Appearing in *Error and Uncertainty in Scientific Practice,*
Marcel Boumans, Giora Hon and Arthur Petersen, eds., 2014
Pickering & Chatto Publishers, London, pp.157–175

Order and Indeterminism: An Info-Gap Perspective

Yakov Ben-Haim
Yitzhak Moda'i Chair in Technology and Economics
Faculty of Mechanical Engineering
Technion—Israel Institute of Technology
Haifa 32000 Israel
yakov@technion.ac.il

Contents

1 Introduction

Order and chaos have impressed themselves on consciousness throughout human history. Modern science attempts to uncover and understand orderliness in creation, leaving tohu vavohu for others to contemplate. In this essay we describe a conception of indeterminism—the unknown, the uncertain, the formless and void—that is relevant both to the scientific endeavor and to the practical attainment of reliable decisions in human affairs. Tohu vavohu can be characterized, even understood in some sense, without dispelling the mystery of the unknown.

Indeterminism—the lack of orderly law-like progression of events—occurs both in the physical world and in human affairs. We will use info-gap theory to characterize this indeterminism and the responses to it.

The prototype of worldly indeterminism is quantum mechanics. Quantum mechanics has enjoyed more than a century of success, explaining black body radiation and the photo-electric effect in the

---

\papers\ErrorAndUncertainty\nlni07.tex 7.4.2013

early years, up to nuclear tunneling, radioactive decay, anti-matter, diode and transistor physics, and more. And yet, quantum mechanics comes at a cost: constriction of the domain of scientific explanation, and weakening of the classical concept of causality. One grain of truth that quantum mechanics captures is that the physical world has an element of irreducible elusiveness, what one might call *natural* or *ontological indeterminism.*

Human affairs are full of surprises, and the most important are those that we bring upon ourselves. People make discoveries and inventions. People have discovered new continents, the structure of the atom, the size of the universe, the processes of biological evolution, and much more. People have invented the kindling of fire, the domestication of wheat, writing, the wheel, the mechanical clock, the printing press, electric motors, the internet, and so on. These inventions and discoveries alter the behavior of individuals and societies in fundamental and far-reaching ways. However, tomorrow's discovery cannot be known today. We cannot know what is not yet discovered, so we cannot predict future behavior in its entirety. The human world has an element of irreducible elusiveness, what one might call *human* or *epistemic indeterminism.*

Ontological and epistemic indeterminism share a basic attribute: they are unstructured and unbounded, they lack order and form and regularity. Rules or laws do not govern or generate patterns of indeterminism. The unknown is unconstrained even by what we know: knowledge, even if true of the past, need not hold in the future.

Info-gap theory provides a framework for understanding both ontological and epistemic indeterminism. We will discuss a conceptual framework for understanding the physical, ontological origin of quantum indeterminism. The info-gap concept of robust-satisficing underlies the probabilistic interpretation of the wave function. We will also explain how info-gap robust-satisficing describes epistemic indeterminism and supports responsible, reliable decisions in the face of epistemic indeterminism.

We begin with a generic conceptual discussion of info-gap decision theory in section 2. We describe ontological and epistemic indeterminism in sections 3 and 4. In section 5 we discuss the concepts of optimizing, satisficing and indeterminism, in preparation for employing the idea of info-gap robust-satisficing as a response to indeterminism in section 6.

2 Info-Gaps

In this section we describe the conceptual basis of info-gap decision theory (see also [11, 12, 48] for further non-technical discussion).

2.1 The Known and the Unknown

The known and the unknown form an exclusive and exhaustive dichotomy, but one that is highly asymmetrical. 'Knowledge' is to the 'unknown' as 'banana' is to 'non-banana' or 'linear system' is to 'non-linear system'. The asymmetry is in both size and quality. For instance, biologists estimate that the number of species that have not yet been identified substantially exceeds the number that have [2, 18, 20], and the ratio of unknown to known is even greater if one includes extinct species. Our immediate physical world—the earth, the solar system—is known in far greater detail than the vast expanses of the universe with its countless stars and planets. The realm of the known is tiny compared to what is not known. The qualitative difference between the known and the unknown is also substantial, as we will see.

"Whereof one cannot speak, thereof one must be silent." [53, section 7]. However, from both practical and speculative points of view, the unknown is not the realm of silence. We can and must say very much about the unknown. I will be discussing the unknown as an epistemic entity, rather than an ontological one. In this subsection I will use anthropomorphisms and sometimes even metaphors, but in section 6 we will extend our discussion beyond exclusively human knowledge and ignorance and attempt to be more precise.

One of the first things to say about the unknown is that its boundaries, if they even exist, are unknown. For example, in the analysis of risk in many fields we don't know which bad event (if any) will occur, and we often cannot realistically specify a worse case. The inventiveness—pernicious or propitious—of the unknown future is boundless. Or, we don't know how many animal phyla existed in the Precambrian age, though in the Burgess Shale of British Columbia scientists identified "eight anatomical designs that do not fit into any known animal phylum .... But this list is nowhere near complete .... The best estimates indicate that only about half the weird wonders of the Burgess Shale have been described." [24, p.212]. We don't know if the total number of Precambrian phyla throughout the world was 20, or 40, or 117. While the number is not infinite, a realistic bound is unknown, even assuming that the concept of a phylum is well defined. Furthermore, we have no idea *why* the preponderance of Burgess phyla became extinct: "we have no evidence whatsoever—not a shred—that losers in the great decimation were systematically inferior in adaptive design to those that survived." [24, p.236]. That this assertion may be hotly disputed by some biologists only strengthens the point.

The unknown is not limited to the physical or biological worlds, but includes the realm of human creativity as well. We don't know the contents of the lost book *Porisms,* nor what Archimedes would have discovered next had he not fallen to that Roman soldier's sword. As a more modern example, "no one knows, really knows, that is, what role nuclear weapons would play in war." [25, p.315] We don't even know how many nuclear weapons have been manufactured or where they all are. The possible answers to these few questions are endless, and they all reside in the realm of the unknown.

The next thing to say about the unknown is that it is in no way limited by the strictures of logical consistency. For instance, the known mass of the earth cannot be both 6×10^{24} kg and *not* 6×10^{24} kg. If knowledge is to be distinguished from error or ignorance, then it cannot contain contradictions (because from a contradiction one can deduce any other statement). However, regarding the mass of an unknown planet that might exist at the other edge of our galaxy, both statements have equal status in the realm of the unknown: one or the other is true, but we don't know which. Furthermore, the unknown nourishes the imagination and is the subject of science fiction as much as it is the subject of science. We can hold in our minds thoughts of 4-sided triangles, parallel lines that intersect, and endless other seeming impossibilities from super-girls like Pippi Longstocking [37] to life on Mars or merciful murder (some of which may actually be true, or possible, or coherent) [9].

Scientists, logicians, and saints are in the business of dispelling all such incongruities, errors and contradictions. Banishing inconsistency is possible in science because (or if) there is only one coherent world. Banishing hypocrisy and deceit is possible if people follow particular moral codes. The unknown is the realm in which scientists and saints have not yet completed their tasks. For instance, we must entertain a wide range of conflicting conceptions when we do not yet know how (or whether) quantum mechanics can be reconciled with general relativity, or Pippi's strength reconciled with the limitations of physiology, or killing with murder. As Henry Adams wrote:

> "Images are not arguments, rarely even lead to proof, but the mind craves them, and, of late more than ever, the keenest experimenters find twenty images better than one, especially if contradictory; since the human mind has already learned to deal in contradictions." [1, p.489]

A theory, in order to be scientific, must exclude something. A scientific theory makes statements such as "This happens; that doesn't happen." Karl Popper explained that a scientific theory must contain statements that are at risk of being wrong, statements that could be falsified [45, pp.133–134]. Deborah Mayo demonstrated how science grows by discovering and recovering from error [40].

The realm of the unknown contains contradictions (ostensible or real) such as the pair of statements: "Nine year old girls can lift horses" and "Muscle fiber generates tension through the action of actin and myosin cross-bridge cycling".

Scientific theories cannot tolerate such contradictions. But it is a mistake to think that the scientific paradigm is suitable to all activities, in particular, to thinking about the unknown. Logic is a powerful tool and the axiomatic method assures the logical consistency of a theory. For instance, Leonard Savage argued that personal probability is a “code of consistency” for choosing one’s behavior [47, p.59]. In contrast, Jim March compares the rigorous logic of mathematical theories of decision to strict religious morality. Consistency between values and actions is commendable says March, but he notes that one sometimes needs to deviate from perfect morality. While “[s]tandard notions of intelligent choice are theories of strict morality ... saints are a luxury to be encouraged only in small numbers.” [38, p.51]. Logical consistency is a merit of a scientific theory. However, logical consistency does not constrain the unknown because the unknown is the reservoir of possibilities and the source of inventiveness.

2.2 Info-Gap Theory

The unknown presents serious challenges to decision makers. Knight, in discussing entrepreneurial decision making, distinguished between ‘risk’ based on known probability distributions and ‘true uncertainty’ for which probability distributions are not known [32]. Wald, in studying statistical decisions, also considers situations of ignorance that cannot be represented by probability distributions [52]. Lempert, Popper and Bankes [35] and many others studying public policy, such as Beck in this volume, discuss severe or ‘deep’ uncertainty that reflects the unknown.

Info-gap theory is a decision theory: a methodology for modeling and managing the unknown and for supporting the formulation of plans, strategies, designs and decisions under severe uncertainty [4]. In this subsection we will describe the basic concepts of info-gap theory. We will be speaking of human decision making, and in section 6 we will refine our language and attempt to remove the anthropomorphism when we apply info-gap theory to both epistemic and ontological indeterminism.

Decision makers are goal-oriented. The engineering designer aims to increase the payload of an aircraft, or to decrease the failure rate of a milling machine. The economic planner seeks to improve profits or to reduce unemployment. The physician seeks to prevent or cure disease. The military strategist seeks to deter the enemy or to win in armed conflict if deterrence fails.

Decision makers do not operate in total darkness. They have data, knowledge, and understanding, which we will collectively refer to as ‘models’. Sometimes those models are confidently known, as in choosing rain gear for a winter trip to London: take an umbrella! Other times the models are imperfect but their uncertainties are highly structured. For instance, the farmer can choose next season’s crops based on reliable statistical data of temperatures and rainfall in recent years. Sometimes the uncertainties are severe and unstructured. Predicting rainfall or storm patterns a century hence is plagued by substantial lack of understanding about the processes of global warming. Info-gap theory is useful in modeling and managing severe uncertainty.

An *info-gap* is the disparity between what one *does know,* and what one *needs to know* in order to make a reliable, responsible decision. An info-gap is part of the disparity between the known and the unknown.

An *info-gap model of uncertainty* is a mathematical device that quantifies an info-gap. An info-gap model is an unbounded family of nested sets of possible realizations of an unknown entity. The unknown entity may be tomorrow’s stock value, or the number of Precambrian phyla, or the functional relation between temperature and fertility of an insect species, or an adversary’s unknown preferences among possible outcomes of a future nuclear war, or the probability that a new veterinary disease will spread to humans, or the probability distribution of useful lifetimes of a new micro-robotic machine. Mathematically, the elements of the sets of an info-gap model of uncertainty are numbers, or vectors, or functions, or sets of such things. An info-gap model is not a single set, but rather an unbounded family of nested sets. The sets become more and more inclusive, reflecting a growing horizon of uncertainty. The sets grow without bound (within the domain of their definition), reflecting the unknown bound on possible realizations of the uncertain entity. An info-gap model represents the unknown degree of error of the decision maker’s models.

Two concepts are central in info-gap theory in supporting the decision maker’s choice of an action in attempting to achieve a goal in the face of severe uncertainty: robustness and opportuneness.

The *robustness* of a decision assesses its tolerance to the unknown for achieving an outcome requirement. A decision is highly robust if the decision maker's goal will be achieved despite vast uncertainty about the decision maker's models. A decision is not robust if low uncertainty jeopardizes the achievement of the goal.

We now consider *opportuneness.* Decision makers are goal-oriented, but sometimes the goal is not a requirement (or demand, or obligation) such as "The useful lifetime of the system will not be less than 30 years". Sometimes the goal is an aspiration for windfall or better-than-anticipated outcomes, like "It would be wonderful—though it's not a design requirement—if the useful lifetime of the system exceeds 50 years". A decision is opportune if wonderful windfalls are possible—though not necessarily guaranteed—even at low levels of uncertainty. A decision is opportune if small deviations of reality from the models can facilitate large windfalls. While the robustness of a decision is its immunity against failure to achieve the goal, the *opportuneness* of a decision reflects its potential for wonderful windfall.

The decision maker who seeks to achieve a goal, despite severe uncertainties in the models, faces two irrevocable trade offs, one regarding robustness against uncertainty, and the other regarding opportuneness from uncertainty.

Robustness to error trades off against the quality of the goal. The decision maker must choose an action (or design, or strategy, etc.) and can use the models to predict whether or not a contemplated action would achieve the goal. The models represent the best available data, knowledge, and understanding, but they are accompanied by severe uncertainty and thus may err greatly. Hence predicted outcomes are not a reliable basis for decision. More specifically, it is not unreasonable to expect that the actual outcome may be worse that the predicted outcome. However, a goal that is somewhat less demanding or somewhat less ambitious than the predicted outcome may be achieved even if the models err a bit. As the goal becomes more modest, the contemplated decision (action) can tolerate greater error in the models and still achieve the goal. In short, the robustness (to error, uncertainty, ignorance or surprise) *increases* as the quality of the required goal *decreases.* This is the first trade off.

Opportuneness from error trades off against the quality of the aspiration. The decision maker's models predict the outcome of a contemplated action. The models represent the best available data, knowledge, and understanding, but they are accompanied by severe uncertainty and thus may err greatly. Hence it is not unreasonable to expect that actual outcomes may be better than predicted. An outcome-aspiration that is somewhat more ambitious than the predicted outcome may be achieved (but cannot be guaranteed) if the models err a bit. As the aspiration becomes more ambitious, the contemplated decision (action) requires greater error in the models in order for achievement of the aspiration to be possible. In short, the opportuneness (from error, uncertainty, ignorance or surprise) *decreases* as the quality of the aspiration *increases.* This is the second trade off.

It is entirely correct that assessment of the robustness or the opportuneness of a decision depends on the decision maker's models. In fact, that's usually the point of performing the assessment. Given the decision maker's current state of knowledge—as expressed by the models—we wish to assess the vulnerability (or opportuneness), of a contemplated decision, to error in that knowledge. The robustness is used to prioritize the available decisions given the existing models: a more robust option is preferred over a less robust option. The opportuneness is used similarly: more opportuneness is preferred over less opportuneness (the robustness and opportuneness prioritizations need not agree). From a practical point of view—and decision makers are usually pragmatic—a choice must be made from a given state of knowledge. What the decision maker needs to know is how robust, or how opportune, each option is, given the available models. A decision theory in which the robustness and opportuneness were *not* sensitive to the decision maker's models would not be useful at all.

It is incorrect to assert that, because the info-gap robustness depends on the decision maker's models, it is local rather than global. Info-gap models of uncertainty are unbounded on the domain of their definition, so they are not constrained "locally" to the region around the decision maker's models. This allows the robustness to be either small or large. If the robustness is large then the goals are achieved even if reality deviates enormously from the models. At the extreme (which sometimes occurs), infinite robustness implies that the goal is guaranteed, regardless of the degree of error. At the other extreme, very small robustness implies that only small deviation is tolerable: the goal will be achieved only if reality is close to the model. That the robustness of some decisions is small does not mean that the info-gap robustness is inherently local. The robustness may be small or large, depending on the decision maker's models and requirements. Small or large robustness implies that the domain of reliability of the contemplated decision is local or not, respectively.

As we have explained, the robustness is used by decision makers to prioritize the available options: the more robust option is preferred over the less robust option. However, the first trade off that we discussed asserted that the robustness of an option depends on the required quality of the goal. Hence the prioritization of the options may depend on the decision maker's goal. It can happen, and often does, that the prioritization of the options is different for different quality requirements: a very demanding requirement might be most robustly achieved with one option, while a less demanding goal may be robust-preferred with a different option. This "reversal of preference" is very important in decision making under uncertainty [14] and in explaining a number of anomalies of human decision making, including the Ellsberg and Allais paradoxes and the equity premium puzzle [4, chap. 11] and the home bias paradox [13]. We will discuss an example in section 6.2.

3 Ontological Indeterminism

3.1 Polarized Photons: An Example of Ontological Indeterminism

Dirac [22] discusses the interaction between polarized photons and the polarizing crystal tourmaline. If a beam of polarized light impinges at an angle α to the crystal axis, then a fraction $\sin^2 \alpha$ will be transmitted and will be polarized perpendicular to the crystal axis upon exiting the crystal. If a single photon impinges, polarized at an angle α to the crystal axis, then it will either be completely absorbed or completely transmitted; in the latter case it will be polarized perpendicular to the crystal axis. A fraction $\sin^2 \alpha$ of such photons, impinging independently, will be transmitted. After describing these observations Dirac writes [22, p.6]:

> Thus we may say that the photon has a probability $\sin^2 \alpha$ of passing through the tourmaline and appearing on the back side polarized perpendicular to the axis and a probability $\cos^2 \alpha$ of being absorbed. These values for the probabilities lead to the correct classical results for an incident beam containing a large number of photons.
>
> In this way we preserve the individuality of the photon in all cases. We are able to do this, however, only because we abandon the determinacy of the classical theory. The result of an experiment is not determined, as it would be according to classical ideas, by the conditions under the control of the experimenter. The most that can be predicted is a set of possible results, with a probability of occurrence for each.
>
> The foregoing discussion about the result of an experiment with a single obliquely polarized photon incident on a crystal of tourmaline answers all that can legitimately be asked about what happens to an obliquely polarized photon when it reaches the tourmaline. Questions about what decides whether the photon is to go through or not and how it changes its direction of polarization when it does go through cannot be investigated and should be regarded as outside the domain of science.

Dirac's discussion illustrates what we are calling ontological indeterminism: "The result of an experiment is not determined … by the conditions under the control of the experimenter." The lack of control is not the experimenter's deficiency, but rather nature's indeterminism. The classical conception of a law of nature is modified by quantum mechanics. When an event occurs it does so in one way and not any other (excluding multiple worlds interpretations). To this extent, some "law" or property inherent in the substances and circumstances may be said to be acting. However, if identical substances and circumstances at different instants or locations result in non-identical outcomes, then "natural law" as understood classically—immutable and universal—does not hold.

There is some variability or indeterminism in those attributes that govern the course of events. What physicists have classically thought of as constant and fully specifiable laws are in fact to some extent indeterminate. It is not the experimenter's competence that is compromised; nature's competence to govern events is incomplete.

3.2 The Success of Science and (Or Despite) Ontological Indeterminism

What can we say about the character of ontological indeterminism? If we follow Dirac (and I will), such questions are "outside the domain of science" and are therefore speculative or philosophical. We are trying to understand how science—the search for laws of nature—works if nature has a non-nomological element.

We require concise qualitative understanding of two concepts from physics: the Lagrangian and the action integral.

The Lagrangian represents the physical properties of the system, and is often the difference between kinetic and potential energy. The physical properties of the system—whether it's a mass attached to a spring, a particle pulled by gravity, a photon in an electro-magnetic field, or whatever—are represented by the Lagrangian. The 'action integral' is the integral over time of the Lagrangian.

The equations of motion of both classical physics and quantum mechanical systems can be derived by finding a stationary point of the action integral ([42] section 3.2; [34] p.60; [22] p.128). Since the stationary point is usually a minimum, this is sometimes called the Principle of Least Action.

Dirac does not dispute or abandon the least-action principle, and Feynman's 1948 derivation of Schrödinger's equation, based on the path-integral method, explicitly exploits the least-action principle [23].

The problem we are facing is this. As we explained in section 3.1, nature is indeterminate to some extent; it lacks a complete universal and invariant set of laws that govern physical processes; nature is non-nomological to some degree. As Dirac writes "[q]uestions about what decides whether the photon is to go through or not and how it changes its direction of polarization when it does go through cannot be investigated and should be regarded as outside the domain of science." And yet, the Lagrangian represents the physical properties of the system. Classically this meant that the Lagrangian is obtained from the relevant laws of nature, such as stress-strain relations or gravitational attraction or Maxwell's equations. How can we reconcile ontological indeterminism and partial lack of law, with a unique and specified Lagrangian?

Classical physics dictated that nature "finds" the unique solution, or path, that obeys the laws of nature. Feynman, in his beautiful derivation of quantum mechanics, recognizes that an infinity of different paths all contribute to the physical process. The 2nd axiom in Feynman's 1948 paper states:

> *The paths contribute equally in magnitude [to the wave function], but the phase of the contribution is the classical action (in units of $\hbar$); i.e., the time integral of the Lagrangian taken along the path.*
>
> That is to say, the contribution $\Phi(x)$ from a given path $x(t)$ is proportional to $\exp(i/\hbar)S[x(t)]$, where the action $S[x(t)] = \int L(\dot{x}(t), x(t))\,\mathrm{d}t$ is the time integral of the classical Lagrangian $L(\dot{x}, x)$ taken along the path in question. [23, p.371].

Herein lies Feynman's reconciliation of quantum indeterminacy with the classical least-action principle. Consider two slightly different paths, x_1 and x_2, whose action integrals, $S[x_1(t)]$ and $S[x_2(t)]$, are slightly different. The phases of their contributions to the wave function will be different, so these paths will tend to cancel one another out. However, there is some path for which the classical action is stationary: for which $S[x(t)]$ does not change as $x(t)$ changes slightly. Consider a bundle of paths near this classical least-action path. These paths all have nearly the same action, and hence nearly the same phase in their contributions to the wave function. These paths will not cancel one another out, but rather re-enforce one another. Consequently, a bundle of paths near the classical path will be more likely to occur than a bundle of paths further from the classical path. Quantum indeterminacy is reconciled with classical least-action.

The classical Lagrangian is incorporated in Feynman's derivation of quantum mechanics, but it has lost its classical sovereignty and no longer determines a unique natural process in response to a unique set of initial conditions. The physicists are happy, and rightly so, because quantum mechanics has enjoyed a long history of success. But from the speculative point of view, how are we to understand the concept of a law of nature? I suggest that natural law is indeterminate; it does not take a specific functional form for any given physical system. Each path has its value of the action integral, based on the classical Lagrangian, but this does not determine the outcome of an experiment by application of the least-action principle, nor does the classical Lagrangian reflect an immutable and universal law of nature. Feynman's path-integral method, exploiting the phase-cancellation phenomenon, is needed to bridge the gap between classical determinism and quantum indeterminism.

Specific events happen specifically, at least when we observe them (quantum mechanics is circumspect about events that are not observed). When a specific event happens in one way rather than another, the classical physicists offered a Lagrangian—and the least-action principle—in explanation. This program collapsed with the advent of quantum mechanics. Ontological indeterminism is the gap between our classical understanding of the world (the classical Lagrangian) and some hypothetical law of nature that governs any specific event. This gap is unbounded in the sense that the "hypothetical law of nature" may have no real existence, no ontological status. The gap is also unbounded in the sense that all path-bundles contribute quantum mechanically to some (unequal) extent. The gap that I'm calling ontological indeterminism results in the constriction of the domain of science to which Dirac attests. In the quantum world, no single path "wins" universally, in a given experimental setup, and no single Lagrangian or law of nature "rules" in the classical sense. The range of paths that participate in the physical event is unbounded. Ontological indeterminism is the gap between what we do know—the classical Lagrangian and the least-action principle—and what the classical physicist would want to know—a complete explanation of specific events. This gap in our understanding is unbounded and, if the quantum mechanicists are right, unbridgeable.

4 Epistemic Indeterminism

The history of humanity is, in large measure, the story of discovery and invention. Our proud self-images as homo sapiens, the Wise Man, and the Tool Maker, convey the central message of deliberation, innovation and change. Despite the deliberative element, discovery and innovation entail surprise, and this has far reaching implications.

Habermas [27] emphasizes the non-nomological nature of social science, and Nelson and Winter stress that in evolutionary economics "things always are changing in ways that could not have been fully predicted" [44, p.370]. We learn things, which Keynes referred to as hearing the 'news', and this sometimes dramatically alters our behavior [31, pp.198, 199, 204]. The idea of indeterminism in human affairs was developed separately and in different ways by Shackle [49, pp.3–4, 156, 239, 401–402] and Popper [46, pp.80–81, 109]. We will refer to this Shackle-Popper indeterminism [5] as epistemic indeterminism because it derives from the limitation of what we know.

The basic idea of Shackle-Popper indeterminism is that the behavior of intelligent learning systems displays an element of unstructured and unpredictable indeterminism. By 'intelligence' I mean: behavior is influenced by knowledge. This is surely characteristic of humans individually and of society at large. By 'learning' I mean a process of discovery or invention: finding out today what was unknown yesterday. Finally, indeterminism arises as follows: because tomorrow's discovery is by definition unknown today, tomorrow's behavior is not predictable today, at least not in its entirety. Given the richness of future discovery, (or its corollary, the richness of our current ignorance), the indeterminism of future behavior is broad, deep and unstructured. The laws of human behavior will change over time as people make discoveries and inventions. These laws cannot be known in their entirety ahead of time, because by definition discoveries cannot be predicted and the laws of behavior depend on the discoveries to be made.

Epistemic indeterminism is the gap between what we *do know* about ourselves, our world, and our future, and what we would *need to know* in order to fully grasp and control that trajectory. This gap is unbounded as long as the universe of possible future discovery and invention is open-ended and accessible. This gap is continually shifting as new knowledge emerges, but the frontier of possibilities is shifting as well.

In section 1 we distinguished between ontological and epistemic indeterminism, and in this and the previous section we elaborated on these two conceptions. Nonetheless, our main claim is the structural similarity between these vastly different phenomena, unified by the concept of info-gap robust-satisficing. We now begin to explore the underlying commonality between epistemic and ontological indeterminism.

5 Optimizing, Satisficing and Indeterminism

5.1 Optimizing

Optimization—finding stationary values of an objective function—is a fundamental concept in physics. The variational principles of mechanics are optimization problems, equilibrium in thermodynamics involves minimizing an energy function, and so on. The "optimization paradigm" asserts that laws describe the behavior of a system and that these laws can be derived by optimizing a physically meaningful objective function. The properties of a system are embodied in its objective function. Examples of objective functions in physics are the action integral or the Gibbs free energy.

The optimization paradigm is prominent in the biological and social sciences as well, where the optimization of substantive or useful outcomes is a normative explanatory tool. For example, a large body of biological literature seeks to explain the foraging behavior of animals as guided by maximizing caloric intake [17]. Fitness or some measure of survival is often optimized in theories of evolutionary competition. In social science, mathematical economists derive equations of motion of an economy from "first order conditions" that specify optimality of economic utility functions [39].

5.2 Satisficing and the Limits of Optimization

Optimization requires that the objective function be determinate and accessible in some sense. This does not mean that the agents involved (projectiles, pigeons or pawn brokers for instance) need conscious knowledge of the objective function. Nonetheless, that function must be stable and discoverable—perhaps implicitly—in some relevant sense. Pigeons need not be conscious of their caloric intake, but measures or expressions of that intake must be accessible to them in some way, for instance by the evolutionary selection process against pigeons whose caloric intake is too low.

The need for stability and accessibility of the objective function limits both the explanatory power and the practical utility of the optimization paradigm in social science. Simon [50, 51] studied what he called the "bounded rationality" of animal, human, and organizational decision makers. Bounded rationality is the behavioral consequence of limited access to information and understanding, and the limited ability to process that information. As a consequence, Simon claimed, organisms "satisfice". Etymologically, "to satisfice" is a variant on "to satisfy", but "satisfice" has come to have a tighter technical meaning in economics, psychology and decision theory. The Oxford English Dictionary [43] defines "satisfice" in this technical sense to mean "To decide on and pursue a course of action that will satisfy the minimum requirements necessary to achieve a particular goal."

The prevalence of the outcome-optimization paradigm in the study of human affairs is due at least in part to the fact that more of a good thing is, by definition, better than less. Of course, things are not so simple. Optimizers are well aware of diminishing marginal utility, of substitution between goods, of budget constraints, and of absurd extremes (chocolate is good; too much chocolate is a belly ache). Theories of optimization account for these subtleties by modifying the objective function and by imposing constraints on the domain of solution. The limitations of the optimization paradigm in the social sciences derive from epistemic indeterminism, not from the subtle and multi-variate nature of human preferences. Optimization is of limited utility or feasibility if either information or information-processing capabilities are too scarce.

The attractiveness of the optimization paradigm in physics derives from various historical and intellectual factors that we cannot explore here. Very briefly, though, one motivation is certainly the beauty and simplicity of deriving all laws of nature from a single concept: the least-action principle. Another speculative motivation may be related to the search for certainty in a confusing world. If God is dead, (or at least dying as in the early stages of modernity), then the optimal perfection of natural law is not a bad surrogate.

The limitation of optimization in physics is subtler than in human affairs. The gist (and the beauty) of Feynman's derivation of the Schrödinger equation, as discussed in section 3.2, is that the least-action principle is preserved, though its impact is muted. The path whose action integral is stationary (this is the unique classical path) is the quantum mechanically dominant (most likely) path. Other (classically forbidden) paths contribute quantum mechanically with diminishing strength as their action integrals become less stationary. The classical optimization, that leads to a unique state of nature, is replaced by an infinity of sub-optimal solutions that all contribute to some degree as possible physical realizations. The discovery of ontological indeterminism led to constriction of the domain of science, and dilution of the concepts of causality and optimality, as discussed in section 3.1.

The success of satisficing—as an alternative to optimizing—in understanding the physical, biological and human worlds, derives from the nature of indeterminism that underlies the course of events in all these domains. The info-gap conception of uncertainty, and the resulting quantification of robustness, provide a framework for characterizing this indeterminism and the responses to it as we will discuss in section 6.

6 Info-Gap Robust Satisficing and the Response to Indeterminism

Knight's discussion of economic behavior under severe uncertainty distinguishes between 'risk' based on known probability distributions and 'true uncertainty' for which probability distributions are not known [32]. The info-gap conception of uncertainty, described in section 2.1, is very Knightian. Info-gap robust-satisficing and opportune windfalling, discussed in section 2.2, are operational responses to Knightian uncertainty. Info-gap theory [4] has been applied to decision problems in many fields, including various areas of engineering [19, 29, 30], biological conservation [16], economics [6, 14, 33], medicine [15], homeland security [41], public policy [28] and more (info-gap.com).

We are now in a position to explain how the concepts of satisficing and robustness to info-gap uncertainty provide a unified framework for understanding both epistemic and ontological indeterminism. This requires us to transcend the anthropomorphic and teleological language that we have used and that often characterizes the discussion of theories of human decision making. Human language, as distinct from mathematics, is embedded in human experience, so this transcendence will be only partial. Nonetheless, section 6.1 introduces the basic definitions that we will use in attempting to ameliorate the anthro-centrism of our discussion. Sections 6.2 and 6.3 discuss epistemic and ontological indeterminism.

6.1 Terminology for Info-Gap Uncertainty and Robust Satisficing

We will talk about *"a world"* as an entire sequence of events, both past and future. There are many possible worlds in the sense that many sequences of events are possible.

We will talk about *"decisions"* without necessarily entailing consciousness or volition. A decision has been made, in the sense we intend, if the course of events in the world goes one way rather than another. A person's decision may result from that person's volitional choices. But we can also talk about the decisions made by an organization, or by a squirrel or a squid, or by nature as a whole.

We will talk about *"satisfying a requirement,"* without necessarily implying either volition or teleology. A requirement is satisfied when a condition is met. A person's condition is met, if we believe Lennon–McCartney, by achieving love [36]. A squirrel satisfies a caloric requirement by eating acorns, and nature needs action integrals that are stationary or nearly so.

A decision is *"robustly satisficing"* if the requirement is satisfied over a wide range of possible worlds. A person's decision to visit a library, (rather than a bar perhaps), is robust if love is achieved for any of a large set of other visitors. A squirrel's decision to eat acorns under this tree (rather than under another tree) is robust if the caloric requirement is satisfied regardless of weather, or predators, or competing squirrels. Nature's decision is robust if its action integral is nearly stationary over a wide range of possible Lagrangians.

6.2 Robust Satisficing and the Response to Epistemic Indeterminism

In human affairs, robust decisions are motivated by epistemic indeterminism. We don't know everything that we need to know, or we don't have the capability to process all the available information, so we want to satisfy our requirements even if we err greatly about the world. Bounded rationality motivates robust-satisficing as a decision strategy.

The motivation for robust-satisficing can also be understood in the context of competition under uncertainty. In human affairs, competition among individuals, organizations, societies and civilizations sometimes comes with high stakes, even survival or demise. An entity whose critical requirements are satisfied despite huge surprises will tend to survive at the expense of competing entities that are more vulnerable to surprise. The same can be said about evolutionary competition of biological species [17]. The squirrel that is evolutionarily programmed to make the right decisions about foraging, resting and mating will tend to survive instead of other squirrels. In many situations, info-gap robust-satisficing is a better bet than any other strategy for satisfying one's requirements [4 section 11.4] [8] [10].

Let's briefly consider a simplified example: public policy for managing an invasive biological species. The Light Brown Apple Moth (LBAM) is native to Australia. It was detected in California in 2007 though some entomologists claimed that it had been present and widespread (though undetected) for a long time. In other words, there was expert dispute over whether the LBAM was in fact an invasive species or not. Furthermore, there was very substantial dispute about the economic damage that could result from the moth, either from direct crop loss or due to inter-jurisdictional trade restrictions on California crops potentially carrying the LBAM. There was also dispute over the efficacy of various means of control or eradication of the LBAM, as well as over possible adverse effects of these interventions on public health. In short, the LBAM created major controversy in both the scientific community and the general public [14].

In our simplified example we will describe how info-gap robust-satisficing would be used to choose between various possible interventions (where doing nothing is one possibility). There are diverse risk factors that are poorly understood and whose adverse impact, in monetary terms, are highly uncertain. The analyst begins by asking what is the largest monetary loss that is tolerable (recognizing that this too may be uncertain). The outcome requirement is that the loss must not exceed this critical value. The analyst then identifies the data, knowledge and understanding—the models—that are relevant to the problem. (The models may contain probabilistic elements.) The models are used to predict the outcome of each intervention. Any intervention whose predicted outcome does not satisfy the outcome requirement is excluded. The vast uncertainties of the models are non-probabilistically represented using unbounded info-gap models of uncertainty. The robustness to uncertainty of each remaining intervention is evaluated. That is, the analyst determines, for each remaining intervention, how large an error in the models can be tolerated without jeopardizing the critical outcome requirement. The analyst adopts the intervention whose robustness is greatest for satisfying the requirement.

Three general characteristics of info-gap robust-satisficing are illustrated by this example.

First, the robustness of each possible intervention decreases as the outcome requirement (maximum tolerable monetary loss) becomes more demanding. This is the robustness trade off discussed in section 2.2.

Second, it could happen that all interventions that are examined have zero or very low robustness at the specified outcome requirement (maximum tolerable loss). This could induce the analyst to seek additional options, or to improve the models or to reduce their uncertainties. Alternatively, the outcome requirement might be declared unrealistic. The robustness trade off might then be used in revising the requirement.

Third, different interventions might be most robust (and hence preferred) for different values of the maximum tolerable loss. This preference reversal among the options is particularly common when choosing between a widely used "state of the art" option, and a new and innovative alternative. The innovation is predicted to provide a better outcome (lower loss) than the state of the art; that's what makes it attractive. However, the innovation is usually less well known and more uncertain than the state of the art which has been widely used. The innovation may have positive robustness for a very demanding requirement, at which the state of the art may even have zero robustness. However, the greater uncertainty of the innovation causes its robustness to fall below the robustness of the state of the art at less demanding requirements. The "innovation dilemma" facing the analyst is to choose between a more promising but more uncertain option (the innovation) and a less promising but less uncertain option (the state of the art). The innovation dilemma is resolved by choosing the more robust alternative for the required outcome [14].

6.3 Robust Satisficing and the Response to Ontological Indeterminism

From the perspective of classical physics, the impact of ontological indeterminism can be summarized by noting that the unique classical Lagrangian no longer determines a unique response (path) to a specified set of conditions. Feynman used the classical Lagrangian in his derivation of Schröderinger's equation, though he did not use it in the way that classical physicists used it. Classically, the Lagrangian determined a unique solution, while Feynman (and quantum mechanics) allows all paths to contribute, each according to the degree of stationarity of its action integral.

The robust-satisficing response to ontological indeterminism begins by hypothesizing that the Lagrangian is indeterminate; it does not take a specific functional form for any given physical system. Consequently, a least-action principle is no longer immediately applicable: there is no specific Lagrangian with respect to which the action can be stationary. Feynman's adoption of the classical Lagrangian in a quantum setting is to be viewed as an ad hoc synthesis.

In the context of robust-satisficing, the requirement that nature must satisfy is that the action integral must be nearly stationary. What Feynman proposed is that, quantum mechanically, path-bundles are favored whose action integrals are relatively stationary with respect to path variation; path-bundles with highly non-stationary action integrals are rare. We interpret this as a response to the ontological indeterminism of the Lagrangian: nature robustly satisfices the stationarity of the action integral.

Feynman's derivation of the Schrödinger equation remains unchanged, except that now it is the magnitude (or, more generically, stationarity) of the robustness, not stationarity of the action integral, that determines the quantum mechanical weight of a wave function. We hypothesize that the classical action depends on the path only through the robustness. Consequently, we obtain the same wave functions as in Feynman's method, but for a different reason. This puts ontological indeterminism in the same conceptual framework as epistemic indeterminism: maximize the robustness to indeterminism while satisfying a requirement. This interpretation of quantum mechanics also sheds light on issues of locality, realism and completeness [7].

Optimization is still central because stationarity of the robustness determines the equations of motion. But the robustness is not itself a property of the system, unlike the Lagrangian in both classical and quantum mechanical physics. The Lagrangian, that embodies specific properties of the system, is indeterminate, and only the robustness to this property-indeterminism is optimized in deriving equations of motion.

7 Conclusion

It is amazing that we can unravel the Rosetta stone, cosmological evolution, and the neurophysiology of primate brachiation. We discern patterns, mechanisms, causes, effects and other relations. Over time, much more will be discovered and understood. Some of our current ignorance, like the location of the Ten Lost Tribes, has been resolved indirectly by understanding the historical processes of conquest and assimilation. Other quests, like a cure for HIV/AIDS, still demand great effort and will be achieved, if at all, by ways unknown today. And some quests, like unifying our understanding of the physical world, seem to evade our best efforts.

It is no less amazing that the formless void of what we don't know—even what we cannot know scientifically—seems to have a handle that we can grasp conceptually. As a start, the indeterminism of the physical, biological, and human worlds is real. Ontologically, the world and its laws are not fixed and finalized. Epistemically, our ability to know is limited by the fact that tomorrow's new knowledge cannot be known today, from which the Shackle-Popper indeterminism establishes strong implications for explanation and prediction in human affairs. Second, both epistemic and ontological indeterminism are (to the best of our knowledge) unbounded: the epistemic gap between knowledge and possible knowledge, and the ontological gap between what has occurred and what will occur, has no limit. The potential for surprise is unbounded.

That these indeterminisms are manageable is evident from the simple fact that we, and nature, manage, either literally or metaphorically. Stars shine and people ramble along despite the limitation of classical causality and of human comprehension. *How* we manage is the amazing thing. Any attempt to explain this management is subject to the inevitable fallibility of all explanation [26] (and subject to the inevitable disputes among explainers). The suggestion in this essay is that an underlying common theme, in response to both epistemic and ontological indeterminism, is to robustly satisfice the achievement of specific requirements. The concept of robust-satisficing reflects what we can—and cannot—say about indeterminism. Nature does not have *one* Lagrangian that determines *the* response, nor do people know future discoveries and inventions. But we know that we don't know the future, and nature allows multiple paths as though each were, for the moment, *the* solution. At the risk (again) of anthropomorphism, we suggest that robust-satisficing "embraces" indeterminism: robustness is the way to "get along" in the face of inevitable surprise.

8 References

1. Adams, Henry, 1918, *The Education of Henry Adams,* edited by Ernest Samuels, Houghton Mifflin Co., Boston.
2. Appeltans, Ward, *et al.,* 2012, The magnitude of global marine species diversity, *Current Biology,* 22: 2189–2202.
3. Bell, David E., Raiffa, Howard and Tversky, Amos, eds., 1988, *Decision Making: Descriptive, Normative, and Prescriptive Interactions,* Cambridge University Press.
4. Ben-Haim, Yakov, 2006, *Info-Gap Decision Theory: Decisions Under Severe Uncertainty,* 2nd edition, Academic Press, London.
5. Ben-Haim, Yakov, 2007, Peirce, Haack and Info-gaps, pp.150–164 in [21].
6. Ben-Haim, Yakov, 2010, *Info-Gap Economics: An Operational Introduction,* Palgrave, London.
7. Ben-Haim, Yakov, 2011a, Optimizing and satisficing in quantum mechanics: An info-gap approach, working paper. http://www.technion.ac.il/yakov/lr05.pdf
8. Ben-Haim, Yakov, 2011b, When is non-probabilistic robustness a good probabilistic bet? Working paper. http://www.technion.ac.il/yakov/prx27a.pdf
9. Ben-Haim, Yakov, 2011c, Jabberwocky. Or: Grand Unified Theory of Uncertainty??? http://decisions-and-info-gaps.blogspot.com/2011/12/jabberwocky-or-grand-unified-theory-of.html
10. Ben-Haim, Yakov, 2012a, Robust satisficing and the probability of survival, *Intl. J. of System Science,* to appear.
11. Ben-Haim, Yakov, 2012b, Doing our best: Optimization and the management of risk, *Risk Analysis,* 32(8): 1326–1332.
12. Ben-Haim, Yakov, 2012c, Why risk analysis is difficult, and some thoughts on how to proceed, *Risk Analysis,* 32(10): 1638–1646.
13. Ben-Haim, Yakov and Karsten Jeske, 2003, Home-bias in financial markets: Robust satisficing with info-gaps, Federal Reserve Bank of Atlanta, Working Paper Series, 2003-35, Dec. 2003, pdf file on the Federal Reserve Bank website. SSRN abstract and full paper at: http://ssrn.com/abstract=487585.
14. Ben-Haim, Yakov, Craig D. Osteen and L. Joe Moffitt, 2013, Policy Dilemma of Innovation: An Info-Gap Approach, *Ecological Economics,* 85: 130–138.
15. Ben-Haim, Yakov, Nicola M. Zetola and Clifford Dacso, 2012, Info-gap management of public health policy for TB with HIV-prevalence, *BMC Public Health,* to appear.
16. Burgman, Mark, 2005, *Risks and Decisions for Conservation and Environmental Management,* Cambridge University Press, Cambridge.
17. Carmel, Yohay and Yakov Ben-Haim, 2005, Info-gap robust-satisficing model of foraging behavior: Do foragers optimize or satisfice?, *American Naturalist,* 166: 633–641.
18. Chapman, Arthur D., 2009, *Numbers of Living Species in Australia and the World,* 2nd ed. Australian Govt., Dept. of Environment, Heritage, Environment and the Arts, Canberra, AU.
19. Chinnappen-Rimer, S. and G.P. Hancke, 2011, Actor coordination using info-gap decision theory in wireless sensor and actor networks, *International Journal of Sensor Networks,* 10(4): 177–191.
20. Costello, Mark J., Robert M. May and Nigel E. Stork, 2013, Can we name earth's species before they go extinct? *Science,* 339: 413–416.
21. de Waal, Cornelis, ed., 2007, *Susan Haack, A Lady of Distinctions: The Philosopher Responds to Her Critics,* Prometheus Books, Amherst, New York.
22. Dirac, P.A.M., 1958, *The Principles of Quantum Mechanics,* 4th ed., Oxford University Press.
23. Feynman, Richard P., 1948, Space-time approach to non-relativistic quantum mechanics, *Reviews of Modern Physics,* vol.20, #2, pp.367–387.
24. Gould, Stephen Jay, 1989, *Wonderful Life: The Burgess Shale and the Nature of History,* W.W. Norton and Co., New York.

25. Gray, Colin S., 1990, *War, Peace and Victory: Strategy and Statecraft for the Next Century,* A Touchstone Book, New York.
26. Haack, Susan, 1998, *Manifesto of a Passionate Moderate: Unfashionable Essays,* University of Chicago Press, Chicago.
27. Habermas, Jürgen, 1970, *On the Logic of the Social Sciences,* trans. by Shierry Weber Nicholsen and Jerry A. Stark, 1990, MIT Press.
28. Hall, Jim W., Robert J. Lempert, Klaus Keller, Andrew Hackbarth, Christophe Mijere, and David J. McInerney, 2012, Robust climate policies under uncertainty: A comparison of robust decision making and info-gap methods, *Risk Analysis,* 32 (10): 1657–1672.
29. Harp, Dylan R. and Velimir V. Vesselinov, 2012, Contaminant remediation decision analysis using information gap theory, *Stochastic Environmental Research and Risk Assessment,* to appear.
30. Kanno, Y. and I. Takewaki, 2006, Robustness analysis of trusses with separable load and structural uncertainties, *International Journal of Solids and Structures,* 43(9): 2646–2669.
31. Keynes, John Maynard, 1936, *The General Theory of Employment, Interest, and Money,* Prometheus Books, 1997 reprint of the original 1936 edition by Harcourt Brace & World.
32. Knight, F.H., 1921, *Risk, Uncertainty and Profit,* Hart, Schaffner and Marx. Re-issued by Harper Torchbooks, New York, 1965.
33. Knoke, Thomas, 2008, Mixed forests and finance - Methodological approaches, *Ecological Economics,* Volume 65(3): 590–601.
34. Lanczos, Cornelius, 1970, *The Variational Principles of Mechanics,* 2nd edition, Dover Publications, New York.
35. Lempert, Robert J., Steven W. Popper and Steven C. Bankes, 2003, *Shaping the Next 100 Years: New Methods for Quantitative, Long-Term Policy Analysis,* RAND Corp. Santa Monica, CA.
36. Lennon, John and Paul McCartney, 1967, All You Need is Love, a song.
37. Lindgren, Astrid, *Pippi Longstocking.*
38. March, James G., 1988, Bounded rationality, ambiguity, and the engineering of choice, pp.33–57 in [3].
39. Mas-Colell, Andreu, Michael D. Whinston and Jerry R. Green, 1995, *Microeconomic Theory,* Oxford University Press, Oxford.
40. Mayo, Deborah G., 1996, *Error and the Growth of Experimental Knowledge,* University of Chicago Press, Chicago.
41. Moffitt, L. Joe, John K. Stranlund, and Barry C. Field, 2005, Inspections to avert terrorism: Robustness under severe uncertainty, *J. Homeland Security and Emergency Management,* vol. 2, no. 3. http://www.bepress.com/jhsem/vol2/iss3/3
42. Morse, Philip M. and Herman Feshbach, 1953, *Methods of Theoretical Physics.* Part I: Chapters 1 to 8, McGraw-Hill Book Co., New York.
43. Oxford English Dictionary, 1989, 2nd ed.
44. Nelson, Richard R. and Sidney G. Winter, 1982, *An Evolutionary Theory of Economic Change,* Belknap Press, Harvard University.
45. Popper, Karl R., 1957, *The Poverty of Historicism,* Harper Torchbooks edition 1964.
46. Popper, Karl, 1982, *The Open Universe: An Argument for Indeterminism.* From the Postscript to *The Logic of Scientific Discovery,* Routledge.
47. Savage, Leonard J., 1972, *The Foundations of Statistics,* 2nd revised edition. Dover
48. Schwartz, Barry, Yakov Ben-Haim, and Cliff Dacso, 2011, What Makes a Good Decision? Robust Satisficing as a Normative Standard of Rational Behaviour, *The Journal for the Theory of Social Behaviour,* 41(2): 209–227.
49. Shackle, G.L.S., 1972, *Epistemics and Economics: A Critique of Economic Doctrines,* Transaction Publishers, 1992, originally published by Cambridge University Press.

50. Simon, Herbert A., 1955, A behavioral model of rational choice, *Quarterly Journal of Economics,* vol.69, #1, pp.99–118.
51. Simon, Herbert A., 1956, Rational choice and the structure of the environment, *Psychological Review,* vol.63, no.2, pp.129–138.
52. Wald, A., 1945. Statistical decision functions which minimize the maximum risk, *Annals of Mathematics,* 46(2), 265–280.
53. Wittgenstein, Ludwig, *Tractatus Logico-Philosophicus,* trans. by D. R. Pears and B.F. McGuinness, 1961, based on the 1921 German original. Routledge.