Decision Making in Times of Knightian Uncertainty: An Info-Gap Perspective[1]

Yakov Ben-Haim, Technion—Israel Institute of Technology[2]
and
Maria Demertzis, De Nederlandsche Bank and European Commission[3]

Contents

Abstract

The distinction of risk vs uncertainty as made by Knight has important implications for policy selection. Assuming the former when the latter is relevant can lead to wrong decisions. With the aid of a stylized model that describes a bank's decision on how to allocate loans, we discuss decision making under Knightian uncertainty. We use the info-gap robust satisficing approach to derive a trade-off between confidence and performance (analogous to confidence intervals in the Bayesian approach but without assignment of probabilities). We show that this trade off can be interpreted as a cost of robustness. We show that the robustness analysis can lead to a reversal of policy preference from the putative optimum. We then compare this approach to the the min-max method which is other main non-probabilistic approach available in the literature.

\people\demertzis\3rd-paper2015\rupi008.tex. 22.6.2015.

[1] All views are our own and do not represent those of any of the institutions we are affiliated with.

[2] Yitzhak Moda'i Chair in Technology and Economics, Technion—Israel Institute of Technology, Haifa, 32000 Israel, Tel: 972-4-829-3262. yakov@technion.ac.il

[3] m.demertzis@gmail.com, +31 20 5242016, De Nederlandsche Bank, PO Box 98, 1000 AB Amsterdam, The Netherlands.

Keywords: Uncertainty vs risk, confidence, robustness, satisficing, info-gap

JEL Classification: C02, C18, D81, G10

1 Introduction

The economic circumstances since the start of the crisis in 2007 to the present are characterized by high levels of uncertainty. What do we mean by high uncertainty and what does it imply for policy design or decision making? High uncertainty can mean one of two things: either high stochastic volatility around known (or well estimated) average future outcomes, or at least partial ignorance about relevant mechanisms and potential outcomes. The first implies that uncertainty can be probabilistically measured (what Frank Knight called 'risk'), whereas the second implies that it cannot (what Knight called 'true uncertainty' and is now known as Knightian uncertainty). We often conflate these two concepts when discussing 'uncertainty' in general. However, it is crucial to distinguish between them for three reasons. First, the relevant methods for decision making depend on which of the two notions of 'high' uncertainty we address. Designing policies under the assumption of probabilistically measurable risk can lead to serious policy mistakes if the underlying uncertainty is non-probabilistic, Knightian. Second, one's measures of confidence differ under risk or Knightian uncertainty. Finally, the use of contextual understanding is different when dealing with risk or Knightian uncertainty. In a probabilistic setting contextual understanding can be used, for example, to select an appropriate probability distribution. In a Knightian setting contextual understanding can be used to intuit a trend or to sense a pending change that is not yet manifested in data.

This paper will make the following points:

- When uncertainty is probabilistically measurable risk, it is possible to design policies that are optimal on average or in some quantile sense. Policy design under risk is based on first principles as expressed by economic theory. The theory underlies policy choices that are designed to optimize specified substantive outcomes (e.g. minimize a high quantile of the inflation, maximize average growth, etc.).

- Under Knightian uncertainty it is not possible to optimize stochastic outcomes because at least some probabilities are unknown. Furthermore, it is unreliable to attempt to optimize substantive outcomes because the underlying models are poorly known. Instead, under Knightian uncertainty one aims to prevent bad results from occurring or at least prepare for them. Building buffers in the financial system, applying unorthodox monetary policies in the monetary system are policies of this type; they aim to provide intervention tools to deal with or prevent bad outcomes from arising, irrespective of how likely they might be.

- A non-probabilistic concept of robustness is used to evaluate the confidence in achieving an outcome under Knightian uncertainty. We will discuss info-gap robustness and compare it with the min-max robustness concept.

Decision making under risk relies on known probability distributions of outcomes. Policy design becomes then a question of identifying the most likely occurrence (or perhaps a quantile of the occurrence) given the underlying models, and applying measures that optimize the outcome. Risks around those most likely occurrences are described probabilistically, and confidence in one's actions in turn is best captured with statistical intervals. An obvious example is the forecasts that central banks present and the confidence intervals around them. The resulting *fan charts* (first used by the Bank of England) are stochastic simulations in underlying variables under assumed probability distributions. Confidence then is defined as the probability of ranges of events.

However, probabilities are measures of frequencies of events that have happened in the past, and therefore, in real time we are not necessarily confident that they represent accurate descriptions of the future. In 2007 most forecasts of, for example growth in most countries, were presenting confidence bands that were quite different from ex post outcomes. The corresponding confidence was no more than simply a false sense of certainty. Naturally, 2007 was the start of two of the most difficult years for forecasting in the past 20 to 30 years. Point estimates were revised both frequently and by substantive amounts and confidence bands were a mechanical tool void of economic significance. If one were to look at other times, one would not necessarily see equal revisions in forecasts and the corresponding confidence intervals would have been useful.

How could we have done it differently? The lesson that the 2007 exercise has taught us is that even though models do serve us satisfactorily most of the time, there will be times that they fail us, and they may even fail us spectacularly. It is on these occasions that probabilities do not provide reliable assessment of, or confidence about, the outcomes. Relying on them provides a false sense of security that can lead to wrong policy decisions. The problem is that the moments at which standard models fail us are moments of crisis and are not known in advance. And importantly, it is difficult to distinguish between times that models serve us well and times that they don't.[4] What does this mean for policy making or more generally for decision making? How can we evaluate confidence in these decisions? In this paper we provide an info-gap approach to decision making under Knightian uncertainty. With the aid of a simplified bank loan allocation example we will describe how the decision problem is handled in the presence of Knightian uncertainty. The info-gap approach will allow the bank to rank different portfolios in a way that it can pick those that provide satisfactory outcomes for the greatest range of adverse future contingencies. Robustness provides a measure of confidence.

The paper is organized as follows. Section 2 briefly reviews some literature in the economics of Knightian uncertainty. It discusses how policies change as we account for Knightian uncertainty. Section 3 uses a simple example of bank loan decisions to illustrate methodological implications of info-gap theory for decisions under Knightian uncertainty. Section 4 compares info-gap and min-max decision methodologies. Section 5 concludes.

[4]See Ahir, H. and P. Loungani, (2014) for a recent article on how difficult it is to predict turning points.

2 Risk versus uncertainty: Implications for policy making

2.1 Risk versus uncertainty

Frank Knight (1921) distinguished between ‘risk’ (for which probability distributions are known) and ‘true uncertainty’ (for which probability distributions are not known). Knightian uncertainty reflects ignorance of underlying processes, functional relationships, strategies or intentions of relevant actors, future events, inventions, discoveries, surprises and so on. Info-gap models of uncertainty provide a non-probabilistic quantification of Knightian uncertainty (Ben-Haim, 2006, 2010). An info-gap is the disparity between what you *do know* and what you *need to know* in order to make a reliable or responsible decision. An info-gap is not ignorance *per se,* but rather those aspects of one’s Knightian uncertainty that bear on a pending decision and the quality of its outcome. Under risk we are confident—at least probabilistically—of the underlying model or combination of models that describe the economy. By contrast, under Knightian uncertainty, the social planner lacks important knowledge of how the system works. The planner starts with a number of models that may be relevant, but cannot identify the likelihood with which they describe the economy. When designing policy under risk, the knowledge of underlying probability distributions permits the identification of policies that are optimal on average or satisfy other quantile-optimality requirements. This is not possible under Knightian uncertainty because one lacks knowledge of the underlying distributions. But if one cannot design policy based on the principle of outcome-optimality, what other principles can one follow and what would these policies look like?

Two approaches have been widely used as alternatives to outcome-optimization based on a reliably known (possibly probabilistic) model: 1) robust control (also called min-max) and 2) info-gap. Neither requires knowledge of probabilities. The overarching principle behind these two approaches is to find policies that are robust to a range of different contingencies.

The literature on robust control relies on identifying and then ameliorating worst outcomes (Hansen et al 2006, Sargent and Hansen 2008 and Williams 2007). The planner considers a family of possible models, without assigning probabilities to their occurrence. Then that model is identified which, if true, would result in a worse outcome than any other model in the family. Policy is designed to minimize this maximally bad outcome (hence ‘min-max’ is another name for this approach). The appeal of this technique is that it provides insurance against the worst anticipated outcome. However, this technique has also been criticized for two main reasons. First, it is unnecessarily costly to assume that the worst will happen all the time (irrespective of how it is defined). Second, the worst may be expected to happen rarely and therefore it is an event that planners know the least about. It is odd to focus the policy analysis on an event that is the least known (Sims 2001). Confidence in robust control is not measured explicitly. It manifests itself in the following form: the planner will have maximally ameliorated the worst that is thought to be possible. The optimization is not of the substantive outcome (growth, employment, etc.) but rather of ameliorating adversity. In this sense min-max is robust to uncertainty.

The second approach is called info-gap (Ben-Haim 2006, 2010) and relies on the principle of robust satisficing.[5] The principle of satisficing is one in which the planner is not aiming at best outcomes. Instead of maximizing utility or minimizing worst outcomes, the planner aims to achieve an outcome that is *good enough*. For example, the planner tries to assure that loss is not greater than an acceptable level, or growth is no less than a required level. When choosing between alternative policies, the robust-satisficing planner will choose the policy that will satisfy the critical requirement over the greatest spectrum of models.[6]

Min-max and info-gap methods are both designed to deal with Knightian uncertainty, but they do so in different ways. The min-max approach requires the planner to identify a range of events and processes that could occur, acknowledging that likelihoods cannot be ascribed to these contingencies. The min-max approach is to choose the policy for which the contingency with the worst possible outcome is as benign as possible: ameliorate the worst case. The info-gap robust-satisficing approach requires the planner to identify the worst consequence that can be tolerated, and to choose the policy whose outcome is no worse that this, over the widest possible range of contingencies. Both min-max and info-gap require a prior judgment by the planner: identify a worst model or contingency (min-max) or specify a worst tolerable outcome (info-gap). However, these prior judgments are different, and the corresponding policy selections may, or may not, agree.[7]

2.2 How do policies change as we account for uncertainty?

A vast literature has analyzed how policies designed to handle risk differ from those designed to handle Knightian uncertainty. In the case of designing policy under risk the most famous result is that of Brainard in his seminal paper (Brainard 1967) in which he showed that accounting for Bayesian uncertainty, in a specific class of problems, implies that policy will be more cautious.[8] In terms of policy changes it therefore means smaller but possibly more persistent steps, and is known as the 'Brainard attenuation' effect. At the limit, as risk becomes very large, the social planner abandons the use of the instrument and is faced with policy inaction.[9] As the social planner is more and more uncertain of the results of policy, it is used less and less. This result has been very popular with policy makers as it appeals to their sense of caution when they lack sufficient information or knowledge.[10]

[5] The technical meaning of "satisficing" as "to satisfy a critical requirement" was introduced by Herbert Simon (1955, 1957, 1997).

[6] Satisficing is a strategy that seems to maximise the probability of survival of foraging animals in adverse conditions (i.e. uncertainty) Carmel and Ben-Haim 2005. There are circumstances for which this can also be proven for economic examples (see Ben-Haim and Demertzis, 2008).

[7] Further discussion of this comparison appears in Ben-Haim, Dacso, Carrasco and Rajan, 2009. See also section 4 here.

[8] This is for uncertainty in the coefficients that enter the model multiplicatively, not the residuals which enter the model additively.

[9] To be fair, this attenuation effect does not hold always but also depends on the cross-correlations of error terms in the assumed model. It is possible therefore, that the policy is more aggressive than that under no uncertainty and Brainard did acknowledge that.

[10] As Blinder (1988, p.12) wrote, there tends to be "a little stodginess at the central bank."

By contrast, policies derived under the principle of min-max (or robust control), and directed against non-probabilistic uncertainty, tend to be comparatively more aggressive. The policy steps taken are typically larger in size by comparison to either risk-based policies or outcome-optimal policies in the absence of uncertainty. The intuition is that under Knightian uncertainty, and when addressing a worst case, there is little knowledge about the transmission mechanisms, and it is therefore important to strongly exercise available tools in order to learn about and manage the economy. It is not surprising that this runs against some policy makers' natural inclination to be cautious and avoid introducing volatility.

It is here that info-gap robust satisficing provides a useful operational alternative. At the heart of the method for dealing with uncertainty lies a fundamental choice: that between robustness against uncertainty and aspiration for high-value outcomes. As we become more ambitious in our aspirations, we need to compromise in the degree of confidence that we can have about achieving these aspirations. Conversely, if we require high confidence in achieving specified goals, then we need to reduce our ambitions. Info-gap is a method developed with the specific aim of capturing this trade-off. Confidence is quantified with robustness to uncertainty. The trade-off quantifies the degree of robustness with which one can pursue specified outcome requirements. Policies therefore are not automatically more or less aggressive. It depends very much on the decision maker's preferences. Furthermore, the decision maker can rank alternative policies: between policies of similar ambitions, those that provide the greater robustness (greater confidence) are preferred. In section 4 we will compare and contrast the policy implications of min-max with robust-satisficing.

3 An informative trade-off: Robustness vs performance

In this section we use a highly simplified example to illustrate how a decision maker deals with the inability to measure uncertainty, to come to informed decisions. We provide a framework, based on info-gap theory, that allows us to derive a trade-off between confidence in outcomes and performance requirements. Decision makers who are ambitious in terms of requiring high-performance outcomes will have to settle for their choices being appropriate only across a small range of events or contingencies (i.e. having low robustness). On the other hand, if the decision maker wants the comfort of knowing that policies chosen will function across a wide range of contingencies (high robustness), then relatively low performance outcomes will have to be accepted.

Consider a bank that aims to give out loans to potential borrowers. Part of the problem that it faces is that the premium it requires depends on the risk type of the recipient agents, where risk here refers to their likelihood to default. However, assessing this probability is subject to Knightian uncertainty and therefore the bank is not in the position to price risk based on well defined underlying distributions. Furthermore, correlations exist between the solvencies of different borrowers which are significant even when they are small. Inter-borrower correlations are typically assumed to be zero though this is quite uncertain, potentially leading to over-optimistic estimates of bank invulnerability (Ben-Haim, 2010, section 4.1). In evaluating or designing the bank's loan portfolio, the following two questions (among others) are pertinent. First, some of the uncertainty in assessing default probabilities can be reduced. How much reduction in uncertainty is needed to substantially increase the bank's confidence? How should uncertainty-reduction effort be allocated among different borrower profiles as characterized by their estimated default probabilities? Second, what loan-repayment programs should be used for clients with different default-probability profiles? We describe how info-gap can help banks to allocate loans and, in section 4, compare it with the robust control (min-max) approach.

3.1 Formulation

Consider a bank that plans a number of loans, all with the same duration to maturity. The potential borrowers are of different risk types but all borrowers of the same risk type are identical.

Let:

N :	number of years to loan maturity
K :	number or risk-types
f_{kn} :	repayment in year n of risk type k
f :	matrix of f_{kn} values
w_k :	number (or fraction) of loans of risk-type k
w :	vector of w_k values
N_d :	number of years at which default could occur
t_j :	year at which default could occur, for $j = 1, \ldots, N_d$
p_{kj} :	probability that a client of risk-type k will default at year t_j
p :	matrix of default probabilities p_{kj}
i :	discount rate on loans

In case of default at t_j, no payment is made in that year and in all subsequent years, for $j = 1 \ldots N_d$. We define $t_{N_d} = N + 1$, so "default" at year t_{N_d} means that the loan is entirely repaid and default has not occurred. We also assume that $p_{k1} \ldots p_{kN_d}$ is a normalized probability distribution, so that the probability that borrowers of risk-type k do not default is:

$$p_{kN_d} = 1 - \sum_{j=1}^{N_d - 1} p_{kj}. \tag{1}$$

The present worth (PW) of the entire loan portfolio, assuming no defaults, is:

$$PW = \sum_{n=1}^{N} (1+i)^{-n} \sum_{k=1}^{K} w_k f_{kn}, \tag{2}$$

The no-default present worth of a single loan of risk-type k is:

$$\widetilde{PW}_k = \sum_{n=1}^{N} (1+i)^{-n} f_{kn}. \tag{3}$$

Eqs. (2) and (3) can be combined to express the total no-default present worth as:

$$PW = \sum_{k=1}^{K} w_k \widetilde{PW}_k. \tag{4}$$

We first formulate the probabilistic expected value of the present worth. We then define the info-gap uncertainty of the probabilistic part of the model. The expected PW of a single loan of risk-type k is:

$$
\begin{aligned}
\mathrm{E}(PW_k) &= \sum_{j=1}^{N_d-1} p_{kj} \sum_{n=1}^{t_j-1} (1+i)^{-n} f_{kn} + \left(1 - \sum_{j=1}^{N_d-1} p_{kj} \right) \sum_{n=1}^{N} (1+i)^{-n} f_{kn} && (5) \\
&= \sum_{n=1}^{N} (1+i)^{-n} f_{kn} - \sum_{j=1}^{N_d-1} p_{kj} \underbrace{\sum_{n=t_j}^{N} (1+i)^{-n} f_{kn}}_{\widetilde{PW}_{kj}} && (6) \\
&= \widetilde{PW}_k - \sum_{j=1}^{N_d-1} p_{kj} \widetilde{PW}_{kj} && (7)
\end{aligned}
$$

where $\widetilde{PW}_k$ is defined in eq.(3) and $\widetilde{PW}_{kj}$ is defined in eq.(6).

From eq.(7) we obtain the following expression for the expected PW of the entire portfolio:

$$
\mathrm{E}(PW) = \sum_{k=1}^{K} w_k \left(\widetilde{PW}_k - \sum_{j=1}^{N_d-1} p_{kj} \widetilde{PW}_{kj} \right). \quad (8)
$$

We note that the expected present worth, $\mathrm{E}(PW)$, depends on the distribution of risk types, expressed by the vector w, and on the repayment plans for the various risk types, expressed by the matrix f, and on the matrix, p, of default probabilities.

3.2 Info-gap uncertainty and robustness

The info-gap model for uncertainty in the default probabilities employs estimated default probabilities, $\tilde{p}_{kj}$. Each estimated probability is accompanied by an assessment of its accuracy, s_{kj}, expressing a judgment such as “The probability could be about $\tilde{p}_{kj} = 0.02$ plus or minus $s_{kj} = 0.07$ or more.”[11] This judgment of the error could come from an observed historical variation but, under Knightian uncertainty, the past only weakly constrains the future and the error estimate does not entail probabilistic information (such as defining a confidence interval with known probability). Or the error estimate could be a subjective assessment based on contextual understanding. The error estimate s_{kj} does not represent a maximal possible error, which is unknown. s_{kj} describes relative confidence in the various probability estimates and does not imply anything about likelihoods.

There are many types of info-gap models for representing Knightian uncertainty (Ben-Haim 2006, 2010). The following info-gap model is based on the idea of unknown fractional error of the estimates, and is applied to the uncertain probabilities of default. This info-gap model is an unbounded family of nested sets, $\mathcal{U}(h)$, of probability distributions p.

Definition 1 *Info-gap model of uncertainty.* For any value of h, the set $\mathcal{U}(h)$ contains all mathematically legitimate probability distributions whose terms deviate fractionally

[11] Subject of course to the probabilistic requirements of non-negativity and normalization.

from their estimates by no more than h:

$$\mathcal{U}(h)=\left\{p:\ p_{kj}\geq 0,\ \sum_{j=1}^{N_d}p_{kj}=1,\ |p_{kj}-\tilde{p}_{kj}|\leq s_{kj}h,\ \forall\ k,\ j\right\},\quad h\geq 0. \qquad (9)$$

The value of the fractional error, h, is unknown, and the range of uncertainty in p increases as h increases, thus endowing h with its name: horizon of uncertainty. The info-gap model of uncertainty in eq.(9) is not a single set. It is an unbounded family of nested sets. This is expressed formally in eq.(9) by the statement "$h \geq 0$". The horizon of uncertainty, h, is unbounded: we do not know by how much the estimated probabilities err. The info-gap model of uncertainty underlies the evaluation of robustness that we discuss shortly.

The performance requirement, at the bank's discretion, is that the expected value of the present worth be no less than a critical value PW_c:

$$\mathrm{E}(PW)\geq PW_c. \qquad (10)$$

Definition 2 *Info-gap robustness.* Robustness is the greatest horizon of uncertainty up to which the expected present worth of the portfolio is guaranteed to be no less than the critical value PW_c, i.e.:

$$\hat{h}(PW_c,w,f)=\max\left\{h:\ \left(\min_{p\in\mathcal{U}(h)}\mathrm{E}(PW)\right)\geq PW_c\right\}. \qquad (11)$$

Robustness is the greatest value of h up to which eq.(10) will be fulfilled for all realizations of p in $\mathcal{U}(h)$.

If probability estimates $\tilde{p}_{kj}$ were accurate (i.e. no Knightian uncertainty), then the bank would be able to give out loans in ways that would maximize the expected present worth of the portfolio. As these estimates become unreliable due to Knightian uncertainty, the bank becomes less confident that the loans would achieve the ex ante expected present worth. Intuitively, the robustness in eq.(11) answers the following question: how wrong can the estimated probability $\tilde{p}_{kj}$ be, in units of s_{kj}, and still achieve outcomes that are no worse than PW_c?[12] It will be evident shortly that, if the bank wants higher confidence in the sense that its choices are robust to a larger range of probability outcomes, then it will have to settle for lower critical present worth PW_c. Appendix A derives the robustness function for the special case where borrowers can default only at the mid point to maturity. Through explicit parameterization we can then compare different portfolios so that the bank can choose between them, to either reduce uncertainty or improve outcomes.

An interim summary of the info-gap robustness idea is as follows. Robustness to uncertainty is good to have, but it is also necessary to ask: how much robustness is sufficient? More robustness is obviously better than less, but the crucial question is: at what cost? It is this cost that is most important to the policy maker and this is where info-gap theory is helpful. Policy makers have views on what policy outcomes they want, and what outcomes they simply cannot tolerate. Quantifying the trade-off between robustness and outcome enables policy makers to make informed decisions, as we illustrate with an example.

[12] Note that the error estimates s_{kj} are somewhat analogous to deviations around the mean in the Bayesian case, but without employing probabilities.

3.3 Numerical Example

3.3.1 Formulation

The bank is designing a loan portfolio for ($N=$)10-year loans and has identified low and high-risk potential borrowers (therefore $K=2$). The bank must decide what fractions of its portfolio to loan to low- and high-risk clients. These fractions are denoted w_1 and w_2 respectively.[13] The bank must also specify the annual repayment schedule for low- and high-risk clients, denoted by $f_{1,1},\ \ldots,\ f_{1,10}$ for low-risk clients and by $f_{2,1},\ \ldots,\ f_{2,10}$ for high-risk clients. That is, client of risk-type k returns the sum $f_{k,n}$ at the end of the n-th year.

If default were not a possibility, then the bank could assess any proposed portfolio by evaluating the discounted present worth (PW) based on a minimal acceptable rate of return. However, default is definitely possible, though assessing the probability of default for each risk type, at each time step, is highly uncertain. The bank has made estimates for default probabilities at the mid-point of the loan maturity (therefore $t_1=5$ is the single potential default time and $N_d=2$). The 10-year repayment plan for the low-risk clients is constant at $f_{1,n}=0.1$ for $n=1,\ldots,10$. We consider two different 10-year repayment plans for the high-risk clients. Both plans decrease in equal increments over time. The first high-risk plan is $f_2^{(1)}=(0.12,\ldots 0.08)$ and the second high-risk plan is $f_2^{(2)}=(0.14,\ldots 0.10)$. The total nominal repayments for f_1 and $f_2^{(1)}$ are the same, while the total nominal repayment for $f_2^{(2)}$ is greater. Further we assume that:

- The discount rate is $i=0.07$.
- The vector of estimated default probabilities is $\tilde{p}=(0.02,0.05)$. Thus the high-risk clients are assumed to be two and a half times as likely to default as the low-risk clients but these are highly info-gap-uncertain: the true values may be much better or much worse.[14]
- We consider two different vectors of error estimates of these probabilities, corresponding to lower and greater precision in the estimated probabilities. The lower-precision case is $s^{(1)}=(0.10,0.15)$ and the higher-precision case is $s^{(2)}=(0.05,0.08)$. Knightian uncertainty accompanies all probability estimates, and these error estimates appear in the info-gap model of eq.(9).
- We consider two different risk-type distributions, expressed by the vector w. The preponderantly low-risk distribution is $w^{(1)}=(0.7,0.3)$, and this will be used in the case where the estimated default probabilities are less well know, as expressed

[13] Note that under no uncertainty, the bank would be able to allocate optimally w_1,w_2 by demanding a repayment that leaves it indifferent between the two types of borrowers. In the presence of uncertainty it cannot do that though and needs to consider alternative portfolios.

[14] There are of course other relevant uncertainties, such as delayed or partial payments, correlations between client defaults, etc. Furthermore, the bank may wish to evaluate a proposed portfolio with a quantile analysis of the PW rather than with the expected PW. This simple example—illustrating the info-gap robust-satisficing methodology and to comparing it with the min-max approach—will ignore these additional issues.

by $s^{(1)}$. The preponderantly high-risk distribution is $w^{(2)} = (0.3, 0.7)$, to be used with $s^{(2)}$.

- We consider two different portfolios, $\mathcal{P}_1 = (w^{(1)}, f_2^{(1)}, s^{(1)})$ and $\mathcal{P}_2 = (w^{(2)}, f_2^{(2)}, s^{(2)})$.

The concept of Knightian uncertainty is quantified, in info-gap theory, with an unbounded family of nested sets of possible realizations of the uncertain entity. In the example discussed in this section, the default probabilities are uncertain and eq.(9) is the info-gap model of uncertainty. The bank has estimates of these probabilities for each client risk-type, as well as error measures of these estimates, though these error measures are insufficient to specify probabilistic confidence intervals, and do not specify maximum possible error. The basic intuition of the info-gap model of uncertainty is that the fractional error of each estimated probability is bounded, but the value of this bound is unknown. That is, the analyst has probability estimates, knows the errors of these estimates are bounded, but does not know the magnitude of the bound. In other words, a worst case cannot be identified.[15]

We now explain the idea of robustness. The default probabilities are unknown and the estimates are highly uncertain. However, we are able to assess any proposed portfolio by asking: how large an error in the estimated default probabilities can occur without causing the PW to fall below an acceptable level? That is, how robust is the proposed portfolio to error in the estimated default probabilities? If a proposed portfolio is highly robust, then acceptable PW will be obtained even if the estimated default probabilities err greatly. Such a portfolio would be more attractive than one whose robustness is low and thus highly vulnerable to error in the estimates. In other words, portfolios are prioritized by their robustness for satisfying a PW criterion, not by their predicted PW.

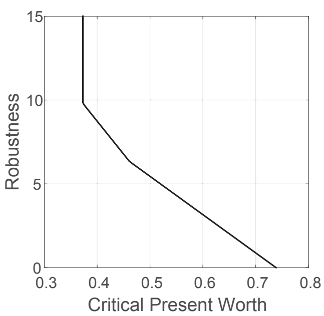

Figure 1: Robustness curve for loan portfolio $\mathcal{P}_1$.

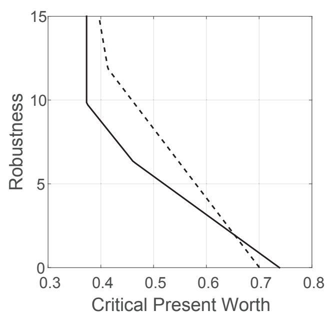

Figure 2: Robustness curves for loan portfolios $\mathcal{P}_1$ (solid) and $\mathcal{P}_2$ (dash).

3.3.2 A robustness curve

Fig. 1 shows a robustness curve for portfolio $\mathcal{P}_1$. The horizontal axis is the critical present worth: the lowest value of PW that would be acceptable to the bank. (The

[15] One could argue that default probabilities all equal to unity is the worst possible case. That is true by definition but does not reflect the bank's knowledge of its specific situation.

PW has the same units as the client repayments, f_{kn}.) The vertical axis is the robustness: the greatest fractional error in the estimated probabilities of default that do not jeopardize the corresponding critical PW. For instance, at a critical PW of 0.6, the estimated default probabilities can err by a factor of 3 without jeopardizing the PW requirement.

Three concepts can be illustrated with this figure: trade off, cost of robustness, and zeroing. The negative slope demonstrates that the robustness decreases as the required PW increases. This expresses a *trade off:* as the requirement becomes more demanding (as critical PW increases) the robustness becomes lower. More demanding requirements are more vulnerable to Knightian uncertainty than lax requirements. This is a generic property of info-gap robustness functions, and is sometimes called "the pessimist's theorem".

The curve in fig. 1 expresses this trade off quantitatively, and the slope can be thought of as a *cost of robustness.* A very steep negative slope implies that the robustness increases dramatically if the requirement, critical PW, is slightly reduced, implying a low cost of robustness. A gradual negative slope implies the opposite and entails large cost of robustness. From fig. 1 we see that the cost of robustness is relatively high when the critical PW is large (lower right). The cost of robustness actually becomes zero at the upper right when the slope is infinite. The robustness rises to infinity at low values of critical PW in fig. 1. Specifically, the robustness is infinite if the required present worth is less than the least possible value (this least possible value occurs when all risk-types default at midterm).

At the lower right end of the graph in fig. 1 we see that the robustness vanishes for large critical values of PW. More precisely, the robustness is zero if the required present worth exceeds the value based on the estimated default probabilities (eq.(18) in the appendix). This is called the *zeroing* property and it states that a required PW that exceeds the estimated PW has no robustness to Knightian uncertainty because default probabilities may exceed the estimated values. While this is perhaps not surprising, it entails two methodological conclusions. First, the estimated PW should *not* be used to prioritize alternative portfolios (because the estimated value has no robustness against Knightian uncertainty). Second, the Knightian uncertainty may in fact motivate a *preference reversal.* We explore these two methodological conclusions in fig. 2, where we plot both portfolios $\mathcal{P}_1$ (solid curve) and $\mathcal{P}_2$ (dashed curve).

3.3.3 Which portfolio for the bank's performance requirement? A preference reversal

Figure 2 shows robustness curves for portfolios $\mathcal{P}_1$ (solid) and $\mathcal{P}_2$ (dash). In $\mathcal{P}_1$ the preponderance of clients are low-risk, while in $\mathcal{P}_2$ the preponderance are high-risk. The estimated default probabilities are the same for both portfolios, but less effort was invested in verifying the estimates for $\mathcal{P}_1$ than for $\mathcal{P}_2$, which might be justified by noting that the preponderance of clients in $\mathcal{P}_1$ are low-risk in any case. The repayment plan for low-risk clients are constant in time and the same in both portfolios. The repayment plans for high-risk clients decrease in time by moving more of the debt to early payments. Furthermore, in $\mathcal{P}_2$ the repayments are greater than in $\mathcal{P}_1$.

Fig. 2 shows that robustness vanishes at a greater value of critical PW for $\mathcal{P}_1$ than for $\mathcal{P}_2$, as seen from the horizontal intercepts of the robustness curves. From the zeroing property, this means that $\mathcal{P}_1$'s estimated PW is greater than $\mathcal{P}_2$'s. If these estimates were reliable (which they are not due to the Knightian uncertainty) then we would be justified in preferring $\mathcal{P}_1$ over $\mathcal{P}_2$. The Knightian uncertainty motivates the first methodological conclusion: do not prioritize portfolios according to their estimated PWs.

The predicted PWs are not a basis for portfolio selection because those predictions have zero robustness. Hence, we "migrate" up the robustness curve, trading critical PW for robustness. At the lower right end of the curves we see that the cost of robustness is greater for $\mathcal{P}_1$ than for $\mathcal{P}_2$ ($\mathcal{P}_2$ has steeper negative slope). The differences in slopes and intercepts result in crossing of the robustness curves. This creates the possibility for a *reversal of preference* between the two portfolios. For instance, suppose the bank requires a PW no less than 0.7. From the solid curve we see that the robustness of $\mathcal{P}_1$ is 1.0 which exceeds the robustness of $\mathcal{P}_2$ which is 0 at this PW requirement. The robust-satisficing decision maker would prefer $\mathcal{P}_1$. However, if the bank can accept a PW of 0.6, then $\mathcal{P}_2$ is more robust against Knightian uncertainty than $\mathcal{P}_1$. The robust-satisficing prioritization would now prefer $\mathcal{P}_2$ over $\mathcal{P}_1$. The robust-satisficing method implies that the prioritization depends on the decision maker's outcome requirement, and thus may change as that judgment changes.

It is important to understand why this reversal occurs. Portfolio $\mathcal{P}_1$ has relatively more low-risk clients than portfolio $\mathcal{P}_2$. Consequently, given the parameterization assumed, $\mathcal{P}_1$ would generate higher expected present worth if there were no Knightian uncertainty and would be the portfolio to choose. However, it is also the portfolio that is less precisely measured. As discussed above, more effort has gone into estimating default probabilities for portfolio $\mathcal{P}_2$ as expressed by the lower s. In other words, while $\mathcal{P}_2$ would be worse than $\mathcal{P}_1$ if there were no Knightian uncertainty, the assessment of $\mathcal{P}_2$ is less uncertain. Thus $\mathcal{P}_2$ has lower estimated expected present worth (intercept further left), but $\mathcal{P}_2$ also has lower cost of robustness (steeper slope). In short, there is a dilemma in the choice between $\mathcal{P}_1$ and $\mathcal{P}_2$. The dilemma is manifested in the crossing of the robustness curves. This crossing has the effect that, for moderate ambitions (anything below $PW_c = 0.65$), portfolio $\mathcal{P}_2$ satisfies these ambitions for a greater range of default probabilities. The choice between the portfolios (and the resolution of the dilemma) depends on the decision maker's choice of the critical present worth. A value less than 0.65 is more robustly achieved with $\mathcal{P}_2$ and this portfolio would be chosen, while a value greater than 0.65 would lead to choosing $\mathcal{P}_1$.

4 Robust satisficing vs min-max

We now use figs. 3 and 4 to compare the min-max and robust-satisficing decision methodologies, identifying situations in which they agree or disagree. We explained earlier that min-max and robust-satisficing require different judgments to be made by the decision maker. Min-max requires specification of worst case probability estimates, which is equivalent to assessing a maximum possible uncertainty (vertical axis). Robust-satisficing requires identification of a worst acceptable outcome (hori-

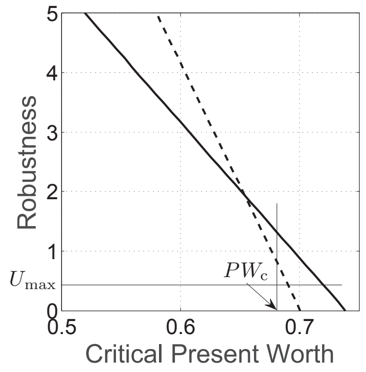

Figure 3: Robust-satisficing and min-max agree.

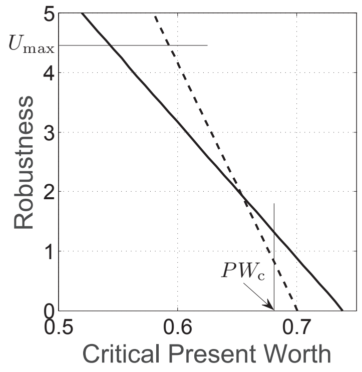

Figure 4: Robust-satisficing and min-max disagree.

zontal axis). Let $U_{\max}$ denote the min-max assessment of the maximum uncertainty, and let PW_c denote the robust-satisficing lowest acceptable PW.

Figs. 3 and 4 shows robustness curves for portfolios $\mathcal{P}_1$ (solid) and $\mathcal{P}_2$ (dash), from the lower-right portion of fig. 2. A robust-satisficing decision maker's least acceptable present worth, PW_c, is labeled on the horizontal axis. The thin vertical line on fig. 3 shows that this analyst would prefer $\mathcal{P}_1$ (solid) over $\mathcal{P}_2$ because $\mathcal{P}_1$ is more robust against Knightian uncertainty for this requirement. A min-max decision maker's maximum possible uncertainty, $U_{\max}$, is labeled on the vertical axis. The thin horizontal line shows that this analyst would also prefer $\mathcal{P}_1$ (solid) over $\mathcal{P}_2$ because the worst outcome at $U_{\max}$ is better with $\mathcal{P}_1$. The min-maxer and the robust-satisficer agree on the prioritization of the portfolios, but for different reasons because their initial judgments differ. The min-maxer tries to ameliorate the maximal uncertainty; the robust-satisficer tries to achieve no less than an acceptable outcome.

Fig. 4 shows the same robustness curves but with a different judgment by the min-max analyst, who now identifies greater maximum possible uncertainty. The min-maxer now prefers $\mathcal{P}_2$ (dash) because, at this larger $U_{\max}$, the worst outcome for $\mathcal{P}_2$ is better than for $\mathcal{P}_1$. The robust-satisficer would probably not dispute that uncertainty could be as great as $U_{\max}$. However, portfolio $\mathcal{P}_2$ is less robust than $\mathcal{P}_1$ for the specified critical outcome PW_c, so the robust-satisficer still prefers $\mathcal{P}_1$ (solid). Now min-max and robust-satisficing prioritize the portfolios differently.

The central ideas illustrated in this example are zeroing, trade off, preference reversal, and the situations in which min-max and robust-satisficing agree or disagree. *Zeroing* states that predicted outcomes (estimated PW in our example) have no robustness against Knightian uncertainty and therefore should not be used to prioritize the options. *Trade off* means that robustness increases as the performance requirement becomes less demanding. Robustness can be "purchased" by reducing the requirement, and the slope of the robustness curve quantifies the cost of robustness. The potential for *preference reversal* between options arises when their robustness curves cross each other. The robust-satisficing analyst's preference between the options depends on the outcome requirement. Finally, min-max and robust-satisficing both attempt to manage non-probabilistic Knightian uncertainty, but they are based on different initial judgments by the analyst, and they may either agree or disagree on prioritization of the options.

5 Conclusion

We have explored some of the implications of Knightian uncertainty for policy selection. Our main claim is that Knight's non-probabilistic "true uncertainty" requires very different management than is required for handling probabilistic risk. We used a simplified bank-loan example to illustrate the method of info-gap robust satisficing, and we compared this with the method of min-max. Both methods are non-probabilistic and both employ concepts of robustness. The choice between these methods hinges on the prior judgments that the analyst can make. Info-gap robust satisficing requires specification of outcome requirements (e.g. minimum acceptable present worth, or maximum acceptable unemployment, etc.). Info-gap robust-satisficing requires the decision maker to specify performance requirements. In contrast, min-max focuses on judgment of uncertainty and requires the analyst to specify the worst contingency. The min-max method then ameliorates this worst case, and does not require specification of an outcome requirement. Info-gap, in turn, does not presume knowledge of a worst case.

The info-gap robust satisficing methodology quantifies an irrevocable trade-off between confidence (expressed as robustness to uncertainty) and performance (embodying the decision maker's outcome requirement). This trade-off can be interpreted as a cost of robustness: robustness can be enhanced in exchange for reducing the performance requirement. The robustness curve characterizes any proposed policy as a monotonic plot of robustness versus performance requirement, where the slope reflects the cost of robustness and the horizontal intercept reflects the putative error-free outcome.

If the robustness curves of two alternative policies do not cross one another, then one policy is more robust than the other for all feasible outcomes. That robust-dominant policy is preferred. In this case, the putative optimum policy (whose estimated outcome is better) is also the robust-preferred policy.

If the robustness curves of two alternative policies cross one another, as seen in fig. 2, then the robustness analysis can lead to a reversal of policy preference from the putative optimum. The policy that is more robust (and hence preferred) at high performance requirement, will be less robust (and hence not preferred) at lower requirement. Info-gap robust-satisficing leads to policy selection that will achieve the performance requirement over the greatest range of Knightian uncertainty.

References

[1] Ahir, H.. and P. Loungani, 2014. There will be growth in the spring: How well do economists predict turning points?, Vox article April 14.

[2] Ben-Haim, Yakov 1999. Set-models of information-gap uncertainty: Axioms and an inference scheme, Journal of the Franklin Institute, 336: 1093–1117.

[3] Ben-Haim, Yakov 2006. *Info-Gap Decision Theory: Decisions Under Severe Uncertainty,* 2nd Edition Academic Press, London.

[4] Ben-Haim, Yakov 2010. *Info-Gap Economics: An Operational Introduction,* Palgrave.

[5] Ben-Haim, Yakov, Clifford C. Dacso, Jonathon Carrasco and Nithin Rajan, 2009, Heterogeneous Uncertainties in Cholesterol Management, Intl J Approximate Reasoning, 50: 1046–1065.

[6] Ben-Haim Yakov and M. Demertzis 2008. Confidence in Monetary Policy, De Nederlandsche Bank Working Paper, No. 192, December.

[7] Blinder, Alan S. 1998. *Central Banking in Theory and Practice,* Lionel Robbins Lecture, MIT Press, Cambridge.

[8] Brainard, W. 1967. Uncertainty and the Effectiveness of Policy, American Economic Review, 57, pp 411–425.

[9] Carmel Y. and Yakov Ben-Haim 2005. Info-gap robust-satisficing model of foraging behavior: Do foragers optimize or satisfice?, American Naturalist, 166: 633-641.

[10] Hansen, L. P., T. J. Sargent, G. A. Turmuhambetova, and N. Williams 2006. Robust control and model misspecification, Journal of Economic Theory 128, 45–90.

[11] Hansen, L.P. and T.J. Sargent 2008. *Robustness,* Princeton University Press, Princeton.

[12] Knight, F. H. 1921 Risk, Uncertainty, and Profit. Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company

[13] Simon, H.A. 1955. A Behavioral Model of Rational choice, Quarterly Journal of Economics, Vol.69, 174-183.

[14] Simon H. A. 1957. *Models of Man,* New York, John Wiley and Son.

[15] Simon H.A. 1997. *An Empirically Based Microeconomics,* Cambridge University Press, Cambridge.

[16] Sims, C.A 2001. Pitfalls of a minimax approach to model uncertainty, American Economic Review, Vol 91, 2, 51054.

[17] Williams, N. 2007. Robust Control: An Entry for the New Palgrave, 2nd Edition.

A A Special Case: One Default Time

We consider a special case for simplicity, $N_d = 2$, meaning that if default occurs then it happens at time t_1. We derive an explicit analytical expression for the inverse of the robustness function, $\hat{h}$, thought of as a function of the critical present worth, PW_c, at fixed loan portfolio (w, f). The analytical expression for the general case is accessible but more complicated and is unneeded to achieve the goals of this example.

Definition 3 *Define a truncation function:* $x^+ = x$ *if* $x \leq 1$ *and* $x^+ = 1$ *otherwise.*

Definition 4 *Let* $m(h)$ *denote the inner minimum in the definition of the robustness function, eq.(11).*

A plot of $m(h)$ vs h is identical to a plot of PW_c vs $\hat{h}(PW_c)$. Thus $m(h)$ is the inverse function of $\hat{h}(PW_c)$. Given that $N_d = 2$, the expectation of the present worth, eq.(8), becomes:

$$\mathrm{E}(PW) = \sum_{k=1}^{K} w_k \left(\widetilde{PW}_k - p_{k1} \widetilde{PW}_{k1} \right). \tag{12}$$

From eq.(12) and the info-gap model of eq.(9) we see that the inner minimum in eq.(11) is obtained, at horizon of uncertainty h, when the probability of default of each risk type, p_{k1}, is as large as possible. Thus:

$$m(h) = \sum_{k=1}^{K} w_k \left(\widetilde{PW}_k - [\tilde{p}_{k1} + s_{k1} h]^+ \, \widetilde{PW}_{k1} \right), \tag{13}$$

and $m(h)$ decreases piecewise-linearly as h increases. Hence, since $m(h)$ is the inverse of the robustness function, $\hat{h}(PW_c)$, we see that $\hat{h}(PW_c)$ decreases piecewise-linearly as PW_c increases.

To explore the significance of this we first define several quantities. Let $\widetilde{\mathrm{E}}(PW)$ denote the expectation of the present worth with the estimated probabilities, from eq.(12) with $\tilde{p}_{k1}$ rather than p_{k1} (recall that $N_d = 2$):

$$\widetilde{\mathrm{E}}(PW) = \sum_{k=1}^{K} w_k \left(\widetilde{PW}_k - \tilde{p}_{k1} \widetilde{PW}_{k1} \right). \tag{14}$$

Let E_0 denote the expectation of the present worth when each probability of default equals unity (eq.(8) with $p_{k1} = 1$ and $N_d = 2$):

$$\mathrm{E}_0 = \sum_{k=1}^{K} w_k \left(\widetilde{PW}_k - \widetilde{PW}_{k1} \right). \tag{15}$$

Note that:

$$\mathrm{E}_0 \leq \widetilde{\mathrm{E}}(PW). \tag{16}$$

Finally,

Definition 5 *Define* $h_{\max}$ *as the value of horizon of uncertainty,* h*, beyond which all the probabilities terms* $[\tilde{p}_{k1} + s_{k1} h]^+$ *in eq.(13) equal unity:*

$$h_{\max} = \max_{1 \leq k \leq K} \frac{1 - \tilde{p}_{k1}}{s_{k1}}. \tag{17}$$

Now we find, from eqs.(13)–(15), that:

$$m(h) = \begin{cases} \widetilde{\mathrm{E}}(PW) & \text{if } h = 0 \\ \text{piece-wise linearly decreasing} & \text{if } 0 \le h \le h_{\max} \\ E_0 & \text{if } h_{\max} < h. \end{cases} \tag{18}$$

From this relation we see that the robustness function has the following form:

$$\hat{h}(PW_c) = \begin{cases} \infty, & PW_c < \mathrm{E}_0 \\ \text{piece-wise linearly decreasing,} & \mathrm{E}_0 \le PW_c \le \widetilde{\mathrm{E}}(PW) \\ 0, & PW_c > \widetilde{\mathrm{E}}(PW). \end{cases} \tag{19}$$

This special case is explored with a numerical example in section 3.3.