad enough to be statistically meaningful but narrow enough to be truly comparable to the project at hand.

2. **Assess the distribution of outcomes:** Once the reference class is chosen, you have to document the outcomes of the prior projects and arrange them as a distribution, showing the extremes, the median, and any clusters.

3. **Make an intuitive prediction of your project's position in the distribution:** Based on your own understanding of the project at hand and how it compares with the projects in the reference class, predict where would fall along the distribution.

4. **Assess the reliability of your prediction:** This step is intended to gauge the reliability of the forecast you made in Step 3. The goal is to estimate the correlation between the forecast and the actual outcome, expressed as a coefficient between 0 and 1, where 0 indicates no correlation and 1 indicates complete correlation. In the best case, information will be available on how well your past predictions matched the actual outcomes. In the absence of such information, assessments of predictability become more subjective. You may, for instance, be able to arrive

at an estimate of predictability based on how the situation at hand compares with other forecasting situations.

5. **Correct the intuitive estimate:** Due to bias, the intuitive estimate made in Step 3 will likely be optimistic – deviating too far from the average outcome of the reference class. In this final step, you adjust the estimate toward the average based on your analysis of predictability in Step 4. The less reliable the prediction, the more the estimate needs to be regressed toward the mean (Ibid, 8).

Not only can reference class forecasting ameliorate the biases arising from affective reasoning, it can also render a more accurate, realistic picture of the prospects of strategic directions. However, it's utility might be limited to cases in which foresight participants are laboring under an overoptimistic view of the future due to affective reasoning. In this case, reference class forecasting can help to re-plant the team's feet firmly on the ground. But it might not do the same for teams who experience option paralysis or buck-passing when confronted with unfavorable future scenarios; reference class forecasting, especially if the results suggest that their solution sits at the bottom end of the distribution of project performance, may further "depress" foresight teams, driving them deeper into affective biases that color decision-making. In this case, our best bet might be to focus on "developing a supportive psychological

climate" aimed at "reducing anxiety, maintaining a future-focus, and alleviating avoidant and dysfunctional behavior" (Healey and Hodgkinson 2007, 581).

We should not preclude the possibility that this might include strategies to dampen the aesthetic presentation of scenarios. Healey and Hodgkinson recommend use of low-fidelity scenarios in situations where individuals might become anchored in one scenario world: "Regularly analyzing multiple scenarios in a fast and simple manner, rather than elaborately and infrequently, is another potentially useful means of reducing scenario anchoring effects" (Ibid, 581). Low-fidelity scenarios might also be useful in situations where fear and anxiety in the face of uncertainty has made a team's thinking rigid and reactionary.

4.6 Findings

Before the revolution of behavioral economics, decision scientists were working under the assumption that human beings were perfectly rational agents that always acted in such a way as to maximize their utility. The work of Kahneman and Tversky demonstrated instead that humans were often led astray by biased thinking process that could lead to decisions that didn't best serve them. Their work also highlighted the fact that humans always had to make decisions under conditions of imperfect knowledge; in this important way, their rationality was bounded by imperfect information, timing considerations, and the

shortcomings of our evolved decision-making processes in the face of complex environments and barely perceived risks.

In similar fashion to the decision scientists who were working under the assumption that humans were perfectly rational agents, the current practice of strategic foresight relies on an overly simplified and essentialized understanding of the ways in which human beings think about the future. This set of assumptions has lead foresight strategists to make blunt assumptions about human future thinking.

One clear example is the unstated assumption among foresight strategists that presenting individuals with multiple possible scenarios of the future naturally loosens the hold of their default mental model of the future by distributing their intuitive probabilistic weightings across a broader range of possibilities. Granted, this is sometimes exactly what happens when people are presented with multiple scenarios of the future, but this outcome is predicated on a very specific set of preconditions: that the alternate futures do not trigger threat rigidity, and that they present causal stories which can compete at parity with that of the default future, for instance. Under strictly defined parameters, we can also expect humans to be perfectly rational actors. But the real world arguably almost never provides us with these perfect decision-making guardrails.

Another is the propensity for humans to overweight certain threats and underweigh others based solely on how *legible* they are to our mental decision making machinery. Perhaps as a result of the important role

forming and keeping track of alliances played in human evolution, we also tend to overweigh the risks associated with threats that are highly causal and agent-centered in nature, and underweigh those risks that do not emanate from agents and cannot be encapsulated into neat narratives. Highly causal, agent-driven future scenarios exploit our tendency to be cognitive misers: because these scenarios are easier to process than highly distributed, non-causal, and non-agent driven scenarios, our minds prefer to think through the former while tabling the latter or treating them with less attention and enthusiasm. The problem is that those non-causal, non-agent driven scenarios of the future – most imminently, the possibility of high disruptive climate changes – that pose humanity's greatest challenge today.

Even without these nuances, the foresight methodology is an important addition to humanity's mental toolkit. The question before us now is how to sharpen this tool. In the concluding chapter, I will examine strategic foresight as a mental tool, and suggest ways strategic foresight can be transformed into a living laboratory designed to investigate human future thinking and formulate new ways with which groups of humans can think about the future together and formulate ways to respond to it's challenges.

5. Conclusion: Strategic Foresight as a Cognitive Toolkit

As biological anthropologist Terrence Deacon points out in his book *Incomplete Nature: How Mind Emerged from Matter*, everything that humans purposefully build – from tools to structures to policy – points to some aspect of the future: a set of stairs anticipates that people will ascend them; a hammer portends the need to affix materials together; a zoning plan preconfigures the space of a future neighborhood development. Humans constantly heed the call of the future. Though immaterial and acausal, it beckons us into engagements with the present and reconsiderations of the past on its behalf.

But just because we design for the future does not mean that we do so with clear *foresight*. Of the many things that future historians will find curious about the first half of the 20th century, one of them may well be the enthusiastic and self-assured way in which powerful professionals sought to restructure whole societies and economies around ideological grand narratives that articulated future states of utopia. Efforts toward improvement within these political regimes were no doubt *future*-oriented, but what was the *quality* of the future they foresaw?

The answer is that it was almost certainly very poor indeed. The mountain of fine-grained detail in the way that humans think about the future – its close coupling with memory, its complex stratification, and its individual variability – is, as a foresight practitioner, disturbing to ponder.

It should force us into a reconsideration of the sophistication of future thinking we should reasonably expect from a creature whose capacities are so plainly truncated as a result of its evolutionary history.

5.1 Case Study: The Long Now

One locus in which we can clearly see debate crystallizing around the extent of our foresight capacities is the discourse surrounding our seeming inability, as a species, to muster real concern and action addressing the long-term viability of our civilization. There is a great deal of worry about the way in which the human mind seizes most readily on those futures that are most adjacent to us in both the spatial and temporal sense. When we're absorbed in and acting on only the absolute newest information available to us – about gas or stock or housing prices, or about what's on television tonight – it becomes nearly impossible to form clear thoughts about and act with consideration for the long term futures of our selves, our communities, and our planet. We no longer *give ourselves much time* to ruminate on the broader themes of our lives: of history, both personal and collective, of truth and responsibility, of where we're all going, of what it all means. The big questions have faded into abstraction and taken on the character of an immature pastime; in an always on culture, careful, considered, expansive thinking is seen as a distraction: something we used to do when we were young and didn't have so many pressing responsibilities. Considerations of our past and future, and what

they mean to us, are put aside; our present states, and the concerns of the short term, seem to dominate our mental landscapes.

And so we behave slavishly toward our daily calendars and appointment alerts in an attempt to be productive and "on it." But it may be those larger forces sneaking up on us, unfolding over years, decades, and generations – those deeper trends driving the state of our environments, the moral fabric of our societies, and the health of our institutions – that are the "it" we should really be "on". Maybe it's not that we have no time. Maybe it's just that we are so wired for the short term that we can't *see* through to more the more distal dimensions of the future.

Alexander Rose worries about this mismatch between native foresight abilities and the modern world a lot. For Rose, our myopia concerning time is due in part to the limitations of our biology. "You can imagine a lot of biological imperatives for staying focused on the short term, especially in a world that was much more dangerous than the one that we in the developed world live in now," says Rose (personal communication, July 18, 2013). But the long term can sneak up on us and insert itself into our day-to-day lives. "Even though the world we live in might not be as life or death as not having crops in the winter," he says, "it can certainly have an impact on our lives when we fail to understand Black Swan (high impact, low probability) events" (Ibid).

Rose's solution is to repurpose those evolved biological capacities for being engaged in short-term timeframes and build institutions, construct artifacts, and make art that stretch our sense of time through creating a visceral and emotional engagement with the distant future. In fact, he is helping to build just such an institution. Rose is the Executive Director of the Long Now Foundation, an organization that hopes to encourage long-term thinking in an age of accelerating culture. Their 10,000 Year Clock – now being built inside a mountain in Western Texas – is designed to operate with minimal human maintenance for millennia, inspiring those who ponder its epochal design to engage in imaginative long-term thinking. The Long Now hopes that the Clock, along with their other projects – like the Rosetta Disk, a repository of 1500 human languages microscopically etched and electroformed onto solid nickel – will lead the human mind to naturally wander across greater expanses of time; to emerge from always-on, short term thinking and adopt a more encompassing sense of what we mean by *Now*.

In some ways, the Long Now's projects can be seen as a museum of the future. Just as visiting a museum can inspire sensations of awe as we are led down deep temporal pathways into the lives and values of ancient people, the 10,000 Year Clock invites us to wander forward in time, speculate about the people who'll live there, and imagine what they might be like. Rose speaks passionately about this possibility:

> Designing and building a large clock, or creating a language archive etched onto a metal disc: those types of acts change the

> conversation about the future. You start wondering very real things about those future people: What aesthetics will they value? Will this be something that gets destroyed in one generation, or will it be something that's valued for tens – or hundreds – of generations? Are the hands of the people that wind this thing going to be the same as ours (Ibid)?

Imagining how the hands of our distant descendants might differ from our own does raise questions about the level of care and responsibility we should feel for them. "In general, the longest that people plan, on a personal level, tends to be for their grandchildren," Rose observes, continuing: "There are certainly mechanisms for changing that. Jonas Salk's sentiment about being a good ancestor is a good one, because it really resonates with people. People can then reflect on their own ancestors, and think about the choices they might've made. Then the question becomes: How will you internalize that in order to be a good ancestor going forward" (Ibid)?

Rose admits that instilling long-term thinking in people is not an easy task, due to our strong inborn biases toward imminent events and current emotional states over future events and states. Just as physical distance dulls the impact of tragedy, temporal distances cause us to discount the value of our future objectives against our immediate ones; my goal to be a good ancestor to future generations seems easily overtaken by my goal to save time in my day by driving to the grocery store two blocks away. And, of course, the Long Now has no proof that their efforts will pay off; barring radical life extension, its likely that no one involved in the Long Now's projects will be around to see whether

their efforts pan out. But it is not part of the Long Now's mission to personally witness the future; cryonicists they are not. Theirs is a project based on hope, an attempt to influence our culture into considering the distant future of humankind.

"If you stand in front of our clock, and are awe-struck about the possibilities of the next 10,000 years, then you might go home and do something that stretches your time frame just a little bit," says Rose.

The projects undertaken by The Long Now Foundation suggest some understanding of the inherent biases of our future thinking. There are, perhaps, good reasons for this: after all, Daniel Kahneman, the father of research into the biases of human decision-making, has been a repeated guest at The Long Now's salons. But, for media theorist Douglas Rushkoff, the struggle The Long Now Foundation is engaged in against our natural orientation toward the present and immediate future in order to save the future might not be without it's own consequences. When the choice of where to toss a burger wrapper is bound up in the mind with considerations of epochal time scales, anxiety and decision paralysis can quickly set in. This decision paralysis is a phenomenon he calls over-winding: making the present responsible for *too much other time*. "Unless we're living in utter harmony with nature," he writes in his newest book Present Shock, "thinking in ten-thousand-year spans is an invitation to nightmarish obsession. It's a potentially burdensome, even paralyzing,

state of mind. Each present action becomes a black hole of possibilities and unintended consequences" (Rushkoff 2013, 135).[24]

For Rushkoff, the Long Now solution is another permutation of what he calls a culture in "present shock", in which we make the present moment the center of our lives and imbue it with inordinate significance and responsibility. In a strange inversion, for him the 10,000-year mindset becomes "less of a Long Now than a Short Forever" (Ibid). Not an escape from the thinking that sends us into present shock, but rather acceleration into that black hole of possibilities.

And, of course, there's the problem of overconfidence: on what grounds do humans, with our pithy foresight capacities, deign to speak for the future? I spoke to Rushkoff as he was travelling by train home from a speaking engagement, a much more comfortable and less temporally overwound method of travel than flying, he opined. "The Long Now People mean very well," he admits, continuing: "What they want to do is help bring to light the long-term implications of stuff that we're doing right now. But we don't have movements of this same sort now. We don't have the narrative expectations of doing things today for

[24] We can clearly see the potential impact of overloading the present with the burden of the future in they way that decision makers sometimes become overwhelmed when presented with foresight scenarios detailing disruptive futures for their organizations. Because exposure to uncertainty and risk can oftentimes stimulate emotions of worry, fear, dread, and anxiety, provocative scenarios can lead decision makers to process information in a narrow and labored manner. As we've seen, visceral emotions can override rational decision-making, leading to impulsive decision-making, states of denial, and a breakdown of strategic conversations (Healy and Hodgkinson 2007, 579). The lesson for foresight practitioners: making scenarios *more* evocative is not always the best course of action.

reward tomorrow. We've now got a more incremental approach to things" (personal communication, July 17, 2013).

The position that Rushkoff advocates in his book is a return to scales of foresight that are consonant with our native abilities. Attempting foresight out of the range of human abilities is not only hubristic, but also potentially detrimental to our mental health as a species. "Your future is less dependent on your 401K plan than it is on how you're connected meaningfully with the community in which you live," he suggests. "The more connected you are, the more *in the present* you are, and the less you're worried about the future" (Ibid). Rushkoff thinks we can safely dispense with long-term future thinking; it's a task for which we're simply not suited. "If we're going to engage in appropriate behaviors," he says, "in the end, the only thing that's going to make it work is if we appreciate in the moment that those methods are superior. Otherwise, we're going to be working against our own natures" (Ibid).

For me, this debate circles around a set of questions that is never really addressed straight on: Given the nature of human foresight, what level of genuine engagement with long term (and multiple) futures can we really expect from individuals and organizations? How can we ensure that we are not unduly reinforcing biased future thinking in both our engagement with short term and long-term horizons? What principles might we formulate, what methods might we devise, and what checks and balances might we create that harness and expand our native ability to

think about the future without stretching that ability beyond its breaking point? I don't think we have good answers to those questions yet, and especially in the context of broad, sweeping claims about either our ability to scaffold a 10,000 sense of *Now* onto human future thinking, or the impossibility of approaching such a project without causing extreme decision paralysis and anxiety, I don't think we'll be able to discover them.

5.2 What's Next?

A. Transform foresight into a living laboratory – The good news is that strategic foresight exercises are a perfect laboratory for exploring these questions in a systematic way, as long as we can set up the experimental parameters and controls properly. The methodology is, at least in some respects, repeatable and measurable. We could, for instance, take a classic foresight engagement – taking participants from signals, to trend analysis, to discovering and prioritizing drivers, to developing scenarios, and finally to devising and windtunnelling strategies – as a 'kernel' for the research, then insert other, less frequently use methods – Delphi polling, experiential futures, futures wheel, etc. – in order to asses the value they add to the engagement.

What kind of observations might we make in the foresight laboratory? The most direct means would be to record video and audio of the group discussions that take place in each phase of the foresight engagement. The data that we could expect from recordings is very rich;

we might discover group polarizing effects amplified by shared worldviews, or a propensity to default to tropes of current blockbuster tropes that feed into availability bias, or the discounting of possible futures based primarily on overly thin causal stories that fail to reach parity with focal scenarios. We might also "seed" signals or trends that are intentionally engineered to engage the subjects' biased future thinking in order to observe how they are carried and developed through the engagement.

B. Reframe "good future thinking" – As I wrote at the outset of this paper, in light of the science of human future thinking the aim of strategic foresight should be to teach participants the skill to *develop clearly authored, reasonably de-biased visions of the future and formulate explicitly rational human-regarding strategies for flourishing within whatever future we happen to get*. Under this lens, strategic foresight is another piece of "mindware" – like probabilistic reasoning and bias-mitigating strategies in decision-making – that augments our normal capabilities and better equips us to deal with the challenges of the modern world. And like these other pieces of mindware, keen foresight might only be possible after protracted exposure to a certain way of thinking; just as probabilistic reasoning is somewhat unnatural to us because we naturally think in terms of possibilities rather than probabilities, robust foresight might have to overcome deeply engrained cognitive biases toward short term thinking, causal explanations, and poor affective forecasting, among other

pitfalls. Robust scenarios and strategies should therefore be regarded at best, as a reflection of whether we've done the work to instill good future thinking in strategic foresight participants. Rather than judging the images of the future created during a foresight engagement on their accuracy, we should judge them on the level of skilled future thinking they indicate; even an implausible image of the future can be judged fairly if the intention in building it was to effect some favorable change in a group's broader future thinking: so-called wildcard scenarios with highly causal explanation, for instance, can force us to re-examine our sense of what *is* and *isn't* possible in the future and open up a space of possible futures around a focal scenario.

C. Build a Foresight Inventory – One possibility for adding a quantitative measure to foresight engagements would be to use a questionnaire similar to the Zimbardo Time Perception Inventory order or Keith Stanovich's in-development Rationality Quotient Test in to establish a pre-engagement baseline time perception for all participants in a foresight exercise. Once the engagement is completed, participants' time perception can be re-polled with the inventory at several intervals – a week later, three months later, a year later – in order to discern the effect and duration (if any) of the foresight engagement on individual time perception. Over enough trials, we might be able to discover something about the way that different configurations of strategic foresight engagements can create greater or lesser effects in altering individual

reports on our inventory. Of course, transforming strategic foresight engagements into living laboratories requires that we develop an inventory that is specific to measuring the preferred outcomes of strategic foresight engagements. A Strategic Foresight Inventory may include questions that investigate an individual's tendency to produce counterfactuals, engage in robust episodic future thinking, examine their biases about the future, create strategies that are robust enough to respond to multiple futures, and so on. My recommendation is for foresight strategist to working together with experimental psychologists that specialize in developing psychological inventories to build a custom inventory to measure foresight acumen, specifically.

D. Create a theoretical grounding for foresight - It's also equally critical that we develop a solid evidence-based theoretical ground to support the practice. My objective in this paper has been to move the conversation around the evidentiary tools and theoretical underpinnings of foresight forward. We have a long way to go before strategic foresight can meet these aspirations, and given the success of foresight, it might be difficult to garner support for fundamental reform. Mark Healey expressed strong views about this in our interview:

> Scenarios are helpful because they're intuitive and easy to work with. But other people would say the reason that they work is that they make people *think harder*, and simply by *thinking harder,* using more information, considering more perspectives, that loosens up some of these fundamental biases that have been written about in the behavioral decision making literature. But once you start to scratch that veneer away and realize there's little evidence to back up these claims, it turns out that, in a sense, foresight is based on

> some fairly limited pieces of evidence (personal communication, November 20, 2013).

The deeper issue is that there hasn't been a concerted effort to build a solid theoretical ground for the practice of foresight. If foresight is to gain legitimacy as an evidence-backed methodology then it needs to test its assumptions, abandon "folk psychological" accounts of future thinking, and transform itself into a discipline built on a sound scientific understanding of the way humans imagine and reason about the future. Peter Bishop, a Professor Emeritus of Foresight at the University of Houston Futures program, points out that the lack of such a solid theoretical foundations is not uncommon in practitioner-led disciplines like foresight. "In any discipline," he notes, "there are at least two levels of conversation going on: the delivery of what the discipline offers, and the "back room" stuff where you might say the ideas and methodologies and research is taking place" (Bishop 2014). For Bishop, foresight lacks this "back room" because most of this type of research happens in universities at the doctoral level:

> There is no organized group of researchers with a cadre of graduate students who spend all of their time doing research on the kinds of questions you're asking. We are a group of practitioners. And I do believe in most disciplines, practice preceded theory. Most people think research happens first and then people apply it. That does happen. But usually, it's the opposite. Architecture is practice driven. The theoretical foundation of any new discipline will be very thin until it becomes successful and then it can afford to have that background available. Without high quality doctoral programs, the growth of our field is inhibited. It would give us much more

> credibility if we had those, but getting those established is a tremendous effort (Bishop 2014).[25]

I am inclined to agree that having doctoral programs dedicated to foresight research would do a great deal to advance the state of the practice. That being said, I also believe that there are research groups doing work that has immediate implications for our practice; in this paper, I have tried to showcase some of this research, which is being conducted in cognitive science, management studies, and psychology doctoral programs throughout the world. There is much that we foresight practitioners could learn from seeking out researchers in these fields who are studying human decision making and foresight in order to test our methods and assumptions against their results. And if the field lacks the resources to establish stand-alone doctoral programs, we can always instead recruit these research professionals into our communities of practice. The roster of The Greatest Good, a Chicago-based consultancy, can serve as a template: it has representatives from the worlds of management consulting as well as psychology and economics, including Daniel Kahneman himself. Especially as institutions of higher education forge closer ties with private enterprise through start-up incubators and other programs that encourage entrepreneurship among young scholars, we should expect to see more organizations that develop both original

[25] During our conversation, Bishop did reference the doctoral program at the University of Manoa as an exception to this observation, but also points out that it is technically a degree in Political Science and so is unlikely to address the kinds of questions posed in this paper.

research and innovative practitioner models in-house. In my opinion, the foresight practice is well positioned to take advantage of these developments.

5.3 Closing Remarks

The influence of the strategic foresight methodology continues to grow. A survey of UK-based organizations conducted in 2006 found that over a third used scenario planning as part of their toolkit for devising strategies (Hodgkinson et al. 2006). A 2009 study by the European Foresight Monitoring Network found that the most widely used foresight methods among 1000 examined exercises were, in order of frequency: literature review (54%), expert panels (50%), scenarios (42%), trend analysis (25%), and futures workshops (24%). The sharp decline between third and fourth place is interesting, but as a whole the top five methodologies would've been at home in a foresight activity conducted in the early seventies. Why have the core methods of foresight remained virtually unchanged for decades?

One likely explanation is that the most commonly used methods can create the impression of a fruitful engagement with possible futures without any need to demonstrate the activity's worth *in the present*. Despite the effort of foresight practitioners to distinguish themselves from futurists, the products of both professions enjoy the benefits of *temporal discounting*, another tendency of the human mind unearthed by

experiments in behavioral economics. In short, we tend to discount the value of those things that are remote from the present; one marshmallow today exerts far more pull on us than two marshmallows a week from now. Savvy futurists place their predictions far enough in the future that their misfires won't (rightfully) damage their forecasting careers. In like fashion, since the horizon of foresight activities is almost always beyond the purview of an organization's operational concerns, the question of whether scenarios are robust enough is necessarily deferred. By the time the investigated horizon *arrives*, organizations have already moved on, and aren't likely to turn their sights back on a five-to-ten year old foresight project to assess its utility.

The temporal remoteness of foresight horizons is, of course, also what makes it difficult to make a case for foresight's actual, as opposed to perceived, utility. But despite the lack of provability, organizations are still willing to engage in foresight activities to assuage their skittishness about the future. The way that the value of foresight is communicated highlights the schizophrenic nature of the foresight practice as it stands today: at once, foresight practitioners insist that foresight isn't about *prediction* per se, but then tout examples of successful foresight like the Shell scenarios, which that organization used to correctly forecast the possibility of an oil shock caused by an embargo originating in the Arab world.

My suspicion is that much of this confusion is due to the fact that foresight practitioners have been overly fixated on the quality of the

outputs of foresight – the scenarios, the strategic recommendations, etc. – and less focused on the effect that the exercises tangibly have on the ability of participants to change their mental models based on clear thinking about the future. One reason for this might be that changes to mental models and toolkits have been difficult to measure, historically. But a maturing science of human decision-making is beginning to change this.

In closing, in this paper I've tried to point the way toward new framing objectives, evidentiary tools, and theoretical foundations for strategic foresight. I've argued that foresight professionals should aim to make organizations aware of the operation of biases in the course of a strategic foresight engagement, or at the very least to lead them through exercises constructed to mitigate the pull that biases and heuristic thinking have on the way we think about the future. The only way to approach this project with any legitimacy is to engage with the sciences that are investigating the ways that features of our minds systematically lead humans astray, and to base our proposed solutions and engagements on evidence-backed methodologies for either temporarily suspending the cognitive barriers to good foresight or, in the preferred scenario, granting clients a deeper sense of control over the way they think about the future.

I believe this means, at a minimum, introducing new checks and balances into the foresight methodology. It seems to me that foresight has a bit of an aversion to overtly quantitative or scientific methods. One

foresight practitioner offered some early feedback on the ideas in this paper, which took the form of reminding me that Peter Schwartz's book is called "The *Art* of the Long View", and Kees van der Heijden's is called "The *Art* of Strategic Conversation".

Of course, I have some sympathy with this view. Historically, it's been important to distinguish foresight from the work of forecasters and futurists, whose liberal use of quantitative trend extrapolation, especially in the social and technological realms, has overemphasized a deterministic view of the future at the expense of credibility and nuance. But if foresight is to lay claim legitimately to being a methodology that *in fact* improves our future thinking, it has to establish the veracity of this claim through painstaking measurement, not mere salesmanship.

And what impact does human foresight have on changing the evolution of culture, in any case? We simply do not have good answers to these more fundamental issues. Anthropologist Alex Mesoudi applies a skeptical lens to the claim that cultural evolution is directed by human foresight in any meaningful way, arguing instead that cultural evolution proceeds in ways that have more in common with biological evolution. The ability for humans to mentally travel forward in time – to employ *foresight* – does not invalidate Mesoudi's claim, since foresight itself "may evolve through a past process of blind evolution" (Mesoudi 2007). Indeed, I have argued that this is the case when we allow the biases of our

foresight to shape our view of the future in ways that are antithetical to constructing and meeting our goals.

It doesn't help that real data on the efficacy of strategic foresight is difficult to come across. In a recent study of 77 multinational firms, top-performing organizations reported that strategic foresight activities delivered value through an enhanced capacity to perceive, interpret, and respond to change, to influence other actors through shared images, and to enhance capacity for organizational learning (Rohrbeck & Schwarz, 4); an interesting finding, though it amounts to self-reporting. As Mark Healey and Gerard Hodgkinson point out in their critique of strategic foresight: "…hard empirical evidence to substantiate these fundamental claims concerning the cognitive benefits of scenario-based techniques is both highly equivocal and limited in scale and scope, comprising in the main descriptive case accounts of apparently successful applications of the techniques in action" (Healy and Hodgkinson 2007, 568).

However, I do not think that all is lost. If strategic foresight is really about improving our *mental models of the future*, as is now commonly understood, then its utility and transformative potential should be immediately demonstrable by re-probing individuals' cognitive toolkits and mental models in the wake of a strategic foresight engagement. "One thing you can always do," says Thomas Suddendorf pointed out during our interview, "is hold a mirror in front of people and make them aware of their potential biases so that they become more alert to them"

(Suddendorf 2014). Developing that "mirror" in the form of protocols and methods for measuring individual and group foresight capacity would help us answer a whole host of questions: Has the exercise decreased the hold of foresight-related biases on the minds of those who went through the exercise, or who have been exposed to the resulting material? Has it made them aware of the origin of these biases in their own minds in a systematic way, so that they can work to avoid their effects? Has it had any impact on the time perception valences of participants to make them more future-oriented? Has it increased their facility with employing semantic knowledge to the future in robust ways? Has it made their images of the future more representative, robust, extended, and contextualized?

Devising ways to measure and report on these changes will serve to legitimize (or perhaps, condemn) strategic foresight. But it's potential exceeds even this. If indeed cultural evolution has largely been guided by blind evolution, as Mesoudi claims, then transforming foresight into a measurable science may be the best chance humans have of wresting control of our future both from blind evolutionary forces and our own inherited biases, and reshaping it to promote human survival and thriving. It is encouraging for the future of robust discourse between the theory and practice of foresight to hear similar sentiments echoed by both the scientists and foresight practitioners who agreed to be interviewed for this project.

BIBLIOGRAPHY

Addis, Donna, and Schacter, Daniel. (2012) The hippocampus and imagining the future: Where do we stand? Frontiers in Human Neuroscience. 5: 173.

Addis, Donna, Wong, Alaina, & Schacter, Daniel (2007). Age related changes in episodic simulation of future events. Psychological Science. 19(1): 33-41.

Ashby, Madeline. (2013). Personal communication. October 16.

Atance, Cristina M. & O'Nell, Daniela K. (2001) Episodic future thinking. TRENDS in Cognitive Sciences. 5(12): 533-539.

———. (2005) The emergence of episodic future thinking in humans. Learning and Motivation. 36: 126-144.

Ariely, Dan (2008) Predictably Irrational: The Hidden Forces that Shape Our Decisions. New York: Harper Collins.

———. (2010) The Upside of Irrationality: The Unexpected Benefits of Defying Logic at Work and at Home. New York: Harper Collins.

Arp, Robert. (2008) Scenario Visualization: An Evolutionary Account of Creative Problem Solving. The MIT Press: Cambridge.

Baines, James. (2013) "The Man with the Thirty Second Memory". Vice Online. June 6.

Baird, B., Smallwood, J., & Schooler, J.W. Back to the future: autobiographical planning and the functionality of mind wandering. Consciousness and Cognition 20(4): 1604-1611.

Bar, M. (2010) Wait for the second marshmallow? Future-oriented thinking and delayed reward discounting in the brain. Neuron 66(1): 4-5

Barrett, D. (2010). Supernormal Stimuli: How Primal Urges Overran Their Evolutionary Purpose. W.W. Norton & Company: New York.

Basso, Anna, Bisiach, Edoardo, and Luzzati, Claudio. (1980) Loss of Mental Imagery: A Case Study. Neuropsychologia 18: 435-442.

Baumeister, R. F., Heatherton, T.F., & Tice, D.M. (1994) Losing control: How and why people fail at self regulation. San Diego: Academic Press.

Bayes, Thomas. (1763). An Essay towards solving a Problem in the Doctrine of Chances. Philosophical Transactions of the Royal Society of London. 53: 370-418.

Berdik, Chris. (2011). Mind Over Mind: The Surprising Power of Expectations. New York: Penguin Group.

Bingham Geoffrey P., Schmidt R.C., & Rosenblum Lawrence D. (1989) Hefting for maximum distance throw: A smart perceptual mechanism. Journal of Experimental Psychology: Human Perception and Performance. 15(3): 507-528.

Bishop, Peter. (2014) Google Hangouts Interview. February 20.

Bradfied, Ronald. (2008). Cognitive Barriers in the Scenario Development Process. Advances in Developing Human Resources. 10(2): 198-215.

Buckner, L.R. & Carroll, D.C. (2007) Self-projection and the brain. TRENDS in Cognitive Sciences 11(2): 49-57.

Busenitz, Lowell W. & Barney, Jay B. (1997) Differences Between Entrepreneurs and Managers in Large Organizations: Biases and Heuristics in Strategic Decision-Making. Journal of Business Venturing, 12: 9-30.

Calvin, W.H. (1982) Did throwing stones shape hominid brain evolution? Ethology and Sociobiology 3(3): 115-124.

Candy, Stuart. (2008). Tangible futures. the skeptical futurist. Retrieved from: http://futuryst.blogspot.ca/2008/11/tangible-futures.html

———. (2014) Interview. March 5.

Carruthers, Peter. (2000) Phenomenal Consciousness: A Naturalistic Theory, Cambridge University Press: Cambridge. 2000.

Carstensen, L.L., Isaacowtiz, D.M. & Charles S.T. (1999). Taking time seriously: A theory of socioemotional selectivity. American Psychologist, 54, 165-181.

Changizi, M. (2011) Humans, Version 3.0. Seed Magazine Online. February 23.

———. (2011) Harnessed: How Language and Music Mimicked Nature and Transformed Ape to Man. Dallas: BenBella Books, Inc.

Clark, Andy. (2013) Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science. Behavioural and Brain Science. 36(3): 181-204.

Cosmides, L. & Tooby, J. (1992) Cognitive Adaptations for Social Exchange. In The Adapted Mind: Evolutionary Psychology and the Generation of Culture. (eds. J. Barkow, L. Cosmides & J. Tooby). New York: Oxford University Press.

Costandi, Mo. (2009) Mental Time Travel. Neurophilosophy. June 15.

D'Argembeau, A. et al (2008) Neural correlates of envisioning emotional events in the near and far future. NeuroImage 40: 398-407.

Danziger, Shai, Levav, Jonathan, & Avnaim-Pesso, Liora. (2011) Extraneous factors in judicial decisions. Proceedings of the National Academy of Sciences. 108(17): 6889-6992.

Dator, Jim. (2009) Alternative Futures at the Manoa School. Journal of Future Studies. November. 14(2): 1-18.

———. (2014) Email interview. February 22.

Dawkins, Richard. (2000) An open letter to Prince Charles. Available at www.edge.org/3rd_culture/prince/prince_index.html.

Deacon, Terry. (2011) Incomplete Nature: How Mind Emerged from Matter. New York: W.W. Norton & Company.

Dennett, Daniel. (1996) Kinds of Minds: Toward an Understanding of Consciousness. Basic Books.

———. (1984) Elbow Room: The Varieties of Free Will Worth Wanting. Cambridge: The MIT Press.

Ebert, Jane E.J. (2001) The role of cognitive resources in the valuation of near and far future events. Acta Psychologica 108: 155-171.

Ecken, Phillip, Gnatzy, Tobias, & von der Gracht, Heiko A. (2011) Desirability bias in foresight: Consequences for decision quality based on Delphi results. Technological Forecasting and Social Change. 78(9): 1654-1670.

Elqayam, Shira and Evans, Jonathan. (2011) Subtracting "ought" from "is": Descriptivism versus normativism in the study of human thinking. Behavioral and Brain Sciences. 34: 233-290

———. (2013). Rationality in the new paradigm: Strict versus soft Bayesian approaches. Thinking & Reasoning. 19(3): 453-470.

Fellows, Lesley and Farah, Martha. (2004) Dissociable elements of human foresight: a role for the ventromedial frontal lobes in framing the future, but not in discounting future rewards. Neuropsychologia 43: 1214-1221.

Ferrante, D., Girotto, V., & Straga, M. (2013) Improving the past and the future: a temporal asymmetry in hypothetical thinking. Journal of Experimental Psychology: General. 142.1. Feb: 23.

Fischhoff, B. (1976) The effect of temporal setting on likelihood estimates: Organizational Behaviour and Human Performance. 15: 180 – 194.

Florovsky, G. (1969). The study of the past. In R.H. Nash (Ed.), Ideas of history (Vol. 2). New York: Dutton.

Foote, A.L. & Crystal, J.D. (2007) Metacognition in the rat. Current Biology. 17(6): 551-555.

Foskett, N. and Hemsley-Brown, J. (2001) Choosing Futures: Young people's decision-making in education, training and careers markets. RoutledgeFalmer: London.

Gardiner, J.M. (2001) Episodic memory and autonoetic consciousness: a first-person approach. Philosophical Transactions of The Royal Society B Biological Sciences 356(1413): 1351-1361.

Gilbert, D.T. & Wilson, T.D. (2007) Prospection: experiencing the future. Science 317: 1351-1354.

Ghamari-Tabrizi, Sharon. (2005). The Worlds of Herman Kahn: The Intuitive Science of Thermonuclear War. Harvard University Press: Cambridge.

Gigerenzer, Gerd and Goldstein, Daniel G. (1996) "Reasoning the fast and frugal way: models of bounded rationality." *Psychological Review* 103(4): 650-669.

Gigerenzer, G. & Brighton, H. (2009) "Homo Heuristics: Why Biased Minds Make Better Inferences. Topics in Cognitive Science 1: 107-143.

Ginsburg, S. & Jablonka, E. (2007) The Transition to Experiencing: I. Limited Learning and Limited Experiencing. Biological Theory 2(3): 218-230.

Haldenby, Trevor. (2013). Bringing the Future to Life: Pervasive Transmedia Scenarios and the World of Worlding. Collection of OCADU: Toronto.

Havas, Attila, Schartinger, Doris, & Weber, Weber, Matthias. (2010) The impact of foresight on innovation policy-making: recent experiences and future perspectives. Research Evaluation. 19(2): 91-104.

Healey, M.P. & Hodgkinson, G.P. (2008) Troubling Futures: Scenarios and Scenario Planning for Organizational Decision Making. In: G. P. Hodgkinson, and W. H. Starbuck, editor(s). The Oxford Handbook of Organizational Decision Making. Oxford, UK: Oxford University Press. 565-585.

Healey, M.P. (2013). Telephone interview. November 20,

Hejazi, Alireza. (2012) Futures Metacognition: A Progressive Understanding of Futures Thinking. World Future Review 4(2): 18-27.

Hodgkinson, G.P., Whittington, R., Johnson, G. & Schwartz, M. (2006). "The Role of Strategy Workshops in Strategy Development Processes: Formality, Communication, Coordination and Inclusion." Long Range Planning, 39: 479-496.

Hodgkinson, G.P., and Wright, G. (2002). 'Confronting Strategic Inertia in a Top Management Team: Learning From Failure.' Organization Studies. 23: 921-945.

Horton, Averil. (2012). "Complexity sciences approaches to disruptive events". Foresight.14(4): 294-303.

Hsee, C.K. and Rottenstreich, Y. (2004). 'Music, Pandas, and Muggers: On the Affective Psychology of Value.' *Journal of Experimental Psychology-General*. 133(1): 23-30.

Janis, I.L. (1972). Victims of Groupthink. Houghton-Miffin: Boston.

Koriat, Asher. (2007) "Metacognition and Consciousness" In P.D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), The Cambridge Handbook of Consciousness. New York: Cambridge University Press.

Kahn, Herman. (1962) Thinking about the unthinkable. New York: Discus Books/Avon.

Kahneman, Daniel. (2013) "Thinking, Fast and Slow". Long Now Foundation. longnow.org. Web. 2 October 2013.

———. (2011) Thinking, Fast and Slow. New York: Farrar, Straus and Giroux.

Kaheman, Daniel & Tversky, Amos. (1973). On the psychology of prediction. Psychological Review, 80, 237-251.

Kosslyn, Stephen M. (1988) "Aspects of a cognitive science of mental imagery." Science 240(4859): 1621-1626.

Kuhn, K.M. & Sniezek, J.A. (1996). "Confidence and Uncertainty in Judgmental Forecasting: Differential Effects of Scenario Presentation" Journal of Behavioral Decision Making. 9: 231-247.

Lakoff, George & Johnson, Mark. (1980) Metaphors We Live By. London: University of Chicago Press.

Levine, Robert. (1997). A Geography of Time: The Temporal Misadventures of a Social Psychologist, or How Every Culture Keeps Time Just a Little Bit Differently. OneWorld Publications: England.

Liberman, Nira, Sagristano, Michael D., & Trope, Yaacov. (2002). The effect of temporal distance on level of mental construal. Journal of Experimental Social Psychology. 38: 523-534.

Lombardo, Thomas. (2008) The Evolution of Future Consciousness. Bloomington: AuthorHouse.

Lovallo, Dan and Kahneman, Daniel. (2003). Delusions of Success: How Optimism Undermines Executives' Decisions. Harvard Business OnPoint: 4279.

Lowenstein, George and Lerner, Jennifer S. (2003) The Role of Affect in Decision Making. In R.J. Davidson, K.R. Scherer & H.H. Goldsmith. Handbook of Affective Sciences. Oxford University Press.

Mackay, R. Bradley and McKiernan, Peter. (2004) "The role of hindsight in foresight: refining strategic reasoning." Futures: 36(2): 161-179

Marcus, Gary. (2008) Kluge: The Haphazard Evolution of the Mind. New York: Houghton Mifflin

Marshall, M. (2011) Your brain chemistry existed before animals did. New Scientist. 10 September 2011. Issue 2828.

Mesoudi, Alex. (2007) "Has mental time travel really affected human culture?" Behavioural and Brain Sciences 30: 326-327

———. (2008) Foresight in cultural evolution. Biology and Philosophy. 23(2): 243-255.

Mesoudi, A., Whiten, A., & Laland, K.N. (2006) Towards a unified science of cultural evolution. Behavioural and Brain Sciences 29: 329-383.

Miller, George A. (2003). The cognitive revolution: A historical perspective. TRENDS in Cognitive Sciences. 7(3): 141-144.

Millett, Stephen M. (2009) Should probabilities be used with scenarios? Journal of Futures Studies. 13(4): 61-68.

Miyashita, Y. (2004) Cognitive memory: Cellular and networking machineries and their top-down control. *Science*. 306:435-440.

Myers, D.G. et al. (1977). Attitude Comparison: Is There Ever a Bandwagon Effect? Journal of Applied Social Psychology. 341: 341-347.

Novacek, M. (2002). Time traveler: In search of dinosaurs and ancient mammals from Montana to Mongolia. New York: Farrar, Straus and Giroux.

Orwell, G. (1945). London letter. Partisan Review. 12: 467–472

Peters, J. and Büchel, C. (2010) Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal mediotemporal interactions. Neuron 66(1): 138-148.

Plummer, T. (2004) Flaked stones and old bones: Biological and cultural evolution at the dawn of the dawn of technology. Yearbook of Physical Anthropology 47: 118-164.

Popper, Rafael. (2009). Mapping Foresight: Revealing how Europe and other world regions navigate into the future. Luxembourg: Publication Office of the European Union.

Porter, M.E. Competitive Advantage – Creating and Sustaining Superior Performance. New York: Free Press.

Polak, Fred. (1973) The Image of the Future. Amsterdam: Elsevier Scientific Publishing Company.

Raby, C.R. & Clayton, N.S. (2009) Prospective cognition in animals. Behavioral Processes 80: 314-324.

Race, E., Keane M.M., Verfaellie M. (2013) Losing sight of the future: Impaired semantic prospection following medial lobe lesions. Hippocampus. 23(4): 268-277.

Ranganath, Charan et al. (2011) Can Cognitive Training Improve Episodic Memory? Neuron 72(5): 688-691.

Reichenbach, H. (1951) The Rise of Scientific Philosophy. University of California Press: Berkeley.

Redish, A.D. & Jackson, J. (2007) Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. Hippocampus. 17(12): 1209-29.

Richard, P.B. (1983) Mechanisms and adaptation in the constructive behaviour of the beaver (C. fiber L.)" Acta Zoologica Fennica. 174: 105-108.

Roberts, William and Feeney, Miranda. (2009) The comparative study of mental time travel. Trends in Cogntive Science. 13(6): 271-277.

Rohrbeck, R. and J.O. Schwarz. (2013). The Value Contribution of Strategic Foresight: Insights From an Empirical Study of Large European Companies. Technological Forecasting and Social Change. 80(8): 1593-1606.

Rouse, Neal J. & Olson, James M. (1996) Counterfactuals, causal attributions, and the hindsight bias: A conceptual integration. Journal of Experimental Social Psychology.

Rushkoff, Douglas. (2013) Present Shock: When Everything Happens Now. New York: Penguin Group.

———. (2013). Telephone interview. July 17.

Rose, Alexander. (2013). Telephone interview. July 18,

Sanna, Lawrence J. & Schwarz, Norbert. (2003) Debiasing the hindsight bias: The role of accessibility experiences and (mis)attributions. Journal of Experimental Social Psychology. 3: 287-295.

Schacter, Daniel et al (2012) The Future of Memory: Remembering, Imagining, and the Brain. Neuron 76: 677-694

Schacter, Daniel et al. (2007) Remembering the past to imagine the future: the prospective brain. Nature Reviews: Neuroscience. 8: 657-661

Schoemaker, P.J.H. (1993). Multiple Scenario Development – Its Conceptual and Behavioural Foundation. Strategic Management Journal. 10: 549-564.

Schroeder, Karl. (2011) Safety Glass: Presenting Foresight Findings as Fiction. Collection of OCADU: Toronto.

———. (2013). Personal communication. November 17.

Schnaars, S.P. and Topol, M.T. (1987). "The Use of Multiple Scenarios in Sales Forecasting – an Empirical-Test." International Journal of Forecasting, 3: 405-419.

Shafir, E., ed. (2003). Preference, belief, and similarity: Selected writings of Amos Tversky. MIT Press.

Sharot, Tali. (2011). The Optimism Bias: A Tour of the Irrationally Positive Brain. Pantheon: New York.

Silver, N. (2012) The Weatherman is Not a Moron. The New York Times Magazine. Sept 7.

———. (2012) The Signal and the Noise. The Penguin Press: New York.

Slaughter, Richard. (2002) Futures Studies as an Intellectual and Applied Discipline. IN Advancing Futures: Futures Studies in Higher Education. Westport: Praeger. 91-108.

Song, X., & Wang, X. (2011) Mind wandering in Chinese daily lives – an experience sampling study. PLoS One 7(9).

Suddendorf, T. and Busby, J. (2005) Making decisions with the future in mind: Developmental and comparative identification of mental time travel. Learning and Motivation. 36: 110-125.

Suddendorf, T. Rose, D, and Corballis, M.C. (2009) Mental time travel and the shaping of the human mind. Philosophical Transactions of The Royal Society B Biological Sciences 364: 1317-1324.

Sunstein, Cass R. (2009) Going to Extremes: How Like Minds Unite and Divide. Oxford University Press: New York.

———. (2013) Simpler: The Future of Government. Simon & Schuster: New York.

Stanovich, K. (2014) Email interview. February 19.

———. (2004) The Robot's Rebellion: Finding Meaning in an Age of Darwin. Chicago: University of Chicago Press.

———. (2009) What Intelligence Tests Miss. New Haven: Yale University Press.

———. (2010) Rationality and the Reflective Mind. New York: Oxford University Press.

Stanovich, K, Toplak, M.E., & West, R.F. (2010) Contaminated Mindware: Thinking biases of the cognitive miser. Rotman Magazine, Winter, 16-21.

Staw, B.M., Sandelands, L.E, and Dutton, J.E. (1981). Threat Rigidity Effects in Organizational Behavior: A Multilevel Analysis. Administrative Science Quarterly. 26(4): 501-524.

Stanovich, K. & West, R.F. (2003) Evolutionary versus instrumental goals: How evolutionary psychology misconceives human rationality. In D.E. Over (Ed.), Evolution and the Psychology of Thinking: The Debate. Psychological Press.

Sterelny, Kim. (2003) Thought in a Hostile World: The Evolution of Human Cognition. Malden: Blackwell Publishing.

Suddendorf, Thomas. (2014). Skype Interview. January 29.

———. (2009) "Episodic memory versus episodic foresight: Similarities and differences" Wiley Interdisciplinary Reviews: Cognitive Science 1(1): 99-107.

Suddendorf, Thomas and Corballis, Michael C. (2007) "The evolution of foresight: What is mental time travel, and is it unique to humans?" Behavioural and Brain Sciences 30: 299-351

Szpunar, Karl and Schacter, Daniel. (2012) Get real: Effects of repeated simulation and emotion on the perceived plausibility of future experiences. Journal of Experimental Psychology: General. 142(2): 323-327.

Szpunar, Karl and McDermott, Kathleen. (2008) Episodic future thought and its relation to remembering: Evidence from ratings of subjective experience. Consciousness and Cognition 17: 330-334.

Tetlock, Philip. (2005) Expert Political Judgment: How Good is It? How Can We Know? Princeton University Press: New Jersey.

Thompson, V.A et al. (2013) The role of answer fluency and perceptual fluency as metacognitive cues for initiating analytic thinking. Cognition 128. 237-251.

Toplak, Maggie. (2014) Email interview. February 19.

———. (2014) Telephone interview. March 6.

Tulving, E. (1985) Memory and consciousness. Canadian Psychologist. 26: 1-12.

Van Alstyne, Greg. (2010). How We Learned to Pluralize The Future: Foresight Scenarios as Design Thinking. In M. Shamiyeh (Ed.), Creating Desired Futures: How Design Thinking Innovates Business (pp. 69–91). Basel: Birkhäuser.

Van der Heijden, K., Bradfield, R., Burt, G., Cairns, G. and Wright, G. (2002). The Sixth Sense: Accelerating Organizational Learning with Scenarios. New York: John Wiley.

Wack, Pierre. (1984) Scenarios: The Gentle Art of Reperceiving. Reprinted by Global Business Network from President and Fellows of Harvard College, Working Paper 9-785-042.

Wasserman, D., Lempert, R.O., & Hastie, R. (1991) Hindsight and causality. Personality and Social Psychology Bulletin, 17: 30-35.

Wang, Shirley. (2012) "The New Power of Memory." The Wall Street Journal. March 14.

Wheeler, Mark A, Stuss, Donald T., and Tulvig, Endel. (1997) Toward a Theory of Episodic Memory: The Frontal Lobes and Autonoetic Consciousness. Psychological Bulletin. 121(3): 331 – 354.

Zimbardo, P. (1999) Putting Time in Perspective: A Valid, Reliable Individual-Differences Metric. Journal of Personality and Social Psychology. 77(6): 1271-1288.

Zimmer, Carl. (2011) The Brain: Memories are Crucial for Looking into the Future. Discover Magazine. April.