The Australian Journal of
Agricultural and
Resource Economics

Journal of the Australian Agricultural and Resource Economics Society

Australian Journal of Agricultural and Resource Economics, 60, pp. 159–176

Ambiguity, learning opportunities and risk-neutral regulation*

Alan Randall†

Two current issues in management of public risks, ambiguity and learning, are addressed in the context of managing ecosystems with thresholds, and regulating treatment safety as might apply, for example, to human and animal health, pesticides and herbicides. Reconsidering a recent claim that, in systems that penalise violation of thresholds, learning opportunities induce riskier decision-making, I find no incentive for ambiguity seeking. But opportunity to benefit from learning may indeed induce riskier decisions, an effect that diminishes and eventually disappears as penalties become larger. A recent claim that a rational regulator of pharmaceutical drugs would be ambiguity preferring – a claim that has obvious applicability to a broader set of treatments – is then examined. Ambiguity-tolerant patients may indeed prefer a menu of ambiguous treatments and opportunity to learn and switch, rather than a single treatment with known risk. But the source of ambiguity matters. Patient heterogeneity, prior to and independent of policy, generates ambiguity for individuals and motivates preference for a menu of treatments. However, expanding the menu does not justify approving treatments that are generally riskier. Finally, I challenge the perennial claim that the regulator of risks to human health and safety should seek to maximise expected value.

Key words: ambiguity, bandit problem, health and safety, learning, risk-neutral regulation, system with threshold.

1. Setting the stage: risk, learning and ambiguity

In what follows, I consider two current issues in management of public risks – ambiguity and learning – and a hardy perennial: whether the regulator should seek to maximise expected value. To establish the context of the inquiry, a brief introduction to the public risks involved and their relevance to agriculture and natural resources is provided, and a preview of the argument is offered. Then, key terms are defined, current thinking is summarised briefly, and some implications and issues that arise are highlighted.

* The present article is drawn from the author's presidential address at the 59th annual conference of the Australian Agricultural and Resource Economics Society, Rotorua, NZ, 2015. He is grateful to Eric Naevdal for an extended conversation on these and related topics over the years, and to seminar audiences at Australian National University, the Productivity Commission, and the Ohio State University for stimulating comments.

† Alan Randall (email: randall.3@osu.edu) is an Honorary Professor of Economics and Agriculture & Environment, The University of Sydney, Sydney, New South Wales, 2006, Australia, and an Emeritus Professor of Agricultural, Environmental & Development Economics, The Ohio State University, Columbus, Ohio, 43210-1067, USA,

doi: 10.1111/1467-8489.12124

1.1 Context and preview of the argument

The two cases at the core of the analysis presented below are prominent among the current and emerging concerns in public risk management. Complexity theory has brought management of systems characterised by thresholds to the fore, and there is widespread awareness of the possibility of regime shifts and the importance of resilience in terrestrial and aquatic ecosystems and the global climate system. The first case involving management of ecosystems characterised by thresholds (Rogers 2013) helps untangle the relationships between ambiguity and learning, and takes a first step towards understanding the value of 'probing the system' in order to learn more about how it works.

Regulation of novel interventions is an issue of long-standing interest and public concern, because there seems always to be room for debate as to the relative priorities of encouraging innovation and protection from harm should things go wrong. Examples include approval of drugs and treatments intended for humans and animals, novel food products (including those that are genetically modified), and pesticides and herbicides, which are regulated in many countries; and have potential application also to nano-particles and synthetic chemicals, which receive little regulatory attention in many jurisdictions. In the second case study, approval criteria for drugs and medical treatments serve as a metaphor for this broad class of public risk regulation problems. Here, the Viscusi and Zeckhauser (2015) claim that the regulator should be ambiguity preferring is examined.

The argument proceeds as follows. First, risk is defined as chance of harm, and three kinds of chance are identified. Only the first kind of chance, in which outcomes are generated by a known random process, offers the possibility of enumerating completely the possible outcomes and specifying their probabilities. Yet most of the public decision problems we encounter involve systems that we do not fully understand, some of which may be nonstationary. In these cases, ambiguity and learning opportunities loom large. Ambiguity is introduced via Ellsberg's (1961) paradox, which treats ambiguity as a kind of uncertainty. To address questions about the impact of ambiguity on decisions, outcomes and welfare, it is necessary to hold the magnitude of risk constant. Compound lotteries provide a way of doing that, albeit for a rather restrictive kind of ambiguity. The first case examined, on systems with thresholds, is mostly about learning – if the threshold is fixed, probing may reveal information about its location that can be employed profitably in subsequent periods – but it is shown also that ambiguity *per se* is never preferred. This result is interesting because the second case reconsidered, regulatory approval of drugs, involves a claim that ambiguity is preferred by consumers and should be valued by regulators. One can indeed construct cases where ambiguity is preferred, because a menu of ambiguous treatments allows patients to stay with a high-performing treatment but switch from an underperformer. However, the value of ambiguity is derived mostly from patient heterogeneity, which is independent of regulatory activity and provides little justification for greater regulatory tolerance of risk. What we learn from the drug approvals case has current and potential application to agricultural and natural resource issues including food safety, herbicide and pesticide safety, and workplace health and safety.

Finally, since the Viscusi–Zeckhauser argument involves appeal to the notion that the drug regulator should adopt a risk-neutral stance, I address that issue head-on, arguing that the public demands for good reasons that the health and safety regulator remain averse to individual risks from treatment.

1.2 Risk as chance of harm

In ordinary usage, *risk* refers to a *chance* of *harm* – a small chance of a large benefit is called a gamble. Harm is a relatively simple concept, an event or change of circumstance that makes one worse-off. Chance is not so simple – developments in complex systems theory and risk management over the last generation have, among other things, broadened and enriched the concept of chance. Now, we recognise three kinds of chance (Randall 2011):

1. Chance, because outcomes are generated by a known random process. The dominant analogy is well-specified games of chance, in which outcomes can be enumerated and probabilities calculated, as in classical *Knightian risk*. Risk management concepts and tools include the expected value and expected utility criteria; the value of delaying a decision if the opponent (often conceptualised as nature) may reveal its hand in the interim, a concept central to real options analysis; and assuming idiosyncratic risk, as we frequently do, insurance and hedging mechanisms.
2. Chance, because we do not understand the system that generates outcomes. We do not know enough to enumerate the complete outcome set and calculate the associated probabilities. Observed outcome frequencies may be invoked to substitute for knowledge of how the system works. But the frequentist approach often seeks to coax additional information out of limited frequency data by imposing structure, commonly by fitting the data to a functional form with known and regular properties. How else do engineers, for example, calculate the 500-year drought and the 500-year flood from (say) ninety years of good rainfall data? But this approach can be misleading. The order imposed on limited data, especially when analysts choose statistical models that provide good estimates of central tendency, may seriously understate the tails of the distribution.

If we have confidence that we can specify the complete set of possible outcomes but we are unsure of the objective probabilities, we face *Knightian uncertainty*. If some of the uncertainties that influence the system remain unidentified, we have *unknown unknowns* – things we do not know that we do not know.

3. Chance, because the system that generates outcomes is itself changing. Complexity theory permits this kind of situation, and we can recognise it in the real world, if mostly in retrospect. This introduces the possibilities of *gross ignorance* (in which the outcome set is unbounded) and really big surprises (where *surprise* is defined as an outcome that was not a member of the *ex ante* outcome possibility set).

It is not easy to distinguish chance types 2 and 3 in real time, a problem that seems to have occurred as the millennium drought dragged on in the eastern parts of Australia. Perhaps we were living through an unlikely but not impossible run of bad luck, or perhaps the system itself had changed ('drought is the new normal') and drastic new remedies therefore were required. Watching the situation unfold in real time, it was difficult for decision-makers to dismiss the latter possibility, and there ensued a rush to build desalination plants that subsequently were mothballed when rainfall returned to more familiar levels.

Often, we are tempted to analyse chance of types 2 and 3 with type-1 apparatus – it is more familiar and more orderly, and we have grown accustomed to justifying chance of type 1 as an idealised but nevertheless serviceable analogy for the risky and uncertain situations we face. Unfortunately, the real world tends to punish such short cuts.

1.3 Learning

Learning is motivated by a perceived disconnect between the world, which is real, and our knowledge of it, which is partial at best and may involve misconceptions as well as imprecision. Learning is the process by which we acquire and process information that we hope will bring us closer to understanding the reality.

The vocabulary of applied economists dealing with risk includes learning – the opportunity for learning drives the value of real options, and Bayesian updating (a stylised kind of learning) has long been used to model how one's expectations respond to new evidence. Economists have estimated the value of learning in a few particular contexts, but we do not have a well-articulated theory of learning. For management of public risks, I distinguish two fundamentally different sources of threats (2011) – from novel interventions that could end badly, and from overstressing familiar systems in the course of business as usual (BAU) – which impose different restrictions on learning opportunities.

1.3.1. Learning before we let the genie out of the bottle
As has become familiar in the processes for approving new pesticides and pharmaceutical drugs, we can invoke prerelease testing protocols to learn about the risks inherent in proposed interventions. Because testing is itself a risky process, it begins in secure laboratories (for example) and containment is relaxed sequentially as results of testing increase confidence that the intervention is benign or its risks are manageable. An idealised sequential testing process is sketched in Randall (2011, figure 13.2).

1.3.2. Learning in BAU
In the business-as-usual context, the threat already is embedded in the economy and society, and the time for prerelease testing protocols has long passed. Instead, learning opportunities fall mostly within the broad rubric of adaptive management, AM (Holling 1978; Stankey *et al.* 2005). Most commonly, AM provides a passive kind of learning: at each round in a continuing process of plan, implement, monitor and adjust, the plan is to implement the action that is most promising based on the best available information at the time, and learning occurs as nature's response is revealed. A more active approach to AM is conceivable – probing the system purposefully in order to increase learning opportunities – but seems to be underused in practice.[1]

In Section 2, learning opportunities figure prominently in the two cases examined: in the first, learning opportunities in systems that penalise violation of thresholds may induce riskier decisions; and in the second, patients may prefer a menu of treatments with opportunity to switch treatments after learning.

1.4 Ambiguity

Like learning, ambiguity is about us at least as much as about the real world we seek to comprehend. Ambiguity can be attributed to probabilities, to the *ex ante* outcome set, and even to our models of the systems we seek to manage – the common element in these applications is a sense of insecurity regarding our knowledge.

Ellsberg (1961) coined the term ambiguity to label a kind of Knightian uncertainty: in the three-colour problem below, the player is provided with insufficient information to calculate some of the probabilities. Klibanoff *et al.* (2005) define ambiguity as, roughly, having to decide on the basis of subjective rather than objective probabilities. They write (and I paraphrase):

[1] In focus groups, local and regional resource managers resisted this more active kind of AM, saying it is important that they can assure their constituents that each management action taken is the most promising based on information available at the time (Randall *et al.* 2012, section 3.2).

> ... we imagine an ambiguity-averse decision maker to be thinking as follows: My best guess is that the probability of outcome x is 20 per cent, but this is based on 'softer' information than knowing that the frequency is 20 per cent. Hence, I apply more caution in the case of the ambiguous risk.

In the discussion of gross ignorance, outcome sets may be called ambiguous if they are specified imprecisely or incompletely. 'Unknown unknowns' suggests that ambiguity applies not only to the answers but also the questions.[2] Analyses of decision making under ambiguity require us to be very clear about what kind of ambiguity is being addressed.

1.4.1. The Ellsberg Paradox: people avoid ambiguity
In Ellsberg's 'three-colour problem' an urn has 30 red balls and 60 that may be either black or yellow. The player is asked to choose one of the following pair of gambles:

A: You receive \$100 if you draw a red ball   B: You receive \$100 if you draw a black ball

Then, with replacement, choose one of the following pair of gambles:

C: You receive \$100 for a red or yellow ball   D: You receive \$100 for a black or yellow ball

Most people choose A and D, which is illogical, because if one believes red is more likely than black (as implied by choosing A over B), then red or yellow must be more likely than black or yellow. The choice of A and D implies that people are averse to something about B and C, presumably the ambiguity that arises from the less complete probability information:

In A, $p(\text{red}) = 1/3$ In D, $p(\text{black or yellow}) = 2/3$
In B, $0 \leq p(\text{black}) \leq 2/3$ In C, $1/3 \leq p(\text{red or yellow}) \leq 1$

This kind of ambiguity aversion has been well corroborated experimentally (Camerer and Weber 1992). Millner *et al*. (2013) provide an influential application to climate policy, given the extant ambiguity in predictions relating global emissions, temperature rise and damage. They define ambiguity, as did Klibanoff *et al*., as making do with probabilities that are subjective to some degree, certainly 'softer' than knowing the relevant frequencies; and they assume ambiguity aversion. Examining the implications of this kind of ambiguity for climate change mitigation, they conclude:

> ... the value of emissions abatement increases as ambiguity aversion increases, and this "ambiguity premium" can in some plausible cases be

[2] Schipper (2014) suggests that the emerging literature on unawareness eventually may provide a more satisfactory analysis of unknown unknowns.

very large. Thus ambiguity aversion seems likely to provide another argument for strong abatement policy.[3]

2. Ambiguity and the value of learning – reconsidering two cases in the literature

The first case reconsidered involves learning in systems characterised by thresholds. Naevdal and Oppenheimer (2007) imagine an individual paid by the step to walk blindfolded towards a cliff an undisclosed distance away. What is the optimal stopping point? Relevant considerations include the payment per step, the individual's expectations re distance to the cliff, and the penalty for breaching the threshold. N and O show that, for at least some plausible values of the relevant parameters, the optimal solution is not a single step.

Rogers (2013) addressed managed ecosystems with undisclosed thresholds. Revenue increases with harvest, but so does the chance of being penalised for breaching the threshold. If the threshold is fixed, each probe into the previously unexplored zone reveals information that the decision-maker can use in subsequent periods. Rogers concluded that explicit consideration of the value of information induces the decision-maker generally to make riskier decisions. This result sounds counterintuitive on first hearing, because we are accustomed to assuming that better information reduces risky behaviours. However, it is consistent with the notion of active learning in adaptive management – the potential value of what might be learned induces probing, that is interventions more intrusive than the best action based on the information available at the time.

The second case reconsidered invokes a context much discussed in economics and policy circles, optimal regulatory criteria for approving new pharmaceutical drugs and medical treatments. Viscusi and Zeckhauser (2015) associate ambiguity with learning opportunities and conclude

> (t)he FDA (US Food and Drug Administration) shows a strong aversion to ambiguous risks. This is the opposite of what is desirable. For any given initial expected risk level, optimal risk-taking decisions involving uncertainty in a multiperiod world should prefer ambiguous risks, and the potential for learning, relative to well-established risks of the same magnitude.

This conclusion seems to be at odds with the notion of ambiguity aversion and its usual implication that ambiguity makes people more cautious.

In preparation for reconsidering the Viscusi–Zeckhauser, V–Z, claim concerning ambiguity and learning, the Rogers claim about learning and risk-

[3] Traeger (2014), taking a different approach, reaches a similar conclusion.

taking is examined in Rogers' ambiguity-free context and in a context where ambiguity is present.

2.1 The problem of risks of similar magnitude

The claims about learning (Rogers) and ambiguity (V–Z) apply to risks of similar magnitude. Fixing the magnitude of risk requires that we know a lot about both harm and chance, ideally the magnitude of each or, taking a familiar short cut, the magnitude of their product, i.e. expected value, EV. Fixing harm is relatively easy at least in principle, but fixing chance in the context of ambiguity presents a fundamental problem – ambiguity (in the senses of Ellsberg and Klibanoff *et al.*) is about not really knowing the probabilities.

In the three-colour problem, the probability of yellow is bounded by 0 and two-third, but within that range is unknown. Ellsberg's ambiguity (and Knightian uncertainty) in their raw forms do not permit calculation of expected values.

Ambiguity, for Klibanoff *et al.*, is about the distinction between objective probabilities and best guesses. We might approach this kind of ambiguity by constructing a compound lottery. The probability of yellow, $p(y)$, is known for the population of balls. Call $p(y)$ the meta-probability. Fill n urns by randomly drawing balls, with replacement, from the population. Then, $E(p_i(y)) = p(y)$ for urn i, but it is unlikely that the frequency of yellow balls in urn i is exactly $p(y)$. The player first draws an ambiguous urn and then draws a ball from the urn chosen. This procedure would allow ambiguity with respect to the urn drawn, but within the context of risks of similar magnitude given by the meta-probability $p(y)$ and the magnitude of harm.

It is well known that compound lotteries reduce to simple lotteries in the case of a single draw (von Neumann and Morgenstern 1944), but not when multiple draws must be taken from the urn drawn in the first stage – drawing an urn in which yellow is under-represented would leave a player at a disadvantage throughout the whole sequence of draws (Berry and Viscusi 1981). The case of multiple draws is appropriate here, since Rogers' decision-maker hopes for a long sequence of harvests and V–Z's patient hopes to survive multiple rounds of treatment.

The compound-lottery framework is serviceable in capturing ambiguity, but it is a restrictive kind of ambiguity – we know the meta-probability and the EV of a single draw. Compare Klibanoff's ambiguity, which allows no more than best guesses about EV, Ellsberg's ambiguity and Knightian uncertainty, which do not permit calculation of EV, and the various formulations that permit ambiguity about outcome possibilities as well as their probabilities.

2.2 Ecosystems with thresholds at unknown locations

Profit increases with intensity of exploitation of an ecosystem, but there is a penalty for breaching a threshold that is not revealed *ex ante*.[4] To keep things simple, imagine a rectangular farm field with the right-side boundary undisclosed. Output and revenue increase with area cultivated, that is with distance D of the threshold from the left boundary, but so does the chance of breaching the threshold. Each production period, the decision-maker chooses the area cultivated and following the harvest is paid for the quantity produced minus the penalty for any breach of the threshold. Breach is a discrete 0, 1 variable (the threshold was breached or not) – there is no feedback about the magnitude of a breach. Thresholds are set, or their distributions are bounded, so that there is some positive level of penalty at which willingness to pay to cultivate the field, WTP, and the area cultivated are positive. Four configurations of the game are specified, and we compare them in terms of EV and WTP to lease the field for multiple periods (to allow profit from learning, if any).

The first two configurations severely restrict the opportunity for learning; they differ in that ambiguity is present in ii:

- i Stochastic threshold, known distribution – pure risk, no learning opportunity. This configuration is denoted (S, K).
- ii Stochastic threshold, undisclosed distribution – ambiguity, potential to learn about the distribution, but not the location of the threshold (S, U).

Configurations iii and iv have fixed thresholds, which sharply increases the value of learning because whatever is learned about the threshold in cultivation period *t* remains true for subsequent periods:

- iii Fixed threshold, drawn from a known distribution – pure risk, opportunity to learn the location of the threshold (F, K).
- iv Fixed threshold, drawn from an undisclosed distribution – ambiguity, opportunity to learn the location of the threshold (F, U).

Note that the magnitude of risk in the undisclosed distribution, U, cases is not fixed unless we add strong restrictions. Suppose we restrict the distribution of threshold distance, D, so that E(D(U)) = D(K), but for U there are no additional restrictions on higher moments of the distribution. This implies that an EV-maximising agent would be indifferent between such prospects, but risk-averse and/or ambiguity-averse agents are unlikely to be indifferent.

[4] Here, I approach Rogers' (2013) problem from a quite different perspective: his contribution was in dynamic modeling of managed ecosystems.

2.2.1. Results

The square brackets below contain the reasoning for the conclusions re relative magnitudes, i.e. less than, equal, or greater than. The letters K, U, F, S refer to the case configurations above, and the bold letters **I** and **P** denote 'is indifferent to' and 'is at least as preferred as', respectively.

First, consider expected values:

$$\text{EV } iii = \text{EV } iv \text{ [K } \mathbf{I} \text{ U]} \geq \text{EV } i \text{ [F } \mathbf{P} \text{ S]} = \text{EV } ii \text{ [K } \mathbf{I} \text{ U]}$$

Expected value comparisons provide a clear result that the decision maker will offer more to lease the field if the threshold distance D is fixed, [F **P** S], and the process of cultivation reveals something about its location. Ambiguity regarding the probability distribution is a matter of indifference under the EV criterion so long as the EVs themselves are equal.

Now, introduce ambiguity aversion:

$$\text{WTP } iii \geq \text{WTP } iv \text{ [K } \mathbf{P} \text{ U]} \geq \text{WTP } ii \text{ [F } \mathbf{P} \text{ S], and}$$

$$\text{WTP } iii \geq \text{WTP } i \text{ [F } \mathbf{P} \text{ S]} \geq \text{WTP } ii \text{ [K } \mathbf{P} \text{ U].}$$

The result that a fixed threshold is preferred [F **P** S] because it provides the greater opportunity for learning still holds. With ambiguity aversion, ambiguity regarding the probability distribution is no longer a matter of indifference. [K **P** U] because the unrestricted and undisclosed higher moments in U generate more risk; for example, the lower bound on D may be lower, the U distribution may have a fatter left tail, etc. There is no general preference ranking of (F, U) and (S, K) – the U distribution is perceived as riskier but F rewards learning, an effect that gets stronger as the penalty decreases and the multiperiod contract gets longer. So with the WTP criterion we have (where learning refers to opportunity to learn the location of the threshold):

(Pure risk, learning) **P** (ambiguity, learning) **P** (ambiguity, no learning)
(Pure risk, learning) **P** (pure risk, no learning) **P** (ambiguity, no learning).
This implies: Pure risk **P** ambiguity
Given risk or ambiguity, learning **P** no learning.

2.2.2. Penalties

The value of learning and the incentive for risky probing hinges on manageably small penalties for breaching the threshold. If the penalty is high enough – perhaps the invisible boundary is guarded by an invisible row of land mines – the Naevdal and Oppenheimer 'not a single step' result comes into play. With a known probability distribution, K, there would be no cultivation beyond the known lower bound of the distribution, and WTP would not reflect any learning opportunity. With an undisclosed distribution, U, a zero-cultivation solution is conceivable. More generally, as the level of penalty increases *ceteris paribus*, WTP to cultivate beyond the (known or guessed) lower bound of the threshold probability distribution falls until it hits a corner solution at zero (Rogers 2013).

One can also imagine penalties that increase with each violation – this commonly is observed in enforcement of regulations and criminal laws, and is plausible in some natural systems, for example repeated breaches of the threshold may weaken resilience. This penalty structure would inhibit probing, especially when there have been previous breach(es), and thus reduce WTP.

It should be noted that the ambiguity addressed above is nontrivial but less challenging than commonly encountered. Typically, there is no assurance that EVs are equal for risky and ambiguous games (c.f. Ellsberg), and in ecosystem threshold cases, there is likely to be ambiguity about the nature and magnitude of harm, as well as its probability.

2.3 Ambiguity, learning and switching: should the regulator be ambiguity preferring?

Viscusi and Zeckhauser (2015) argue that, for optimal risk-taking decisions involving uncertainty in a multiperiod world, ambiguous risks should be preferred. This raises an obvious question: if, in the case of systems with thresholds, there is value in learning but not in ambiguity *per se*, why would ambiguity be valued in drug regulation?

The V–Z formulation features stylised cases with binary outcome possibilities, the patient dies or not. Treatment may kill; and the affliction may kill if untreated, but is less likely to kill if treated. The game is multiperiod in the sense that, should the patient survive a round of treatment, s/he may proceed to a further round. Cure is defined as surviving *n* rounds of treatment, and the game ends when the patient dies, or is considered cured.

2.3.1. V–Z, and balls in urns

It is useful to explore the sorts of games V–Z play in the context of urns and red and yellow balls, but with two helpful modifications. The games specified below involve more rounds to increase the scope for learning, and stakes that are substantially reduced compared to V–Z's life-and-death stakes.

2.3.2. Game 1. Pure risk, no switching (S, K)

An urn has 50 red or yellow balls with known $p = 0.5$ of drawing yellow. The player gets 10 draws, with replacement after each, and is paid \$10 for each yellow ball drawn. Using notation introduced above (Section 2.2), this game can be characterised as (S, K), because it provides no learning opportunity and the probability distribution is known.

2.3.3. Game 2

A compound lottery is introduced to generate ambiguity that can be characterised as U′ (c.f. U in Section 2.2). In the games below, p for the meta-distribution is known but p_i is not; however, the binomial meta-distribution places some limits on the higher moments of the distribution in the particular urn. Since opportunity to benefit from switching after learning is crucial to the V–Z argument in favour of ambiguity, three variants of Game 2 are considered, each with a different rule concerning switching. Games with learn-and-switch opportunities are classified as F′ because a learn-and-switch opportunity here is analogous to a fixed threshold in 2.2, in that both provide opportunity to gain from learning.

2i Ambiguity, no switching (S, U′). Draw an ambiguous urn with individual p_i from a set of urns with $E(p_i) = p$, and draw 10 balls with replacement from the urn drawn. The payoff is $10 per yellow ball drawn.

2ii Ambiguity, can switch to another ambiguous lottery (F′, U′). After 2 draws from an ambiguous urn, players can reject their urn and (with replacement) draw another ambiguous urn for the remaining 8 draws. This game is characterised as F′ because it provides opportunity to learn and switch after the first two draws.

2iii Ambiguity, can switch to a simple risk lottery (F′, K′). After 2 draws from an ambiguous urn, players can switch to the pure-risk game, 1, for the remaining 8 draws. This rule is denoted as F′ as in game ii and K′ because there also is opportunity to switch to K after two rounds of U.

Games 2ii and 2iii are ‘bandit problems’, a reference to slot machines: first-generation machines were nicknamed one-armed bandits, alluding to the single lever used to activate each play. Believing there might be some variation in expected payout among the machines in a room, players sought to develop methods of identifying the more fruitful machines. Trialling a machine chosen at random and ‘staying with a winner but dumping a loser’ made intuitive sense. Berry and Viscusi (1981) demonstrate a method for up-dating a machine’s payout probabilities as the trial proceeds.

2.3.4. Results

As in Section 2.2, the four games (game I and three variants of game 2) are compared in terms of EV and, to allow ambiguity aversion, WTP to play; and the square brackets contain the reasoning for conclusions regarding relative EV and WTP.

Expected Values.

$$2iii = 2ii\ [\mathrm{K}'\ \mathbf{I}\ \mathrm{U}'] > 2i\ [\mathrm{F}'\ \mathbf{P}\ \mathrm{S}] = 1\ [\mathrm{K}'\ \mathbf{I}\ \mathrm{U}'].$$

To interpret these results:

- Opportunity to learn is worthless without the opportunity to switch.
- Opportunity to switch is valuable, but only if there is a chance to learn before switching.
- A learn-and-switch game is more valuable than an ordinary risk game – it provides opportunity to stay with a winner, and dump a loser.
- While the analogy with the systems with thresholds case (Section 2.1) is imperfect, the EV results are similar.

Ambiguity aversion matters. With any positive aversion to ambiguity

$$\text{WTP } 2iii > \text{WTP } 2ii \; [\text{K}' \; \mathbf{P} \; \text{U}'] \text{ and WTP } 1 \geq \text{WTP } 2i \; [\text{K}' \; \mathbf{P} \; \text{U}'].$$

That is, ambiguity *per se* is never preferred: switching to a pure risk game is preferred to switching to another ambiguous game, and with no switching opportunity a pure risk game is preferred to an ambiguous game. Nevertheless, ambiguous compound lotteries with a switching opportunity may be preferred to a simple risk lottery – they provide opportunity to stay with a winner but dump a loser.

With great enough aversion to ambiguous distributions

$$\text{WTP } 1 > \text{WTP } 2iii > \text{WTP } 2ii \geq \text{WTP } 2i.$$

That is, the games are preferred in strict order of exposure to pure risk versus ambiguity; but, given inescapable ambiguity (games 2i and 2ii), a switching opportunity (as in 2ii) may be preferred.

2.3.5. The nature of the probabilities and the sources of ambiguity

A menu of ambiguous treatments and the opportunity to learn and switch is preferred to a pure risk situation, unless ambiguity aversion is great enough to reverse this ranking, a result that V–Z would find congenial. However, this result does not imply that the regulator (e.g. FDA) should be ambiguity seeking.

Suppose the p_i are treatment characteristics, that is variability in response to treatment is attributable entirely to heterogeneity in treatments. In time, with trial and error, patients will tend to converge on the treatment(s) that provide the preferred combination(s) of safety and effectiveness. However, given the strong public good nature of the p_i, a structured program of prerelease testing is likely to be preferable to leaving patients to do their own testing; and this program would reduce ambiguity and reject proposed treatments that failed tests of safety and effectiveness. New proposed treatments that survive prerelease testing would be welcomed to the market. But, quite explicitly, there would be no virtue in ambiguity-seeking regulation, especially if it involved relaxing standards of safety and effectiveness.

Now suppose the probabilities are really p_{ij}, that is, patients are heterogeneous and the characteristics of individuals, j, as well as characteristics of treatments, i, influence treatment outcomes. Then individuals face personal ambiguity with respect to the safety and effectiveness of treatments. They potentially can gain from a choice of treatments with opportunity to learn more about their $p_{ij.}$. It follows that a menu of safe and effective treatments is preferred to a single approved treatment – especially for afflictions that are recurring and seldom fatal – because it would allow patients to try treatments, learn and switch, converging eventually on the treatments that work best for them. But this does not establish that the regulator should promote ambiguity by relaxing standards of safety and effectiveness. Nature provides, in the form of patient heterogeneity, the ambiguity that makes learn-and-switch opportunities for individual patients valuable.

3. The risk-neutral regulator

Consistent with a long-standing tradition in economists' critiques of drug safety regulation, V–Z argue strongly for risk-neutral regulation – that is, for a maximise expected quality-adjusted life years, E(QALY), criterion that makes no distinction between sources of risk (from treatment or from exposure to the affliction itself), and recommends treatment up to the point where expected net marginal risk reduction is zero. Yet, economists find themselves rather lonely in this point of view – there seems to be broad public and political support for FDA's 'safe and effective' criteria, and public indignation when adverse reactions to approved drugs (ADRs) exceed acceptable norms.

3.1 The arguments for risk neutrality are unconvincing

The claim that government should be a risk-neutral regulator often is framed as an extension of Arrow and Lind's (1970) conclusion that government should be risk-neutral as an investor, because each citizen has a trivially small stake in each of the many items in the government's investment portfolio. See, for example, Zivin and Bridges (2002). But Arrow and Lind claimed no such extension, and Fisher (1973) pointed out (with encouragement from Arrow and Lind) that many public investments are locally and/or regionally focused (e.g. disaster relief, local public works, and many environmental improvements), others focus on particular sectors of the economy, and yet others focus on providing local and national public goods (including public health). In all of these cases, individuals have nontrivial stakes in outcomes of public investments, which undermines any appeal to the Arrow-Lind argument.

A foundation for risk-neutral regulation of health and safety might be sought, instead, in Benthamite welfarism. However, Rawlsian reasoning would question the justification for such a criterion (Rawls 1971). A first-cut conjecture might be that, if individuals could retreat behind a veil of ignorance (VoI) concerning their own circumstances, they would agree to a criterion that would maximise expected quality-adjusted life years, E(QALY), for the population and leave individuals to take their luck. But arguments also might be heard for a 'maximise median QALY' criterion (which would equalise the numbers of people who do better and worse than the median), and for special concern for those with afflictions or exposures that make their individual default QALY prospects unusually poor. There also would be arguments addressed to individual differences in ability to pay for prevention and treatment. Finally, and analogous to Rawls' recognition that distributional guarantees might undermine incentives and generate adverse behavioural consequences, an E(QALY) or similar criterion may well have adverse behavioural consequences.

Suppose, as is reasonable, that the public interest is served when people take prescribed treatments and preventative measures. But to do so, they must overcome inertia. A risk-neutral regulatory posture would leave individuals at the margin indifferent between taking prescribed measures and treatments and letting the exposure or affliction run its course. The behavioural consequences of indifference suggest a public interest in ensuring that prescribed measures and treatments offer substantially better prospects than the no-action default (Carpenter *et al.* 2010).

So, while it may be tempting to invoke a VoI argument in support of an E (QALY) criterion, it is likely that a VoI process would pay more than ordinary attention to the concerns of those at greatest risk, and to the behavioural consequences of failure to ensure that prescribed measures and treatments offer markedly better prospects than just letting exposure and affliction take their course. Even if the conclusion had been different, that is that VoI reasoning supports an E(QALY) criterion, there remains the issue that individual health and safety offer poor prospects for VoI thinking – people simply have too much information about their personal circumstances, and personal health and safety are just too important, for plausible retreat behind a veil of ignorance. So analysts concerned with real-world policy must start with a presumption that the public demands *ceteris paribus* that treatment risk be kept low, typically lower than the risk of harm from exposure or affliction.

3.2 How might death as a kind of harm affect preferences regarding treatment risk?

Reasoning and behavioural evidence suggest the following:

3.2.1. Treatment risk for minor afflictions

Suppose worst-case harm from treatment, however unlikely, substantially exceeds harm from the affliction. For example, suppose there is known to be a small chance that an acne remedy may cause serious illness. We would expect very strong aversion to treatment risk, for at least four reasons: the desire to avoid worst-case harm, the heightened sense of responsibility for any harm that results from volitional action (better to take my chances with the affliction than to proactively undergo a treatment that might really hurt me), there are cost impediments to treatment, and there is inertia when the net EV of treatment is fairly small.

As worst-case harm from treatment becomes large (e.g. death), the Naevdal and Oppenheimer 'not a single step' result becomes relevant. On the other hand, when the chance of serious harm from treatment is miniscule, it likely is ignored entirely.

3.2.2. When worst-case harm from treatment and from no-action are modest and similar

We might expect a degree of risk aversion to treatment risk – three of the four reasons offered immediately above apply to this case. For common ailments that do not threaten great harm, it seems people want the authorities to assure that approved treatments are as risk free as possible. As they see it, there is already enough risk in their world, and they seldom are offered a direct trade of a familiar risk for a new and certifiably lesser risk.

3.2.3. When people know their prospects under no-action are truly grim

The 'minimise EV of harm from all sources' criterion, adjusted for patients' cost of treatment, is more likely to be adopted by people who know their prospects are truly grim in the no-action case. Viscusi and Chesson (1999) report experiments showing that, as the probability of serious harm under no-action gets high, say >0.5, ambiguity-aversion diminishes and ambiguity-seeking choices may be observed. Desperate patients are less ambiguity averse and may even become ambiguity seeking regarding treatment.

3.3 Do drug regulators really show strong ambiguity avoidance?

Given that regulators are tasked with approving only those drugs and treatments shown to be safe and effective, is not easy to distinguish regulatory ambiguity aversion from reluctance to approve treatments that may exceed acceptable rates of adverse reactions. Nor is it easy to devise an alternative strategy that would tolerate more ambiguity without increasing the rates of ADRs deemed acceptable.

However, there has been an observable evolution in regulatory practice. The elimination of treatments that entail unacceptable risks remains the driving goal. But within that rubric, regulators have reduced the 'drug lag', the time required to approve new treatments. Aware of the trade-off involved – there is longitudinal and cross-jurisdictional evidence that reducing the drug lag has been associated with modest but measurable increases in ADRs (Olson 2013; Frank *et al*. 2014) – they have done it cautiously, but they have done it. They have approved accelerated trials of new treatments for people with devastating afflictions and made it easier for such patients to participate in these trials. They also have approved use of treatments known to be risky, under rigorously controlled conditions, for devastating afflictions – for example, *Thalidomide* (notorious for its association with birth defects in the early 1960s) is now approved for treatment of leprosy under rigorous controls. One can always debate the optimality of drug regulatory regimes, but I am not inclined to dispute the general direction in which they are evolving.

4. Conclusions

The foregoing analysis supports several conclusions.

4.1 Re ambiguity and learning

Ambiguity is never preferred *per se*, except perhaps by people in desperate circumstances. Learning opportunities are beneficial and will be sought by decision-makers when there is opportunity to gain from what has been learned. Risky probing of ecosystems with thresholds may be induced, a result that is firmest when the threshold is fixed and the penalty for breaching it is manageably small. This analysis may suggest ways to improve models of testing novel interventions and probing as a kind of active learning in adaptive management.

When there are opportunities to learn and switch, and the cost of a bad outcome from the attempt to learn is modest, a menu of treatments can be beneficial. The benefit is greatest when individual patient characteristics influence treatment outcomes, and the responsibility of discovering the best match of treatment and patient falls naturally upon the patient. This benefit from a menu of treatments arises from patient heterogeneity and the compound-lottery multiple-round structure that facilitates learn and switch, rather than from ambiguity associated with treatment characteristics; and we might expect it to dissipate quickly if menu expansion increases risk.

4.2 Re the regulator of public health and safety as a risk-neutral EV maximiser

For many reasons – the desire to avoid worst-case harm, the heightened sense of personal responsibility for harm that results from volitional action, cost impediments to treatment, and limited motivation to undergo treatments that improve one's chances only marginally – the public demands that the regulator be averse to treatment risk in all but the limiting case where expected outcomes in the default situation are truly grim.

References

Arrow, K. and Lind, R. (1970). Uncertainty and the evaluation of public investment decisions, *American Economic Review* 60, 364–378.

Berry, D. and Viscusi, W.K. (1981). Bernoulli two-armed bandits with geometric termination, *Stochastic Processes and Their Applications* 11, 35–45.

Camerer, C. and Weber, M. (1992). Recent developments in measuring preferences: uncertainty and ambiguity, *Journal of Risk and Uncertainty* 5, 325–370.

Carpenter, D., Grimmer, J. and Lomazoff, E. (2010). Approval regulation and endogenous consumer confidence: theory and analogies to licensing, safety, and financial regulation, *Regulation and Governance* 4, 383–407.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms, *Quarterly Journal of Economics* 75, 643–669.

Fisher, A. (1973). Environmental externalities and the Arrow-Lind public investment theorem, *American Economic Review* 63, 722–725.

Frank, C., Himmelstein, D., Woolhandler, S., Bor, D., Wolfe, S., Heymann, O., Zallman, L. and Lasser, K. (2014). Era of faster FDA drug approval has also seen increased black-box warnings and market withdrawals, *Health Affairs* 33, 1453–1459.

Holling, C.S. (1978). *Adaptive Environmental Assessment and Management*. John Wiley & Sons, New York.

Klibanoff, P., Marinacci, M. and Mukerji, S. (2005). A smooth model of decision making under ambiguity, *Econometrica* 73, 1849–1892.

Millner, A., Dietz, S. and Heal, G. (2013). Scientific ambiguity and climate policy, *Environmental and Resource Economics* 55, 21–46.

Naevdal, E. and Oppenheimer, M. (2007). The economics of thermohaline circulation – a problem with multiple thresholds of unknown locations, *Resource and Energy Economics* 29, 262–283.

von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.

Olson, M.K. (2013). Eliminating the US drug lag: Implications for drug safety, *Journal of Risk and Uncertainty* 47, 1–30.

Randall, A. (2011). *Risk and Precaution*. Cambridge University Press, Cambridge, UK and New York.

Randall, A., Capon, T., Sanderson, T., Merrett, D. and Hertzler, G. (2012). Understanding end-user decisions and the value of climate information under the risks and uncertainties of future climate. Report to National Climate Change Adaptation Research Facility. The University of Sydney, Sydney. Available from URL: http://www.nccarf.edu.au/publications/making-decisions-under-risks-and-uncertainties-future-climates [accessed 12 August 2015].

Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge, MA.

Rogers, D.J. (2013). Managing ecological systems with unknown threshold locations. PhD dissertation, The University of Sydney, Australia.

Schipper, B. (2014). Unawareness – a gentle introduction to both the literature and the special issue, *Mathematical Social Sciences* 70, 1–9.

Stankey, G., Clark, R. and Bormann, B. (2005). Adaptive Management of Natural Resources: Theory, Concepts, and Management Institutions. US Department of Agriculture, Forest Service, Portland (OR), Pacific Northwest Research Station General Technical Report no. PNW-GTR-654.

Traeger, C. (2014). Why uncertainty matters – discounting under intertemporal risk aversion and ambiguity, *Economic Theory* 56, 627–664.

Viscusi, W.K. and Chesson, H. (1999). Hopes and fears: the conflicting effects of risk ambiguity, *Theory and Decision* 47, 153–178.

Viscusi, W.K. and Zeckhauser, R. (2015). Regulating ambiguous risks: the less than rational regulation of pharmaceuticals, *Journal of Legal Studies* 44 (forthcoming).

Zivin, J. and Bridges, J. (2002). Addressing risk preferences in cost-effectiveness analysis, *Applied Health Economics and Health Policy* 1, 135–139.