Journal of Public Economics 145 (2017) 94–102

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

The insurance value of medical innovation☆

Darius Lakdawalla [a,c], Anup Malani [b,c], Julian Reif [d,*]

[a] University of Southern California, United States
[b] University of Chicago, United States
[c] NBER, United States
[d] University of Illinois at Urbana-Champaign, United States

ARTICLE INFO

Article history:
Received 25 January 2016
Received in revised form 16 November 2016
Accepted 17 November 2016
Available online 24 November 2016

JEL classification:
I13
I18
J17
O30

Keywords:
Cost-benefit analysis
Health insurance
Medical innovation
Value of health

ABSTRACT

Economists think of medical innovation as a valuable but risky good, producing health benefits but increasing financial risk for consumers and healthcare payers. This perspective overlooks how innovation can lower *physical* risks borne by healthy patients facing the prospect of future disease. We present an alternative framework that accounts for all these sources of value and links them to the value of healthcare insurance. We show that any innovation worth buying reduces overall risk and generates positive insurance value on its own. We conduct a stylized numerical exercise to assess the potential empirical significance of our insights. Our calculations suggest that conventional methods meaningfully understate the value of historical health gains and disproportionately undervalue treatments for the most severe illnesses, where physical risk to consumers is the costliest. These calculations also suggest that the value of physical insurance from new technologies may exceed the financial spending risk that they pose.

1. Introduction

Economists traditionally measure the benefit of a medical innovation as the improvement in health it produces in a person who is already sick (Drummond et al., 2005, Murphy and Topel, 2006). Likewise, economists measure the benefit of healthcare insurance by valuing the reduction in financial risk associated with lower out-of-pocket spending for medical care (Finkelstein and McKnight, 2008, Abaluck and Gruber, 2011, Engelhardt and Gruber, 2011). Studying either innovation or insurance in isolation, however, overlooks fundamental connections between them. As a result, the true economic benefit of medical technology has been inaccurately characterized and measured.

It is certainly true that a medical technology can improve the health of the sick, and that it can raise financial risk for the healthy by imposing the burden of paying for expensive medical technologies in the event of illness. But it also does two other things that affect its value. First, a technology can reduce physical risk for healthy consumers who might get sick.[1] New treatments make illness less unpleasant and thus effectively raise utility in the bad state of the world, just like standard insurance contracts. Failure to account for this feature understates the value of medical technology, particularly when it comes to treating the most severe illnesses in the most risk-averse consumers. Second, medical technology does not merely create financial risk. Rather, it expands insurance possibilities by converting a previously uninsurable physical risk into a potentially insurable financial risk.

We present a framework for valuing a morbidity-reducing medical innovation that brings together all these benefits and costs. To illustrate our key points, consider a healthy consumer facing the risk of

☆ We are grateful to Jeff Brown, Marika Cabral, Amitabh Chandra, Tatyana Deryugina, Don Fullerton, Gautam Gowrisankaran, Nolan Miller, David Molitor, Heidi Williams, and participants at the 2014 American Economic Association meetings, the 2015 American Health Economics Conference, the 2014 ASHEcon meeting, Harvard Medical School, Harvard Law School, the Indiana University Health Economics Workshop, the NBER Joint Healthcare and Insurance Program meeting, the NBER Health Economics Program meeting, the Pritzker School of Medicine, the University of Chicago Applications workshop, the University of Chicago Health Economics Workshop, the University of Illinois CBPP seminar series, UCSB, UIC, and USC for helpful comments. Charles Zhang provided outstanding research assistance.
* Corresponding author.

[1] Our paper focuses on medical innovations that reduce morbidity and their relationship to healthcare insurance markets. See Lakdawalla et al. (2016) for an economic analysis of mortality-reducing innovations, which are related to annuity insurance markets. Both these types of innovation are important, although Murphy and Topel (2006) estimate that the historical reduction in morbidity is more valuable than the accompanying reduction in mortality.

http://dx.doi.org/10.1016/j.jpubeco.2016.11.012
0047-2727/

developing Parkinson's disease,[2] a neurological disorder that reduces patients' quality of life. As is the convention in health economics, we measure the quality of a life-year as a proportion of a year spent in a perfectly healthy state. For instance, a severely ill person might derive 10% of the value from one year of life that a perfectly healthy person would, while a fairly healthy person might derive 90% of this value. Using this construct, Parkinson's might reduce quality of life from, say, 80% of a perfectly healthy year to 40%.[3] If a perfectly healthy life-year is worth $50,000, Parkinson's imposes a cost of $20,000 per year $((80\% - 40\%) \times \$50,000)$. Imagine a new medical treatment that costs roughly $5000 per year and increases quality of life for Parkinson's patients from 40% to 70%. This increase in quality of life is worth $15,000 annually $((70\% - 40\%) \times \$50,000)$ to patients with Parkinson's but costs only $5000 annually. The traditional approach in health economics compares the benefit of $15,000 to the cost of $5000, and computes the net value of the treatment as $10,000.

This calculation neglects, however, the way the medical treatment's introduction also compresses the *variance* in the quality of life between the Parkinson's and non-Parkinson's states. Prior to the availability of treatment, the risk of Parkinson's represented a gamble that could lower quality of life by 40% of a perfectly healthy year, or about $20,000. The treatment transforms this risk into a new gamble that could lower quality of life by just $5000. This reduction in the variance of quality of life outcomes generates value for consumers who dislike risk. Of course, this risk-reduction value is mitigated by the arrival of a new risk, namely the potential $5000 per year expenditure. However, if the treatment is priced to generate consumer surplus, the ex post improvement in health outcomes will outweigh its financial cost. For instance, prior to the development of treatment, Parkinson's imposes a risk of losing $20,000 in reduced health. After its development, the risk of disease is transformed into a $5000 financial risk plus a $5000 health risk. The new medical treatment cut the total risk of Parkinson's in half. Furthermore, the new financial risk created by the treatment can be mitigated or even eliminated by healthcare insurance.

We conduct a stylized empirical exercise to explore the practical relevance of these insights. We calculate the extent to which conventional economic studies such as Murphy and Topel (2006) have underestimated the benefit of new medical technologies by ignoring their insurance value. The physical insurance value associated with aggregate quality-of-life improvements over the past 50 years may add as much as 50% to the conventional value of quality-of-life improvements, depending on how those gains are distributed throughout the population. Our calculations also suggest that the physical insurance value offered by new technologies greatly exceeds the financial spending risks that these technologies pose and that health insurance ameliorates.

Our theoretical framework clarifies the relationship between the value of medical innovation and of healthcare insurance. First, our model implies that medical technology itself acts as insurance. Even if a consumer has no healthcare insurance, technology can reduce the physical risk she faces. In the Parkinson's example, she faced a health risk of $20,000 prior to the technology but just a $10,000 risk after it, even if no healthcare insurance is available. This insight has important implications for health policy. For example, providing consumers with access to better medical technology by encouraging medical innovation may reduce risk more efficiently than providing them with healthcare insurance.

Second, our framework provides insights into the economic relationship between healthcare insurance and medical technology. The existing literature has argued that these two products are complements by showing that the provision of healthcare insurance can drive medical technology (Goddeeris, 1984, Newhouse, 1992).[4] Our approach highlights the possibility of reverse causality. Medical technology converts a physical risk (sickness) into a financial risk (payment for treatment) that can be mitigated by healthcare insurance. Thus, medical technology, by making healthcare insurance more useful for smoothing health-related risk, generates demand for healthcare (and more generally, income) insurance (Weisbrod, 1991).

Third, our framework allows economists to incorporate risk-reduction into existing estimates of the value from medical technology. This correction has the greatest empirical impact on treatments for severe diseases, where risks to consumers are greatest. This insight reconciles the conventional economic approach to valuation with the findings of population surveys suggesting that people prefer to allocate resources to treating severe diseases rather than milder ones, even holding fixed the cost-effectiveness of treatment across the two types of diseases (Nord et al., 1995, Green and Gerard, 2009, Linley and Hughes, 2013). Conventional approaches are hard-pressed to account for this finding.

Our paper unites two large literatures. The first, which estimates the consumer surplus value of health and longevity, has found that advances in medical technology generate enormous value for consumers (Shepard and Zeckhauser, 1984, Rosen, 1988, Murphy and Topel, 2006). Because these studies all operate within riskless environments, their estimates do not reflect any potential benefits accruing from risk reduction. A second, more recent literature has documented that healthcare insurance delivers significant value to consumers (Engelhardt and Gruber, 2011, Verguet et al., 2014, Barcellos and Jacobson, 2015). This is an important finding because it justifies the cost of public healthcare insurance programs, even if they do not generate significant increases in overall health as several studies have found (Finkelstein and McKnight, 2008, Baicker et al., 2013). The framework used in these studies, however, is unable to compare the value of financial healthcare insurance to the value of physical insurance provided by medical technology.

Our paper is related to Philipson and Zanjani (2014), a contemporaneously written paper that, like the present paper, notes that innovation converts non-insurable physical health risks into insurance financial payment risk. Unlike this paper, it uses this point to frame medical R&D expenditures as insurance payments that help protect against physical risks, though ones that only pay off if R&D results in new treatments. By contrast, this paper uses the same point to revise the literature's estimate of innovation to account for its insurance value. Our paper is also related to Hendren (2015), who demonstrates that individuals obtain value from unemployment insurance not only as insurance against unemployment, but also as ex ante insurance value against the risk of losing a job. He estimates that this ex ante value comprises over 35% of the total value of unemployment insurance.

The remainder of this paper is organized as follows. Section 2 provides a model that describes the different components of value of medical innovation that reduces morbidity in the sick state. Section 3 presents the results of our empirical exercises. Section 4 concludes.

2. Framework for valuing medical treatments

Consider an individual who faces a health risk. We are interested in calculating the value of a new medical technology that improves health in the sick state and is priced to generate non-negative consumer surplus even in the absence of healthcare insurance.[5] Our baseline

[2] Parkinson's disease is a progressive disorder of the nervous system that degrades a patient's movements. It typically manifests as a hand tremor but can also cause slowing of movement and slurring of speech, and later dementia. Its most famous patient is the boxer Muhammad Ali. Parkinson's symptoms can now be treated with medications such as Levodopa or MAO-B inhibitors that raise the level of dopamine in the brain.

[3] To be clear, the numbers on the impact of Parkinson's on quality of life in this example are made up.

[4] In general, healthcare insurance is treated as an outward shift in the demand for medical technology (Acemoglu et al., 2006, Blume-Kohout and Sood, 2013, Clemens, 2013). However, Malani and Philipson (2011) observe that healthcare insurance can reduce the supply of human subjects for the clinical trials required for medical innovation. Lakdawalla and Sood (2013) demonstrate that healthcare insurance and medical innovation are complementary in the sense that healthcare insurance reduces the static inefficiency from patents and thus reduces the cost of using patents to encourage innovation.

[5] Lakdawalla et al. (2016)examine the value of technologies that increase longevity by reducing mortality in the sick state.

framework values the technology from the third-party payer's perspective and therefore we equate the marginal "cost" of a technology with its unit price. In reality, the marginal cost of producing many medical innovations lies below its unit price. At the end of the theoretical presentation, we revisit our analysis from the social perspective, which would focus on this lower marginal cost of production rather than price.

The individual derives utility from non-health consumption and from health according to $u(c,h)$. She is either sick with probability π, or well with probability $1-\pi$. Absent medical treatments, health is h^w when well and $h^s<h^w$ when sick. The individual is endowed with income y^w when well and $y^s\leq y^w$ when sick. Let u_j^i denote the marginal utility of good $j\in\{c,h\}$ in state $i\in\{s,w\}$.

We examine a medical treatment that promises an increase in health—specifically a reduction in morbidity—of Δh in the sick state at a price of p to be paid in the sick state. Our theoretical analysis will focus on valuing marginal doses of the technology, i.e., dh and dp, because that will yield the most intuitive expressions for the different components of value.[6] Our numerical exercise, presented in the next section, allows technologies to have discrete benefits and prices. In addition, our theoretical approach here calculates a consumer's ex ante willingness to pay for a new technology. In the Appendix we discuss how to value technology using certainty equivalents, a related approach.

A key assumption that we maintain throughout our paper is that consumers have positive demand for income insurance, i.e., that the marginal utility of consumption is higher in the sick state than in the well state ($u_c^s>u_c^w$). This holds if one or both of the following are true: illness raises the marginal utility of consumption, by affecting the curvature of utility directly; or illness reduces consumption in the sick state by, for example, necessitating the purchase of medical care or reducing earnings, thereby increasing marginal utility as a result. The first condition is sometimes referred to as "positive state dependence". While there is no consensus among economists as to whether consumers exhibit positive state dependence, there is little doubt that the demand for income insurance is positive. Thus our theoretical analysis maintains the weaker assumption about insurance demand, without imposing a specific assumption around state dependence. We note that if this weaker assumption is violated, so that $u_c^s<u_c^w$, our results still obtain, but the sign of the value of insurance flips from positive to negative. In this case, both medical technology and healthcare insurance exacerbate risk.

To clarify our terminology, we conceive of medical technology as providing "physical" benefits—e.g., health improvements—along with financial benefits and costs. Its value thus consists of the "physical value" plus the "financial value". As described in Table 1, each of these values can be further decomposed into the value of: (1) changes in mean physical and financial outcomes; and (2) changes in the variance of physical and financial outcomes. We call the sum of the mean effects the "conventional value" of the technology, since this is what conventional economic analysis (e.g., cost-effectiveness) estimates.

The variance consists of two effects, which we refer to as "physical risk" (lower health outcomes risk), and "financial risk" (greater healthcare spending risk). We call the sum of these two variance effects the "insurance value" of technology. This component is not accounted for by conventional health economic analysis. If the consumer has access to full healthcare insurance, which eliminates spending risk, then it is obvious that medical technology must have positive insurance value. Our analysis will show that the insurance value is always positive, even when healthcare insurance is unavailable.

Table 1
Elements in the value of medical technology.

	Mean	Variance
Physical value	Improvement in health outcomes	Lower health outcomes risk
Financial value	Increase in healthcare spending	Greater healthcare spending risk
Full value	Conventional value	Insurance value

Notes: Traditional cost-effectiveness analysis calculates the conventional value of medical technology. The spending risk component of insurance value is absent if the consumer has access to comprehensive healthcare insurance.

Before deriving our main results, we pause to note that our theoretical framework is quite general and could be applied to any "innovation" that generates high value in a bad state. Health improvements are a natural application because there is a broad consensus among economists that they are valuable, but our analysis may also be enlightening for other applications outside of this context.

2.1. The conventional value of medical technology

The standard approach to valuing medical technology typically proceeds by quantifying how much patients are willing to pay for the technology in the sick state (Drummond et al., 2005). That value, V, is defined implicitly according to the following expression:

$$u(y^s-p-V,h^s+\Delta h)=u(y^s,h^s)$$

Taking the full derivative of this expression with respect to components of technology (Δh and p) and willingness to pay (V) shows the ex post marginal value of technology (dV) to sick patients is:

$$dV=\frac{u_h^s}{u_c^s}dh-dp$$

This expression is the difference between the technology's marginal benefit ($u_h^s dh$) and its marginal price ($u_c^s dp$), normalized by the marginal utility of income (u_c^s).[7] The "conventional" value of health technology, dV_0, is simply this marginal value to sick patients multiplied by the probability of being sick:

$$dV_0=\pi\left(\frac{u_h^s}{u_c^s}dh-dp\right) \tag{1}$$

The conventional value measures what a sick person would pay for the health improvement.

The conventional approach is very often used to determine whether and to what extent third-party health insurers should cover the new technology. Yet, in a world with healthcare insurance, the more salient question is the value of innovation as perceived by healthy people paying premiums or taxes that finance insurance coverage. Unfortunately, the valuations of sick and healthy people coincide only when consumers are risk neutral or possess first-best indemnity insurance against health shocks.[8] Neither of these conditions seems realistic or appropriate.

The remainder of this section provides an alternative approach that values medical technology ex ante, before the health state is realized, and illustrates how medical technology influences the size and nature of health risks borne by consumers. We show how this approach relates to the conventional framework. The "ex ante" value of technology

6 Allowing for endogenous investments in prevention does not affect our analysis. Consider, for example, a new therapeutic treatment for an infectious disease, which can be prevented by avoiding infected individuals. Assuming that prevention is chosen optimally, the envelope theorem implies that the choice of prevention level will have no impact on the value of a new treatment on the margin.

7 As mentioned earlier, our theoretical analysis models the marginal value of the introduction of a new technology. Mathematically, this means we evaluate the derivative at the point $u(y^s,h^s)$.

8 By "indemnity insurance", we mean contracts that pay consumers when a health shock occurs, regardless of how much healthcare they use. In contrast, real-world "healthcare insurance" pays consumers only when they consume healthcare. See Section 2.4 for additional discussion.

consists of the conventional value plus "insurance value", which corresponds to the incremental value placed on medical technology by risk-averse consumers. Under risk neutrality, this insurance value component goes to zero, and the conventional approach coincides with the ex ante value to premium-paying consumers. Thus, in keeping with the terminology shown in Table 1, we will decompose total value into "conventional value" (dV_0) and "insurance value" (dV_I). We further break down insurance value into a "physical" (dV_{I_P}) and a "financial" (dV_{I_F}) component.[9] "Physical" insurance value accrues to risk-averse consumers even when they do not have any financial healthcare insurance. It represents the value of reducing the physical risks of ill health. "Financial" insurance value is the incremental gain to risk-averse consumers from gaining access to financial healthcare insurance.

We derive the total value under three different settings: "no healthcare insurance" (dV^{NHI}), "with healthcare insurance" (dV^{WHI}), and "complete indemnity insurance" (dV^{CII}).

2.2. The total value of technology in the absence of healthcare insurance

We first assume consumers do not have access to healthcare insurance in order to show that technology reduces risk even in the absence of financial healthcare insurance. The willingness to pay for a technology under "no healthcare insurance", V^{NHI}, from the perspective of all consumers who face the relevant health risk, is implicitly defined by:

$$\pi u\left(y^s - p - V^{NHI}, h^s + \Delta h\right) + (1-\pi)u\left(y^w - V^{NHI}, h^w\right) = \pi u(y^s, h^s) + (1-\pi)u(y^w, h^w)$$

The marginal value in this case is given by the difference between the expected marginal benefit ($\pi u_h^s dh$) and the expected marginal price ($\pi u_c^s dp$). Since we are taking the ex ante perspective of the healthy consumer, we normalize by the ex ante marginal utility of income:

$$dV^{NHI} = \frac{\pi(u_h^s dh - u_c^s dp)}{\pi u_c^s + (1-\pi)u_c^w}$$

Rearranging this expression shows that the value of technology with no healthcare insurance, dV^{NHI}, can be expressed as the conventional value, dV_0, plus an additional component that reflects the insurance value of the technology, dV_I^{NHI}:

$$dV^{NHI} = \overbrace{\pi\left(\frac{u_h^s}{u_c^s}dh - dp\right)}^{\text{Conventional value, } dV_0} + \overbrace{\pi(1-\pi)\left(\frac{u_h^s}{u_c^s}dh - dp\right)\left(\frac{u_c^s - u_c^w}{\pi u_c^s + (1-\pi)u_c^w}\right)}^{\text{Insurance value with no health insurance}, dV_I^{NHI}} \quad (2)$$

Insurance value in the absence of insurance, dV_I^{NHI}, is always positive, provided that the technology is priced such that its conventional value is positive, and provided the individual has positive demand for income insurance against the health risk (i.e., $u_c^s > u_c^w$). This important result bears repeating: even absent healthcare insurance, any medical technology that is worth purchasing ex post reduces overall risk ex ante, because the reduction in physical risk more than offsets the increase in financial risk.

The insurance value of technology in the absence of healthcare insurance can be written explicitly as the reduction in physical health risk minus the increase in financial risk:

$$dV_I^{NHI} = \overbrace{\pi(1-\pi)\left(\frac{u_c^s - u_c^w}{\pi u_c^s + (1-\pi)u_c^w}\right)\frac{u_h^s}{u_c^s}dh}^{\text{Reduction in physical risk } (dV_{I_P})} - \overbrace{\pi(1-\pi)\left(\frac{u_c^s - u_c^w}{\pi u_c^s + (1-\pi)u_c^w}\right)dp}^{\text{Increase in financial risk } (dV_{I_F})}$$

The reduction in health risk gets larger as the value of the health improvement, $\frac{u_h^s}{u_c^s}dh$, gets larger. The increase in spending risk gets larger as the technology's price, dp, gets larger. Our numerical exercise will quantify the size of these two insurance components and compare them to the conventional value.

2.3. The total value of technology with healthcare insurance

Healthcare insurance mitigates the spending risk created by new technology and thus boosts the overall insurance value created when new technologies are introduced. Consider an actuarially fair, fee-for-service healthcare insurance contract that pays the consumer $I(p)$ when she falls sick.[10] When $I(p) = p$, the individual has complete fee-for-service (FFS) healthcare insurance; when $I(p) < p$, the individual has incomplete FFS insurance due to, e.g., deductibles, co-payments, annual caps, or other patient cost-sharing features.

In this environment, the individual solves the problem:

$$\max_\tau \pi u\left(y^s - p + \frac{1-\pi}{\pi}\tau, h^s + \Delta h\right) + (1-\pi)u(y^w - \tau, h^w) \text{ subject to } \tau \le \pi I(p)$$

In practice, most consumers are not completely insured against health risks. Therefore, the transfer constraint will typically bind, and the consumer will choose $\tau^* = \pi I(p)$.

Define V^{WHI} as the total value of technology "with healthcare insurance". Using the expression for the optimal transfer τ^*, we can implicitly define this value as:

$$\pi u\left(y^s - p + (1-\pi)I(p) - V^{WHI}, h^s + \Delta h\right) + (1-\pi)u\left(y^w - \pi I(p) - V^{WHI}, h^w\right) = \pi u(y^s, h^s) + (1-\pi)u(y^w, h^w)$$

The corresponding marginal value of technology with healthcare insurance (dV^{WHI}) is given by:

$$dV^{WHI} = \frac{\pi\left[u_h^s dh - u_c^s dp + (1-\pi)(u_c^s - u_c^w)\frac{dI}{dp}dp\right]}{\pi u_c^s + (1-\pi)u_c^w} \quad (3)$$

We can relate dV^{WHI} to the earlier expression for conventional value, dV_0, and insurance value without healthcare insurance, dV_I^{NHI}, according to:

$$dV^{WHI} = \overbrace{dV_0 + dV_I^{NHI} + \pi(1-\pi)\left[\frac{u_c^s - u_c^w}{\pi u_c^s + (1-\pi)u_c^w}\right]\frac{dI}{dp}dp}^{\text{Value of health insurance}, dV_I^{WHI}} \quad (4)$$

The value of technology with healthcare insurance is equal to its conventional value, plus the insurance value that accrues without any healthcare insurance available, plus a component that reflects the incremental value of financial healthcare insurance made possible by technology.

If healthcare insurance is complete, so that $I(p) = p$, then it will perfectly offset and eliminate the financial risk introduced by the new

[9] Using the language of Ehrlich and Becker (1972), one could call dV_I "self-insurance value" when healthcare insurance is absent because it measures the ability of technology alone to reduce risk, and call dV_{I_F} "market insurance value" because it reflects the ability of financial insurance to mitigate spending risk introduced by a new technology.

[10] Because the contract is actuarially fair, the insurance premium is equal to $\pi I(p)$. This means a consumer in the sick state will receive a net transfer of $I(p) - \pi I(p) = (1-\pi)I(p)$ when sick.

technology. Mathematically, if $\frac{dI}{dp}=1$, then $dV_I^{WHI}=dV_{I_r}$. In this special case, the total value of technology is equal to the conventional value plus the value of physical risk reduction:

$$dV^{WHI}=dV_0+dV_{Ip}$$

2.4. The total value of technology under complete indemnity healthcare insurance

What happens if the consumer has access to perfect indemnity insurance, as opposed to healthcare insurance covering only the cost of medical care? While rarely observed in practice, indemnity insurance is frequently assumed in economic models of health for analytical convenience (e.g., Murphy and Topel, 2006). Because the consumer faces no constraints on the amount of money she can transfer across states, she will choose an amount $\tau\sim$ that equalizes the marginal utility of wealth across states, even when she does not have access to medical technology:

$$u_c\left(y^s+\frac{1-\pi}{\pi}\tilde{\tau},h^s\right)=u_c(y^w-\tilde{\tau},h^w)$$

Full indemnity healthcare insurance is fundamentally different from the healthcare insurance we considered earlier, because indemnity insurance operates even in the absence of medical technology. This means that the marginal value of a new technology is measured at the point $u(y^s+\frac{1-\pi}{\pi}\tau\sim,h^s)$, not $u(y^s,h^s)$. We therefore denote the indemnity-insured marginal utility of good $j\in\{c,h\}$ in state $i\in\{s,w\}$ as $u\sim_j^i$. Because $u\sim_c^w=u\sim_c^s$, it is straightforward to show that the value of a new medical technology under "indemnity insurance" is equal to

$$dV^{CII}=\pi\left(\frac{u\sim_h^s}{u\sim_c^s}dh-dp\right)$$

Notice that this expression is substantially similar to Eq. (1), the expression for the conventional value. In fact, if the marginal utilities for dV_0 are calculated in the indemnity insured state, then $dV^{CII}=dV_0$. In that case, the conventional value of medical technology is equal to the value in a setting where consumers face no risk thanks to perfect indemnity insurance. In principle, some differences arise, because conventional approaches typically fail to calculate marginal utilities in the indemnity-insured state. Nonetheless, the structure of the conventional value calculation is identical to that of the indemnity insurance case. In other words, the conventional approach to valuing medical technology coincides with the ex ante value of technology when individuals are fully insured against health shocks by perfect indemnity insurance.

2.5. Implications for valuing health gains

If individuals are risk averse and have positive demand for income insurance—which empirical evidence suggests is true—then our model shows that the conventional valuation of medical technology underestimates the true value. This has important implications for cost-effectiveness analysis, which is widely employed by healthcare systems across the world to determine which medical treatments qualify for insurance coverage. Moreover, economic studies such as Murphy and Topel (2006) abstract away from the insurance value of technology and thus exclude an important source of value associated with health improvements.[11] We return to this point in our numerical exercise.

Our results also have important implications for the relative values of different types of medical technologies. This can be seen by examining the effect of health status in the sick state, h^s, on our analytical expressions for the value of a marginal technology. This is of particular interest because low values of h^s reflect diseases with high "unmet need", e.g., Parkinson's disease, hepatitis C, or amyotrophic lateral sclerosis (ALS). There is much contemporary debate concerning how much insurers should pay to treat these diseases. Our framework suggests that both the conventional and insurance values of medical technology may be higher for more severe diseases. Thus, the total value of a medical technology is higher for diseases with a higher degree of unmet need (i.e., diseases associated with low values of h^s). Moreover, the difference between the conventional value and the total value grows as the degree of unmet need rises. This suggests that errors in the use of the standard approach are most likely for severe diseases with a poor current standard of care. This prediction is consistent with stated preference evidence that people who are sicker have higher willingness to pay for medical technology, even when the health benefits remain constant. Survey evidence suggests that people value a given level of health investment more highly when provided to sicker patients (Nord et al., 1995, Green and Gerard, 2009, Linley and Hughes, 2013).

To understand these arguments, differentiate Eq. (2) with respect to health in the sick state:

$$\frac{dV^{NHI}}{dh^s}=\pi\left(\frac{u_c^s u_{hh}^s-u_h^s u_{ch}^s}{(u_c^s)^2}dh\right)+\pi(1-\pi)\left(\frac{u_c^s u_{hh}^s-u_h^s u_{ch}^s}{(u_c^s)^2}dh\right)\left(\frac{u_c^s-u_c^w}{\pi u_c^s+(1-\pi)u_c^w}\right)$$
$$+\pi(1-\pi)\left(\frac{u_h^s}{u_c^s}dh-dp\right)u_{ch}^s\left(\frac{(u_c^s-u_c^w)\pi-(\pi u_c^s+(1-\pi)u_c^w)}{(\pi u_c^s+(1-\pi)u_c^w)^2}\right)$$

Suppose that, as we maintain throughout the paper, there is positive demand for insurance ($u_c^s-u_c^w>0$). Also assume a slightly stronger version of this assumption: $u_{ch}^s<0$. Then a sufficient condition for $\frac{dV^{NHI}}{dh^s}<0$ is that $u_c^s u_{hh}^s-u_h^s u_{ch}^s<0$ (which holds necessarily for certain classes of utility functions, e.g., Cobb-Douglas). Moreover, since each individual component term of V^{NHI} also falls with h^s, it follows that both the conventional value and the insurance value fall with h^s as well.

From a practical standpoint, our analytical framework can be implemented empirically with knowledge of a few key parameters. Standard cost-effectiveness analyses rely on estimates of the value of health gains ($\frac{u_h^s}{u_c^s}$), the marginal health gain from the new technology (dh), and the marginal cost of the new technology (dp). The two additional parameters in our framework are the probability of illness (π), and the marginal rate of substitution between the sick and well states (u_c^s/u_c^w). The probability of illness can be estimated as the incidence of disease in the population of interest. In principle, the marginal rate of substitution can be recovered by estimating the demand for healthcare insurance. For instance, shocks to the loading cost of healthcare insurance can be used to trace out the marginal rate of substitution as a function of the quantity of healthcare insurance purchased. The advantage of this approach is the use of revealed preference to recover the quantity of interest directly. The drawback is the difficulty of finding suitable experiments for every population of interest. An alternative would be to parameterize a utility function and recover the implied estimates of $u_c^s(y^s,h^s)$ and $u_c^w(y^w,h^w)$. The income in the sick and well states could be obtained from population surveys of health and income, such as the US Medical Expenditure Panel Survey (MEPS). The health levels could be estimated from similar sources. For example, the MEPS measures health-related quality of life using a five-question instrument known as the EuroQol five dimensions questionnaire (EQ-5D). Quality of life is measured using this scale, which ranges from zero to one.

The approach above took the perspective of the third-party payer, for whom marginal cost is equal to price. The social planner's perspective takes marginal cost to be the marginal cost of production. In many cases, e.g., pharmaceuticals, marginal production costs are quite small

[11] Although consumers in Murphy and Topel (2006) have uncertain lifespans, their quality of life at any given age is known with certainty and they have access to perfect credit markets. Thus healthcare insurance has no value in their model.

relative to prices. For example, some estimates suggest that pharmaceutical prices are about 5 times marginal cost (Caves et al., 1991). From the social planner's perspective, therefore, dp is relatively small, because it reflects only the marginal cost of production. As such, the social value of technology is determined primarily by the conventional value plus the value of physical risk-reduction. The value of financial risk-reduction is much smaller.

This also raises the question of innovation incentives. The classic Nordhaus model (Nordhaus, 1969) implies that the social value of the technology should optimally be equal to innovator profits. “Patent race” models (Loury, 1979) and other alternative models take issue with this result, although none of them provides a clear prediction on how to relate profits to value. Under the simpler Nordhaus formulation, our theory predicts that rewards for innovation should rise with higher insurance value.

Our theoretical analysis focused on technologies that treat disease, because the distinction between ex post and ex ante valuation is most relevant here. Technologies that prevent disease are always assessed from the ex ante standpoint of healthy people. A subtler issue arises with treatment technologies that affect incentives for preventive behavior. For example, a new treatment for an infectious disease might affect incentives to prevent infection. In this case, however, the behavioral response is second-order to the introduction of the treatment technology. Thus it has no first-order welfare consequences and does not affect the value of the treatment technology on the margin. It may, however, have some inframarginal value that could be considered in special cases like that of infectious diseases where behavior substantially affects risk.

3. Illustrating the value of medical innovation

There is substantial evidence that the average quality of life has improved dramatically over the past fifty years. The proportion of elderly who are disabled has decreased, and the proportion who are active has increased (Cutler, 2005). Previous work has estimated that the increase in quality of life may be more valuable than the accompanying increase in life expectancy between 1970 and 2000, which itself is valued at roughly \$95 trillion (Murphy and Topel, 2006). We conduct a numerical exercise to understand how Murphy and Topel's estimated value from improvements in quality of life might change when one accounts for the insurance value of medical innovation.

Because data on the marginal costs of medical technologies to consumers (i.e., prices) or to producers (including the cost of innovation and manufacturing) are generally unavailable, our value estimates account only for consumer surplus and assume prices are a fixed percentage of the willingness to pay for health improvements. Focusing on consumer surplus does not affect the policy implications of our findings: most health technology assessment agencies ignore producer surplus in their calculations (Jena and Philipson, 2008), and our estimates of the relative importance of the insurance value of medical innovation are not sensitive to reasonable changes in the price of technology.

We assume that utility is additively separable in consumption and health, and that it takes a CRRA form:

$$u(c,h) = \frac{c^{1-\sigma}-1}{1-\sigma} + \frac{h^{1-\sigma}-1}{1-\sigma}$$

The theoretical model presented in the first half of this paper allowed for only one possible sick state. Here we employ a generalized version that allows for an arbitrary number of sick states. The health status and probability of a particular state i are given by h^{s_i} and π_i, respectively. A medical technology can improve the health of any state i by an amount Δh_i for price p_i.

The proper value of risk aversion among real-world populations remains controversial, with estimates ranging from <1 to >10.[12] We adopt $\sigma=2$ as our preferred estimate. Employing a larger (smaller) value would increase (decrease) our estimates of insurance value.[13]

We assume throughout that income in the well state, y^w, is equal to \$120,000, which is approximately the value of full income for a typical individual (Murphy and Topel, 2006). Full income here embeds all sources of non-health consumption, including leisure. We assume that income decreases by 20% in the sick state ($y^s=0.8y^w$), which incorporates the documented empirical finding that poor health tends to decrease income (Smith, 1999).

Our numerical exercise requires us to measure health in some manner. We pursue this in a manner consistent with the prior literature. We estimate the lifetime benefits of an increase in quality of life, Δh, similar to the measure considered in Murphy and Topel (2006), using data from a nationally representative sample of individuals from the Medical Expenditure Panel Survey (MEPS). We follow Philipson and Jena (2006), who find that average prices for medical technology are about 20% of ex post willingness to pay plus producer surplus. Assuming that producer surplus is roughly equal to price, since production costs are quite small, this would imply that price is roughly 25% of ex post consumer surplus.

We report all estimates of insurance value from an ex ante perspective. Thus, they should be regarded as the values accruing to an individual who is facing a risk of illness rather than to an individual who is already ill. The appendix provides details on how we implement our calculations, and also generalizes our main model to accommodate an arbitrary number of sick states.

Our numerical exercise estimates the conventional and insurance values of technology. We also decompose insurance value into its two subcomponents: physical insurance value and financial spending risk (see Table 1). Absent healthcare insurance, the total value of technology is equal to the sum of conventional value, physical insurance value, and the offsetting financial spending risk. If complete fee-for-service healthcare insurance is available, financial spending risk is equal to 0 and the total value of technology is then equal to the conventional value plus physical insurance value.

Let $f(h^s)$ represent the distribution of health risks. We measure h^s as “quality of life”, employing a widely used and well-validated tool for measuring quality of life known as the EQ-5D (or EQ-5D-3L). The EQ-5D measures quality of life on a scale from zero to one, using answers from five survey questions regarding the extent of the respondent's problems in mobility, self-care, daily activities, pain, and anxiety/depression. All these questions are asked of respondents in the 2000–2003 MEPS, which serves as our host database.

We use the EQ-5D measure to estimate baseline health state status by age group and gender. The measure is summarized in Table 2 for every 10th percentile. The table shows that the 10th percentile of health status for 18–34-year-old males corresponds to an EQ-5D score of 0.726. For each quantile, health status declines with age, as expected. Conditional on age, males are estimated to have a higher quality of life than females. In every group, the 90th percentile enjoys perfect health. In our analysis, we assume that each health status displayed in Table 2 represents health status in an untreated sick state, h^{s_i}. For each gender and age group, there are nine possible states, each occurring with probability $\pi_i=1/9$. Throughout our analysis, we conceive of the health state as a percentage of perfect health, consistent with the literature on quality-adjusted life-years. However, the numerical calculations require us to choose some absolute scale for health. For convenience, we rescale it to range from 0 to 120,000, so that health and consumption are scaled similarly.

Next, we estimate how much a consumer facing the health risk distribution described in Table 2 would be willing to pay, ex ante, for a hypothetical average increase in her quality of life. We assume an average increase in the quality of life of 6000 (or 5% of perfect health), and concentrate all of the gain in the two poorest health states. This average

[12] A less than comprehensive list includes Barsky et al. (1997), Chetty (2006), Cohen and Einav (2007), Kocherlakota (1996), and Mehra and Prescott (1985).

[13] Employing a larger value of σ also increases our estimate of the conventional value, but by less than the increase in insurance value.

Table 2
Average health status for selected quantiles, by age group and gender.

Group	Observations	Quantiles								
		10	20	30	40	50	60	70	80	90
Males (18–34)	11,382	0.726	0.796	0.886	1	1	1	1	1	1
Males (35–49)	11,424	0.681	0.743	0.796	0.835	1	1	1	1	1
Males (50–64)	7998	0.62	0.691	0.74	0.796	0.796	1	1	1	1
Males (65–79)	4344	0.569	0.681	0.704	0.727	0.796	0.796	0.962	1	1
Males (80+)	1120	0.208	0.56	0.638	0.699	0.725	0.761	0.796	0.916	1
Females (18–34)	13,049	0.717	0.787	0.835	0.895	1	1	1	1	1
Females (35–49)	13,351	0.62	0.725	0.787	0.8	0.857	1	1	1	1
Females (50–64)	9210	0.534	0.689	0.725	0.761	0.796	0.826	1	1	1
Females (65–79)	5567	0.332	0.62	0.691	0.721	0.743	0.796	0.814	0.971	1
Females (80+)	2077	0.116	0.427	0.62	0.681	0.696	0.731	0.796	0.844	1

Notes: Table presents pooled, weighted estimates from the 2000–2003 MEPS. Each cell represents the average EQ-5D index for that quantile and group. The EQ-5D index is a measure of quality of life that ranges from 0 (poor health) to 1 (perfect health).

Table 3
Value of a modest health improvement that is concentrated in the two lowest health states.

Group	Conventional	Insurance value		Total
		Physical insurance	Financial spending risk	
Males (18–34)	$1386	$217	$28	$1575
Males (35–49)	$1950	$393	$45	$2299
Males (50–64)	$2629	$683	$69	$3243
Males (65–79)	$3461	$1186	$106	$4541
Males (80+)	$7421	$10,128	$525	$17,024
Females (18–34)	$1785	$336	$39	$2081
Females (35–49)	$2541	$647	$66	$3122
Females (50–64)	$3355	$1143	$102	$4396
Females (65–79)	$5592	$4084	$269	$9408
Females (80+)	$9494	$22,420	$1034	$30,880

Notes: This table displays the annual value of a modest hypothetical increase in quality of life that is concentrated in the two poorest health states, for an individual facing the health risk profile displayed in Table 2. The total value is equal to the conventional value plus the physical insurance value minus the financial spending risk.

health increase is in line with the hypothetical increase considered by Murphy and Topel (2006).[14]

Our results are displayed in Table 3. They show that the total annual value of the health increase is equal to $1575 for males between the ages of 18 and 34. This total can be broken down into $1386 of conventional value and $189 (=$217–$28) of insurance value. Both the conventional and insurance values increase with age because the elderly are less healthy and thus have more to gain from health improvements. Young individuals, by contrast, already have a high probability of enjoying perfect health, which cannot be improved. The conventional value is responsible for the bulk of the health gain when individuals are young, but the fraction of the gain due to insurance increases steadily with age. For the oldest age groups, the insurance value of the health gains significantly exceeds the conventional value. This is due to the large dispersion in health states for the elderly, as shown in Table 2. Because the elderly face the most health risk, they enjoy the highest insurance value from an increase in the quality of life.

We calculate the aggregate per capita lifetime value of these health gains for an 18-year-old by aggregating over age groups. We discount

Table 4
Aggregate lifetime value of the health improvement from Table 3.

Gender	Conventional	Insurance value		Total	Value added by insurance
		Physical insurance	Financial spending risk		
Male	$61,915	$25,378	$1957	$85,336	38%
Female	$84,582	$55,796	$3578	$136,800	62%

Notes: Estimates are weighted to reflect discounting and survival probabilities. The total value is equal to the conventional value plus the physical insurance value minus the financial spending risk.

our calculations by the probability of survival and by a real rate of discount equal to 3%.[15] The results, displayed in Table 4, show that the hypothetical health increase we consider generates about $60,000 and $80,000 in conventional value for an 18-year-old male and female, respectively.

We also compute that the insurance value adds 38 to 62% to the conventional value, as shown in the last column of Table 4. This suggests that the value of advances in the quality of life may be significantly higher than has previously been recognized. The results from Table 3 also suggest that the magnitude depends greatly on how those gains were distributed across the population.[16] The value is greatest if it accrues to the oldest and sickest individuals (those with a high degree of "unmet need"), and lowest if it accrues to those who were already relatively healthy. This reflects our earlier theoretical result that the difference between the conventional value and the total value of health gains grows with the degree of unmet need. To date, relatively little attention has been paid to the *distribution* of historical health gains, and how it influences the total social value of those gains.

Finally, we note that the physical insurance values reported in Table 4 greatly exceed the financial spending risks. Prior studies have argued that healthcare insurance programs provide enormous value by reducing financial spending risk (Finkelstein and McKnight, 2008, Engelhardt and Gruber, 2011, Barcellos and Jacobson, 2015). Our numerical exercise suggests that implementing policies to encourage

[14] Murphy and Topel (2006) assume that advances in quality of life are related to the declines in mortality from 1970 to 2000. Life expectancy for 18-year-olds increased by about 5% during that period.

[15] Survival probabilities are obtained from www.mortality.org. Discount rates are calculated for the midpoint of the age group. For example, the expected conventional value for an 18-year-old male for the period covering ages 18–34 is equal to $\$1,479 \times 17 \times 0.99886/(1+0.03)^{17/2}$, where the first term comes from Table 3, $17 = 34 - 18 + 1$, the third term is the probability of surviving from age 18 to age 35, and the last term is the discount rate.

[16] Like most insurance studies, we implicitly assume that health shocks are uncorrelated over time. Relaxing this assumption would increase the value of insurance because correlated shocks imply even greater physical risk (Kowalski, 2015).

medical innovation may generate even greater insurance value than expanding access to healthcare insurance.

4. Conclusion

When real-world healthcare insurance markets are imperfect, risk-averse consumers derive value from medical technologies that limit the consequences of bad events and expand the reach of financial healthcare insurance.

Our numerical exercise suggests that these theoretical observations are empirically meaningful. New medical technologies provide substantial insurance value above and beyond standard consumer surplus. Under plausible assumptions, the insurance value adds about 50% to the conventional value. Notably, the physical insurance value of technology is generally much larger than the financial insurance value created by healthcare insurance. The latter point suggests that medical technology on its own may do more to reduce risk than healthcare insurance.

Our argument also suggests that the academic literature, which tends to focus exclusively on the standard consumer surplus value of medical technology, may have failed to capture a significant portion of its value. For example, Murphy and Topel (2006) value health increases over the past century at over $1 million per person. Our results suggest that accounting for insurance value could significantly increase their estimates.

The ability of medical innovation to function as an insurance device influences not just the level of value, but also the relative value of alternative medical technologies. The conventional framework understates the value of technologies that treat the most severe illnesses, compared to technologies that treat mild ailments. This helps explain why health technology access decisions driven by cost-effectiveness considerations alone often seem at odds with public opinion. For example, survey evidence suggests that representative respondents evaluating equally "cost-effective" technologies strictly prefer paying for the one that treats the most severe illness (Nord et al., 1995).

From a normative point of view, our analysis also implies that the rate of innovation functions in a manner similar to policies or market forces that complete or improve the efficiency of insurance markets. Increases in the pace of medical innovation reduce overall physical risks to health, and thus function in a manner similar to expansions in healthcare insurance. As a result, policymakers concerned about improving the management of health risks should view the pace of medical innovation as an important lever to use in their efforts. US policymakers have focused their efforts on improving healthcare insurance access and design. While these are worthy goals, medical innovation policy may have an even greater impact on reducing risks associated with bad health shocks.

More practically, our analysis informs the contemporary debate over how new medical technologies should be reimbursed. The United Kingdom provides an instructive example, as the UK health authorities hew closely to the use of ex post consumer surplus as a measure of value for a new technology, and thus a guide to how generously it should be reimbursed. Perhaps as a result, the UK performs poorly in the reimbursement of drugs to treat cancer, which has motivated legislators there to provide exceptional reimbursement for such products, above and beyond what the UK health authorities dictate (Anon, 2010). Controversy has erupted over the appropriateness of this approach, and the legislation has drawn a great deal of criticism (Anon, 2010). Yet, our analysis illuminates how the severe nature of cancer might contribute to the major misalignment between the standard economic approach to valuing medical technology and the preferences of legislators and voters. The policy lesson is that more attention needs to be paid by third-party payers and other health policymakers to covering treatments for severe diseases in order to align payment policies with the values of consumers. Moreover, the standard economic approach to valuing health technology should itself work towards alignment with the preferences of healthy consumers and sick patients.

Appendix A. Accommodating multiple sick states

The model presented in the main text allowed for two health states, one sick and one well. Here we generalize the model to allow for an arbitrary number of sick states. Let the probability of each sick state be π_i, where $i = 1 \ldots N$. Define the probability of the well state as $1-\pi = 1-\sum_{i=1}^{N}\pi_i$. Suppose that, for each sick state, there is a medical technology available that increases health by an amount Δh_i for a price p_i. The conventional value is equal to

$$V_0 = \sum_{j=1}^{N} V_{0_i}$$

where each V_{0_i} is defined implicitly as

$$u\left(y^{s_i} - p_i - V_{0_i}/\pi_i, h^{s_i} + \Delta h_i\right) = u\left(y^{s_i}, h^{s_i}\right)$$

The full ex ante willingness to pay for a technology under "no healthcare insurance" is defined implicitly as

$$\sum_{i=1}^{N}\left[\pi_i u\left(y^{s_i} - p_i - V^{NHI}, h^{s_i} + \Delta h_i\right)\right] + (1-\pi)u\left(y^w - V^{NHI}, h^w\right) = EU$$

where EU, expected utility absent medical technology, is defined as

$$EU = \sum_{i=1}^{N}\left[\pi_i u\left(y^{s_i}, h^{s_i}\right)\right] + (1-\pi)u\left(y^w, h^w\right)$$

The ex ante willingness to pay "with healthcare insurance" is implicitly defined as

$$\begin{aligned}&\sum_{i=1}^{N}\left[\pi_i u\left(y^{s_i} - \sum_{j=1}^{N}\pi_j \overline{p}\left(p_j\right) - V^{WHI}, h^{s_i} + \Delta h_i\right)\right]\\&+(1-\pi)u\left(y^w - \sum_{j=1}^{N}\pi_j \overline{p}\left(p_j\right) - V^{WHI}, h^w\right)\\&= EU\end{aligned}$$

when $N = 1$, we have $(1-\pi) = 1-\pi_1$ and all of the above expressions simplify to the two-state case presented in the main text.

We solve for V_0, V^{NHI}, and V^{WHI} using standard numerical methods. The insurance value is equal to the incremental willingness to pay when accounting for risk: $V_I = V^{NHI} - V_0$. Financial spending risk is equal to the incremental willingness to pay when an individual gains access to financial insurance markets: $V_{I_F} = V^{WHI} - V^{NHI}$. The physical insurance value can then be easily computed as $V_{I_P} = V_I + V_{I_F}$.

Appendix B. Employing certainty equivalents

Define a certainty equivalent as the maximum amount that a consumer is willing to pay to completely insure against risk. For an individual without access to medical technology or financial insurance markets, the certainty equivalent, CE_0, is defined implicitly as:

$$u\left(y^w - CE_0, h^w\right) = \pi u\left(y^s, h^s\right) + (1-\pi)u\left(y^w, h^w\right)$$

Following the introduction of a new medical technology that generates positive consumer surplus, the certainty equivalent, $CE_1 < CE_0$, is defined implicitly as:

$$u\left(y^w - CE_1, h^w\right) = \pi u\left(y^s - p, h^s + \Delta h\right) + (1-\pi)u\left(y^w, h^w\right)$$

The new medical technology reduces the certainty equivalent for two distinct reasons. First, the new technology generates consumer surplus for sick individuals. This is what we call the "conventional value" of technology. Second, the technology generates what we call "insurance value" because the consumer now faces less risk. Note that the first source of value comes from a reduction in the mean and the second comes from a reduction in the variance.

Finally, consider the case where the consumer has access to fee-for-service healthcare insurance. The certainty equivalent, CE_2, is now defined implicitly as:

$$u(y^w - CE_2, h^w) = \pi u(y^s - p + (1-\pi)p, h^s + \Delta h) + (1-\pi)u(y^w - \pi p, h^w)$$

The conventional value (V_0), net insurance value (V_I) and financial spending risk (V_{I_F}) associated with a new technology are equal to the incremental reductions in uncertainty associated with the introduction of the technology and the availability of health care insurance:

$$V_0 + V_I = CE_0 - CE_1$$

$$V_{I_F} = CE_1 - CE_2$$

The physical insurance value can then be defined as $V_{I_P} = V_I + V_{I_F}$. A shortcoming of employing certainty equivalents is that it does not separately identify V_0 (a mean shift) and V_I (a variance shift). This is not a problem for most studies that value insurance because the mean shifts are typically already measured in dollars. For example, Finkelstein and McKnight (2008) subtract out changes in mean medical spending following the introduction of Medicare so that their welfare estimates can be attributed solely to risk reduction. We cannot do that in our setting because changes in health, unlike changes in medical spending, are not measured in dollars.

One might be tempted to estimate V_0 using the willingness to pay method we present in the main text. Doing so can generate nonsensical estimates, however, because willingness to pay is calculated from the perspective of the sick state while certainty equivalents are always calculated from the healthy state, and those states employ different marginal utilities of income. For example, it is possible to generate scenarios where the insurance value is negative even though a consumer has positive demand for income insurance in the sick state. Nevertheless, we obtain similar results overall if we estimate our model using certainty equivalents rather than the willingness to pay method we present in the main text.

References

Abaluck, J., Gruber, J., 2011. Choice inconsistencies among the elderly: evidence from plan choice in the Medicare part D program. Am. Econ. Rev. 1180–1210.

Acemoglu, D., Cutler, D., Finkelstein, A., Linn, J., 2006. Did medicare induce pharmaceutical innovation? Am. Econ. Rev. 96 (2), 103–107.

Anon, 2010. New 50 million cancer fund already intellectually bankrupt. Lancet 376 (9739), 389.

Baicker, K., Taubman, S.L., Allen, H.L., Bernstein, M., Gruber, J.H., Newhouse, J.P., Schneider, E.C., Wright, B.J., Zaslavsky, A.M., Finkelstein, A.N., 2013. The Oregon experiment—effects of Medicaid on clinical outcomes. N. Engl. J. Med. 368 (18), 1713–1722.

Barcellos, S.H., Jacobson, M., 2015. The effects of Medicare on medical expenditure risk and financial strain. Am. Econ. J. Econ. Pol. 7 (4), 41–70.

Barsky, R.B., Juster, F.T., Kimball, M.S., Shapiro, M.D., 1997. Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement study. Q. J. Econ. 112 (2), 537–579.

Blume-Kohout, M.E., Sood, N., 2013. Market size and innovation: effects of Medicare part D on pharmaceutical research and development. J. Public Econ. 97, 327–336.

Caves, R.E., Whinston, M.D., Hurwitz, M.A., Pakes, A., Temin, P., 1991. Patent expiration, entry, and competition in the U.S. pharmaceutical industry. Brookings Papers on Economic Activity. Microeconomics 1991, pp. 1–66.

Chetty, R., 2006. A new method of estimating risk aversion. Am. Econ. Rev. 96 (5), 1821–1834.

Clemens, J., 2013. The Effect of US Health Insurance Expansions on Medical Innovation. National Bureau of Economic Research.

Cohen, A., Einav, L., 2007. Estimating risk preferences from deductible choice. Am. Econ. Rev. 97 (3), 745–788.

Cutler, D.M., 2005. Your Money or your Life: Strong Medicine for America's Health Care System. Oxford University Press, USA.

Drummond, M.F., Sculpher, M.J., Torrance, G.W., O'Brien, B.J., Stoddart, G.L., 2005. Methods for the economic evaluation of health care programmes. OUP Catalogue.

Engelhardt, G.V., Gruber, J., 2011. Medicare part D and the financial protection of the elderly. Am. Econ. J. Econ. Pol. 3 (4), 77–102.

Ehrlich, I., Becker, G.S., 1972. Market insurance, self-insurance, and self-protection. J. Polit. Econ. 80 (4), 623–648.

Finkelstein, A., McKnight, R., 2008. What did Medicare do? The initial impact of Medicare on mortality and out of pocket medical spending. J. Public Econ. 92 (7), 1644–1668.

Goddeeris, J.H., 1984. Medical insurance, technological change, and welfare. Econ. Inq. 22 (1), 56–67.

Green, C., Gerard, K., 2009. Exploring the social value of health-care interventions: a stated preference discrete choice experiment. Health Econ. 18 (8), 951–976.

Hendren, N., 2015. Knowledge of Future Job Loss and Implications for Unemployment Insurance. National Bureau of Economic Research.

Jena, A.B., Philipson, T.J., 2008. Cost-effectiveness analysis and innovation. J. Health Econ. 27 (5), 1224–1236.

Kocherlakota, N.R., 1996. The equity premium: it's still a puzzle. J. Econ. Lit. 34 (1), 42–71.

Kowalski, A.E., 2015. What Do Longitudinal Data on Millions of Hospital Visits Tell Us About the Value of Public Health Insurance as a Safety Net for the Young and Privately Insured? National Bureau of Economic Research

Lakdawalla, D., Sood, N., 2013. Health insurance as a two-part pricing contract. J. Public Econ. 102, 1–12.

Lakdawalla, D., Reif, J., Bauer, D., 2016. Mortality Risk, Insurance, and the Value of Life, Working Paper.

Linley, W.G., Hughes, D.A., 2013. Societal views on NICE, cancer drugs fund, and value based pricing criteria for prioritising medicines: a cross-sectional survey of 4118 adults in Great Britain. Health Econ. 22 (8), 948–964.

Loury, G.C., 1979. Market structure and innovation. Q. J. Econ. 395–410.

Malani, A., Philipson, T.J., 2011. Can Medical Progress be Sustained? Implications of the Link Between Development and Output Markets. National Bureau of Economic Research.

Mehra, R., Prescott, E.C., 1985. The equity premium: a puzzle. J. Monet. Econ.

Murphy, K.M., Topel, R.H., 2006. The value of health and longevity. J. Polit. Econ. 114 (5), 871–904.

Newhouse, J.P., 1992. Medical care costs: how much welfare loss? J. Econ. Perspect. 6 (3), 3–21.

Nord, E., Richardson, J., Street, A., Kuhse, H., Singer, P., 1995. Maximizing health benefits vs egalitarianism: an Australian survey of health issues. Soc. Sci. Med. 41 (10), 1429–1437.

Nordhaus, W.D., 1969. Invention Growth, and Welfare; a Theoretical Treatment of Technological Change. M.I.T. Press, Cambridge, Mass.

Philipson, T.J., Jena, A.B., 2006. Who benefits from new medical technologies? Estimates of consumer and producer surpluses for HIV/AIDS drugs. Forum Health Econ. Policy 9 (2).

Philipson, T.J., Zanjani, G., 2014. Economic analysis of risk and uncertainty induced by health shocks: a review and extension. Handbook of the Economics of Risk and Uncertainty, p. 453.

Rosen, S., 1988. The value of changes in life expectancy. J. Risk Uncertain. 1 (3), 285–304.

Shepard, D.S., Zeckhauser, R.J., 1984. Survival versus consumption. Manag. Sci. 30 (4), 423–439.

Smith, J.P., 1999. Healthy bodies and thick wallets: the dual relation between health and economic status. J. Econ. Perspect. 13 (2), 145–166.

Verguet, S., Laxminarayan, R., Jamison, D.T., 2014. Universal public finance of tuberculosis treatment in India: an extended cost-effectiveness analysis. Health Econ.

Weisbrod, B.A., 1991. The health care quadrilemma: an essay on technological change, insurance, quality of care, and cost containment. J. Econ. Lit. 523–552.