Unknowns, Black Swans and Keynesian Uncertainty

Phil Faulkner, Alberto Feduzi* and Jochen Runde

Cambridge University and SOAS University of London*

Tony Lawson's work on probability and uncertainty is both an important contribution to the heterodox canon and an early strand of his ongoing enquiry into the nature of social reality. In keeping with both mainstream and heterodox discussions of uncertainty in economics, however, Lawson focuses on situations in which the objects of uncertainty are imagined and can be stated in a way that, potentially at least, allows them to be the subject of probability judgments. This focus results in a relative neglect of the kind of uncertainties that flow from the existence of possibilities that do not even enter the imagination and which are therefore ruled out as the subject of probability judgments. This paper explores uncertainties of the latter kind, starting with and building on Donald Rumsfeld's famous observations about known and unknown unknowns. Various connections are developed, first with Nassim Taleb's Black Swan, and then with the Keynes-inspired interpretation of uncertainty offered by Lawson.

Key words: Risk, Uncertainty, Unknowns, Black Swans, Rumsfeld, Taleb, Keynes, Knight, Lawson

JEL classifications: B21, D81

1. Introduction

Over the course of the second half of the 1980s Tony Lawson published a series of papers on probability and uncertainty (Lawson 1985, 1987, 1988). Written against the backdrop of the rational expectations revolution in macroeconomics, and drawing heavily on the writings of J.M. Keynes's (1921, 1936, 1937), one of their principle aims was to restore and provide philosophical foundations for a conception of uncertainty not reducible to numerically definite probabilities. The papers were widely cited and went on to become part of the canon of heterodox approaches in economics that often insist on distinguishing situations of risk, in which numerical probabilities can be determined, from situations of uncertainty, in which they cannot.

Important as these papers were, there is an important form of uncertainty that they did not consider, namely that which arises in the face of eventualities that are not even imagined as possibilities before being revealed. Although this form of uncertainty is an old and familiar part of the human condition, it has recently been attracting particular attention in the wake of a slew of apparently widely unexpected financial crises, industrial accidents, technological shifts, wars and political uprisings.

Two ideas are currently enjoying considerable currency in this connection, so much so that they have entered the popular lexicon. The first is the distinction between "known unknowns" and "unknown unknowns", already familiar in engineering circles in the 1950s (Wideman, 1992), but which became famous via an enigmatic aside in a 2002 news briefing by the then US Secretary of Defense, Donald Rumsfeld:

> Reports that say that something hasn't happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are **known unknowns**; that is to say we know there are some things we do not know. But there are also **unknown unknowns** − the ones we don't know we don't know (Rumsfeld, 2002).

The second is the Black Swan, popularized by Nassim Taleb's (2007) best-selling book of the same name. "Black Swan" is Taleb's term for an event characterized by three attributes:

> First, it is an outlier, as it lies outside of the realm of regular expectations, because nothing in the past can convincingly point to its possibility. Second, it carries an extreme impact (unlike the bird). Third, in spite of its outlier status, human nature makes us concoct explanations for its occurrence *after* the fact, making it explainable and predictable (Taleb, 2007, pp. xvii–xviii).

Both Rumsfeld and Taleb have their critics, not least Rumsfeld whose remarks won the annual "foot in the mouth award" awarded by the British Plain English Campaign for the most nonsensical remark made by a public figure (http://news.bbc.co.uk/1/hi/3254852.stm). But there is little doubt that they are onto something, and a good deal has now been written on known and unknown unknowns and Black Swans, both individually and in tandem, across disciplines including economics and finance (Barberis, 2013), operations management and management science (Pich, Loch & De Meyer, 2002), probability and statistics (Chichilnisky, 2010), decision theory (de Palma et al., 2014), applied psychology (Feduzi & Runde, 2014), security studies (Mitzen & Schweller, 2011), ecology and environmental studies (Wintle, Runge & Bekessy, 2010), political science (Blyth, 2010) and sociology (Beck, 2006). There are also many contributions that explore similar ideas

under different headings such as "unforeseen contingencies" (Kreps, 1992), "unawareness" (Modica & Rustichini, 1994) and "unknowledge" (Shackle, 1983) in economics, the "small worlds problem" in decision theory (Binmore, 2009; Savage, 1954), "state space uncertainty" in philosophy (Bradley & Drechsler, 2014), "post decision surprises" and "bolts from the blue" in management and organization theory (March, 1994; Weick & Sutcliffe, 2007), and "Knightian uncertainty" or "ignorance" in entrepreneurship (Kirzner, 1979; Sull, 2004).

Taken as a whole, this literature has yielded many insights into the uncertainties it considers. But as yet it contains little in the way of attempts to systematize the relationship between known and unknown unknowns and Black Swans. The purpose of the present paper is to make a start on doing so, and then to investigate how this relates to Lawson's contribution on the nature of uncertainty. We will do so with reference to the particular conceptions proposed by Rumsfeld and Taleb themselves. While the Black Swan has probably made the larger splash in terms of the attention it has received in recent years, in our view the Rumsfeldian categories are the prior and more fundamental of the two. We accordingly begin with Rumsfeld's dichotomy, before moving on to Taleb and the Black Swan, and finally to Tony Lawson's work on probability and uncertainty.

2. Known unknowns and unknown unknowns

To understand what Rumsfeld has in mind with the distinction between "known unknowns" and "unknown unknowns" it is helpful to begin with the more fundamental distinction he draws between "knowns" and "unknowns". According to Rumsfeld knowns are those "things…we know", while unknowns are those "things we do not know" (Rumsfeld 2002) or, equivalently, "gaps in our knowledge" (Rumsfeld 2011, xiv).

The basic idea here—that at any given point in time there are some things that are known and some things that are not known to an individual—is unobjectionable. Yet beyond giving examples of the kinds of things that may be knowns or unknowns, such as the role of gravity in making objects fall, for example, or the extent of a country's nuclear weapons program (Rumsfeld 2011, xiv), Rumsfeld provides little in the way of further elaboration of these two categories. In particular, he sidesteps difficult philosophical questions regarding the nature of knowledge, details that we now briefly address in order to flesh out the meaning of knowns and unknowns.

Rumsfeld's examples indicate that when he talks of a person knowing or not knowing something it is generally factive knowledge, in particular descriptive or propositional knowledge, that he has in mind. Such knowledge concerns what we will call "features of the world", by which we mean facts about past, present and future reality, including the existence, properties and so forth of any (kind of) entity (ranging from animate and inanimate objects, to ideas, theories, opinions and the like), event or state of affairs. Examples of things that may be knowns or unknowns to an individual thus range from quite broad or abstract features of the world such as the regulations that currently govern company mergers within a particular jurisdiction, the new kinds of consumer devices that will emerge as part of the "Internet of Things" or the extent of Iran's existing nuclear weapons program, to much more specific features such as the closing level of the Dow Jones Industrial Average on June 22 2016, the existence of a group of individuals intending to hijack commercial airliners and employ them as weapons or the winner of the next FIFA World Cup.

Perhaps the most difficult issue Rumsfeld leaves unexamined is what it means for an individual to possess knowledge of some feature of the world. According to the standard account within Epistemology an individual is deemed to have knowledge if they possess beliefs about some feature of the world and where those beliefs are both true and justified. We will accept this account for the purposes of what follows, subject to a few qualifications.[1]

Truth is relevant in the present context since false beliefs about features of the world cannot constitute knowledge. But this does beg the question of just how accurate an individual's beliefs have to be to count as knowledge. Evidently our beliefs rarely correspond to the world in a simple one-to-one way, and are often significantly affected by our particular attitudes, interests and ways of seeing things. And even where they are accurate in some respects, they are usually partial and often fragmentary. There is no easy way around these complications. If we set the standard for what counts as knowledge too high, we will exclude many of the beliefs that serve us very well. By the same token, if we set the standard too low we run the risk of including beliefs that verge on the plain false.

Justification is relevant in this context because truth alone is insufficient for

[1] The Gettier (1963) problem provides the most prominent challenge to the idea that true and justified belief is sufficient for knowledge. We will ignore these complications and assume that truth and justification are both necessary and sufficient.

beliefs to be considered knowledge. In addition, in order to rule out beliefs that, while true, are the product of invalid reasoning or are based on false evidence, we also require them to be based on sound reasoning and reliable evidence. But a problem similar to the one raised in the preceding paragraph arises here too: just how well justified do our beliefs have to be in order to count as knowledge? On the one hand if we restrict knowledge to only those true beliefs arrived at on the basis of flawless reasoning and watertight evidence, we are again likely to exclude from counting as knowledge many of the beliefs that serve us well in our decision making. On the other hand, with true beliefs generated on the basis of limited evidence or little reasoning beyond guesswork, we risk including beliefs that are true solely by luck.

Over and above the considerations raised in the preceding paragraphs, it is questionable whether judgments about whether beliefs are sufficiently true and justified to count as knowledge can sensibly ever be made in absolute terms, or whether such judgments shouldn't themselves depend on context. For example, we might reasonably say that we have knowledge of the time for most purposes even when our watch is three minutes slow, even though our beliefs about the time wouldn't qualify as knowledge if we were estimating how long we have to escape a time bomb that is about to explode. Similarly, in a situation that requires immediate action the standard of justification by which one judges a true belief to be sufficiently justified so as to constitute knowledge may well be lower than in a situation where there is the time for a more considered decision.

In light of these three considerations, then, for the purposes of what follows we will regard beliefs as knowledge when they are approximately true and reasonably well justified relative to the context in which they play a part. On this basis we can then define a "known" as any feature of the world that an individual has knowledge of in this sense, and an "unknown" as any feature of the world that an individual lacks knowledge of. So understood, three forms of unknown can be identified. In some cases, they will be features of the world about which we have no beliefs whatsoever. In other cases, they will be features about which we have false (or only very vaguely accurate) beliefs. And in yet other cases, they'll be things about which we have accurate beliefs but where these beliefs lack adequate justification.

Having clarified the meaning of the term unknown we now turn to the distinction between known unknowns and unknown unknowns, which is usually expressed in terms of whether or not an individual knows about a particular gap in

their knowledge. Thus Rumsfeld defines known unknowns as those "gaps we know exist", while unknown unknowns are those "gaps we don't know exist" (Rumsfeld 2011, xiv). Once again the basic idea seems clear enough on the surface: while there may be many features of the world of which an individual is ignorant, there are likely some (the known unknowns) that she knows she is ignorant of while there are others (the unknown unknowns) that she does not even know she does not know.

There is, however, a degree of ambiguity here, since at times Rumsfeld and others appear to regard a known unknown as a gap in knowledge that an individual not only knows about but is actually consciously aware of at the relevant time (e.g. when deliberating over the best course of action, undertaking scenario analysis and so on). And from this perspective the definitions above are imprecise, since knowledge of a thing need not imply the presence of that thing in the conscious mind. Our own view is that the key distinction to be drawn here is indeed concerned with an individual's awareness or unawareness of their own gaps in knowledge, since this is crucial to whether or not an individual takes into account the things they do not know. We will therefore proceed on the basis that a known unknown is a gap in knowledge that an individual knows about and is aware of (at the relevant time), while an unknown unknown is a gap in knowledge that an individual does not know about and by definition therefore cannot be aware of.

3. Hypothetical values and the subjective space of possibilities

Building on the basic Rumsfeldian framework outlined above we now introduce a number of ideas concerned with the beliefs a person holds about what we call the "hypothetical values" associated with a gap in knowledge. By a hypothetical value we mean any value—outcome, state of affairs, result, quantity and so on—that *could conceivably be thought* to be a candidate for the actual or true value of the unknown under consideration. Thus in relation to our earlier examples hypothetical values associated with an individual's ignorance of what may emerge as part of the Internet of Things include such devices as smart expresso makers and smart bicycles; hypothetical values associated with the winner of the next FIFA World Cup include Brazil, Germany and Algeria; and hypothetical values associated with the extent of Iran's nuclear weapons programme would include a range of different stages of development. By the "set" of hypothetical values associated with an unknown we mean the entire collection of values that could conceivably be regarded as the true

value of that unknown by any person within the group or community concerned. Thus while an unknown refers to a gap in the knowledge of a *particular* individual, the set of hypothetical values associated with that unknown is defined for the group as a whole.

Three distinctions can usefully be drawn between members of the set of hypothetical values. The first is between those hypothetical values that are *genuinely possible* and those that are not, where the former are values that could actually turn out to be the case while the latter are values that could not. Two points are worth highlighting about this distinction. The first is that whether some hypothetical value is a genuine possibility or not is something that depends on the way the world is rather than what is believed to be the case. Thus where the unknown concerns the outcome of the next spin of a roulette wheel, then whether or not the outcome 00 is a genuine possibility depends on the type of roulette being played. If it is American roulette, with pockets numbered from 1 to 36 plus two additional pockets numbered 0 and 00 respectively, then 00 is genuinely possible. And this is so irrespective of any individual's beliefs about the type of roulette table being used. If European roulette is being played, where the wheel has only one additional pocket numbered 0, then 00 is not a genuine possibility, again irrespective of beliefs.

The second point is that where an unknown relates to something that is already determined there can be only one genuinely possible value. Thus where the unknown relates to a past event or feature of a currently existing entity there is a definite way the world is and therefore only one genuinely possible value the unknown can take. The situation is different in the case where the unknown concerns something that is not yet determined such as an event which is yet to occur. In the case of the next spin of a roulette wheel for example, provided the wheel is fair and the outcome of any spin is indeed a matter of pure chance, then there are many genuinely possible values.

The two remaining distinctions relate to the beliefs of the individual to whom the unknown relates. The second distinction is between those hypothetical values that individual has consciously imagined and those she has not. To illustrate, consider the case in which the unknown concerns the next spin of an American roulette wheel. Further, suppose the individual is familiar with European roulette but unaware of its American cousin. In this case the outcomes 1 to 36, as well as 0, are likely to have been imagined by the individual, while the outcome 00 is not. Notice it is only in the case of a *known* unknown that an individual can have consciously imagined

hypothetical values at all. For if the gap in knowledge is an *unknown* unknown to the individual then, being unaware of the gap in knowledge in the first place, she cannot have contemplated values that might turn out to be the true value.

Our final distinction is between those hypothetical values the individual has consciously imagined and regards as possible and those that individual has imagined but regards as impossible. By "regards as possible" we mean that there is nothing the individual can think of that she would regard as an insurmountable barrier to the hypothetical value being true or proving to be the case. Returning to our roulette example, if the individual is aware of both American and European variants but believes that a European wheel is being used, it is likely she has consciously thought of outcomes 1 to 36, 0 and 00 as hypothetical values, but regards the latter as impossible given the nature of the wheel being used.

Taken together these three distinctions imply that all elements within the set of hypothetical values associated with a given unknown must fall into one of six categories. These categories and the relationship between them are depicted in Table 1.

			Hypothetical value is…	
			a genuine possibility.	not a genuine possibility.
Hypothetical value is…	consciously imagined by the individual and…	regarded as possible.	1. Hypothetical values imagined and correctly regarded as possible.	2. Hypothetical values imagined and incorrectly regarded as possible.
		regarded as impossible.	3. Hypothetical values imagined and incorrectly regarded as impossible.	4. Hypothetical values imagined and correctly regarded as impossible.
	not consciously imagined by the individual.		5. Hypothetical values not imagined.	6. Hypothetical values not imagined.

Table 1: Classifying the hypothetical values associated with an unknown

The six categories depicted in Table 1 provide a useful framework for describing an individual's beliefs about what may turn out to be true in relation to some unknown. The first column of the table contains hypothetical values that are genuine possibilities. Those in Cell 1 are consciously imagined and correctly regarded as genuine possibilities, while those in Cell 3 are consciously imagined but mistakenly regarded as impossibilities. Hypothetical values in Cell 5 are ones that, despite being genuine possibilities, are not consciously imagined. The second column contains hypothetical values that are not genuine possibilities. Cell 2 contains hypothetical values that are imagined and mistakenly regarded as genuine possibilities, while Cell 4 contains hypothetical values that are imagined and rightly regarded as impossibilities.[2] Cell 6 contains unimagined hypothetical values that are anyway not genuine possibilities.

For a given unknown the distribution of hypothetical values between the six cells in the table depends on two sets of factors. Whether a particular hypothetical value occupies a cell located in the first or the second column of the table is something that, as we noted earlier, is determined by the nature of the world and the constraints this places on which outcomes are genuinely feasible. Within each column the distribution of values between rows is then determined by factors specific to the individual concerned, notably their knowledge of the circumstances relevant to the unknown and their cognitive abilities in relation to reasoning, imagination and so forth. If we take the hypothetical values in the first column of the table for example, those that are genuine possibilities, for someone with strong cognitive abilities and a good understanding of the situation surrounding the unknown, most or all of the values will reside in Cell 1, with relatively few or even none residing in Cell 3 and Cell 5. An individual with a more limited or mistaken grasp of the situation, or who is less able to reason to, or imagine, genuine possibilities, is likely to make more errors, with fewer elements in Cell 1 and more therefore in Cell 3 and Cell 5. Returning to our roulette example, if American roulette is being played and the individual knows only the European variety, then 1-36 and 0 would likely be located in cell 1 and 00 in Cell 5. Cell 3 in this case is likely to be an empty set. If, however, the individual wrongly believed that 12 could not come up on the next spin (it perhaps having come up on the three previous spins), then in this case 12 would be located in Cell 3. Similar considerations apply to the distribution of values among the cells in the second column, where for someone knowledgeable of the situation and with adequate mental powers Cell 2 should be empty, or contain very little, with most or all of the

[2] Examples of hypothetical values in Cell 4 include what might be called "fairy tale" instances such as encountering Superman or sighting a unicorn.

hypothetical values residing in either Cell 4 or Cell 6.[3]

Before we move on it is worth noting that the analysis here, at least as we have presented it so far and as summarized in Table 1, provides only a static framework for thinking about hypothetical values. That is to say, it characterizes a person's beliefs at a given point in time and says nothing about how these beliefs may change over time. We will keep to this static analysis for the remainder of the paper. In doing so however let us note that it would be possible to bring in learning in response to new information by showing how hypothetical values move *within* a column from one row to another as a person changes their beliefs in response to new information. For example, returning to our roulette example, if the individual had thought they were playing European roulette but subsequently learns that an American wheel is being used, the 00 will move from either Cell 3 or Cell 5, to Cell 1.

The hypothetical values located in Cells 1 and 2, those consciously imagined and regarded as genuine possibilities, are the kind dealt with in conventional decision theory and which an individual would list explicitly in the framing of a decision problem. We will call this set of values the individual's *subjective space of possibilities* in respect of an unknown (henceforth SSP). Notice that the SSP in respect of an unknown unknown must necessarily be empty, for as we noted earlier if an individual is unaware of a gap in knowledge that individual cannot even begin to contemplate candidates for its true value. Thus the SSP is of particular relevance to known unknowns, where its contents reflect an individual's beliefs about what might turn out to be the case with respect to some gap in knowledge.

The contents of the SSP in any particular case will of course always be relative to the individual's cognitive abilities and background beliefs in connection with the relevant unknown. Thus someone with a background in nuclear weaponry and a knowledge of conditions in Iran will likely come up with a different space of possibilities regarding the state of that country's nuclear weapons programme than someone who does not. The key issue for our purposes is the fit between an individual's beliefs about the possible values an unknown might take and the way the world actually is in this regard. In some cases, the SSP may coincide perfectly with the objective situation, such that the SSP contains all genuinely

[3] In some cases, having a good understanding of the scenario may imply that an individual is likely to consciously imagine, and then dismiss, particular values (Cell 4) rather than not imagining them at all (Cell 6). An example of this is an experienced gambler who, playing European roulette, may be expected to consciously rule out 00 as a possible outcome rather than not imagine it at all. In general, however, for a knowledgeable individual it will be a matter of that individual's background beliefs as to whether something that is impossible is thought of and then dismissed, or simply not considered at all.

possible values (i.e. Cell 3 and Cell 5 of Table 1 are empty) and no others (i.e. Cell 2 is empty).Thus where the gap in knowledge concerns the outcome of the next spin of a fair roulette wheel and the individual knows that American rather than European roulette is being played, the individual's SSP is likely to be objectively accurate, that is, containing the outcomes 1–36, 0 and 00, and nothing else.

Such examples represent something of a special case, however, since roulette and games of chance more generally usually feature well-defined, reasonably simple, closed systems, where it is relatively straightforward for an individual to determine to a high degree of accuracy what may, and may not, turn out to be the case. Yet for unknowns that arise in more practical, real life, scenarios things are often quite different. The complexity of the situations concerned will usually be higher here, making it more difficult for individuals to accurately determine the set of genuinely possible values. In such cases, the SSP is likely to be deficient in some way. There are two possible types of error here. The first is that the SSP may omit genuine possibilities, either as a result of a hypothetical value having been contemplated but wrongly judged impossible or because the individual failed to even imagine that value at all. The second source of error is that the SSP may wrongly include hypothetical values that are i*mpossibilities*, outcomes or states of affairs that cannot actually arise.

4. Taleb's Black Swan

According to Taleb (xvii–xviii), a Black Swan is an event distinguished by the following three properties:

P1: "It is an outlier, as it lies outside of the realm of regular expectations, because nothing in the past can convincingly point to its possibility";

P2: "it carries an extreme impact"; and

P3: "in spite of its outlier status, human nature makes us concoct explanations for its occurrence after the fact, making it explainable and predictable."

We will focus primarily on P1 in what follows, as it is here that the issues addressed in the present paper surface most clearly in Taleb's account. P2 and P3 concern things that for the most part lie beyond our immediate concerns, and we will therefore have little to say about them directly.

As we have already noted, the property of a Black Swan highlighted in P1 is that it "lies outside the realm of regular expectations" (Taleb, 2007, p. xvii). The same kind of idea is expressed towards the end of the book where Taleb writes: "Remember that for an event to

be a Black Swan, it does not just have to be rare, or just wild; it has to be unexpected, has to lie outside our tunnel of possibilities" (Taleb, 2007, p. 213). Unfortunately Taleb does not appear to define terms such as "realm of regular expectations" or "tunnel of possibilities", and, evocative as these phrases may be, they are lacking in precision. We will accordingly try to add some by formulating what he has to say in terms of our earlier framework and the notion of the SSP in particular.

In order to do so notice that, in virtue of P2 and P3, an event can only be a Black Swan on Taleb's schema from an *ex post* perspective, since it is only once it has occurred that it can have an extreme impact and be the subject of *post hoc* rationalisations. Black Swans, in other words, are necessarily events that have already occurred. P1, however, stipulates an *ex ante* requirement for an event to be a Black Swan, being concerned with an individual's expectations about the possibility of that event actually occurring in the future. In particular, when Taleb writes of events that lie outside the "realm of regular expectations" or the "tunnel of possibilities", we will interpret him to mean events that, prior to their occurrence, had simply not been foreseen even as possibilities by the individual concerned.

To capture this within our framework it is first necessary to say something about the context within which an individual's expectations about the future arise. In particular, we need to specify the relevant unknown, namely the aspect of future events about which the individual is unsure. In some cases, this unknown may be quite specific, such as when a financial analyst is contemplating possible scenarios in the wake of Brexit or a gambler the outcome of the next spin of a roulette wheel. In other cases, the unknown may be considerably vaguer, reflecting uncertainty about the course of events in general rather than with respect to a particular, and well-defined, part of the world.

Viewed this way Taleb's P1 can then be given more precise formulation by relating it to an individual's SSP. For events that are not foreseen even as possibilities—those that lie outside the "realm of regular expectations" or the "tunnel of possibilities" in Taleb's words—are hypothetical values that lie outside the SSP in relation to the relevant unknown. Table 1 above includes four cells whose members fall outside the SSP, namely cells 3, 4, 5 and 6. From the point of view of Taleb's first property, then, hypothetical values in these four cells represent candidates for Black Swans, events that if they subsequently occur would constitute Black Swans provided P2 and P3 are also satisfied.

The analysis here can be taken one step further by incorporating the fact that the events Taleb has in mind are necessarily ones that go on to occur, since it then follows that a Black Swan must correspond, *ex ante*, to a hypothetical value that was a genuine possibility.

Making this requirement explicit Taleb's first condition for an event to be a Black Swan is then that:

> P1': *Ex ante*, the corresponding hypothetical value is a genuine possibility and lies outside the SSP in respect of the relevant unknown, either because it was not imagined or was wrongly deemed impossible.

In relation to Table 1 this implies we can disregard as candidate Black Swans those hypothetical values located in the right-hand column; only those hypothetical values located in Cell 3 and Cell 5 satisfy P1'. We summarize this in Table 2, a cut down *ex post* version of Table 1, in which the right-hand column now provides a classification of hypothetical values that have gone on to occur.

<table>
<tr><td></td><td></td><td></td><td>Revealed event was ...</td></tr>
<tr><td rowspan="3">Hypothetical value was ...</td><td rowspan="2">consciously imagined by the individual and…</td><td>regarded as possible.</td><td>1a. correctly regarded as possible ex ante.</td></tr>
<tr><td>regarded as impossible.</td><td>3a. incorrectly regarded as impossible ex ante.</td></tr>
<tr><td>not consciously imagined by the individual.</td><td></td><td>5a. not imagined ex ante.</td></tr>
</table>

Table 2: Classifying hypothetical values that go on to occur

Members of Cells 3a and 5a are ones that satisfy Taleb's first property, corresponding to events that have occurred and that, *ex ante*, were hypothetical values that fell outside the SSP and were genuine possibilities. Cell 5a, which contains hypothetical values not even imagined *ex ante* but that were in fact experienced or found to be true, is the more straightforward of the two. This case is illustrated by the example in which American roulette is being played by someone who knows only the European variety and accordingly does not even imagine the hypothetical value 00 and include it in their SSP. Cell 3a is illustrated by the case of hypothetical value 00 dismissed as impossible by the roulette player who knows the difference between American and European roulette, but who is operating under the mistaken assumption that European rather than American roulette is being played. It is interesting that the source of the title of Taleb's book, the discovery in the late 17th century of the existence of black swans in Western Australia by a party led by Dutch explorer William de Vlamingh, falls into this second category. For the existence of black swans was up to that point regarded as an impossibility in Europe, a source of the simile that can be traced back as far as the *Satires* of the Roman poet Juvenal writing at the turn of the first century (Juvenal, Satire 6).

5. Uncertainty

Our aim in this section is to explore the relationship between the ideas discussed in the preceding sections and the conception of uncertainty advanced by Tony Lawson in a string of papers published over the course of the 1980s (Lawson 1985, 1987, 1988). In particular, since Lawson never addressed the subject of unimagined possibilities in these papers, we will attempt to show how our earlier discussion of unknowns, hypothetical values and the SSP may be used to augment his account in this regard.

The primary influence on Lawson's writings on uncertainty is J.M. Keynes, a rich source in virtue of the significant contribution Keynes made to the subject of probability in his 1921 *A Treatise on Probability* (Keynes 1921/1973) and, in some ways perhaps even more so, for the glimpses of this work in his later economic writings on uncertainty in his *General Theory* and beyond (Keynes 1936; 1937). Much of Lawson's contribution in the three papers cited above lies in unpacking and building on Keynes's ideas on probability—and also, but to a lesser extent, those of Knight (1921)—which went on to form an important plank of his wider critique of, and alternative to, mainstream economics (Lawson, 2003, 2007).

Understanding Lawson on the subject of uncertainty therefore requires going back to *A Treatise on Probability*. This is not an entirely straightforward matter as the conception developed there is rather different from the modern view of probability as a branch of mathematics and its more familiar frequentist or subjectivist interpretations. What Keynes proposes instead is a view of probability as an indicator of the strength of the argument, or what he calls the *probability-relation*, between some conclusion and the evidence bearing on it. He formalizes this idea as follows:

> Let our premisses (sic) consist of any set of propositions *h*, and our conclusions consist of any set of propositions, *a*, then, if a knowledge of *h* justifies a rational belief in *a* of degree α, we say that there is a *probability-relation* of degree α between *a* and *h* (Keynes 1921/1973, p. 4)

The probability-relation is written:

$$a/h = \alpha \qquad (1)$$

where, for example, *a* might be the proposition "inflation will be 4% next year", *h* a body of evidential propositions concerning recent changes in the money supply, interest rates, aggregate demand and so on relevant to *a*, and α the rational degree of belief *h* justifies in *a*. Keynes makes a useful distinction between "primary" and "secondary" propositions here,

where the latter involve assertions about probability-relations while the former do not. With respect to (1) above, *a* is the primary proposition and the probability relation (1), the secondary proposition. The secondary proposition is thus both a proposition in its own right and a statement about the primary proposition. The set of evidential propositions, *h*, may include (other) secondary as well as primary propositions.

Keynes holds that the value of α in any probability relation ranges between 0, where not-*a* is a logical consequence of *h* and *a* is therefore impossible, and 1, where *a* is a logical consequence of *h*. The intermediate values represent cases in which *h* only partially entails *a*. An important and, from a modern point of view rather idiosyncratic feature, of Keynes' theory is his insistence that the rational degrees of belief represented by α are generally not numerically definite, nor necessarily even pair-wise comparable in terms of more than, less than, or equal to (Keynes 1921/1973, p. 37).

To flesh out this scheme it is helpful to consider Keynes's views on knowledge and the cognate categories of certainty and truth. Knowledge, for Keynes, corresponds to true and certain rational belief, where certainty denotes the highest degree of rational belief in a proposition.[4] He distinguishes between direct and indirect knowledge, reflecting the two ways in which he believes knowledge of propositions can be acquired. The first is by "direct acquaintance", whereby "we are able to pass from direct acquaintance with things to a knowledge of propositions about the things of which we have sensations or understand the meaning" (Keynes 1921/1973, p. 13). Although alive to the difficulties and possible objections involved, he proceeds on the basis that direct acquaintance always leads to *certain* rational belief.[5] The second route is "indirectly, *by argument*, through perceiving the probability-relation of the proposition, about which we seek knowledge, to other propositions" (Keynes 1921/1973, p. 12), where these other propositions correspond to the body of known evidential premises *h*.

Returning to (1), Keynes holds that the probability relation is only ever known directly, and, when known in conjunction with the body of evidential propositions *h*, provides *indirect* knowledge *about* a (when $a/h < 1$), or *of a* (when $a/h = 1$). In the former case this knowledge is said to entail a *probable* degree of rational belief of degree α in the proposition

[4] "… knowledge of a proposition always corresponds to certainty of rational belief in it and at the same time to actual truth in the proposition itself. We cannot know a proposition unless it is in fact true" (Keynes, 1921/1973, p. 11).

[5] "… I have assumed that all direct knowledge is certain. All knowledge, that is to say, which is obtained in a manner strictly direct by contemplation of the objects of acquaintance and without any admixture whatever of argument and contemplation of the logical bearing of any other knowledge on this, corresponds to *certain* rational belief and not to a mere probable degree of belief." (Keynes 1921/1973, p. 17).

a, while the latter case involves certain belief in *a*.[6]

Keynes's theory therefore represents a form of epistemic probability, concerned as it is with degrees of belief, and to this extent has affinities with the subjectivist or personalist interpretation associated with Ramsey (1926), de Finetti (1937) and Savage (1954). The main difference between Keynes and these authors, besides his view that probabilities are generally not point valued, is that the degrees of belief he has in mind are not merely subjective or personal to the individual concerned, but the degrees of belief it is rational to hold in any proposition given the evidence bearing on it. The only subjectivity Keynes admits in his scheme is with respect to the contents of *h*, which he sees as depending on individual circumstances. Given *h*, however, he believes there is only one objective probability-relation between *a* and *h*.

Keynes does not define uncertainty in *A Treatise on Probability*, and one of the aims of Lawson's 1985 paper is to provide an interpretation of uncertainty consistent with the framework sketched above. Building on the idea that uncertainty must involve a lack of certainty of some kind, Lawson proposes that uncertainty corresponds to the situation in which direct knowledge of the secondary proposition is absent (Lawson 1985, p. 913). On this basis he argues that uncertainty can arise in two ways: (1) where the relevant probability relation is unknown due to an individual's inability to argue from given evidence to the degree of rational belief it justifies in some proposition, and (2) where there exists no method for determining a numerical measure of the probability relation, namely where probabilities are numerically immeasurable or indeterminate (Lawson 1985, p. 913).[7]

On the basis that the idea of unknown probabilities does not seem to feature elsewhere in Keynes' work, Lawson argues for the second interpretation as being most in keeping with Keynes' use of the term uncertainty in his economic writings. A difficulty with this interpretation, however, is that on Keynes's own account in *A Treatise on Probability*, it is quite possible to have direct knowledge of the secondary-proposition and for the probable

[6] There is tension between the places in which Keynes describes knowledge as corresponding to true and certain belief, and the places in which he talks about indirect knowledge where "knowledge" falls short of certainty. Keynes seems to sense something of this problem where he writes: "I assume then that the (sic) only true propositions can be known, that the term 'probable knowledge' ought to be replaced by the term 'probable degree of rational belief', and that a probable degree of rational belief cannot arise directly but only as the result of an argument, out of the knowledge, that is to say, of a secondary proposition asserting some logical probability-relation in which the object of belief stands to some known proposition" (Keynes, 1921/1973, p. 18). It is therefore helpful to read Keynes' references to indirect knowledge as referring to degrees of belief that are only probable, save of course for the limit cases in which $a/h = 1$ or $a/h = 0$.

[7] An alternative interpretation of (2), noted but rejected by Keynes primarily on the grounds of maintaining as general a notion of probability as possible, is to say that in such circumstances there exists no probability relation at all (Keynes, 1921/1973, p. 11).

degree of belief it justifies nevertheless to be numerically immeasurable or numerically indeterminate. That is to say, for Keynes, numerically immeasurable or indeterminate probabilities are not necessarily a sign of an absence of direct knowledge of the secondary proposition.

This problem is avoided in a later paper in which Lawson returns to the subject of Keynesian uncertainty, but this time moves away from associating uncertainty with the absence of direct knowledge of the secondary proposition. Instead uncertainty is simply equated with numerically immeasurable or indeterminate probability-relations. The crucial passage is a long footnote, the first half of which runs as follows:

> In *A Treatise on Probability* and later in the *General Theory*, Keynes essentially distinguishes three types of probability-relation: the first where a probability-relation is numerically indeterminate and possibly not even comparable to (in terms of less than, equal to, or greater than) other probability relations; the second where probabilities are numerically determinate but less than unity (and greater than zero); and the third where probabilities take the value of unity (or zero). The first type of probability, for Keynes, corresponds to a situation of uncertainty (see Lawson, 1985) and the third to a situation of certainty. Moving from the first type through the second towards the third, Keynes talks … of the argument in question being less 'uncertain' … (Lawson 1987, p. 953, f. 2)

The conception of uncertainty captured in this passage is strongly reminiscent of that expressed in Knight's (1921) famous distinction between risk and uncertainty, where the former corresponds to situations in which numerically definite probabilities can be determined and the latter to situations in which they cannot. In a slightly later paper in which he compares various interpretations of probability and uncertainty, Lawson again comes down explicitly in favour of this Keynesian / Knightian view of uncertainty as "corresponding to situations wherein numerically measurable probabilistic knowledge is not available" (Lawson 1988, p. 62).

We are now in a position to consider how these ideas fit with our earlier discussion of unknowns, hypothetical values and the SSP. The site we use to explore this question is the famous passage from Keynes's 1937 *QJE* defense of *The General Theory*, also quoted and discussed by Lawson (1985, p. 914), in which, sixteen years after *A Treatise on Probability* was published, Keynes outlines what he means by uncertainty:

> By "uncertain knowledge", let me explain, I do not mean merely to distinguish

> what is known for certain from what is merely probable. The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn. Or, again, the expectation of life is only moderately uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth-holders in 1970. About these matters there is no scientific basis on which to form any calculable probability whatsoever. We simply do not know (Keynes 1937:213-214).

While Keynes does not explicitly contrast risk with uncertainty in this passage, his message is clearly in keeping with the Knightian view of uncertainty outlined above. Keynes illustrates his view with a spectrum of examples aimed at demonstrating the extent to which numerically definite probabilities can be determined in each case. We reproduce Keynes' spectrum diagrammatically in Figure 1.

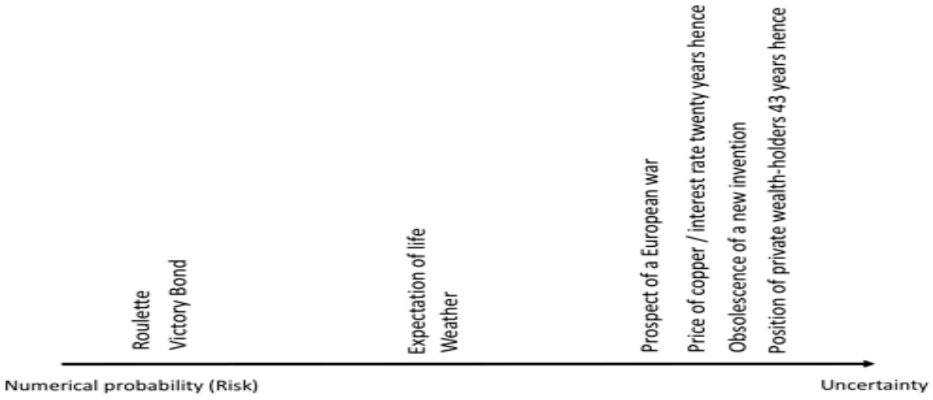

Figure 1. Keynes' spectrum of probability situations

The examples on the left, the cases of roulette or a lottery draw, are ones in which numerical probabilities can be calculated on an *a priori* basis by assuming equal probabilities of the elementary outcomes. Those in the middle, the expectation of life and the weather, are ones in which numerical probabilities can sometimes be determined, at least within limits, by way of an empirical estimation of underlying frequencies.[8] The cases to the right, according to Keynes, are ones in which there is no scientific method for the calculation of numerical probabilities. Moving from left to right across the spectrum, then, we move progressively

[8] Knight, like Keynes, distinguishes between a priori probabilities and frequencies when talking about numerical probabilities, which he then counterposes against uncertainty (Runde 1998).

from situations in which numerical probabilities can be determined, situations of "risk" (Dequech, 2000; Dow, 2015; Knight, 1921; Lawson, 1988), through to situations of uncertainty in which they cannot.[9]

As well as illustrating Keynes' notion of uncertainty, the spectrum of examples in Figure 1 also provides a useful bridge to the various concepts discussed in earlier parts of the present paper. Each of the examples corresponds to an unknown in Rumsfeld's sense, since they all concern gaps in knowledge (in this case pertaining to the future, such as the outcome of a lottery, the weather over some period or the price of copper twenty years' hence). Each of these gaps in knowledge have a set of hypothetical values associated with them, each member of each set corresponding to an outcome or state of affairs that could be thought to be the actual or true value associated with the unknown under consideration. And the subset of each set of hypothetical values that the individual consciously imagines and regards as genuinely possible with respect to any unknown corresponds to the SSP of that individual, where each member of that subset may be the subject of probability judgments by that individual ("primary propositions" in the language of *A Treatise on Probability*). Returning to our earlier example of the proposition that inflation will be 4% next year, this proposition is one hypothetical value associated with the unknown that is next year's inflation rate. This proposition will be subject to uncertainty in Keynes's sense, insofar as it does not have a numerically measurable probability.

One of the advantages of locating Keynes' and Lawson's ideas within our earlier framework in this way is that it brings into focus aspects of uncertainty—here understood in its most generic sense as referring to an absence of certainty of some kind—neglected in their accounts. The remainder of this section is devoted to one of these aspects, what we will call unimagined possibilities. What we have in mind here are the kinds of outcomes or eventualities we began with at the start of the paper: industrial accidents, terrorist actions, scientific discoveries, political uprisings and the like that many people seem not to have entertained even as possibilities prior to their occurrence.

Unimagined possibilities of this sort are hypothetical values, associated with an unknown, that while being genuine possibilities have not been consciously imagined as such

[9] Proponents of the subjectivist Bayesian interpretation of probability have sometimes argued that, if subjective degrees of belief are "rational" in the sense of conforming to the strictures of something like the Savage axioms (Savage 1954), all probabilities become numerical and the distinction between risk and uncertainty evaporates (Leroy and Singell XX). However, there is considerable empirical evidence that people often do not behave "as if" they are assigning numerically definite probabilities to the things they are uncertain about even in relatively straightforward situations (Ellsberg, XXX XXX), and a large amount of theoretical work, much of it by people working within the subjectivist Bayesian tradition, aimed at modelling these cases (Bewley, XXX).

by the individual concerned. They therefore correspond to the hypothetical values located in Cell 5 of our earlier Table 1. It could be argued that hypothetical values in Cell 3, those that have been contemplated but wrongly deemed impossible, should also be considered unimagined possibilities since they too concern things the individual has never imagined actually occurring. While the existence of such values may be significant, as a potential source of Black Swans for example, we will however exclude them from the present analysis in order to focus on those genuine possibilities that have not entered the individual's conscious mind at all.

We touched on the notion of unimagined possibilities earlier in the paper, in relation to the question of how closely an individual's SSP is likely to match objective reality. In particular, we noted that the SSP may be deficient in two ways: first, it may omit values that are genuine possibilities, and second, it may include values that are not genuine possibilities. Unimagined possibilities are one source of the first kind of error, where an individual has failed to come up with a value that the unknown might actually take. Values that have been imagined but rejected as impossible—those in Cell 3—represent a second source of this first kind of error. The extent to which the SSP exhibits a close fit with objective reality, and so avoids such errors, depends jointly on the nature of the unknown and the cognitive abilities and background knowledge of the individual concerned. And it was in this context we argued that unknowns associated with games of chance represent something of a special case, for unlike many of the unknowns that arise in real life, such games typically involve well-defined, relatively simple, closed systems in respect of which it is comparatively straightforward to determine the set of genuinely possible outcomes.

From this perspective it is significant that most of the examples Keynes refers to in the earlier *QJE* passage are similar to games of chance in terms of the limited scope for unimagined possibilities. To see this, consider what the SSP is likely to look like for a typical individual in each of the examples in Figure 1.[10] In the cases of roulette and the drawing of a Victory bond we are in the realm of games of chance of the sort just described and for which the set of genuinely possible hypothetical values is typically known from the outset or else easily determined. In these cases, we would expect the SSP to comprise only, and all, the genuinely possible outcomes of the spin or draw and for there to be no unimagined possibilities.

[10] By a typical individual here we mean someone with an average level of cognitive ability and relevant background knowledge. Restricting things in this way allows us to focus solely on the impact the nature of the unknown has on the scope for unimagined possibilities.

In cases in which the unknown concerns the value of a single variable—the price of copper or the interest rate in twenty years' time, and even the expectation of life—the set of genuinely possible hypothetical values is restricted to a range on the real line. Here again, then, we would expect the SSP to match closely the set of genuinely possible values, where the only way in which the latter might fall outside the SSP is where an individual sets the upper or lower bounds of the range too conservatively. In such a scenario, however, the genuinely possible values omitted from the SSP are unlikely to constitute unimagined possibilities. Rather they are likely to reside in Cell 3 rather than Cell 5 of Table 1, having been imagined and considered impossible, rather than not having been imagined at all. The cases of the prospect of a European war or the obsolescence of a new invention illustrate a further class of unknown where the set of genuinely possible hypothetical values is easily deduced, for in these cases there are only two genuinely possible hypothetical values, namely that the war or obsolescence either occurs or does not. In cases of this kind, where the possible outcomes are binary and restricted to *x* and not-*x*, the SSP will plainly already include the full set of hypothetical values and nothing else. Once more, unimagined possibilities are unlikely.

For the most part, then, Keynes' examples impose fairly definite and easily understood constraints on the genuinely possible hypothetical values, either because the unknown relates to a relatively simple part of the world or because of the (e.g. binary) way in which the unknown is framed. There is then minimal scope in these examples for the kind of deficiencies in the SSP mentioned above and for the existence of unimagined possibilities in particular. The only two of Keynes' examples which differ in this regard are those concerning the weather and, perhaps most clearly, the position of private wealth-holders 43 years hence. What is special about these cases is that the unknowns relate to a complex part of reality—systems characterized by numerous components, their interactions and then, in the case of the position of private-wealth holders, viewed over a long time horizon—and that they are framed in a relatively loose and open-ended way. As a result, in determining the SSP an individual is faced with coming up with hypothetical values in relation to scenarios where key factors and their possible interactions may be difficult to identify, and perhaps even unknowable at the time, and where the framing of the unknown imposes far fewer constraints than the earlier examples. The scope for deficiencies in the SSP in these cases is accordingly much greater than in Keynes' other examples, and where this, particularly with respect to the position of private wealth-holders more than four decades into the future, will likely include the existence of unimagined possibilities.

The absence of an explicit treatment of unimagined possibilities in the Keynes / Lawson account of uncertainty is perhaps not surprising in view of it starting with the Keynesian probability relation. For the focus then falls quite naturally on the probability of, or uncertainty in respect of, what Keynes calls the primary proposition—as opposed to the uncertainty in regard of perhaps not even being able to come up with one or more primary propositions that might be relevant to the problem at hand. Or, to put it in our terms, the focus from the outset is on imagined hypothetical values.[11]

However, and as writers such as Rumsfeld and Taleb have become famous for pointing out, unknown unknowns and unimagined hypothetical values may be important too. This is evident simply from considering examples in which this type of uncertainty looms large, for example when thinking about the kind of products a major scientific advance might lead to in 10 year's time, about potential improvised weapons that might be used by a terrorist organization, or about the future consequences of something like Brexit. In these kinds of scenarios traditional approaches to probability that presuppose an exhaustive list of mutually exclusive hypothetical values ("outcomes") simply do not apply, and accordingly, the approaches to decision-making on which they are based.

Up to this point in the discussion we have focused on uncertainty regarding membership of the SSP, particularly in relation to unimagined possibilities that lie outside the SSP. We close this section by drawing attention to an associated, though distinct, form of uncertainty. The kind of uncertainty we have in mind here concerns a lack of awareness about features of the world—what we will henceforth refer to generically as "influences"—that may go on to affect which of the hypothetical values in the SSP turns out to be the true or actual one. Note that uncertainty here is not about membership of the SSP, but about influences that, if they were known and contemplated, might enter probability judgments concerning individual members of the SSP. Keynes' example of the prospect of European war nicely illustrates the distinction. While the binary nature of the unknown means that there is little scope for uncertainty in relation to the contents of the SSP in this case, the complexity of the scenario means there is likely to be considerable uncertainty over influences relevant to the determination of probabilities for either of the two hypothetical values.

This kind of uncertainty is closely related to what Keynes calls the weight of evidence. As presented in Chapter 6 of *A Treatise on Probability* (Keynes, 1921/1973, pp. 77-85), the

[11] The contrast between unimagined possibilities and those hypothetical values located in Cell 3 of Table 1 is again instructive here, for unlike the former values which are ones that have not even been contemplated, the latter values are ones the individual is aware of but believes that the evidence points to their impossibility.

weight of evidence is an indicator of the size in some sense, or the degree of completeness or spread, of evidence on which a probability judgment is based.[12] While Keynes reveals a measure of ambivalence about the concept (1921/1973, p. 77, 345), he never fully rid himself of the idea that "the degree of completeness of the information upon which a probability is based does seem to be relevant … in making practical decisions" (Keynes, 1921/1973, p. 345) and it resurfaces in his later economic writings in the *General Theory* (XXX) and subsequent correspondence (Runde 1990, 2003; O'Donnell 1989, 1991).

Lawson does not have much to say about the weight of evidence, but touches on the idea in the second part of the long footnote on uncertainty from his 1987 paper quoted above:

> Moving from [numerically indeterminate probabilities to numerically definite probabilities], Keynes talks … also of the weight of the argument increasing. For Keynes the weight of an argument appears to be the 'degree of completeness of the information upon which a probability is based' (Keynes, 1973, p. 345) or, equivalently, the 'balance... between the absolute amounts of relevant knowledge and of relevant ignorance respectively' (Keynes, 1973, p. 77). Thus as relevant evidence increases so too does the weight. Keynes seems to suggest that certainty can only be reached when weight is at its highest, and this in turn appears to correspond to a situation where the relevant evidence is complete (Lawson 1987, p. 953, f. 2)

We agree that low weight and Keynesian uncertainty are linked and with the idea that uncertainty of this sort often decreases as evidence becomes more complete (Runde 1990). But we disagree with the suggestion, if it is indeed even being made by Lawson here, that increases in weight *necessarily* equate to coming closer to being able to determine numerically definite probabilities. For there are many situations in which the evidence may be highly complete but in which it is not possible to determine probabilities, whether this be on the basis of equal probabilities, a knowledge of stable frequencies, or indeed even by looking at peoples' propensities to bet on different outcomes (Bewley 1986; Gärdenfors and Sahlin 1988; Levi 1974). While increase in weight may well reduce uncertainty in some sense they do not necessarily lead any closer to numerical probabilities.

That aside, the ideas expressed by Keynes and Lawson on this front are comfortably accommodated within the framework developed in the first part of the present paper. Putting

[12] Keynes is circumspect about the extent to which weight can be measured, and his views in this regard mirror what he has to say about probability: "Only in a restricted class of cases can we compare the weights of two arguments in respect of more or less" (Keynes, 1921/1973, p. 77).

it in slightly more general terms than Keynes does in *A Treatise on Probability*, we propose that the weight of evidence on which a probability judgment is based is usefully interpreted as an indicator of the extent to which relevant influences are contained in the evidence on which that judgment is based. This could be described as the weight of evidence reflecting the degree of completeness of the evidence, or from a Rumsfeldian perspective, as the balance between influences that are knowns and those that are unknowns. That is, when weight is low and the evidence highly incomplete, many of the relevant influences are unknowns and so absent from the set of evidential propositions. Conversely when weight is high and the evidence relatively complete, most of the relevant influences are knowns.

Thinking about the weight of evidence in this way leads us to two closing observations. The first, if low weight is indeed a reflection of unknown influences, is that unknown unknowns—influences that an individual is not even aware of being ignorant of—are likely to be especially pernicious in this regard. An individual who is not even aware of not knowing about some influence cannot begin to factor these things into their probability judgments, whereas, if the influences are known unknowns, they can at least begin to form a view as to the hypothetical values associated with those influences. In either case, of course, the analysis here takes us back to our earlier discussion concerning the accuracy of the SSP, with a given unknown influence contributing most to low weight of evidence when the accuracy of the associated SSP is poor.

The second observation concerns how the nature of the situation in which probability judgments are being made affects evidential weight. In general, the more open or complex such situations, particularly where the unknown concerns outcomes that lie in the distant future, the greater the scope for unknown influences. Keynes' examples again serve to illustrate this point. Thus as we move across the spectrum from left to right in Figure 1, the situations involved generally become progressively more open in the sense of providing increasing scope for the existence of as yet unknown influences that may go on to affect which of the hypothetical values in the SSP turns out to be the true one. And as before, this openness may also contribute to the inaccuracy of the SSP in relation to known unknowns, exacerbating further the problem of low weight. In contrast, the situations in which it is possible to derive probabilities on an *a priori* basis, or to measure stable frequencies, situations that approximate what Lawson in his later work calls closed systems, are likely to be ones in which weight is that much higher.

5. Conclusion

Most of the discussion of uncertainty in economics over the last 30 years or so, both in mainstream and heterodox contributions such as Tony Lawson's, has focused on situations in which the objects of uncertainty can at least be imagined and stated in a way that allows them to be the subject of probability judgments (even if such judgments tend mostly to be qualitative rather than quantitative in nature). This focus has resulted in a relative neglect of the kind of uncertainties that flow from the existence of possibilities not even imagined and which therefore cannot be the subject of probability judgments.

We have attempted to address this neglect by developing a conceptual framework that takes as its starting point Donald Rumsfeld's famous distinction between known unknowns and unknown unknowns. The resulting framework provides a precise interpretation of Rumsfeld's categories, terms that while widely employed are nonetheless under-theorized and prone to inconsistent usage. Building on these ideas we then put forward an exhaustive categorization of the different kinds of hypothetical values that might be associated with any unknown, and a means of describing what is contained in any individual's subjective space of possibilities.

One of the principal benefits of frameworks of this kind is that they provide useful vehicles for making connections with other accounts. We did so in the present paper, concentrating on Nassim Taleb's Black Swan and Tony Lawson's writings on Keynesian uncertainty. The payoff is that we have been able to distinguish four distinct varieties of uncertainty:

(1) Uncertainty that arises in connection with unknown unknowns in Rumsfeld's sense, where even the gap in knowledge is unknown and the SSP is accordingly an empty set.

(2) Uncertainty that arises in connection with whether or not the membership of the SSP is "correct". The unknown or gap in knowledge in Rumsfeld's sense is known in this case, but where one or more genuinely possible hypothetical values associated with that gap in knowledge may be unknown and therefore wrongly excluded from the SSP, or where one or more impossible hypothetical values may be wrongly included in the SPS.

Neither Lawson nor Keynes have much to say directly about these two varieties of uncertainty. However, it is not difficult to think of examples in which they loom large, e.g. when thinking about the kind of products a major scientific advance might lead to in 10 year's time, about potential improvised weapons that might be used by a terrorist organization, or about the future consequences of something like Brexit.

There remain two further possibilities:

(3) Uncertainty that arises in connection with which of the hypothetical values in the SSP is the true one; and

(4) Uncertainty that arises in connection with unknown possible influences that might affect which of the members of the SSP will turn out to be the true one.

The third variety of uncertainty is the one that Lawson, Keynes and Knight have most to say about, namely that reflected in the examples in Figure 1 and ranging from cases involving numerically measureable risk on the left of the spectrum to cases involving numerically immeasurable uncertainty on the right. Finally, the fourth variety of uncertainty, we have argued, has affinities with Keynes' evidential weight, and is where we see an at least subliminal intersection between the work of Lawson and Keynes on the one hand and that of Rumsfeld and Taleb on the other.

Bibliography

Barberis, N. 2013. The Psychology of Tail Events: Progress and Challenges, *American Economic Review*, vol. 103(3), 611-16

Beck, U. 2006. Living in the world risk society, *Economy and Society*, vol. 35(3), 329-345

Bewley, T.F. 1986. Knightian Decision Theory: Part I, New Haven, Cowles Foundation Discussion Paper, number 807. Reprinted in 2002, *Decisions in Economics and Finance*, vol. 25, pp. 79–110

Binmore, K. 2009. *Rational decisions*. Princeton, NJ, Princeton University Press

Blyth, M.M. 2010. What if Most Swans are Black? The Unsettling World of Nassim Taleb. *Critical Review*, vol. 21(4), pp 447-465

Bradley, R., & Drechsler, M. 2014. Types of uncertainty, *Erkenntnis*, vol. 79(6), pp. 1225-1248

Chichilnisky, G. 2010. The Foundations of Probability with Black Swans, *Journal of Probability and Statistics*, ID 838240, http://dx.doi.org/10.1155/2010/838240

de Palma, A., Abdellaoui, M., Attanasi, G., Ben-Akiva, M., Fehr-Duda, H., Erev, I., Fok, D., Fox, C.R., Hertwig, R., Picard, N., Wakker, P.P., Walker, A.L., Weber, M. 2014. Beware of black swans: Taking stock of the description–experience gap in decision under uncertainty, *Marketing Letters*, vol. 25(3), pp. 269-280

Dequech, D. 2000. Fundamental uncertainty and ambiguity, *Eastern Economic Journal*, vol. 26:1, pp. 41-60

Dow, S.C. 2015. Uncertainty: A Diagrammatic Treatment, *Economics*.

http://www.economics-ejournal.org/economics/journalarticles/2016-3

Faulkner, P. & Runde, J. 2009. On the identity of technological objects and user innovations in function, *Academy of Management Review*, vol. 34(3), pp. 442-462

Feduzi, A. & Runde, J. 2014. Uncovering unknown unknowns: towards a Baconian approach to management decision-making, *Organizational Behavior and Human Decision Processes*, vol. 124(2), pp. 268-283

Gärdenfors, P. & Sahlin, N.E. 1988. Unreliable Probabilities, Risk Taking and Decision Making, in P. Gärdenfors & N. E. Sahlin (eds.), *Decision, Probability and Utility*, Cambridge, Cambridge University Press, pp. 313–34

Gettier, E.L. 1963. Is Justified True Belief Knowledge? *Analysis*, vol. 23, pp. 121–123

Juvenal. Satire 6: The ways of women, Translated by G.G. Ramsey, http://www.poetryintranslation.com/PITBR/Latin/JuvenalSatires6.htm#anchor_Toc282858852

Kirzner, I.M. 1979. *Perception, opportunity, and profit*, Chicago, University of Chicago Press.

Lawson, T. 1985. Uncertainty and Economic Analysis, *Economic Journal*, vol. 95, pp. 909-27

Lawson, T. 1987. The Relative/Absolute Nature of Knowledge and Economic Analysis., *Economic Journal,* vol. 97, 951-70

Lawson, T. 1988. Probability and Uncertainty in Economic Analysis, *Journal of Post Keynesian Economics*, vol. 11, pp. 38-65

Lawson, T. 2012. Ontology and the study of social reality: emergence, organisation, community, power, social relations, corporations, artefacts and money. *Cambridge Journal of Economics*, 36, 345–385.

Levi, I. 1974. On Indeterminate Probabilities, *Journal of Philosophy*, vol. 71, pp. 391–418

Keynes, J. M. 1921 [1973]. A Treatise on Probability, The Collected Writings of John Maynard Keynes, vol. VIII, London, Macmillan

Keynes, J. M. 1936 [1973]. The General Theory, The Collected Writings of John Maynard Keynes, vol. VII, London, Macmillan

Keynes, J. M. 1937 [1973]. The general theory of employment, Reprinted from the *Quarterly Journal of Economics*, in The General Theory and After: Part II Defence and Development, The Collected Writings of John Maynard Keynes, vol. XIV, London, Macmillan

Knight, F.H. 1921. *Risk, uncertainty and profit*, Chicago, University of Chicago Press

Kreps, D.M. 1992. Static Choice in the Presence of Unforeseen Contingencies, in Dasgupta,

P., Gale, D., Hart, O. and Maskin, E. (eds.) *Economic Analysis of Markets and Games: Essays in honor of Frank Hahn*, Cambridge, MIT Press, 258-281

March, J.G. 1994. *A primer on decision making: How decisions happen*, New York, Free Press

Mitzen, J. & Schweller, R.L. 2011. Knowing the Unknown Unknowns: Misplaced Certainty and the Onset of War, *Security Studies*, vol. 20(1), pp. 2-35. DOI:10.1080/09636412.2011.549023

Modica, S. & Rustichini, A. 1994. Awareness and partitional information structures. *Theory and Decision*, vol. 37(1), pp. 107-124

O'Donnell, R.M. 1989. *Keynes: Philosophy, Economics and Politics*, *The Philosophical Foundations of Keynes's Thought and their Influence on his Economics and Politics*, London, Macmillan

O'Donnell, R.M. 1991. Keynes on Probability, Expectations and Uncertainty, in O'Donnell R.M. 1991. (ed.) *Keynes as Philosopher-Economist*, Macmillan, London

Pich, M.T., Loch, C.H., & De Meyer, A. 2002. On uncertainty, ambiguity and complexity in project management. *Management Science*, vol. 48, pp. 1008-1023

Rumsfeld, D. 2011. *Known and Unknown: A Memoir*, New York, Sentinel

Runde, J. 1990. Keynesian uncertainty and the weight of arguments, Economics and Philosophy, vol. 2, 275–92

Runde, J. 1998. Clarifying Frank Knight's Discussion of the Meaning of Risk and Uncertainty, *Cambridge Journal of Economics*, vol. 22, pp. 539–46

Runde, J. 2003. On some explicit links between Keynes's A treatise on probability and The general theory, in Runde, J. and Mizuhara, S. (eds.), *The philosophy of Keynes's economics: probability, uncertainty and convention,* London, Routledge, pp.46-54

Savage, L. J. 1954. *The foundations of statistics*, London, John Wiley & Sons

Schoemaker, P. 2002. *Profiting from uncertainty. Strategies for succeeding no matter what the future brings*, New York, The Free Press

Shackle, G.L.S. 1979. *Imagination and the nature of choice*, Edinburgh, Edinburgh University Press

Sull, D. 2004. Disciplined Entrepreneurship. *Sloan Management Review*. 46(1), 71-77.

Taleb, N.N. 2007. *The black swan: The impact of the highly improbable*, London, Penguin books

Weick, K.E., & Sutcliffe, K.M. 2007. *Managing the unexpected. Resilient performance in an age of uncertainty* (2nd ed.), San Francisco, John Wiley & Sons

Wideman, R. M. 1992. *Project and program risk management: A guide to managing project risks and opportunities*, Newton Square, PA, Project Management Institute

Wintle, B.A., Runge, M.V., & Bekessy, S.A. 2010. Allocating monitoring effort in the face of unknown unknowns, *Ecology Letters*, vol. 13(11), pp. 1325–1337