Health Security and Risk Aversion

Jonathan Herington (Kansas State University)
Bioethics (2016), 30(7): 479-489
Available at: http://onlinelibrary.wiley.com/doi/10.1111/bioe.12255/abstract

Abstract
Health security has become a popular way of justifying efforts to control catastrophic threats to public health. Unfortunately, there has been little analysis of the concept of health security nor the relationship between health security and other potential aims of public health policy. In this paper I develop an account of health security as an aversion to risky policy options. I explore three reasons for thinking risk avoidance is a distinctly worthwhile aim of public health policy: (i) that security is intrinsically valuable, (ii) that it is necessary for social planning and (iii) that it is an appropriate response to decision-making in contexts of very limited information. Striking the right balance between securing and maximising population health thus requires a substantive, and hitherto unrecognized, value judgement. Finally, I critically evaluate the current health security agenda in light of this new account of the concept and its relationship to the other aims of public health policy.

The concept of 'health security' has become a popular justification for policy interventions against rare, unusual and potentially catastrophic threats to public health. Newly emerging and re-emerging diseases (e.g. SARS, Ebola), pandemic influenzas, and the deliberate and accidental misuse of biotechnology have all been identified as potential threats to health security.[1] The phrase appears to be especially successful at mobilising large allocations of resources and attention to the establishment of global surveillance networks, pharmaceutical stockpiles, and expansions in biodefense research.

[1] T.R. Frieden et al. Safer Countries through Global Health Security. *The Lancet* 2014; 383: 764–6; US Department of Health and Human Services. 2009. *National Health Security Strategy of the United States of America.* US Department of Health and Human Services; World Health Organisation. 2007. *World Health Report 2007 - A Safer Future: Global Public Health Security in the 21st Century.* Geneva: World Health Organisation; World Health Organisation. 2010. *Responsible Life Sciences Research for Global Health Security: A Guidance Document.* Geneva: World Health Organisation.

Despite its success within the policy community, there has been precious little analysis of the characteristics which define a threat to health security nor the relationship between health security and other aims of public health policy.[2] Definitions by policymakers have been vague and overly expansive. The 2007 World Health Report, for instance, defined 'global public health security' as 'the activities required, both proactive and reactive, to minimize vulnerability to acute public health events.'[3] The US government defines 'national health security' as 'a state in which the nation and its people are prepared for, protected from, and resilient in the face of incidents with health consequences.'[4] These offer little to distinguish health security from the ordinary business of public health protection. Moreover, the growing literature on the ethics of pandemic preparedness,[5] dual-use research,[6] and public health emergencies,[7] has had little to say about the concept of health security. This conceptual lacuna is especially troubling because it is often claimed that health security threats receive an inordinate amount of attention and resources relative to the risk they pose to population health.[8] Whilst protection from catastrophic threats has *prima facie* importance, we lack the theoretical tools to weigh the competing values at stake, and to adjudicate the

---

[2] pace W. Aldis. Health Security as a Public Health Concept: A Critical Analysis. *Health Policy Plan* 2008; 23: 369–75.
[3] World Health Organisation, *op. cit.* note 1, p. 1.
[4] US Department of Health and Human Services, *op. cit.* note 1, p. 1.
[5] M. Selgelid et al. Infectious Disease Ethics: Limiting Liberty in Contexts of Contagion. *J Bioethical Inq* 2009; 6: 149–52; M. Verweij. Health Inequities in Times of a Pandemic. *Public Health Ethics* 2009; 2: 207–9; M.K. Wynia. Ethics and Public Health Emergencies: Restrictions on Liberty. *Am J Bioeth* 2007; 7: 1–5.
[6] T. Douglas & J. Savulescu. Synthetic Biology and the Ethics of Knowledge. *J Med Ethics* 2010; 36: 687–93; G. Samuel, M.J. Selgelid, & I. Kerridge. Back to the Future: Controlling Synthetic Life Science Trade in DNA Sequences. *Bull At Sci* 2010; 66: 9–20; World Health Organisation, *op. cit.* note 1.
[7] J. Herington, A. Dawson, & H. Draper. Obesity, Liberty, and Public Health Emergencies. *Hastings Cent Rep* 2014; 44: 26–35; M.K. Wynia, *op. cit.* note 5.
[8] W. Aldis, *op. cit.* note 2; S. Elbe. Should HIV/AIDS Be Securitized? The Ethical Dilemmas of Linking HIV/AIDS and Security. *Int Stud Q* 2006; 50: 119–44; C. McInnes & K. Lee. Health, Security and Foreign Policy. *Rev Int Stud* 2006; 32: 5–23.

balance between protecting population health from catastrophic threats and improving population health overall.

In this paper I develop an account of health security which helps to explain the special importance of catastrophic risks to population health. I argue that two standard models of public health policy decision-making do not adequately capture the full range of plausible attitudes towards risky policy options. To capture these attitudes, I introduce the notions of health security and health maximisation. Briefly, "health security" involves distributing security (i.e. cumulative probability) towards initial increments of population health; thereby reducing, *ceteris paribus*, the probability of public health catastrophes. This can be contrasted with "health maximisation", which involves distributing security towards later increments of population health; thus increasing, *ceteris paribus*, the maximum possible values of population health. Our judgement about the trade-off between these two considerations models our degree of risk aversion. I give three reasons placing special emphasis on health security: (i) security may be valuable for its own sake, (ii) securing health may facilitate individual and social planning, and (iii) securing health may be an appropriate response to non-probabilistic uncertainty regarding future health threats. Finally, I briefly discuss the potential policy implications of pursuing health security, and how these inform a critique of current 'global health security' initiatives.

Before I begin, there are three essential preliminary remarks. First, I will not be addressing the debate over the concept of health. Whilst the debate between functionalist, biostatistical, and capability accounts of health is important, it is largely orthogonal to the question of whether we ought to maximise population health or secure it. Likewise, I place to one side the debate

over the appropriate measure of population health. In what follows I will assume a broadly functionalist model of population health, as measured by the Health-Adjusted Life Expectancy (HALE) metric described in the 2010 Global Burden of Disease Study.[9] The HALE of a population represents the average number of healthy life years that an individual born into that population could expect to live, given currently observed rates of morbidity and mortality.[10] Whilst the HALE measure is controversial its use here is relatively benign, since regardless of how we define or measure health we will be faced with a choice between maximising possible population health and minimizing the variability of population health across possible outcomes.

Second, I will place to one side distributional considerations within populations. Because HALE is an average of the health outcomes within a population, it elides the distribution of health within that population. Thus a population with highly variable individual health outcomes may have an equivalent HALE to a population where individual health outcomes cluster around the expected value. To place this consideration to one side, I will proceed as if population health is distributed perfectly equally, so that the HALE for the population is the actual number of years of healthy life lived by *all* individuals within that population.

Third, I place to one side the question of how we ought to balance an appropriate pursuit of population health, against other health-related goals including, *inter alia*: (i) a fair distribution of opportunity for health (ii) the

---

[9] J.A. Salomon et al. Healthy Life Expectancy for 187 Countries, 1990–2010: A Systematic Analysis for the Global Burden Disease Study 2010. *The Lancet* 2012; 380: 2144–62.

[10] Note that "years of healthy life" does not mean the number of years lived at full health, but rather the total number of years of life discounted by the health-related quality of life (HRQL, measured on a scale of 0-1.0, where 0 is death and 1 is full health) of each year. C.J.L. Murray et al. Comprehensive Systematic Analysis of Global Epidemiology: Definitions, Methods, Simplification of DALYs, and Comparative Results from the Global Burden of Disease Study 2010. *The Lancet* 2012; 380: 1–141.

promotion of health autonomy, and (iii) the protection of health-related rights.[11] As I will argue below, the balance between securing and maximising population health is a value judgement which is internal to our goal of promoting population health. How an appropriate policy of promoting population health ought to be constrained by these other considerations is an open question which I will not address here.

I. PUBLIC HEALTH UNCERTAINTY

Whilst there is sometimes disagreement over the precise ends of public health policy, there is broad agreement that at least one of those goals is *promoting* population health. Given two possible public health interventions (P and Q) and a summary measure of population health, *H*, promoting population health involves pursuing the policy which will bring about the greatest value of *H*. Of course, this task is complicated by the enduring presence of uncertainty: born of the limits of population health data and the inherent unpredictability of the future. Whilst P may maximise *H* in one state of affairs, Q may maximise *H* in some other state of affairs. We thus cannot straightforwardly know which policy will *actually* promote population health.

To illustrate this concern, consider two different policies aimed at promoting the HALE of a population. The "Risky" policy involves investing in treatment and prevention schemes which specifically target only a select set of highly prevalent diseases (e.g., malaria, tuberculosis, etc.). The "Robust" policy involves investing in multi-purpose public health projects – such as universal primary care, health education programs and robust access. If the

[11] C. Munthe. The Goals of Public Health: An Integrated, Multidimensional Model. *Public Health Ethics* 2008; 1: 39–52.

aetiological environment stays roughly the same, the Risky interventions will improve the HALE of the population to a much greater degree than the Robust intervention. Of course, there is some non-negligible risk that a novel, devastating disease will emerge in the intervening years: with the potential to severely diminish population health in the absence of a robust public health system. To place distributional considerations to one side, assume that regardless of which outcome occurs, individual health is distributed equally within the population. Let us suppose that there are three equally likely states of affairs, with the following impact on population health given implementation of either the Robust or Risky policy:[12]

	Health-Adjusted Life Expectancy (HALE)	
State of Affairs	**Risky Policy**	**Robust Policy**
A	31	50
B	55	55
C	80	60

Table 1: Health-adjusted life expectancy (HALE) for *Risky* and *Robust*, given three equally likely states of affairs.

Whilst the Risky intervention is likely to promote health to a greater degree than the Robust intervention, it carries with it the risk of a catastrophic collapse in population health. Which policy will best promote population health is therefore uncertain.

Two models of decision-making are frequently invoked in an attempt to resolve this uncertainty. The *maxi-min* model adopts the policy which maximises the minimum possible value of population health. On this view, P is a better public health intervention than Q just in case the minimum possible value of H given implementation of P is greater than the minimum possible

[12] These values are purely illustrative, but they are plausible. Consider that a HALE of 60 years is equivalent to middle income countries such as Bolivia, Belarus and Indonesia, whilst a HALE of 31 years is roughly equivalent to Haiti. J.A. Salomon et al., *op. cit.* note 9.

value of *H* given implementation of Q. Because the worst case scenario given Risky is 31 HALE, and the worst-case scenario given Robust is 50 HALE, a decision maker who adopts the maxi-min model is rationally required to adopt the Robust intervention. Unfortunately, like the precautionary principle,[13] maxi-min is often criticized as an irrationally conservative approach to health policy. Since maxi-min focuses all its decision-making attention on the value of *H* which a policy would preserve in the 'reasonable worst case scenario',[14] it disregards the value of H in all other states of affairs. By doing so, it irrationally requires foregoing almost certain improvements in health in order to avoid the slimmest possibility of catastrophe.[15]

The *expectational* model adopts the policy which maximises expected population health. Popular amongst health economists,[16] this model is concerned with *all* of the possible values of *H* given implementation of a particular policy. It therefore avoids the complaint levelled at the maxi-min model, since both catastrophic and stupendous health outcomes are considered in strict proportion to their probability of occurring. Thus, P is a better public health intervention than Q just in case the probability-weighted average of the possible values of *H* given implementation of P is greater than the probability-weighted average of the possible values value of *H* given implementation of Q. Applied to the case above, the expected HALE of the Risky policy is $55\frac{1}{3}$ years,

---

[13] As usually interpreted, the precautionary principle suggests that: we may implement policy, P, if and only if a catastrophic decrease in H is not a possible outcome of P. Famously, this version of the principle can rule all courses of action as impermissible, if both implementing and *not* implementing P could possibly lead to catastrophic decreases in H. On the contrary, the maxi-min principle will always permit at least one policy to be implemented. See C.R. Sunstein. 2009. *Worst-Case Scenarios.* Cambridge: Harvard University Press.

[14] Science and Technology Committee. 2011. *Scientific Advice and Evidence in Emergencies.* London: UK House of Commons.

[15] See J.C. Harsanyi. Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls's Theory. *Am Polit Sci Rev* 1975; 69: 594–606.

[16] P. Dolan et al. QALY Maximisation and People's Preferences: A Methodological Review of the Literature. *Health Econ* 2005; 14: 197–208.

whilst the expected HALE of the Robust policy is 55 years. Thus a decision-maker who adopts the expectational model is rationally required to adopt the Risky policy.

Intuitively, however, I think many would prefer the caution of the Robust policy to the grave gamble required by Risky.[17] Ensuring a population HALE of at least 50 years, rather than 31 years, seems more important than a 1/3 probability of enjoying a HALE of 80 years. Thus, whilst we are sensitive to the probability of health improvements, we might nonetheless place a disproportionate emphasis on avoiding public health catastrophe, even if it involves choosing a policy which does not maximise expected health.

In the parlance of decision theory, such a judgement exhibits *risk aversion* with respect to population health. As I shall argue below, we can think of the degree of risk involved in a policy as the variance in distribution of probability amongst the possible values of H. Thus, though Risky and Robust have almost equivalent expected values for H, the possible values of H given Robust are relatively stable (between 50 and 60 HALE), whilst the possible values of H given Risky are highly variant (between 31 and 80 HALE). Maxi-min, because it considers *only* the worst possible value of H, is blind to variance. On the other hand, whilst the expectational model is sensitive to all the possible values of H, it is insensitive to any change in variance which preserves the expected value of H.[18] Moreover, by taking each possible value of H into account *in strict proportion* to its probability, the expectational model does not allow for trade-

[17] Indeed, health economists have long noted that our preferences display an aversion to risks of this kind. D. Gyrd-Hansen & I.S. Kristiansen. Preferences for "life-Saving" Programmes: Small for All or Gambling for the Prize? *Health Econ* 2008; 17: 709–20; M. Hoel. Allocating Health Care Resources When People Are Risk Averse with Respect to Life Time. *Health Econ* 2003; 12: 601–8; A. Oliver. A Quantitative and Qualitative Test of the Allais Paradox Using Health Outcomes. *J Econ Psychol* 2003; 24: 35–48.

[18] Assuming all outcomes are equally probable, the following health policies are equivalent according to the expectational model: (A) 10 or 90 years, (B) 30 or 70 years, (C) 50 years.

offs between minimizing variance and maximising expected value. Thus, both the maximin and expectational models are blind to a salient consideration: that some policy actions involve far greater *risk* than others.

II. RISK AVERSION

It is my conjecture that the concept of health security can play an important role in helping us capture this concern for avoiding risk policy options. In order to argue for this connection between health security (as I understand it) and risk aversion, I will examine three potential explanations of risk averse preferences: (i) loss aversion, (ii) diminishing marginal value, and (iii) sensitivity to the distribution of probability. Ultimately I believe the first two accounts fail to fully explain our preferences across cases, and so should be rejected in favour of the third account.

(A) Loss aversion

The first explanation for the judgement that Robust is preferable to Risky is that we implicitly treat losses from the current level of population health as more salient than equivalent gains in population health.[19] This phenomenon, known as loss aversion, is typically invoked to explain why many people would prefer missing out on an opportunity to gain $100 as opposed to losing $100 they already possessed.[20] If we assume that the current level of population health is 55 health-adjusted life years, then the preference for Robust over Risky can be explained by the fact that we disvalue the potential loss of health in Risky (24 years) to a greater degree than we value the potential gain (25

[19] I thank an anonymous reviewer for pushing me to explicate this point.
[20] A. Tversky & D. Kahneman. 2000. Judgment under Uncertainty: Heuristics and Biases. In T. Connolly, H.R. Arkes, & K.R. Hammond, eds. *Judgment and decision making: An interdisciplinary reader* Cambridge: Cambridge University Press p. 35–52.

years). If this were an adequate explanation of risk aversion with respect to population health, then we might conceptualise health security as a special emphasis on ensuring that we do not fall below our current level of population health. Indeed, the security of a good is sometimes characterised simply as "protection against (the) loss" of that good.[21]

We ought to reject this as an explanation of our judgements between Risky and Robust, however, since it fails to explain risk averse preferences with respect to policies which are guaranteed to improve population health. Consider the position of a policymaker in present day Haiti, who is considering whether to implement Risky or Robust (with precisely the same payoffs). Current population health is dismal – a HALE of 30 years – such that both Risky and Robust guarantee at least some improvement. In such a situation, some people's preference for Robust might evaporate, since all possible outcomes are better than the status quo. Nonetheless, Risky still guarantees far less than Robust (an improvement of 1 vs 20 years of HALE), and thus many of us will still prefer the Robust policy.[22] This suggests that loss aversion cannot wholly explain risk averse preferences, and thus that our concept of health security ought not simply be about protecting the current level of population health.

---

[21] M.J. Selgelid. 2012. The Value of Security: A Moderate Pluralist Perspective. In C. Enemark & M. Selgelid, eds. *Ethics and Security Aspects of Infectious Disease Control: Interdisciplinary Perspectives* Burlington: Ashgate p. 35.

[22] To push the point, suppose the administrator of a welfare program had the choice of either allocating $100 directly to a recipient, or allocating $1 to the recipient and placing $99 on a 2:1 pony at Eagle Farm Raceway, with the winnings to be awarded to the recipient. Assuming a fair bookie, the expected utility of both options is equivalent, and so a person concerned *merely* with improving the status quo ought to be indifferent between these two options. Yet almost all of us would fault the administrator for taking such a grave gamble.

(B) Diminishing Marginal Utility

A second explanation for the judgement that Robust is preferable to Risky is that population health has diminishing marginal utility. The "utility" of average population health is simply a measure of its all things considered contribution to aggregate wellbeing.[23] Population health has "diminishing marginal utility" if each additional increment of population health promotes aggregate wellbeing to a lesser degree than the last. For example, an increase in HALE from 50 to 55 would, all over things being equal, increase utility to a greater degree than an increase in HALE from 55 to 60 years. This could explain our preference for Risky rather than Robust, in so far as the increments of HALE between 31 and 50 might have greater utility than each increment between 60 and 80 years. If this is so, then adopting the Robust policy will maximise expected *utility*, even if it does not maximise expected *health*. Nonetheless, there are two reasons to resist this as a full explanation of risk averse preferences.

First, the claim that population health has diminishing marginal utility is not straightforwardly vindicated by studies of individual preferences for population health. To begin with, most studies of health policy preferences focus on the preferred *distribution* of a fixed health gain to different populations.[24] Whilst these studies do suggest that improving the health of badly-off populations is preferred to an equivalent improvement in the health of populations with good health, this is likely because of a preference for socially equitable distributions of health.[25] To avoid contamination with equity

[23] I place to one side whether 'wellbeing' is to be understood as pleasure, preference satisfaction, or the enjoyment of objective goods.
[24] P. Dolan et al., *op. cit.* note 16.
[25] *Ibid.*; P. Dolan & A. Tsuchiya. It Is the Lifetime That Matters: Public Preferences over Maximising Health and Reducing Inequalities in Health. *J Med Ethics* 2012; 38: 571–3; L.

considerations, and given our previous assumption that average population health is distributed equally amongst individuals, we might therefore appeal to data on the marginal utility of *individual* health as a proxy for judgements about the marginal utility of population health. Unfortunately, whilst many people do seem to exhibit diminishing marginal utility for individual health, a significant proportion (between 15-25%) appear to value health linearly.[26] Moreover, there appear to be threshold effects, such that small gains in healthy life have *increasing* marginal utility for individuals at any age.[27] Finally, there is some reason to suspect that these results are an artefact of the model of individual decision-making used in these studies, and instead of being due to the diminishing marginal utility of health are actually the result of an aversion to probabilistic variance (discussed below).[28] In sum, whilst the empirical corpus provides some support for the claim that population health has diminishing marginal utility, it is by no means definitive.

Second, if there is disagreement over the degree to which individual health has diminishing marginal utility, then we ought not assume that average population health has diminishing marginal utility. Consider that the value of additional healthy life is particular to the ends an individual chooses to pursue.[29] For a philosopher, the value of health may diminish quickly, whereas for a triathlete the value of health is more linear. At best, then, we might claim that above the threshold necessary for the pursuit of any life plan,

---

Echazu & D. Nocetti. Priority Setting in Health Care: Disentangling Risk Aversion from Inequality Aversion. *Health Econ* 2013; 22: 730–40.

[26] J.-M. Abellán-Perpiñán et al. Towards a Better QALY Model. *Health Econ* 2006; 15: 665–76; H. Bleichrodt & J.L. Pinto. The Validity of QALYS under Non-Expected Utility. *Econ J* 2005; 115: 533–50.

[27] D. Gyrd-Hansen & I.S. Kristiansen, *op. cit.* note 17; M.K. Kvamme et al. Increasing Marginal Utility of Small Increases in Life-Expectancy?: Results from a Population Survey. *J Health Econ* 2010; 29: 541–8.

[28] J.N. Doctor et al. A New and More Robust Test of QALYs. *J Health Econ* 2004; 23: 353–67.

[29] F.M. Kamm. 1998. *Morality, Mortality: Death and Whom to Save from It.* Oxford University Press p. 242–243.

healthy life has diminishing marginal utility for some individuals but is very near to linear for others. Thus, given a plausible principle of liberal neutrality with respect to the ends individuals seek, we ought not assume that average population health has diminishing marginal utility. These empirical and philosophical considerations therefore suggest that the diminishing marginal utility of population health cannot wholly explain risk averse preferences.

(C) Sensitivity to the distribution of probability

The final explanation for the judgement that Robust is preferable to Risky is that the distribution of probability amongst the possible values of population health is an independently important criterion of decision-making. Given any health policy, P, that policy will distribute probability across the possible values of H, making some values more likely than others (e.g. the Robust policy distributes a probability of 1/3 to the values 55, 60, & 65, and a probability of zero to all other values of HALE). Our judgements about risk may derive from our sensitivity to particular features of such a distribution. For example, policies with a wider dispersal of probability across a larger range of possible values of H have greater "variance" than policies where the distribution clusters tightly around the expected value of H.[30] Thus, our preference for Robust over Risky might be explained by a preference for less probabilistic variance in population health, since Robust distributes probability towards values more tightly clustered around the expected value (i.e. to HALEs of 55, 60, and 65 years), whilst Risky distributes probability more widely (i.e., to HALEs of 41, 60 and 80 years).

[30] Thus, a policy P which is *certain* to result in a HALE of 45 years, no matter the state of the world, has no variance. Likewise, a policy Q which is equally likely to result in any value of HALE between 0 and 90 years is highly variant.

There are strong reasons to take such an explanation of risk aversion seriously. First, there is some empirical evidence that risk averse health policy preferences are the result of an aversion to policies with a high degree of probabilistic variance.[31] Second, this account of risk aversion is compatible with a range of ways in which population health might contribute to aggregate wellbeing. Plausibly, even if the utility of additional increments of population health does not diminish, we might still be averse to policies which involve grave risks (i.e. policies which have a high degree of variance). Thus, if we allow variance aversion to be an explanation of risk aversion, then we can remain neutral with respect to the utility of additional increments of average population health. Finally, this concern for the variance in population health between outcomes, shares an important parrallel with our concern for the variance in individual health *within* a population. In determining public health policy we are typically concerned not only to maximise average individual health within a population, but also to ensure the gap between the best off and the worst off is not too great. Likewise, we might be concerned, not only to maximise expected population health, but also to ensure that the gap between population health in bad outcomes and good outcomes is relatively small.[32] Indeed, there is some evidence that individual preferences for equality in social distributions track individual preferences for risk aversion.[33] Given these three considerations, it is at least plausible that risk aversion with respct to public health policies might be a function of an aversion to probabilistic variance.

---

[31] J.N. Doctor et al., *op. cit.* note 28; A. Oliver, *op. cit.* note 17.

[32] For an argument along these lines at the level of individual health, see A. Voorhoeve & M. Fleurbaey. 2013. Decide As You Would With Full Information! An Argument Against Ex Ante Pareto. In O. Norheim, S. Hurst, N. Eyal, & D. Wikler, eds. *Inequalities in Health: Concepts, Measures, and Ethics* Oxford University Press.

[33] F.A. Cowell & E. Schokkaert. Risk Perceptions and Distributional Judgments. *Eur Econ Rev* 2001; 45: 941–52; L.L. Lopes. Risk and Distributional Inequality. *J Exp Psychol Hum Percept Perform* 1984; 10: 465–85.

III. SECURING AND MAXIMISING HEALTH

How then does the concept of health security help us capture our sensitivity to particular distributions of probability? Whilst there is little contemporary philosophical work on the concept of security,[34] I have argued previously that the core of security is a concern with the goods which we will enjoy regardless of how the future turns out.[35] One way of capturing this concern is through the probability of enjoying *at least* these goods in the future. For example, in determining our security with respect to a liveable income, we are concerned with the probability of enjoying *at least* a liveable income, regardless of whether we are retrenched, change jobs, or receive a raise. Security is a measure of the robustness of our safety net.

Whilst this account suffices for talk of discrete goods, slightly more subtlety is required for the security of incremental goods such as population health. Consider that for each particular increment of population health, there is some probability that the actual value of population health will be greater than or equal to that increment.[36] This is the security of that *increment* of population health. Note that different policies will distribute different degrees of security to different increments of population health. For instance, given implementation of the Robust policy, increments of HALE less than or equal to 50 years are totally secure (see Table 2). Contrast this with Risky, where only those increments of HALE below 31 years are totally secure, and

---

[34] pace S. John. Security, Knowledge and Well-Being. *J Moral Philos* 2011; 8: 68–91; J. Waldron. Safety and Security. *Neb Law Rev* 2006; 85: 454–507; J. Wolff. 2012. Health Risk and Health Security. In R. Rhodes, M.P. Battin, & A. Silvers, eds. *Medicine and Social Justice: Essays on the Distribution of Health Care 2nd ed.* Oxford: Oxford University Press p. 71–8.

[35] J. Herington. forthcoming. Liberty, Fear and the State: Philosophical Perspectives on Security. *Security: dialogue across disciplines* Cambridge: Cambridge University Press.

[36] Thus, for each possible value, h_i, of a population health measure, H, there is some probability, p_i, that the actual value of H will be greater than or equal to h_i. In this sense the security of h_i is the *complementary cumulative probability* of h_i.

increments between 31 and 50 years are secure to degree 2/3. Likewise, whilst the Robust policy provides no security to increments of HALE between 60 and 80 years, Risky provides 1/3 measure of security. In this respect, policies with the same expected value for population health may distribute security differently amongst the increments of population health.

	Probability of HALE ≥ *h*	
h	**Risky**	**Robust**
31	1	1
50	2/3	1
55	2/3	2/3
60	1/3	1/3
80	1/3	0

Table 2: Distribution of security over increments of HALE for *Risky* and *Robust*.

Given this characterisation of the security of incremental goods like population health, is there some all-things-considered sense in which a policy increases health security? Officially, a concern for "health security" is best described as a preference for a particular *distribution* of security amongst the increments of population health. In particular, increasing "health security" involves increasing the security of the *initial* increments of H; potentially at the expense of the security of later increments of H. Thus, Robust provides greater health security (all-things-considered) than Risky, because it provides more security to the initial increments of HALE than Risky. Importantly, by increasing the amount of population health which it is highly probably we will enjoy (at the very least), Robust decreases, *ceteris paribus*, the probability of public health catastrophe. Whilst talk of "health security" is therefore imprecise, it is a shorthand for placing a special emphasis on the security of initial increments of population health.

We can now understand the sense in which a concern for health security captures a concern for avoiding risk. In particular, risky policies "spread" probability more widely, such that they diminish the security of initial increments of HALE (e.g. <55) in order to increase the security of later increments of HALE (e.g. >65). The most risky policies are willing to risk total catastrophe (e.g. population extinction) in order to pursue small probabilities of enormous gains (e.g. immortality). The distribution of security by a particular policy therefore captures the degree of risk it entails.

Making decisions with health security

How then does a concern for health security fit into our overall framework for promoting population health in the context of uncertainty? To begin with, pursuing health security is distinct from the maxi-min model of public health decision-making. Recall that the maxi-min model ignores probabilistic information entirely and rejects any option whose worst-case outcome is less favourable than another option. Increasing health security involves distributing probability in a particular way and does not require that we pay attention only to the 'reasonable worst case scenario' (nor a threshold of population health which we deem 'basic' or 'tolerable'). Instead it places special, but not lexical, priority on the probability of enjoying initial increments of population health. In this respect it is not a decision-rule, but a defeasibly important consideration in decision-making under conditions of uncertainty.

Second, health security can be contrasted with health *maximisation*, which aims at increasing security of later increments of population health, and hence increasing the maximum value of population health which is possible. This often involves pursuing risky health policies, which redistribute security away from initial increments of H in order to improve the security of later

increments of H. In this respect, health security and health maximisation are separate considerations in decision-making under conditions of uncertainty. It is true, of course, that securing and maximising health will often be furthered by the same policy, since increasing the probability of enjoying a very high value of H often increases the *cumulative* probability of enjoying initial increments of H. Nonetheless, as Risky and Robust illustrate, maximisation and security are often in tension with one another. Promoting population health in the context of uncertainty thus requires a value judgement regarding the balance between our interest in health security and our interest in health maximisation. This judgement models our degree of risk aversion with respect to population health.

Third, we can incorporate this judgement into a risk-sensitive calculation of expected value. The technicalities of this are complex, but in broad outline, they involve weighting the importance of each degree of security in the calculation of a policy's expected value.[37] At one extreme, a hyper risk-averse agent is solely concerned with those increments of H which are totally secure (i.e. the minimum possible outcome). At the other extreme, a hyper risk-seeking agent is solely concerned with those increments of H with the least amount of security (i.e. the maximum possible outcome). A risk-neutral agent will be equally concerned with the increment of H at each degree of security, and thus weights each degree of security equally). Finally, a security-conscious (moderately risk-averse) agent places special weight on the increments of H which are highly secure, but is still somewhat sensitive to the increments of H which are least secure. The weighting applied to each degree of security is then

[37] For a discussion of "risk-weighted" decision rules, which typically involve a rank-dependent decision procedure, see L. Buchak. 2013. *Risk and Rationality*. New York: Oxford University Press p. 36–47.

used to discount the contribution of the increments of H at that degree of security in an overall calculation of a policy's "risk-weighted" expected value.[38] On this "security-conscious" model of public health decision-making, P is a better public health intervention than Q just in case the risk-weighted expected value of *H* given implementation of P is greater than the risk-weighted expected value of *H* given implementation of Q.

IV. THE VALUE OF HEALTH SECURITY

There are, I think, three reasons we might consider placing an emphasis on the security of health, and hence being risk averse with respect to health policy. First, security may be intrinsically valuable. Second, health security may be instrumentally valuable in so far as having a set of robust background conditions enables individual and social planning. Finally, health security may be an appropriate response to the ambiguity of epidemiological evidence regarding future health threats.

1. The Intrinsic Value of Security

To begin with, the security of health may be distinctly valuable because, as Michael Selgelid has recently argued,[39] security may be valuable for its own sake. To illustrate, Selgelid's asks us to imagine two possible worlds, World 1 and World 2:

> Both worlds contain equal numbers of exactly similar people who live exactly similar lives The only difference between the two worlds is that World 2 is constantly surrounded by distant orbiting meteors that could randomly (i.e., by chance) collide with and destroy it at any given time…No such meteors surround World 1 ….Let us suppose that, by chance, no meteor ever strikes and destroys the society in World 2.[40]

---

[38] L. Buchak, *op. cit.* note 37.
[39] M.J. Selgelid, *op. cit.* note 20.
[40] *Ibid.*

Selgelid claims that if we believe that the society in World 1 is *ex post* better off than the society in World 2, even knowing that they will actually experience equivalent lives, then we appear to value security as an end in itself.

Many will find such an intrinsic commitment to security baffling, but there are some tentative reasons to endorse the intuition. Consider that some individuals may have an intrinsic *desire* not to be subject to risk, even they are unaware and unaffected by such risks.[41] Or perhaps risk directly diminishes an individual's wellbeing by diminishing their *opportunity* to function in ways which are valuable to them.[42] Whilst these focus on the value of security to individuals, they also explain the value of security in social decision-making. If the presence of risk directly diminishes individual wellbeing, then public health policies which gravely risk population health will likely diminish the wellbeing of (most) individuals within that population.

Moreover, even if grave risks do not directly diminish individual wellbeing, security might still be independently important if we have prioritarian duties to possible, as well as actual, members of a population.[43] This line of reasoning claims that, just as we may have a duty to prioritise the worst-off in the actual circumstances (even at the expense of overall wellbeing), we ought to prioritise the worst-off across all *possible* states of affairs (even at the expense of expected wellbeing).[44] If we place an emphasis on the wellbeing of the worst off possible persons (conditional on their

---

[41] C. Finkelstein. Is Risk a Harm? *Univ Pa Law Rev* 2003; 151: 963–1001.
[42] J. Wolff & A. de-Shalit. 2007. *Disadvantage*. Oxford: Oxford University Press.
[43] L. Bovens. Concerns for the Poorly Off in Ordering Risky Prospects. *Econ Philos* 2015; 31: 397–429; L. Buchak, *op. cit.* note 37, p. 55; M. O'Neill. Priority, Preference and Value. *Utilitas* 2012; 24: 332–48.
[44] D. Parfit. Another Defence of the Priority View. *Utilitas* 2012; 24: 399–440; D. Parfit. 2002. Equality or Priority? The Lindley Lecture. In M. Clayton & A. Williams, eds. *The Ideal of Equality* Basingstoke: Palgrave Macmillan p. 81–125.

existence),[45] we will shun policies which involve grave risks, since these involve diminishing already low levels of wellbeing in the "unlucky" outcomes in order to improve already high levels of wellbeing in the "lucky" outcomes. This kind of "modal" prioritarianism is controversial,[46] but if true then it provides a compelling reason to favour health security. By distributing probability towards initial increments of population health, risk averse policies prioritise the health of populations (conditional on their existence) in the unlucky outcomes (i.e. they ensure that all possible populations have a low probability of public health catastrophe).

These explanations of the intrinsic value of security are admittedly tentative, and much more work needs to be done to establish the value of security to individuals and the connection to security as a population-level good. Nonetheless, if we believe that security directly contributes to individual wellbeing, or that we have duties to prioritise the worst-off (whether they be possible or actual persons), then this may help explain our aversion to risky public health policies.

2. Security and Planning

Second, health security may be instrumentally valuable in so far as it limits the scope of the risks individuals and social planners must take into account when pursuing other goals. To begin with, risks to population health may undermine the ability of individuals to form complex life plans. Living with grave risks imposes a cognitive load, and thus undermine the capability of some agents to

[45] I thank an anonymous reviewer for pushing me to clarify that the assessment of a possible person's wellbeing must be conditional on their existence (since different states of affairs will bring different individuals into existence, and individuals have no complaint if they do not exist). Similar issues are explored in A. Voorhoeve & M. Fleurbaey, *op. cit.* note 45.

[46] M. Otsuka. Prioritarianism and the Separateness of Persons. *Utilitas* 2012; 24: 365–80; A. Voorhoeve & M. Fleurbaey, *op. cit.* note 45.

formulate complex life plans.[47] Indeed, Stephen John has recently suggested that securely enjoying ones 'vital needs' is necessary in order to be able to form 'reasonable plans'.[48] His claim is that whatever ones long-term ends may happen to be, pursuing them requires that an individual presuppose that they will survive long enough to realise that plan. Yet being able to assume that you will continue to *survive* hardly seems sufficient. No decision is taken in isolation, and almost every decision will be swiftly followed by another. Our goals are often tightly connected, such that risk in the pursuit of one end affects the risks involved in pursuit of all others. Indeed, being able to treat the enjoyment of some goods beyond survival as provisionally settled seems to also be a prerequisite for being able to pursue complex ends.[49] In so far as, *ceteris paribus*, risks to population health also increase risks to individual health, then risky public health policies may diminish the capacity of individuals to make complex plans.

Moreover, grave risks to population health may undermine the ability of policymakers to pursue other social goals. Consider that very grave risks to population health do not simply divert public resources away from other projects, but also impose a kind of "planning blight", such that we can no longer assume a stable minimum of population health when deliberating between options in other policy areas.[50] This is worrisome in so far as policymakers must typically take at least some facts as non-probabilistic background assumptions (the absence of war, population size remaining within reasonable bounds, etc.) if they are to reach reasonably determinate conclusions on the proper course of action in other domains (such as education, pension or

---

[47] A. Mani et al. Poverty Impedes Cognitive Function. *Science* 2013; 341: 976–80.
[48] *op. cit.* note 34.
[49] R.E. Goodin. 2012. *On Settling*. Princeton, N.J.: Princeton University Press.
[50] J. Wolff & A. de-Shalit, *op. cit.* note 42.

taxation policy). Grave risks to population health can therefore undermine our ability to make reliable predictions of the effect of particular educational, military and economic policies. Given the centrality of population health to other social and individual goals, it is therefore plausible to suggest that very risky public health policy undermines individual and social planning.

3. Security and the Limits of Epidemiological Knowledge

A final reason to believe that health security may be an independently valuable desideratum of public health policy is that risk-aversion may be an appropriate decision-making strategy in the context of limited probabilistic information. Recall that the expectational model of population health requires very precise information about the probability *of each and every* possible outcome. Often, however, we lack such information, such that in order to operate the expectational model effectively we must make highly subjective, unwarranted, judgements about the probability of particular states.

To begin with, merely determining the *current* state of population health is fraught with difficulty because of a lack of precise data about the prevalence of disease within a population and so must generate estimates of the likely prevalence based on data from past years or similar populations.[51] Moreover, the presence or absence of infrequent but high mortality disease events (such as the 2010 Haiti Earthquake or the 2004 Indian Ocean tsunami) within a data period can pose additional challenges.[52] It should be no surprise, therefore, that

[51] For example, the recent Global Burden of Disease study estimated the HALE of the male population of the UK in 2010 by gathering data on mortality and the prevalence of disease within the UK between the years 1997-2012. J.A. Salomon et al., *op. cit.* note 9; T. Vos et al. Years Lived with Disability (YLDs) for 1160 Sequelae of 289 Diseases and Injuries 1990-2010: A Systematic Analysis for the Global Burden of Disease Study 2010. *The Lancet* 2012; 380: 2163–96: app. 6 & 7.

[52] see R. Lozano et al. Global and Regional Mortality from 235 Causes of Death for 20 Age Groups in 1990 and 2010: A Systematic Analysis for the Global Burden of Disease Study 2010. *The Lancet* 2012; 380: 2095–128.

estimates of current population health often generate very large confidence intervals: i.e. ranges of values within which the data allows us to be, for instance, 95% confident that the true value lies somewhere within that interval. Thus, the 2010 GBD study suggests that we can be 95% confident that the HALE of the male population in the UK is some value between 63.6 – 67.5 healthy years.[53]

The problem is even more acute when it comes to calculating the expected health of populations in the future. Consider that the evidence available to us about the future prevalence and severity of diseases is limited and admits of multiple interpretations. Suppose, for instance, that I know that there has been an influenza pandemic in 1918, 1957, 1968 and 2009, and that their relative health impact ranged between severe and very mild. What is the probability that a pandemic will occur in the next twenty years? How severe ought we expect it to be? What is the probability of each possible impact on population health being the actual impact? Whilst we may be very confident that there is *some* probability that a pandemic of *some* strength will occur, there is a very large range of expected health impact which is compatible with the evidence.

Policies which place an emphasis on increasing the security of initial increments of health, rather than maximising expected health, may be an appropriate response to such probabilistic ambiguity. Indeed, in situations of total ignorance of probabilistic information, it is often argued that we decide by consulting the maximally risk-averse maxi-min principle.[54] Thus, where we have some probabilistic information about future population health, but not

---

[53] J.A. Salomon et al., *op. cit.* note 9, p. 2152.
[54] M.D. Resnik. 1987. *Choices: An Introduction to Decision Theory.* U of Minnesota Press p. 40–45.

particularly precise information, we might think that we should adopt policies that place some (but not all) emphasis on increasing the security of initial increments of health.[55] By ensuring that population health is not too fragile to unpredictable threats, we can hedge against the possibility that (objective) expected population health is at the lower end of the range suggested by our evidence.

V. HEALTH SECURITY POLICY

Given this characterisation of the concept and value of health security, we can make some tentative critiques of the present structure of the health security agenda. Whilst most of these criticisms are highly dependent on the empirical details, the conceptualisation of health security provided in this paper does suggest that we health security should potentially move beyond the focus on "acute", "rapid" and "novel" infectious disease threats.

Health Security Threats

As was suggested above, the distinctive character of securing, as opposed to maximising, health becomes evident when we consider health policy decisions which are fragile to small changes in the circumstances. The decision, for instance, to concentrate solely on the predictable global burden disease, might maximise expected population health at the expense of vulnerability to catastrophic and unpredictable health threats such as pandemics or bioterrorism. Whilst there is no strict characterisation of which changes are relevant to the security of health, I think that a highly relevant set of changes are those related to unforeseeable or highly unpredictable threats to health.

[55] I. Levi. Ignorance, Probability and Rational Choice. *Synthese* 1982; 53: 387–417.

This includes some of the disease threats captured by the current health security agenda, but also expands the list to include complex social, political, and environmental catastrophes. There are three major types of health risks which have particular relevance to the aim of securing health.

First, novel threats to health. As HIV/AIDS and SARS illustrate, novel diseases may have a dramatic impact on population health. Likewise, novel technologies (such as synthetic microbes, 3D printing and nano-medicines), even when they are developed to solve health problems, often increase the risk of health catastrophe. Whilst it is rare that catastrophic changes in population health occur when a new disease emerges, as HIV/AIDS illustrates, it is clearly plausible, and reducing vulnerability to the emergence of such a disease might therefore seem prudent.

Second, infrequent and/or unpredictable health events. Some very serious threats, such as the emergence of novel influenza viruses, occur perhaps once or twice a generation, with no discernible pattern. Likewise, the intentional or unintentional release of a biological agent (e.g. smallpox) is unpredictable, but could have a catastrophic impact on population health. Whilst the aetiological agents behind such threats are known, the impact upon expected population health is difficult to predict. In this sense, our limited epistemic situation with regard to future health seems to suggest that there might be virtue in improving population health across all plausible futures, including those were we unluckily and unforeseeably experience a health emergency.

Third, systemic changes in the environment. Consider, for instance, the health impact of changes in the climate. Whilst there have been some attempts to model climate related health changes, such as the geographical distribution

of mosquito-borne disease (e.g., malaria, Dengue fever), these are rarely factored into long term forecasts of overall population health. Changes in diet, exercise, farming practices and economic growth are unpredictable, but could have a catastrophic impact on the health of vulnerable populations. Whilst it might be the case that most health-related impacts of climate change will be mild, it is plausible that radical changes in the climate (perhaps caused by so-called 'runaway' climate change) will impact upon fragile health systems (i.e. those propped up by vertical health interventions) in catastrophic ways. In this sense, placing some emphasis on improving health across the range of plausible changes in the climate, even if this means failing to maximise expected health, might be a reasonable reaction to the risk of catastrophe.

Health Security Policy

A full account of the value of health security does not merely explain the set of threats which should be constitute the agenda. It also suggests ways in which our response to these threats ought to be conducted. Once again, whilst the empirical details matter greatly, the concept of health security does suggest some potential directions.

To begin with, we can now begin to justify the focus of health security policy on rare, but catastrophic diseases, even when this is not the most efficient use of limited public health resources. Critiques of current health security policy have tended to claim that it misallocates scarce resources towards rare diseases (e.g. Ebola) at the expense of programs that could more efficiently improve population health (e.g. malaria netting). Whilst addressing the likely burden of disease will maximise expected health, it also seems important to protect communities from potential catastrophe, even if that catastrophe is so unlikely that its effect on our overall expected health is

limited. In this respect, preferring policies which secure health to those which increase overall expected health is a prima facie reasonable value judgement.

Second, increasing health security will sometimes involve interventions which are broad-based and not targeted to specific aetiological agents. Consider that two populations might have the same degree of expected population health given the actual distribution of disease, but be vulnerable to unknown or unpredictable health threats in radically different ways. Policy A may improve expected population health through the introduction of a measles vaccination program, whereas Policy B's may improve expected population health through improved access to basic primary care services. The former remains highly vulnerable to unpredictable or novel health threats, whilst the latter is more resilient to the emergence of these threats.

In this sense, my characterisation of health security, decouples it from a strict focus on emergency 'vertical' interventions against specific diseases, and suggests that broad-based 'horizontal' initiatives which improve community resilience may be an integral part of health security policy. Whilst the current response to threats to 'health security', such as severe influenza pandemics, has been to stockpile influenza-specific antivirals and develop pre-pandemic vaccines based upon putative candidate strains (e.g. such as H5N1), these will have very limited impact given even small changes in the aetiology of a pandemic. Indeed, my account of securing health-related goals seems to recommend health policy which focuses on the broad social and economic determinants of health, including access to healthcare, adequate sanitation and the alleviation of poverty.

CONCLUSION

This paper has developed the concept of health security in a bid to help explain why aversion to risky policy options should be a distinctly important aim of public health policy. In particular, I explored how health security ought to be understood, its relationship to health maximisation, and the implications for public health policymaking in the context of uncertainty. Much remains to be done, however, in explicating a full theory of health security. In particular, understanding precisely how such a value ought to be weighed against the other ends of public health, and whether it may serve as a useful organizing principle for public health are interesting questions worth exploring. By examining this concept we ought to have a much clearer understanding of the kinds of value judgements involved in protecting populations from catastrophic health events.