Extending the Communicative Possibilities of Foresight Through Formal Artifacts

John Benjamin Cassel

Strategic Foresight and Innovation

OCAD University

john.benjamin.cassel@gmail.com

Suggested running head: “Formal Futures Artifacts”

Abstract

Foresight tools make representational choices that impact understanding futures correctly. This paper explores these choices though formal representations. Formal representations promise dual possibilities for foresight communications: foresight through formal artifacts and artifacts of formal foresight. The foresight of formal artifacts speculates through detailed specifications, including legal documents, financial instruments, and software interfaces. Similarly, this paper explores the artifacts of formal foresight (such as how time, causality, risk, and regret can be represented) through a domain-specific scenario-description language. Research in formal representation addresses the psychological content of foresight, which is critical for effective foresight communication.

Keywords: cognitive biases, domain-specific languages, formal artifacts, representation methodology

Introduction

As foresight now has practitioners prepared to engage each other and a broader public through a wide array of platforms and media, foresight tools face choices in how to represent possible futures. These tools are enabled by the variety of content they can juxtapose and recombine, as well as a tremendous scope for the communicative and collaborative engagement. Perhaps we would like to design these platforms so that they represent futures in a way that is cognitively natural, promoting an understanding of contingency that is at least as rich as real events.

Recently, the foresight community has seen a great deal of innovation in the mediums though which potential futures are told, shared, and experienced, as well as in terms of new methods of collaboration that allow different kinds of political and social engagement about potential futures (see the Journal of Futures Studies special issue on the Communication of Foresight for a sample of these new developments). This essay will concern itself with a different communication problem: how futures are represented explicitly in computer systems and other formal mediums. This kind of explicit representation is important for three reasons: many kinds of foresight platforms will have to make representational choices in the structure of their content; foresight communications that do not need to make such choices explicitly (such as films) may be stored and retrieved by information systems that do; and there are many artifacts which contain important implications about potential futures but for which there are not yet associated narratives or platforms contextualizing them, including methods for eliciting fragmentary or predictive perceptions of the future, as well as many everyday formal documents and systems.

Formal representations promise exciting dual possibilities for the futures of foresight communications: foresight through formal artifacts and the artifacts of formal foresight. The foresight of formal artifacts looks at extending speculation into realms with detailed specifications about future possibilities, including legal documents, financial instruments, application programming interfaces (APIs), and other such specifications. In a similar vein, examining the artifacts of formal foresight is to look at how various means of representing futures portray aspects of the formal content of foresight, including representations of time, causality, risk, and regret. The development of either of these contributes to the other, and how there can be a kind of conceptual liquidity and concept reuse when foresight concepts are captured formally. To investigate this interface we will introduce aspects of a "domain-specific language" specifically designed for doing formal foresight work in high-risk infrastructure projects.

However, not only is the formal content of foresight a means for people to communicate, it is also critical to get right for foresight content to be communicated effectively. Simply asking about potential outcomes does not do objectively well at all, and may in fact bring the worst ideological biases to the surface (Tetlock, 2005), even when the questions themselves are designed to avoid ideological content. Foresight should take these biases seriously and follow parallel avenues of response: to create methodologies for addressing them and to cultivate objective standards for assessment. This paper will describe how the formalisms of the scenario-description language mentioned earlier are designed to respect this psychological understanding. By the end of this article, the reader should be convinced that it is possible to structure the design of a wide array of foresight-bearing artifacts in ways that avoid some of the known deficits in causal understanding.

Formal Foresight Artifacts: What and Why?

A formal foresight artifact is a complete specification of a system or subsystem used within a foresight work. To be formally specified, such an artifact should, under specified conditions, with a specified supporting context, entail particular self-standing interventions with temporal constraints or semantics. It does not need to specify the mechanisms that implement it. Consider, for example, a criminal law: on a particular violation, and given a set of enforcement and adjudication mechanisms with sufficient political support to operate, it will entail certain consequences for the individual having undertaken the violation to be applied over a given period. A valid specification should meet the formal standards of the artifact's user community as it would stand at the time the artifact would be employed.

The relationship between scenarios or narratives and artifacts is simple: a scenario will mention an artifact to demonstrate that 1) in this scenarios, the supporting conditions are right to allow that artifact to function and 2) this artifact has meaningful consequences for actions within this scenario.

The reason why we go to the trouble of constructing these artifacts and model their behavior with structural formalisms instead of merely making narrative use of them to make sure our understanding of them is clear (Tilly, 2004). Artifact construction and modeling are design research activities aimed at discovering if what we know coheres, and if we have thought through its implications. The utility of formalism is in its capacity for constructive self-skepticism.

Design Concerns for Understanding Futures

Even in a discipline as diverse as foresight, there is conceptual bedrock from which a professional norms can be prescribed. We would like to introduce individuals to plausible futures, determine which are desirable, and come to understand the dynamics by which they may effective pursue their desired futures. The degree to which we help accurately discover and convey futures to those ends is the degree to which we have communicated effectively at a professional level. Therefore, what we do not want to do is to misunderstand or distort their preferences, nor preclude reachable futures, nor confuse their understanding of their potential paths such that they act against their own interests. Foresight is professionally limited by personal needs and cognitive limits, or in short, psychology.

It is appropriate to question if the methods currently used within foresight provide the right cognitive toolkit for addressing future challenges (Verdoux, 2010). Unfortunately, the answer appears to be negative, as experiments in cognitive psychology combined with longitudinal studies of expert prediction (Tetlock, 2005) have cast foresight in a negative light (Gardner, 2011). Classic psychological studies (such as Tversky and Kahneman (1974)) have shown that individuals have systematic tendencies toward certain mistakes (these tendencies are called biases) in judgment about the likelihood of future events. Asking for possibilities and forming scenarios may actually be harmful, bringing these biases to the surface and making spurious connections, even when the methods themselves are designed to avoid ideological content.

To properly frame the relationship between foresight practices and psychological understanding, we can look to another field which has made a similar transition: design. Today, designers often serve as psychologists by proxy, in that they are aware that psychological factors strongly affect usability through factors such as cognitive load, attention, and memory. It is not that designers did not produce material improvements before becoming aware of these psychological factors, but instead once understanding them were able to introduce new usability testing and reduce designs that would produce mistakes. Similarly, foresight practitioners use processes that often help their clients come to new insights about their potential futures, but will be able to improve that guidance if their methods are able to reduce processes which generate biases. Let us start by looking at simple guards against mistakes.

Several biases are introduced by priming, framing, and implicitly evaluating input by presenting prior information. These included support-theoretic bias, which gives preference to events with more potential causes or variants, even if these causes and variants are themselves of extremely low probability[1]. Support-theoretic biases can be guarded against by not producing supports, using methods such as open-ended interviews (Jarratt, 1996), sense-making interviews, and other protocols that do not directly introduce content. Simply put, if you want to know what people think without your input, ask them first.

Another bias for against which there are methodological guards is base-rate neglect, in which the likelihood of common, but uninteresting, outcomes are underestimated. We know that people understand probabilistic results

[1] Support-theoretic biases also make arguments with more justifications seem stronger, even if those justifications are weak or spurious.

as the competing occurrence of deterministic paths of events, including paths that disrupt disruptions, such that explicitly describing these different paths can help defeat base-rate neglect (Krynski and Tenenbaum, 2007). If we can elicit potential paths in a non-directed way, there could be a chance to simultaneously avoid base-rate and support theoretic biases.

However, some biases create great methodological challenges. One biases generated by pursuing paths of probability is representation bias. In representation bias, an antecedent with availability bias may distort the assessment of a consequent. For example, after the 9/11 terrorist attacks, randomly sampled airplane passengers would pay more for flight insurance against terrorism than flight insurance against all causes, even though the second would cover more cases (Sunstein, 2007). In general, representation bias implies that individuals will assess $P(A|B)P(B) > P(A)$, even though $P(A) = P(A|B)P(B) + P(A|\neg B)P(\neg B) \geq P(A|B)P(B)$, if B is salient. This paper will propose a line of attack against representation bias later.

However, there is a second, more challenging, concern when considering the role of bias in foresight communications: we need to make sure that the objective standards we choose are themselves right for the environment we face. In Tetlock's work, even the best experts are outclassed by very simple autoregressive predictors (Tetlock, 2005), yet these estimators were taken over commonly recognized variables. It is as though humans unanimously agree on the very hard solutions of feature selection problems, yet disagree on their saliency and interconnections, falling flat just yards before the marathon finish line.

Perhaps it is the case that, in the real causal environments we face, the variables we have to work with are not fixed. Even stronger than this, for the cost of some biases, we gain incredible inductive power of what we can know about our world, and these are not only to our benefit in our evolutionary past. It seems to be true that we can form highly reliable structures of causal inference from very few examples and these inferences are better explained by choosing between models rather than picking the parameters of models (Griffiths and Tenenbaum, 2009). This theory explains the effectiveness of path-based methodologies in guarding against base-rate neglect, as they construct a more sensible structural model.

Some in foresight have suggested that Bayesian analysis is inappropriate as it can only evaluate the appropriateness of belief for known variables, and has nothing to say about the unknown (Fidler, 2011). While this is true for models that characterize our uncertainty about known variables, it fails to characterize Bayesian models that characterize what we believe is unknown. We can characterize what we know about the unknown because of what is already known about our discovery processes. Bayesian models that are capable of incorporating the results of stochastic discovery processes are said to be non-parametric (surveyed nicely in Jordan (2010)).

The consequence using non-parametric methods with intervention-based causal representations is that, given the inductive constraints of the foresight modeling language described below, causal effects can be learned effectively (Kemp et al., 2010). It is now an open program for futures studies to take this knowledge about how the mind learns and represents causal effects back to the operational setting of foresight practice (see Figure 1).

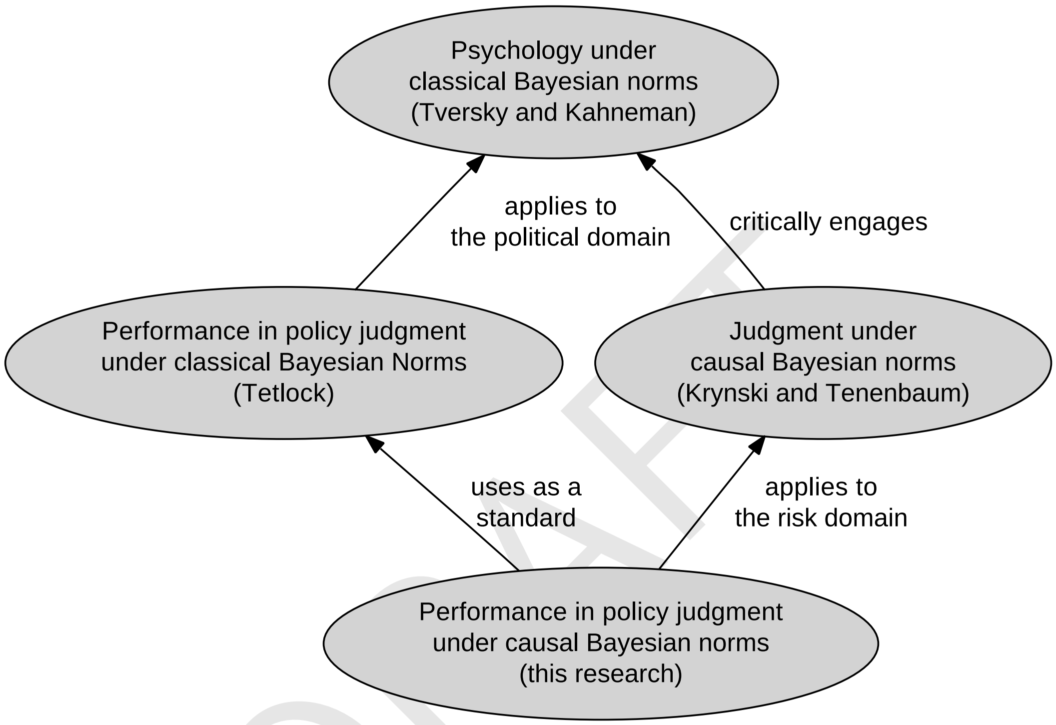

Figure 1: An Open Research Program

Foresight through Formal Artifacts

This section looks at a very tiny sampling of the kinds of specifications that can be modified for foresight purposes, as well as the formal aspects that they contain that are relevant for foresight. These are selected due to saliency, tradition, and availability, but without doubt more types of artifacts can be found and the reader is strongly encouraged to draw from their own research. Although this section is divided into three areas (Policy and Law, Finance and Economics, and Engineering and Technology), we will find that these distinctions are crossed frequently and that foresight addresses these crossings well.

Policy and law

Policy and law are rich with demand for solid foresight, with judicial futures being one example. Changes in both content and process are possible to imagine in both fields, such that a scenario about a future trial may change the process of that trial, the laws applied, the rules of evidence, the information technology of trial records, and even the process of deliberation; and likewise for the stages of developing and administering policies. Every one of these stages can currently be described by a document describing its procedure formally, and therefore each stage also admits a document describing different procedures.

One might wonder about why it is appropriate to speculate about policies and laws. After all, systems of government are not easily constructed, but instead represent long-suffering adaptations to capture and maintain cultural principles (Fukuyama, 2011). However, even within the bounds of those principles, the appropriateness of particular laws and policies change as underlying conditions change. For example, new technology is always challenging the boundaries of existing statutes. Consider recent rulings (United States District Court for the District of Vermont, 2009) examining if a person be compelled to disclose a private encryption key to decrypt files upon their computer that may incriminate them. In the law of the United States, this presents a gray area. On one hand, combinations to safes are considered no different than keys, and thus may be requested as part of a reasonable search. On the other hand, it is a constitutional amendment that an individual may refuse to testify against themselves. If we imagine future devices developed to extend our memory, then we can imagine that the personal refusal to disclose information would be compromised by reasonable searches of such devices (Dunagan, 2011). Of course, new technological arrangements have always produced boundary issues in legal precedents. Consider *Adams vs. New Jersey Steamboat Co.* (New York Court of Appeals, 1896) in which the room of a steamship was burglarized. Should the steamship company be held to strict (as observed in hotel rooms) or negligence (as in trains) liability. Although superficially one might expect that modes of transport share similar liability, the case proved to turn on the steamship offering individual locked rooms, as opposed to shared berths with multiple passengers, and thus were subject to the same expectations of care as hotels.

New technology is not the only kind of change that can alter the intended consequence of an existing policy. For example, federal flood insurance, in addition to other policy criticisms (Holladay and Schwartz, 2010), may mask the actuarial assessment and corresponding price signal for climate change.

As the future can change the impact of regulations, let us now look at policies and laws directly. We can roughly describe the overall purpose of policies and laws to transfer benefits and harms from where they occurred to their originating causes. Policy and law involve different causal considerations, as policy seeks sufficient causes to bring about a desired effect, while law is concerned with determining which the necessary causes in producing an event. Both laws and policies are formed in contexts that shape their acceptability. To violate the acceptability of a risk is to cause a second-order harm, namely the mis-allocation of its prevention. Conversely, if a risk is acceptable, but treated otherwise, then mitigation activities will be seen as wasteful, and the cost of which will be seen as a loss by stakeholders. Beyond the losses of mis-allocation, offenses to acceptability lead to other harms including the social amplification of risk and the loss of institutional credibility and support.

Let us look at legal regulations specifically. These are often concerned with an action for which a legal entity may be held liable (called the tortuous action), a harm caused by that action in fact, and the circumstances by which an individual may ultimately deemed liable regardless of actually having caused the harm.

At first, it would seem that the key issue in law is to determine the facts of the *actual cause*. For example, supposing that there were two fires, and although one of them advanced in such a way that it would have burned down a house, it occurred after the other fire, such that the other fire was the actual cause of burning down the house. As a logical guideline, one can refer to the actual cause as a necessary member of a sufficient set (NESS) (Wright, 1988) or, slightly more elaborately, a possibly insufficient but definitely necessary member of a possibly unnecessary but definitely sufficient cause (INUS) (Mackie, 1965). Yet, logical representations are not entirely appropriate, as standard logical implications do not necessarily hold (the disease is sufficient to cause the symptom, the contrapositive, namely that without the symptom there is no disease does not hold under interventions that treat the symptom) (Pearl, 2000). However, it is often the contextual circumstances which mitigate responsibility for an actual cause. Consider a gun fired into the air at midnight to ring-in a new year. If the bullet struck another individual, although the act itself was negligent, we may assess their liability differently if the gun was fired in a rural area with no nearby neighbors versus it being fired in an urban area.

Furthermore, the way in which conditions alter liability is not strictly actuarial. An assessment of liability may be more likely if the factors are known, are professional standards, or are involuntary instead of an unusual combination of factors that cross-cut categories or involve voluntarily undertaken activities, especially if those activities are deemed socially questionable. Consider a pharmaceutical company that produces a drug that produces a crippling side-effect in a very small percentage of the population. Therefore, it may be more likely to find liability if this side-effect is due to a rare allergy, as opposed to a join exposure to the fumes of metalwork and methamphetamine consumption, even if the second is more likely than the first. To this end, disparities with an actuarial assessment give a window into the beliefs and values of a legal culture.

Regulations have several temporal variables. First of all, what is the relationship between the time a statute was created and the time which it covers. We say that a given statute imposes retroactive liability [2] if a legal entity is liable for an activity even if it took place before that statute was undertaken. Regulations may also be limited by statues of limitation and repose, which limit the time after the tortuous action and the harm, respectively.

How could we discover the plausible outcomes of temporal parameters on policies? Although it may not be the complete story, an economic analysis is a good baseline. First of all, to do an economic analysis would require realistically characterizing the economic interests of all of the participants. To assume that these interests guide them accurately is to neglect other possibilities, but that is also true of ignoring the potential consequences of economic analysis. One example where economic analysis sheds light on the potential outcomes of retroactive policies is in the regulation of underground storage tanks (USTs) (Boyd and Kunreuther, 1995). In this analysis, small gas stations operating at tight margins will avoid maintenance preventing future leaks if retroactive costs and upgrading are greater than the firm value minus the retroactive cost and the expected leakage cost for future leaks. The general lesson is that if a cost that was pooled is suddenly experienced directly, this change in expectation may reduce preventative risk measures of all kinds as other risk pooling may not be judged as effective. On the other hand, for some risks a credible threat of retroactive liability may have consequences even without regulation being passed, due to uncertainty it would pass to investors, insurers, and other financiers, and it is difficult to know how credible such a threat would be without clear precedent.

Changes in regulatory precedent not only include what is regulated and when, but also the legal arrangements that are permitted. In recent times there have been great innovations in open-source and creative-commons copyright licenses, providing a regulatory framework to match the technical capabilities of digital distribution. These changes in contracts and contract frameworks may spread to other contexts. An example of this kind of speculation is the Urban Versioning System (Fuller and Haque, 2008), which proposes licensing, zoning, and property schemes for building and architecture similar to those found in software licenses such as the GNU license. These schemes have become more relevant with the incorporation of manufacturing technology into construction, and will become of key importance if design-for-disassembly, 3D printing, and other technologies bring information technology and its legal culture into the built environment.

One area of policy that squarely addresses not which is prohibited, but which is normal, is insurance law. The ability to buy insurance on a given activity is an indication that this activity is normal, widespread, and generally safe, such that pooling risk is feasible. As an example of insurance innovation, consider that Monitor Liability Managers insures against various legal concerns stemming from the use of social media (Monitor Liability Managers, LLC,

[2]Retroactive liability is also sometimes referred to as superfund liability, after the common name of the Comprehensive Environmental Response, Compensation, and Liability Act of 1980 (CERCLA) regulation.

2011). In some cases, to clarify liability concerns is arguably the largest obstacle facing new technologies.

In general, the intersection between finance and policy is a great place to look for boundaries not between the permitted and prohibited, but risks normalized and spread among the public versus those undertaken at risk to oneself. An area of that bends this distinction to an extreme degree is the granting of institutional structure and legal rights, responsibilities, and liabilities that can be held by companies, trusts, sovereign nations, and other non-human legal entities. This is played to great effect in Charles Stross's Accelerando(Stross, 2005) in which the protagonist pursues his ventures though a distributed network of corporations, each of which serve as officers of each other[3].

Having seen the first intersection of policy and finance, let us look at the formal opportunities of finance itself.

Finance and economics

Finance is thought to deal with a different kind of *futures* than that referred to in futures studies. Before exploring how finance and foresight contribute to each other, it is worth underlining these differences. First, a prediction in foresight is not an objective assessment, but also a speech act. It may serve as a rallying point from which to pursue or avoid that projected future. Further, finance presupposes a context in which benefits and risks can be pooled or transmitted, while foresight does not necessarily take such contexts to be a given.

Futures studies occasionally concerns itself with who makes predictions, what the content of those predictions are, and why they are made. A wide variety of financial matters can be interpreted as predictions in an actuarial sense. Further, in theory, it is possible to construct a financial security for any verifiable prediction. Simply choose a value for the security to pay out if the prediction is true (the strike price). Then, the percentage between the price of the security (the premium) and the strike price is the odds given the prediction. In practice, one must also specify an arbiter for verifying the prediction and how to insure payment. These prices need not be represented in any particular real or imaginary currency, but as the risk and rewards must be respected by all traders, real currencies provide the greatest liquidity. This kind of securities can be traded in prediction markets, of which there are many varieties, some of which support combinatorial bids (Hanson, 2003).

A key artifact for representing future in finance is the financial option (Banks and Siegel, 2007). For a premium, it allows one the option to buy (call) or sell (put) a particular commodity for a set price during some specified later times. The institutional context is such that the holder of the option believes that the writer will remain solvent

[3] Revealing perhaps too much of the plot (i.e. spoiler alert), the weakness of this financial strategy is revealed when a rival begins generating a large number of spurious lawsuits aimed throughout the network, creating a distributed denial-of-service attack.

or will have otherwise insured the contract, and will either voluntarily honor the contract or be legally compelled to do so.

The purposes to which such contracts can be employed are manifold, but one of their most straightforward is insurance. An individual buying a contract may wish to hedge the risk against them, such as a farmer buying an put option for a particular crop to ensure a minimum price, or a factory buying a call option for a particular raw material or commodity to ensure a maximum price. On the other side, the seller may act as an insurance agency, able to offset the occasional exercise of an option with premiums gained from the normal behavior of other assets, for which the options are not exercised. The portfolio of options held by an given party tells us what the assess as risk and whom they trust to provide countermeasures, and thus what they believe about the context of the instrument.

At first glance, the instruments of finance appear to take a crude first-cut in representing future, trading only in specific predictions and not concerning itself with the causal connectivity between events, which would only manifest themselves in trading strategies. However, the options described so far are only the most basic and much more exotic formulations (appropriately called exotic options) are possible. For example, consider compound options, which provide the buyer the right to buy another option, and the digital option, which grants a fixed payout when the underlying asset is within a particular range. Therefore, a conditional prediction could be implemented as a digital option that yields another option as its payout. Composing these options further, one could embed the better part of a trading strategy in a single artifact. This is not to say that such instruments are common, such that trading strategies embody the richest views of the future in the market.

As we have seen how finance is addressed by policy and law, let us see how policy and law are addressed by finance, beyond economic analysis. It is possible to think of a particular set of regulations as a portfolio of hedging measures with the premium being the cost of enforcement. Such an interpretation has now been applied to disaster law (Chen, 2012), despite the admitted impossibility of assessing tail events like disasters probabilistically (Taleb, 2011).

What would a speculative financial instrument look like? Here is one that could be relevant for foresight practitioners interested in trend work. Suppose that one has access to the number of academic papers or news outlets that mention a particular word, and that this indexed is independently maintained. Then, one could run a market using a fake reputation currency using securities tied to those counts. Such a market requires the context that market participants believe that firms capable of generating such numbers accurately and can effectively avoid manipulation.

When one reads finance, we sometimes understand that work institutions of finance often frame their work as information technology, framing their work in terms of price signals. This is certainly how prediction markets frames its own practice. Now let us look to the fields of engineering and technology directly.

Engineering and technology

New technologies are one of the most frequent artifacts encountered in foresight. Even some of the loosest and well-worn stereotypes (say flying cars) can produce interesting analysis (what kind of regulatory barriers would need to be crossed to make flying cars feasible). There are a profusion of new technologies that have yet to cross efficiency, reliability, cost, safety, maintainability, and deployability thresholds; as well as those that are making the policy transition from experimental to private liability and financing within a clarified regulatory environment. I cannot speak directly to how effectively this kind of technological forecasting is done. I have little doubt that technology road-maps maintained in international standards bodies, government agencies funding scientific and engineering research, and industrial leaders do track these concerns with varying degrees of effectiveness. At the same time, I will admit to not knowing of formal standards and tools suitable for maintaining and manipulating manifests desired performance characteristics, the populations that they serve, the body of technologies that are making progress progress toward these objectives, and how funding is allocated given the combination of need and track record. Like finance, it is worthwhile to see what views of the future are already present in their understanding of research programs and if any further formalization can help with disseminating this understanding.

In many cases one can construct useful yet speculative formal specifications without coming to a full understanding of current technology pipelines. One particular artifact useful for foresight development is the application programming interface (API). An API is a set of commands by which another computer program may make requests of a particular software component. It may seem that designing application programming interfaces is an obscure way to go about any particular design problem; after all, why should we use an artifact of computer programming in foresight? However, there is a common-sense reason for thinking through the API design in the earliest possible stages of a project, namely that it serves as the seam between infrastructure and artifact: the API has to be right for both. Before an infrastructure is imposed, we have to have some reassurance that the artifacts interacting with it can support the demands. Simultaneously, before designing an artifact, we must understand infrastructures under which the artifact will be expected to perform. The design of the API is a simultaneous exploration of both preconditions.

How does one go about constructing a speculative API? One can provide a relatively formal specification by describing each command with only a handful of components. First, one gives the syntax of how to issue a command to that application. Next, one describes the effect of the command, including any necessary preconditions, the outcome, conditions guaranteed to be left unaffected (the invariants), and possibly how errors are handled if the preconditions do not hold. Then, one shows an example of its use, and finally describes its effect in that specific example.

Here is an example from an API supporting spimes (Sterling, 2005) that would allow one to annotate, critique, and search for objects. This interface supported tagging, or the association of short text descriptions, with objects.

```
item.tag(tag, [tag,]* tagger)
```

Given this command, the item would then be associated with those provided labels, as well as who provided them. Here is an example of a loaf of bread being designated as being low in sodium.

```
thisLoafOfBread.tag('low sodium', jane_doe)
```

As a result of this action, another user who had access to Jane's tags would be able to see that a particular brand of bread had a loaf that was designated to be low in sodium, and so could find another loaf of that variety of bread, if they trusted Jane and if sodium was a concern, with a call such as `find('low sodium')` (which would be described elsewhere in the specification).

Engineering and technological systems must take into account the future in a broad number of ways, which can only briefly be surveyed here. Control systems must handle a wide variety of potential future conditions in their regulation of systems, often under sharp performance conditions if they are to avoid accident. The capacity of many systems must anticipate arrivals in a wide variety of distributions. Mechanical and civil systems must anticipate the loads upon them for their expected lifetimes. The components of a software system are configured so that they most likely extensions and maintenance can proceed easily and without system interruption. Database systems must remain coherent under concurrent usage. Engineering systems must also cope with a wide variety of performance criteria, often with multiple conflicting objectives, and occasionally with unknown future distributions over their period of life. It is in this last point that engineering and foresight have their greatest overlap, as they are both committed to discover and act upon potential design requirements. Foresight also has a great deal to say about the boundary conditions of when an engineering solution is appropriate given expected cost, cultural, and environmental constraints.

As a matter of process, model developers should also seriously consider the potential advantages afforded by test-driven development (Beck, 2003). These tests might describe whether various background conditions hold. Foresight modelers can manage the complexity of scenarios by checking that, for all artifacts employed, the necessary background conditions are met at the time of their deployment in scenarios. These background conditions also serve as directions to watch for change in present conditions, which may suggest that introducing particular legal, financial, technical, and other instruments may be appropriate.

Information technology also provides the means for representing aspects of foresight formally. Let us now turn to consider formal representations in foresight.

Artifacts of Formal Foresight

Suppose now that foresight practitioners, in addition to the formal representations that they can use from other fields, would like their own formalizations. As foresight practices are very diverse, we may want any number of specifications for different purposes. Fortunately, the development of formalizations has been substantially easier through the increasingly widespread use of domain-specific programming languages. Domain-specific programming comes from a long-held tradition in computer science, perhaps best summarized by the following quote from a well-regarded textbook: *"...a computer language is not just a way of getting a computer to perform operations but rather that it is a novel formal medium for expressing ideas about methodology. Thus, programs must be written for people to read, and only incidentally for machines to execute."* (Abelson et al., 1996). This tradition is stronger than ever as new pedagogy (Fowler and Parsons, 2011) and tools (Bovet and Parr, 2008) have brought the development domain-specific programming languages within the skill-set of ordinary programmers.

With many different appropriate formalizations for foresight, it is appropriate to understand the design criteria for the language features presented here. This language was developed for risk governance problems (Renn, 2008) in which the stakeholder population are likely to have varying understanding of the issues involved, the mental models of other stakeholders, the appropriate criteria to be satisfied, and other's perceptions of the appropriate criteria; and which all of these factors must be managed to avoid deficits in risk governance (Graham et al., 2009). One of the great challenges of developing a formalisms for this particular foresight area is that it must not only capture the descriptions of its professionals, but also accurately capture the perceptions of clients, participants, and sources with fidelity, and yet not introduce any background information beyond those distinctions, so that we get a picture the judgment of similar individuals whom we do not get the opportunity to elicit. At the same time, we wish to not irritate our stakeholders by asking about common knowledge, nor to produce formalizations that are incomprehensibly large, but largely composed of trivialities. This kind of formalisms necessary to meet this challenge are exceptionally challenging to represent, elicit, and analyze, and therefore a more detailed treatment is left elsewhere (Cassel, 2011). This background should now allow one to understand this language's choices for the formal representation of states-of-affairs, criteria, time, causality, risk, and regret.

Before we begin, let me offer a warning. Resist the temptation to consider what is represented to be carved in stone! The strength of a particular model comes from the ability to discover its own incompleteness[4].

A structuring for states-of-affairs

An overall factual state-of-affairs can be thought of as the aggregation of all that happens to be true at a given time. Let us call some of the these individual configurations of true facts structures, as they represent the way some components of the state-of-affairs could be structured. Here is an example which demonstrates two different ways to describe the conditions in which Bob owns a basketball.

```
structure bobOwnsABasketball
          (basketball (owner bob)) .
structure bobOwnsABasketball
          (bob (owns basketball)) .
```

This kind of expression is a nested tag cloud. By default, a computer knows nothing of the equivalence between these expressions, which might as well read (a (b c)) and (c (d a)). Instead, the meaning of these tags is supplied through additional operations done to these structures and regularities in the use of tags, as well as whatever additional interpretation is provided by the reader.

In addition, suppose that by default each of these tags is associated with a decimal number, such that (basketball (owner bob)) can be read as (bob:1.0 (owns:1.0 basketball:1.0)). These numbers may prove to be useful later, as we could interpret them to mean that the first Bob mentioned owns-in-full exactly one basketball. Overall, nested sets of weighted-tags are a powerful representational choice for a condition.

Whether or not a given structure is actually the case is a matter of testimony. It will commonly be the case that a participant may simply assert that a condition is true now. Any time a structure is asserted to be true, or having been true, we must declare who asserted it.

```
structure usGreatPlainsAtDesertRisk
  (usGreatPlains atDesertRisk)
  current according to CanadianCACTechExpert.
```

[4]The dynamics of the models themselves is a complex subject, beyond the scope of this paper.

However, individuals may be uncertain, but not entirely uncertain, about what is true currently and in the past. Consider the phrase "The keys are either in my purse or on the kitchen table." We can say that either structure is current with likelihood 0.5, and that these conditions are mutually exclusive, which we can specify like so.

```
depending keysPurseOrKitchenTable
on keysPurse
mutually exclusive
        keysKitchenTable
according to KeysOwner.
```

These examples are certainly overkill for expressing the ideas shown, but this representation can flexibly adjust to more detail and relationships.

Systems and interventions

A critical stage in many foresight projects is the presentation of scenarios, among which a relatively small number are selected. This number is justified not merely by the fact that a small number of scenarios is easy to understand and compare, but also due to the fact that scenarios should not represent every outcome or contingency, and far from it. For example, although it is useful for foresight to discover, present, and monitor cases that can be detected and acted upon in the ordinary course of events, scenarios should generally not represent them, as these do not necessarily alter how the scenario participants should structure their ventures and institutions. Other cases that scenarios should generally not include are cases between which all stakeholders are indifferent, cases posing events for which there is no reasonable prior position that alters the outcome, or cases in which the distribution of events is known such that any risks can be priced accordingly. Instead, scenarios should present possible sets of general state trajectories initiated by particular mixes of reinforcing and balancing trends. We can see this 'systems' view at work in much of the terminology used in the most common resources for understanding scenarios (Schwartz, 1991), such as drivers which determine particular trajectories, critical uncertainties as factors that can determine one path versus another, and weak signals which serve as early indicators that a particular path seems to be underway.

Yet, is not the case that these networks of phenomena necessarily determine the actual trajectory of the future, for there are many events which change the topology of how different factors influence each other. Often incorporated as "wild-card" scenarios, these events represent "gray swans", imaginable yet unpredictable shifts (Taleb, 2007). At the same time, not all topological changes in the dynamics of a system are best represented as such. To compare the difference, consider a movie theater. Whether each seat is taken should be considered discretely, unlike the fluid level of a gas tank or bathtub, which is continuous. Yet, the rate of arrivals to see a particular movie would probably be best represented as a steady flow, until all of the movie theater seats are spoken for, at which time the movie theater will have to handle new arrivals differently. At that point, it is clear a discrete cut-off is reached and that change is better represented not as a change in dynamics, but a change in the topology of dynamics. Following the interventionist school of causality (Pearl, 2000), we will call such a change in topology an 'intervention'.

It is possible to start by making a map of some set of dynamics, and then develop rules for how the topology of that map should be modified under various conditions. One challenge is that further considerations change the rules on how the topology is changed. In order to generate maps that correctly handle dynamic topologies, it may mean generating a combinatorial number of them, and being careful to maintain them under new conditions.

There is an alternative to explicit mapping to describe dynamic systems. This method would allow for every change to the state-of-affairs to change both the structural and numerical character of the dynamics. If we describe events as changing the composition and weights within the kind of tagged structures we have described before, we can discover system dynamics that are only visible once accommodating shifts on topology.

The possibility of this kind of macro-dynamics appearing is compounded when dependencies between events, similar to those described above, are introduced. Dependencies state what will also happen or not, and with what likelihood, given the occurrence of other events. Dependencies allow the representation of reliable theories about the result of interventions, even if these interventions have not been taken, and remain 'counterfactual'. Developmental psychologists have posed the possibility that children, at a very young age, learn mental representations called 'causal maps' (Gopnik et al., 2004), which allow them to make sensible inferences about the results of possible actions without actually having to undertake them. These psychological theories and their experimental results are grounded by new development in the philosophy of causality and its applications in statistics and computer science (Pearl, 2000) (Spirtes et al., 2000). This work finally puts common intuitions about the difference between causality and correlation into a workable framework. For example, we should be able to capture such clear intuitions that tampering with your barometer does not mean you can control the weather, nor does accidentally setting off your burglar alarm imply someone is robbing your house.

Subjectivity and regret

Foresight is clearly not only concerned with states-of-affairs, and how they may come about, but which of those outcomes are preferred and by whom.

First of all, we must have some criteria for evaluation. The minimum description for each criteria is to note that we are describing a criteria, to name it, and to give it a direction in which to change it is to improve the overall experience of the individual evaluating it. As much as possible, we would like each of the criteria to be independent of each other, such that a single action could improve the conditions according to one criteria while worsening it in another. Here are some criteria.

```
criteria nourishedPopulationPercentage maximize.
criteria nourished maximize.
criteria recognition maximize.
criteria pain minimize.
criteria agricultureProduction maximize.
```

When looking at this particular criteria, one might wonder if agriculture production is appropriate as an end in itself. Remember however that we are often seeking to capture the subjective judgments of stakeholders, however they might be truncated.

Criteria are assessed by stakeholders. We can identify different parties as having similar stakes if the ordering of their preferences by different criteria is maintained across analogical states-of-affairs.

```
stakeholder UnitedStates.
stakeholder CanadianCACTechExpert.
```

To be a particular stakeholder is a state-of-affairs. That is to say, being a stakeholder is a condition in time and a condition that can be changed by events. A farmer who loses their farm to a flood likely has a different set of concerns about agricultural policies afterwords. It is also possible to become a different stakeholder for sociological (Hedstrom, 2005), structurational (Orlikowski, 1992), or other reasons than sharp changes in circumstance.

Finally, we can tie stakeholders and their criteria to states-of-affairs by specifying how each condition impacts them.

```
impact cropLoss
       high declining agricultureProduction
       to agricultureRegion if (drought)
       according to CanadianCACTechExpert.
```

With the combination of states, and their dynamics, and their impacts, we can form regrets. Regret is the disparity between what did happen under what actions were undertaken and what would be projected to happen under a different course of action. However, regrets are curious things, for regrets are relative, not absolute, harms. For there to be a regret, there must be another state-of-affairs that did not occur, but that we take so seriously as a possibility that we judge ourselves in reference to it (Gopnik, 2009). These other states-of-affairs are counter to the facts of what happened, and thus are called counterfactuals. As there are no future facts, but there are many future states of affairs that could be. Therefore, it is appropriate that foresight should concern itself with counterfactuals, how they are understood, and how they are applied.

We would like to avoid regrets, but is impossible. Imagine that you are walking down the street and are asked for spare change. Whether or not we comply with the request, we experience regret, at either being taken advantage of and deprived of what is ours, or for not showing mercy to fellow person with sufficient needs as to make that demand. For this reason, there is no such thing as perfect foresight: we are, and will be, in circumstances that require conflicting trade-offs between multiple criteria that have no objective answer.

Conclusion

Those building foresight models have a grand opportunity to collaborate with individuals from a wide variety of fields using common instruments of their disciplines. Understanding how people understand the future is far from complete if we do not understand the tools people regularly use to interact with the future in their working lives, including the functioning, metaphors, and consequences of those tools. **The cognitive power of artifacts is that, by acting outside of the line of narrative within given scenario, they present independent streams of intervention, capable of derailing causal paths made too salient by representation bias**.

Overall, the structure of how we communicate about futures has a tremendous effect in designing how we understand and incorporate the content of that communication. The belief that we can see and capture whole systems is merely a device our mind. At times, this belief allows us to perform literally better than we know through the powers of speech acts and self-deception. At other times, it leads us to see connections that all too clearly hide other evidence. May we have the wisdom to create methods that will hold us to the mast when the sirens of our fancies lead us to the rocks.

Bibliography

Abelson, Harold, Gerald Jay Sussman, and Julie Sussman. *Structure and Interpretation of Computer Programs*. MIT Press, Cambridge, MA, 2nd edition (1996).

Banks, Erik and Paul Siegel. *The Options Applications Handbook: Hedging and Speculating Techniques for Professional Investors*. McGraw-Hill, New York (2007).

Beck, Kent. *Test-Driven Development by Example*. Addison-Wesley, Boston, MA (2003).

Bovet, Jean and Terence Parr. Antlrworks: an antlr grammar development environment. *Software: Practice and Experience*, 38:1305–1332 (2008). ISSN 0038-0644. doi:10.1002/spe.v38:12.

Boyd, James and Howard Kunreuther. *Retroactive Liability and Future Risk: The Optimal Regulation of Underground Storage Tanks*. Resources for the Future, http://ageconsearch.umn.edu/bitstream/10768/1/dp960002.pdf (1995).

Cassel, John Benjamin. *Addressing Risk Governance Deficits through Scenario Modeling Practices*. Master's thesis, OCAD University (2011). Available from http://john-benjamin-cassel.com/FinalProject.pdf.

Chen, James Ming. Modern disaster theory: Evaluating disaster law as a portfolio of legal rules. *Emory International Law Review*, 26 (2012).

Dunagan, Jake F. *Politics for the Neurocentric Age*. Ph.D. thesis, University of Hawaii at Manoa (2011).

Fidler, Devin. Foresight defined as a component of strategic management. *Futures*, 43:540–544 (2011).

Fowler, Martin and Rebecca Parsons. *Domain Specific Languages*. Addison-Wesley, Boston, MA (2011).

Fukuyama, Francis. *The Origins of Political Order: from prehuman times to the French Revolution*. Farrar, Straus, and Giroux, New York (2011).

Fuller, Matthew and Usman Haque. Urban versioning system 1.0. *Situated Technologies Pamphlets*, 2:1–61 (2008).

Gardner, Dan. *Future Babble*. Dutton, New York (2011).

Gopnik, Alison. *The Philosophical Baby: What Children's Minds Tell Us About Truth, Love, and the Meaning of Life*. Picador, New York (2009).

Gopnik, Alison, Clark Glymour, David M Sobel, Laura E Schulz, Tamar Kushnir, and David Danks. A theory of causal learning in children: causal maps and Bayes nets. *Psychological review*, 111(1):3–32 (2004). ISSN 0033-295X. doi:10.1037/0033-295X.111.1.3.

Graham, John D., Beat Habegger, Belinda Cleeland, and Marie Valentine Florin. *Risk Governance Deficits: An analysis and illustration of the most common deficits in risk governance*. International Risk Governance Council, http://irgc.org/IMG/pdf/IRGC_rgd_web_final.pdf (2009).

Griffiths, Thomas L. and Joshua B. Tenenbaum. Theory-based causal induction. *Psychological Review*, 116:661716 (2009).

Hanson, Robin. Combinatorial information market design. *Information Systems Frontiers*, 5:107–119 (2003).

Hedstrom, Peter. *Dissecting the Social: On the Principles of Analytical Sociology*. Cambridge University Press, Cambridge (2005).

Holladay, J. Scott and Jason A. Schwartz. *Flooding the Market: The Distributional Consequences of the NFIP*. Institute for Policy Integrity: New York University School of Law, http://policyintegrity.org/files/publications/FloodingtheMarket.pdf, policy brief no. 7 edition (2010).

Jarratt, Denise. A comparison of two alternative interviewing techniques used within an integrated research design: a case study in outshopping using semistructured and non-directed interviewing techniques. *Marketing Intelligence and Planning*, 14:6–15 (1996).

Jordan, Michael I. Bayesian nonparametric learning: Expressive priors for intelligent systems. In *Heuristics, Probability, and Causality: A Tribute to Judea Pearl*, chapter 10. College Publications, http://www.collegepublications.co.uk (2010).

Kemp, Charles, Noah D. Goodman, and Joshua B. Tenenbaum. Learning to learn causal models. *Cognitive Science*, 34 (7):1185–1243 (2010). ISSN 03640213. doi:10.1111/j.1551-6709.2010.01128.x.

Krynski, Tevye R. and Joshua B. Tenenbaum. The role of causality in judgment under uncertainty. *Journal of Experimental Psychology*, 136:430–450 (2007).

Mackie, J. L. Causes and conditions. *American Philosophical Quarterly*, 2 (1965).

Monitor Liability Managers, LLC. Social media endorsement. http://www.monitorliability.com/ fileadmin/docs/pdfs/ Social_Media_Overview.pdf (2011).

New York Court of Appeals. *Adams v. New Jersey Steamboat Co.* 45 N.E. 369, New York (1896).

Orlikowski, Wanda. The duality of technology: Rethinking the concept of structure in organizations. *Organization Science*, 3(3):398–427 (1992).

Pearl, Judea. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York (2000).

Renn, Ortwin. *Risk Governance: Coping with Uncertainty in a Complex World*. Earthscan, London (2008).

Schwartz, Peter. *The Art of the Long View: Planning for the Future in an Uncertain World*. Currency Doubleday, New York (1991).

Spirtes, Peter, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, Cambridge, MA, 2nd edition (2000).

Sterling, Bruce. *Shaping Things*. MIT Press, Cambridge, Ma. (2005).

Stross, Charles. *Accelerando*. Ace Books, New York (2005).

Sunstein, Cass. *Worst-case Scenarios*. Harvard University Press (2007).

Taleb, Nassim Nicholas. *The Black Swan: The Impact of the Highly Improbable*. Random House, New York (2007).

Taleb, Nassim Nicolas. *The Future Has Thicker Tails than the Past: Model Error As Branching Counterfactuals*. Social Science Research Network, http://ssrn.com/ abstract=1850428 (2011).

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

Tilly, Charles. Observations of social processes and their formal representations. *Sociological Theory*, 22:595602 (2004).

Tversky, Amos and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185:11241131 (1974).

United States District Court for the District of Vermont. *In re Grand Jury Subpoena to Sebastien Boucher*. No. 2:06-mj-91, 2009 WL 424718, Vermont (2009).

Verdoux, Philippe. Risk mysterianism and cognitive boosters. *Journal of Futures Studies*, 15:1–20 (2010).

Wright, Richard W. Causation, responsibility, risk, probability, naked statistics, and proof: Pruning the bramble bush by clarifying the concepts. *Iowa Law Review*, 73:1001–1078 (1988).