One Medicare for All? The Economics of a Uniform Health Insurance Program

Mark Shepard
Harvard Kennedy School and NBER

Katherine Baicker
University of Chicago and NBER

Jonathan Skinner
Dartmouth College and NBER

September 5, 2019

Abstract

There is increasing interest in expanding Medicare coverage in the U.S., but it is not clear that the current Medicare program is the right foundation on which to build. Traditional Medicare provides uniform coverage to all enrollees. Combined with supplemental plans, Medicare coverage and payments are more generous than public plans in nearly all other developed countries. In this paper, we develop an economic framework to study the tradeoff between a uniform public plan structure like Medicare and one that is more like some other countries' systems: A basic insurance plan that can be supplemented with additional benefits. We argue that three major shifts make a uniform design less efficient today than when Medicare began in 1965. First, rising income inequality makes it more difficult to design a single plan that serves the needs of both higher- and lower-income people simultaneously. Second, the dramatic increase in the availability of expensive medical technology means that funding a generous uniform program is likely to crowd out other redistributive spending. Finally, as medical spending rises, the tax-financing of the system creates mounting economic costs and increasingly untenable policy constraints.

∗ We are grateful to Victor Fuchs for helpful comments, and to the National Institute on Aging (T32-AG000186, P01AG19783, P01AG032952) for financial support. Shepard gratefully acknowledges funding from the National Science Foundation Graduate Research Fellowship. Disclosures: Skinner has consulted for Sutter Health, and Baicker is a Director of Eli Lilly and HMS and a member of the CBO and NIHCM Advisory Boards.

1 Introduction

The United States spends almost 20% of GDP on health care spending, and, at 6% of GDP, federal health spending already represents a quarter of the federal budget. With baby-boomers reaching retirement age and the continuing development of expensive new medical treatments, the Congressional Budget Office (CBO) projects a more than doubling of Medicare spending in the next decade, from $711 billion in 2018 to $1.5 trillion in 2029 (CBO, 2019). Recent proposals to extend Medicare to new populations – ranging from lowering the eligibility age to age 50, or "Medicare for All" (Kliff and Scott, 2019) – further raise the stakes of using the current Medicare program structure for government-funded health insurance.[1]

Most countries face similar pressures arising from an aging population and medical technology growth, but the traditional Medicare program has specific features that distinguish it from many other insurance plans. It provides a uniform benefit with nearly unlimited scope of coverage, even for unproven technologies. By contrast, many countries such as England limit access to new treatments and technologies based explicitly or implicitly on estimates of cost-effectiveness (e.g., Thorlby and Arora, 2019). Unlike many commercial plans for higher-income Americans, the traditional fee-for-service Medicare also has no network restrictions on providers, nor (with the wrap-around plans held by most enrollees) significant deductibles or copayments. And while other countries such as Sweden and Norway provide universal coverage and eschew narrow networks, their reimbursement rates for services are considerably lower than those paid by the Medicare program.[2]

In this paper, we consider the long-term implications of expanding a uniform health insurance system like Medicare to a larger universe of enrollees. We first ask whether the Medicare program currently provides the optimal generosity for a uniform public insurance program, in the sense of capturing tradeoffs between benefits and costs. Second, we address the question of how the benefit structure should change in light of rising income inequality, ongoing medical technology innovation, and the budget pressures imposed by an aging population – even in the absence of any

[1] Frakt and Pearson (2019), for example, has posed the question of whether expanding Medicare will simply spread inefficiency inherent in the current Medicare program to a larger number of patients, given the unwillingness of Medicare coverage decisions to consider cost-effectiveness.
[2] See https://www.oecd.org/health/country-health-profiles-eu.htm.

eligibility expansion. Finally, we examine the effects of an alternative, non-uniform benefit structure on economic efficiency and equity.

To study these questions, we build on a rich literature in health economics and social insurance design to develop a simple economic model of Medicare that incorporates income inequality, medical technology growth, and distortionary taxes. The model allows us to assess how the welfare consequences of Medicare's uniform benefit structure have evolved, as well as the welfare effects of potential alternative public insurance designs. We derive predictions using both a stylized graphical framework and simulations from a calibrated version of the model.

The model suggests that while Medicare's uniform benefit has the advantages of simplicity and lower administrative costs, it comes with a "cost of uniformity" – the fact that high-income households would likely prefer a more generous plan while low-income households would likely prefer lower health care spending and higher take-home pay or more generous non-medical benefits such as food stamps or housing assistance (Baicker, 2000).

We argue that three macro trends have increased this cost of uniformity appreciably since Medicare's creation in 1965. First, income inequality has risen substantially. Rising inequality leads to growing divergence between rich and poor in willingness (and ability) to pay for generous medical care. Second, there have been dramatic innovations in medical technology: there was much less health care available to buy in the 1960s, and even advanced technologies of the day were relatively inexpensive. Third, average marginal tax rates have increased from less than 25% in 1965 to 30% in 2012 (Mertens and Olea, 2018), commensurately increasing the deadweight loss (or economic cost) associated with publicly financed benefits – a trend that will likely continue with the budget pressures from population aging (Baicker, Shepard, and Skinner, 2013).

These changes imply that demand among the rich for generous medical care increasingly diverges from what a uniform public system can afford to fund. While a universal, generous Medicare program was sensible in 1965 when options for treatment were both limited and relatively inexpensive, tax rates were lower, and income more evenly distributed, the efficiency costs arising from maintaining uniform coverage have grown over time. The current benefit design thus may not be a sustainable foundation upon which to expand public health insurance.

We describe an alternative insurance benefit design in which the government provides basic insurance but allows higher-income households to "top up" by purchasing additional coverage for additional services. (Medicare does have in place supplemental "Medigap" plans, but these are

primarily designed to cover copayments and deductibles, rather than cover additional services, and are heavily regulated.)[3] These types of supplemental plans are common in other countries, where governments underwrite a basic insurance plan or mandate the purchase of subsidized, regulated plans, but then allow households to add private supplemental insurance (or even opt-out of the national insurance plan). For example, roughly one-half of Danish individuals hold private insurance plans (Vrangbaek, 2019). In Australia, private insurance offers access to a wider choice of hospitals and faster access to discretionary services (Glover, 2019), while in England 10.5 percent of the population opt for private health insurance (Thorlby and Arora, 2019). Similarly, higher-income households in Germany can opt-in to private health insurance plans with better amenities and coverage (Blümel and Busse, 2019).

Our calibrated model shows that switching from a uniform Medicare benefit to a top-up structure would generate substantial welfare gains in the long term. The distributional implications of such a policy change would depend on the alternative uses to which the resources saved on public insurance would be devoted. Many European countries spend substantially more on other social insurance programs than the US, and some of those non-health programs may themselves yield health benefits (e.g. Bradley and Taylor, 2013; Havighurst and Richmond, 2007; Baicker et al., 2012; although see Papanicolas et al., 2019).

While our model considers benefit design as applied to people age 65+, the implications for Medicare benefit design are clearly amplified under proposals to *expand* the eligible population. For example, while the cost of "Medicare for All" proposals depends crucially on the details of eligibility, coverage, and provider payment rates, the 2016 Bernie Sanders tax proposal to pay for expanding Medicare included a top federal rate of 52%, with state and local taxes on top of that.[4] As medical technology expands, these budget pressures will grow, as will the welfare cost of a uniform benefit. The implication of our simple model is that a more basic public benefit – closer to "Medicaid for All" than to "Medicare for All" – with the option for individuals to opt-in to more generous private coverage could prove to be a higher value, more sustainable alternative.

[3] Medicare Advantage (managed care) plans add additional benefits to the existing package of Medicare, thus making Medicare even more generous; however, in turn they may impose other types of limitations on enrollees.
[4] For example, see https://www.bernietax.com/. His current proposal does not include a specific marginal tax rate.

2 Health Care Demand and In-Kind Transfers

The public sector plays an outsized role in the financing and provision of health care in nearly all developed countries. Health care is often considered a merit good that ought to be available to all – although often without consensus on the quantity or quality of care to which all are entitled. But health care is often privately produced, with economic forces driving supply and heterogeneous patient preferences and needs driving demand. These forces have important implications for the optimal design of public insurance programs.

In this section, we present a stylized model of demand for health care and show how it varies across income groups based on ability to pay. We then consider how this heterogeneous demand interacts with the design of a public health insurance system in a graphical analysis, and consider how the tradeoffs interact with three key macro forces: rising inequality, expensive new health care technology, and rising tax rates.

2.1 A Simple Model of Health Care Demand with Public Benefits

Consider a stylized model of individuals' consumption (*c*) and medical care (*m*) choices, as in past work such as Hall and Jones (2007). A population of individuals $(i = 1,...,N)$ vary in their private income y_i. Individuals derive flow utility from their consumption $u(c_i)$, and medical spending determines their life expectancy, $\lambda(m_i)$, where we assume diminishing returns to more health care spending. The government provides each person with a cash benefit, *R*, and an in-kind medical benefit, *M*, funded by taxes collected according to a progressive tax schedule $T(y_i)$. Individuals choose whether to buy additional top-up care m_i subject to constraints imposed by the system. We write the individual choice problem as:

$$\begin{aligned} &\max_{c_i, m_i} \lambda(M + m_i) \cdot u(c_i) \\ &\text{s.t.} \quad c_i + m_i = y_i - T(y_i) + R \end{aligned} \tag{1}$$

The constraint on m_i depends on the design of public insurance. We consider two possibilities:

1. **Uniform benefit:** $m_i = 0$ *(individuals cannot top-up the public benefit)*
2. **Top-up benefit:** $m_i \geq 0$ *(individuals can top-up above M with private funds)*

Note that this setup also allows for a purely private health system, which is equivalent to a top-up benefit with $M = 0$.

Now consider demand for health care under a completely private system ($M = 0$). The model predicts, not surprisingly, that privately chosen medical spending rises with income. That is, health care spending equates the marginal willingness to pay (or "value") to extend life by an additional year, $WTP_{LY}(c)$, with the marginal cost of doing so, $MC_{LY}(m)$.[5] A key insight of Hall and Jones (2007) is that willingness to pay to extend life rises steeply with income; for the wealthiest in society, the marginal value of another Lamborghini is low, but an additional year of good health in which to enjoy it is nearly priceless.[6] As a result, privately chosen medical spending also increases rapidly with income. Relative to the poor, the rich proceed further up the marginal cost curve, spending more on less valuable services until $MC_{LY}(m)$ equals their much higher value of a life-year.

Figure 1 Panel A graphically depicts an example of this private health care-income relationship in a simple case with two income types, rich (H) and poor (L). The x-axis is medical spending (m), and the green curve is the marginal cost of a life-year (MC_{LY}), as defined above. For a given level of medical technology, marginal cost first rises gradually but then steeply as the limits of medical technology are reached. Unlike a physical production technology, at some point one can no longer "produce" additional units of health simply by spending more.[7]

---

[5] This can be seen from the first-order condition for m, which can be expressed as:

$$WTP_{LY}(c_i) \equiv \frac{u(c_i)}{u'(c_i)} = \frac{\lambda(m_i)}{\lambda'(m_i)} \equiv MC_{LY}(m_i)$$

[6] This is evident from the expression for WTP of a life year in Footnote 4: higher income people have higher consumption, implying a larger $u(c_i)$ and smaller $u'(c_i)$.

[7] An anecdote illustrates this point. An ICU physician, Goetz (2004) stated in a letter to Health Affairs: "Here is an example I have used when teaching medical students and residents: You are taking care of a patient in the ICU. You have done every test and procedure you know to do and have done everything that all the consultants have recommended. I now tell you that you must spend another $5,000 (originally I used $1,000) to improve the patient's quality of care. What would you do with the money? By this point the student or resident is in a bit of a quandary because they are not quite sure how to use the additional money. If there were a continuing positive linear relationship, it should be reasonably easy to suggest more things that result in improved patient care. Generally, the suggestions are more, or repeated, tests and procedures. I respond to the common answers with a statement that if you do more tests or procedures, you could in fact make the patient worse. How? If you do more tests, all tests have false positives and negatives. How will you use results that contradict earlier tests? With again more tests, and the subsequent potential for much more confusion. If you repeat or do another procedure, how do you interpret the results? Also, procedures generally have potential side effects or complications, so again you have a very high risk of NOT improving quality or outcome with more money."

At some point, when all possible treatments are exhausted, marginal cost would be vertical. The WTP for an additional life-year is depicted by the downward sloping blue curve ("WTP_{LY} (*Rich)*") for the rich type and the orange curve ("WTP_{LY} (*Poor)*") for the low-income type. (We defer discussion of the dashed orange curve to the next subsection.) Because willingness to pay is much higher for the rich, the level of medical spending is greater for the rich (m_H) than the poor type (m_L).

While this medical utilization differential is the natural outcome of an unrestricted market for health care, the income-based disparity may be viewed as inequitable if the level of care the poor can afford is below what society has deemed adequate. Cash transfers can narrow the gap, but most countries also use in-kind health insurance benefits as a more direct way of doing so – most likely because of a deep-seated concern about the health of their fellow citizens, or what we call an *egalitarian social preference* for health. We turn next to study how such egalitarian preferences affect the optimal benefit design for health insurance.

2.2 Egalitarian Social Welfare and Uniform Medical Benefits

Absent political constraints, the government has considerable latitude in designing both how transfer programs are funded, and the composition of benefits between in-kind medical care versus cash (or near-cash) benefits The government can levy taxes to fund cash and in-kind medical transfers. But why provide in-kind health insurance benefits at all? The argument in favor of cash benefits is familiar from Econ 101: for the same cost to taxpayers as an in-kind benefit, cash benefits raise recipients' utility more because they do not constrain their choice sets. But there are multiple arguments in favor of in-kind benefits, such as better differentiation between targeted and non-targeted beneficiaries, and taxpayer preferences about recipients' use of resources (see Currie and Gahvari, 2008).

We focus on a rationale based on *egalitarian social preferences* for an equitable distribution of health care.[8] A simple way to capture the idea of egalitarian social preference is to add to private utility an additional term, $e(m_i)$, capturing the value to society of an individual *i*'s access to

[8] Other rationales suggested in the literature include direct externalities from the in-kind good, indirect benefits through the tax/transfer system (e.g., subsidizing goods that complement work, or providing goods differentially attractive to low-ability types), and insurance market failures such as adverse selection. We focus on equity, which seems the most straightforward way to understand a uniform program like Medicare.

medical care – or the extent to which taxpayers care about the health care of others. We assume that $e'(m) > 0$ for low levels of spending but that it declines and reaches zero at some "adequate" level of health spending – such that for the high-income person in our graphical analysis, it has dropped to zero. Returning to Figure 1 Panel A, this egalitarian preference shifts out the social value of health care for the poor type from WTP_{LY}^{Poor} up to the dashed orange line. Society would like the *L* type to obtain medical care $m_L' > m_L$.

In the case of a uniform benefit, the government must decide on a level of care that balances the different demands for the lower income households and for the high income households. One seemingly egalitarian solution would be to make the uniform level of coverage equal to m_H, so that few if any high-income households would want to opt-out into a more generous private plan. This appears to be what we observe with the U.S. Medicare program. Yet the downside of providing the generous plan is the high cost of doing so. Generous medical benefits come at an opportunity cost, including less ability to fund cash transfers or other social programs that lower-income households might prefer.

The economic model presented above suggests an alternate way of setting an optimal uniform public medical benefit: finding a middle ground between the ideal points for the rich and poor. Figure 1 Panel B shows this middle-ground benefit, M_{Unif}^{*}, which is less generous than what the rich would like (m_H) but more generous than the social optimum for the poor (m_L').[9] This result is analogous to the Samuelson condition for an optimal public good: it is set where the average WTP of beneficiaries is equal to the average marginal cost of providing the service.

Even when chosen optimally, there is an implied efficiency loss as shown in the red highlighted areas in Figure 1 Panel B. High-income households would prefer more than what was provided, while low-income households demand less, meaning that they might prefer the cash to high-amenity health care or access to unproven treatments. We call this loss the "cost of uniformity" involved with a uniform benefit system. We discuss next the reasons that this cost of uniformity has likely risen over time.

[9] Technically, the optimum occurs where the curve is vertically equidistant from the blue curve and dashed orange "social value for poor" curve. Note that in our model, lower benefit levels involve more restrictive coverage of specific treatments and amenities, rather than simply imposing higher deductibles and copayments (as one sees under U.S. "bronze" plans on health insurance exchanges).

3 Rising Cost of Uniformity: Medical Technology, Inequality, and Taxes

Uniform benefits provide equal access to health care, but this uniformity comes at a cost when income groups differ in their optimal tradeoff between health spending and other needs. In this section, we ask how this "cost of uniformity" has changed over time. The medical and economic world was quite different in 1965 when Medicare was created. How should we think about the impact of forces like improved medical technology, rising income inequality, and changing tax rates? Using the graphical framework developed above, we argue that these forces tend to imply *rising costs of uniformity*. We present a more formal argument with equations in Appendix A.[10]

3.1 Baseline: Medicare in 1965

The four panels of Figure 2 walk through the basic logic. Panel A starts by applying the model to the environment at Medicare's founding in 1965. There was much less income inequality than today, reflected in the narrower gap between rich and poor's WTP for medical care. Medical technology was much less advanced (and expensive) than today. In 1965, few options for treatment – and particularly for pharmaceutical treatments – were available for common health care conditions. Rather than a production function with gradually diminishing returns, there was relatively little to do beyond relatively few low-cost interventions for common diseases such as cancer and cardiovascular disease.[11] Graphically, this is reflected in the steep asymptote in the green marginal cost curve per life-year at a relatively low level of medical spending.

The key implication of this 1965 environment is that optimal medical spending for rich and poor were not too divergent. Just as today, the rich could certainly afford more and better health care than the poor. But after accounting for the egalitarian social value of providing basic health care to the poor, the optimal spending for rich and poor (labeled as m_H and m_L') are quite close. As a result, optimal uniform benefits (M^*_{Unif}) are also not too different from either group's private optimum, and the loss from the uniform constraint (red area) is small. This logic suggests that when Medicare was established, a single uniform program for rich and poor seniors made good sense. The cost of uniformity was low, so even a small complexity cost from a more flexible top-

[10] To be added in final version of the paper.

[11] Bypass surgery had been developed by 1965, but this was a high-risk and quite rare intervention that was not appropriate for the vast majority of people with cardiovascular disease; statins, stents, and trans-catheter aortic valve replacements (TAVR) were still far in the future.

up system would be enough to tip the scales towards a uniform program. Moreover, this program would optimally be quite generous, since health care was relatively cheap and the government budget (and associated taxes) relatively small.

3.2 Improving Medical Technology

Now consider the impact of improving medical technology from new and more expensive treatments (Cutler et al. (2006), Chandra and Skinner (2012), and Howard et al. (2015)). Treatments for nearly all conditions – but especially heart disease and cancer – have vastly advanced since 1965, but these new treatments are also very expensive. The heart attack that would have resulted in death in 1965 can now be treated and a life extended – but at a cost per hospital admission of $20,000 or more.

Figure 2, Panel B shows how improving medical technology is reflected in our graphical model. The result is a large outward shift and flattening of the green marginal cost curve. This shift reflects the idea that at a given level of *m*, the marginal returns to medical spending are much higher than 1965. The marginal cost curve (= dollars per life-year) is the reciprocal of the marginal returns (= life-years per dollar), so the marginal cost curve is lower and flatter. Of course, the marginal cost curve eventually steepens at a much higher level of *m* – and these are the marginal technologies that are both high-cost and low-value.

The outward shift and flattening of the MC curve have an important implication: privately demanded medical spending for rich and poor are now much further apart. This follows from the simple logic of demand that higher-income households are more willing (and able) to pay the huge bills associated with modern medicine. Indeed, Medicare's annual costs of about $10,000 per elderly enrollee is 58% of the average Social Security benefit ($17,000 in 2018), which is the only source of income for many lower-income seniors. The poor simply cannot afford U.S. modern medicine so will not obtain it without government help.

The egalitarian social value narrows but does not close this growing gap between rich and poor. As a result, the optimal uniform benefit (M^*_{Unif}) is much further apart from the group-specific optima for rich and poor (m_H and m_L') than in the baseline in Panel A. This divergence implies a larger cost of uniformity, reflected in the bigger red shaded area in Panel B.

3.3 Rising Income Inequality

One of the major economic trends of the past 50 years is rising income inequality. While low and middle-income earners have seen modest economic growth since about 1975, the rich have experienced rapid gains. Rising inequality should also be reflected in rising inequality in demand for medical care.

We depict rising inequality in Panel C of Figure 2 via a large upward shift in the willingness to pay (WTP) for a life-year for the rich. (Note that we retain the marginal cost curve for 2015 from Panel B.) The rich therefore proceed much further up the marginal cost curve up to the point that it starts becoming quite steep – reflecting the marginal low-value care. For example, proton beam therapy for prostate cancer is far more expensive (~$50,000) than conventional surgery or radiation therapy ($25,000), but there is little or no evidence of better effectiveness (Schroeck et al., 2017). Higher-income people may prefer it because of convenience, but the additional cost wouldn't be worth it to low-income people.[12] Indeed, the example of proton beam therapy shows how rising inequality works together with expensive new medical technologies to produce the patterns we describe, since rising income inequality would not have much impact on medical spending gaps without the existence of expensive care to use.

Therefore, rising income inequality – in conjunction with improving technology – results in a growing divergence in optimal medical spending between rich and poor and a growing divergence between the optimal top-up and uniform benefits. The red shaded "cost of uniformity" grows larger versus Panel B. Further, the rich are increasingly dissatisfied with the care provided by the uniform medical benefit, which falls short of the generosity they would privately choose.

The poor, while not dissatisfied with their generous medical care itself, would increasingly prefer less generous medical care (perhaps closer to Medicaid in cost) and more help with non-medical needs like their mortgage and their grandkids' college tuition bills. This is the paradox of the egalitarian motive to provide equitable access to health care. While leveling the health care playing field, it comes at the opportunity cost of forgoing other public assistance that the poor and middle class might prefer. This opportunity cost becomes sharper over time as medical costs rise and inequality grows, making a basic top-up program increasingly attractive even to the non-rich.

[12] Other examples of tradeoffs between costs and convenience/quality come from comparing Medicaid to Medicare. Medicaid is lower cost than Medicare but provides less choice of physicians (about 30% of doctors do not accept Medicaid). Medicaid is more aggressive about managing care, denying claims, and requiring prior authorization for expensive treatments (Gottlieb, Shapiro, and Dunn, 2018).

3.4 Rising Deadweight Loss of Taxes

Although top income tax rates have fallen since the 1960s, average overall marginal tax rates are higher today than in 1965 (Mertens and Olea, 2018). Moreover, the large federal debt and impending cost of Social Security and Medicare for an aging population suggest that tax rates will likely rise further, which suggests rising deadweight loss (DWL) of taxation. What effect does this have in our framework? The idea, which we show more formally in our model in Section 4, is that when the marginal tax rate is 50 percent (for example), the efficiency costs (whether through reduced labor supply, capital accumulation, or tax avoidance) of raising an extra dollar of revenue is higher. Thus, the publicly provided health insurance must be that much more valuable to society. We depict this in Panel D of Figure 2 via a leftward shift (i.e., reduction) in the optimal uniform benefits.

Although higher DWL of taxes would also reduce optimal benefits under a top-up system, the welfare impact is larger under a uniform system. With a top-up system, the rich can purchase additional care above the public benefit and obtain their privately optimal generosity (m_H). But with the uniform system, the rich must consume only the public benefit – unless they completely opt out and pay for all care on their own. But the uniform benefit has moved further away from what the rich would privately choose. This increases the total size of the red shaded area, which was previously set to optimally balance needs of rich and poor but is now more tilted towards the poor. Intuitively, higher DWL of taxes requires cutting the public benefit, which in turn makes the constrain on topping-up in the uniform system more costly.

We have presented a simple graphical model that describes a variety of factors that may affect the desirability of a uniform health insurance plan like Medicare, but one might reasonably ask whether these factors are important in more than a theoretical sense. For this reason, we turn next to a simulation model of health insurance demand across the income distribution (rather than just for rich and poor), using plausible parameters based both on the empirical literature and calibrated to fit the observed patterns of health care utilization during the past 50 years.

4 A Simulation Model of Medicare Benefit Design

Our goal is to study two key design choices for any public health insurance: (1) how generous should the public benefit be, and (2) should it be a fixed *uniform benefit* for all recipients or a *basic*

benefit that recipients can *top-up* using their own money? To do so, we first discuss how top-up benefits are optimally set in our graphical model in comparison to uniform benefits. We then specify a simple simulation model with parameters drawn from the literature to simulate our predictions of how secular trends in income inequality, medical care technology, and the marginal cost of taxation affect the relative advantages of a uniform insurance benefit compared to a top-up system similar to those used in other countries.

4.1 Tradeoffs between Uniform vs. Top-Up Medical Systems

To consider the tradeoffs between the uniform vs. top-up medical benefit designs, we first derive the optimal policy for each design. To do so, we focus on the simple two-type (rich and poor) case that can be shown graphically, as introduced in Section 2. While stylized, this case conveys nearly all the intuition of the more general case. Appendix A[13] presents formal mathematical expressions defining optimal policy using the full model above.

Figure 3 plots the optimal determination of top-up benefits. The top-up design retains a medical floor to promote equity for the poor, but it does not impose a medical ceiling on the rich. However, it also creates additional administrative/complexity costs (κ^{Top}), shown by the upward shift of the green marginal cost curve and the highlighted losses in the green area. The optimal top-up benefit (M^*_{Top}) is set based on the social optimal level for the poor – where the marginal cost curve intersects the social WTP for a life-year. The rich top-up by purchasing health care privately up to their desired value.

Notice several implications of optimal top-up benefits in comparison to the uniform design:

- ***Public* health insurance benefits are lower under top-up benefits** – (i.e., $M^*_{Top} < M^*_{Unif}$)

  This result follows from the logic of the public goods problem. The uniform benefit determines medical spending for everyone, while the top-up benefit determines medical spending only for the (poorer) constrained group who choose not to top-up. Eliminating uniformity frees up the government to set up a low-cost, basic public benefit based on demand by poorer households, rather than having to balance the desires of the rich and poor in a single system, or (as Medicare appears to be designed) a uniform program targeted to the needs of higher income households.

[13] To be added.

- **However, *total* medical spending (public and private) may be higher or lower in the top-up design.** It will be lower for poor households (who receive a smaller benefit) but larger for the rich (who top-up), so the overall change depends on the shape of the marginal cost curve and the size of each group in the population.
- While not evident from the graphs, the math in Appendix A shows that under the parameters we assume in the model, **optimal *cash transfers* can be higher under a top-up medical system.** Intuitively, cash and health insurance benefits are substitutable forms of redistribution. As noted by Baicker (2000) and Bradley and Taylor (2013), when the government spends less on health care, there is more to redistribute as cash (or use for other purposes). This fact is important for interpreting the equity implications of the two designs. While the poor get less generous health care under top-up benefits, they may also get more cash income.[14]
- **The top-up system (by design) allows for greater medical inequality.** This health care disparity for rich vs. poor is often seen as inequitable – even if it originates from private choices and allows the government to fund more cash transfers.

Whether a uniform vs. top-up design is preferred depends on the relative losses from uniformity (red areas) versus losses from administrative complexity (green area). We turn next to building a simulation model that can capture these tradeoffs empirically.

4.2 Model Setup

To judge the magnitude of these effects in practice, we calibrate a model matched to U.S. data with income heterogeneity, income redistribution, technology growth, and an egalitarian preference that values longevity for those with low income levels; the details are included Appendix B.[15] The model provides first for the distortionary effects of taxation necessary to pay for in-kind transfers. The marginal excess burden of these taxes will be higher when either current tax rates are high (as in many European countries), or when tax rates are projected to rise in the future (as in the U.S.). Second, we include a more realistic health production function. The

[14] This result depends on the preferences of taxpayers; if they only care about health care and nothing else, they may not choose to distribute more in cash.
[15] To be added.

assumption of unbounded production functions (as in Hall and Jones, 2007), where any extension to longevity can be purchased with sufficient spending, generates health spending as the ultimate luxury good – with the rich devoting a higher share of their income to health care than the poor. A bounded health production function capturing more closely both the empirical evidence (Acemouglu et al., 2013). Finally, we allow for changes over time in income inequality (Piketty and Saez, 2014), which we hypothesize will affect the efficiency costs of in-kind transfers.

In the model, the government solves the following public policy problem. It provides cash transfers (*R*) and a medical benefit *(M*), both of which are financed by taxes. To allow for progressivity in a simple framework, we model taxes as a scaled version of a baseline progressive schedule $T_0(y_i)$, so that $T(y_i) = \tau \cdot T_0(y_i)$ where τ is the scale factor. Rather than model the full labor supply problem, we model deadweight loss in a simple way. We assume that levying taxes that reduce individuals' consumption by \$1 only raises tax revenue of $\$(1-\chi(\tau))$, where $\chi(\tau)$ captures revenue leakage due to the excess burden of taxes. As taxes rise, $\chi(\tau)$ will increase due to the rising marginal excess burden of higher tax rates.

Putting these pieces together, the government chooses the tax scalar (τ), cash transfer (*R*), medical benefit (M), and design of medical benefits to maximize the social welfare function:

$$SW = \sum_i \left[\lambda(m_i) u(c_i) + e(m_i) \right] \tag{2}$$

subject to public budget constraint:

$$\sum_i \left(R + M \cdot (1+\kappa) + E \right) = \sum_i \tau \cdot T_0(y_i) \cdot (1-\chi(\tau))$$

and where $\{m_i, c_i\}$ are set by individual choices from the problem in (1) and κ captures the excess administrative costs of providing in-kind medical benefits as opposed to cash. This extra term is included because in practice, a top-up system creates additional complexity, including greater administrative costs and adverse selection problems associated with letting consumers choose among generosity levels; we model this additional cost by assuming that the incremental costs of providing a top-up system are positive: $\kappa^{Top} > \kappa^{Unif}$. (These costs were another reason why a uniform Medicare program in 1965 made sense.) Finally, *E* is additional "extra" government spending, which captures other public programs that are not transfers.

4.3 Model Calibration

The simulation model is based closely on the public policy problem in equation (3) and the individual decision problem in equation (1). We calibrate the model's parameters to capture the growth in income, income inequality, and medical spending over the 1968-2045 period.[16] (We begin in 1968 to allow the phasing in of Medicare benefits and enrollment.) We then solve the model under different policies at 10-year intervals (1968, 1975, 1985, …, 2045) for a population of 100 individuals who each represent one percentile of the income distribution.

To match the growth of medical technology, we specify a functional form of the health production function, $\lambda(m)$ (see Appendix B for details). We calibrate its parameters in each simulation year to match two moments: (1) the historical or projected Medicare spending per beneficiary, drawing on projections from the Medicare trustees, and (2) historical or projected average life expectancy at age 65 from the Social Security Administration (Bell and Miller, 2005). By flexibly calibrating $\lambda(m)$ in each simulation year, we capture the way that technology improves to both extend life and to induce people to spend more on medical care. Calibrated medical spending rises from 4% of average income in 1968 to 15% in 2015 and to 24% in 2045.

Two key specification choices are for the egalitarian social externality of health care ($e(m)$ in the model) and the excess burden of taxes ($\chi(\tau)$). We specify the egalitarian externality as an additional social value of longevity, $\lambda(m_i)$, valued at a multiple of the flow utility of consuming the year's median income, $u(\bar{y})$. By using a fixed $u(\bar{y})$ for all individuals, it effectively places more weight on improving the health of the poor relative to their private value of health. We specify the excess burden term as $\chi(\tau) = \varepsilon_{TI} \cdot \tau \cdot MTR_0$, where ε_{TI} is the elasticity of taxable income (assumed to be 0.5) and $\tau \cdot MTR_0$ is the average marginal tax rate for tax scalar τ.[17]

We model government policies as follows. For simplicity, we model a fixed exogenous cash transfer of R = 5% of the year's average income. We also assume "extra" government spending (on non-transfer programs like defense) of E = 10% of average income. With these two policies set, the government makes a single policy choice of the generosity of health insurance benefits

[16] We draw historical income data from the Census/ACS, with future values projected forward using an assumed 1.5% real annual growth rate. See Appendix B for details. All dollar values are inflation-adjusted to 2015 dollars.

[17] This is derived from a calculation of average revenue leakage (starting from an initial tax of 0) assuming a constant elasticity of taxable income. Appendix B presents a derivation. Note that this differs from the more familiar formula for the *marginal* revenue leakage of a tax increase (see Saez, Slemrod, and Giertz, 2012).

(M). Taxes (τ) are set at the level needed to balance the government budget. We consider three policies for health insurance benefits:

1. **Generous uniform:** A uniform benefit (no top-up allowed) with the benefit level exogenously set the level of medical care a rich individual (95th income percentile) would privately choose to purchase; we believe this is a reasonable approximation to Medicare's generous plan.
2. **Optimal uniform:** A uniform benefit (no top-up allowed), with the level of M set to maximize the egalitarian social welfare function. This would correspond to a Medicare-style uniform plan but with likely more restrictive benefits.
3. **Optimal top-up:** A basic benefit M that individuals can privately top-up (with $m_i \geq 0$) if they choose. The level of M is set to maximize egalitarian social welfare, recognizing that some individuals will choose to top-up.

We think of the generous uniform policy as roughly analogous to Medicare in that it provides extensive coverage and provider choice so that even the rich are happy with it. Because its generosity is set mechanically – not to maximize social welfare – it continues to grow even in the face of budget pressures and changing economic fundamentals. Optimal uniform and top-up benefits, by contrast, reflect the model's optimal tradeoff between equity and efficiency given the respective benefit structures. For top-up benefits, we assume an additional administrative cost of $\kappa = 2\%$ (with $\kappa = 0$ for the uniform systems), which ensures that the top-up policy need not always be socially preferable.

4.4 Simulation Results

Medical Spending Inequality

A key premise of our analysis is that inequality in *private demand* for medical care across the income distribution has widened over time. Although it is difficult to observe real-world measures of the income elasticity of private medical demand – not influenced by public programs or employment-based health insurance[18] – we can test this idea in our simulation model.

[18] Perhaps the best source of information is individual market health insurance choices in exchanges. Demand estimates from these settings are consistent with higher-income people being less price-sensitive – and therefore more sensitive to quality – in plan choices (Shepard 2016, Jaffe and Shepard 2019, Tebaldi 2018). However, there are few measures on whether or how this income elasticity has changed over time.

Figure 4 plots privately chosen medical spending for various income percentiles (25^{th}%, 50^{th}%, 75^{th}% and 95^{th}%) in a simulation with cash redistribution (R = 5% of average income) but no public health insurance benefits (M = 0). All values are reported as a share of the average income in each year; this adjusts for general economic growth but makes levels comparable across groups for a given year. Medical spending rises sharply for all groups over time, reflecting the improved medical technology in our calibrated health production function.

But in addition to general growth, *inequality* in health spending across incomes widens. While the rich always buy more health care than the poor, the gap grows as health care becomes more expensive relative to income. In 1968, the 95^{th} income percentile spends 4.1% of mean income on health care (about $2,500) versus 1.4% of mean income (about $880) for the 25^{th} percentile. By 2045, these numbers grow to 23.9% of mean income ($31,000) for the 95^{th} percentile versus 9.6% of income ($12,500) for the 25^{th} percentile. The spending gap between rich and poor grows from 2.7% to 14.3% of mean income. Most of this widening gap reflects the simple fact that medical spending grows as a share of income. As a result, similar *proportional* gaps – the 95th percentile spends about 2-3 times the 25^{th} percentile in all years – implies much larger *absolute* gaps in desired spending.

Public Health Insurance Benefits

Under a uniform insurance system, the government must choose a single level of health insurance benefits for everyone. But because of rising inequality in desired spending, any single choice will increasingly diverge from the private optimum for many income groups. This presents a dilemma for policymakers. If they choose benefits to satisfy the rich (as in the "generous uniform" policy), this choice will be increasingly inefficient for the middle class and poor who would prefer other forms of support. But if they choose benefits to satisfy the middle class, the rich will become increasingly dissatisfied.

Figure 5 shows how our simulated policymaker balances this tradeoff under the three policies we consider: generous uniform, optimal uniform, and top-up benefits. It plots the chosen level of the public medical benefit (M) over time, both in levels as a share of average income (Panel A) and in terms of its percentile in the privately medical distribution from Figure 4 (Panel B). Recall that the generous uniform policy is mechanically set at the 95^{th} percentile of private medical spending, so it grows accordingly. The optimal uniform and top-up benefits also grow but at a slower rate so that they diverge from the generous uniform policy over time. In 1968, all three

policies are similarly generous: 4.2% of mean income for generous uniform, 3.8% for optimal uniform, and 3.6% for top-up – or a gap of 0.6% of income from most to least generous policies.

By 2015, the gap has widened, with benefits of 14.7%, 12.4%, and 11.7% of mean income for the three policies – or a gap of 3% of income. By 2045, it has still further widened to 24.0%, 18.5%, and 16.7% of the mean income – or a gap of over 7% of mean income. These numbers also indicate that the gap between optimal uniform and top-up benefits grows over time (from 0.2% of income to 1.8% of income), another prediction of our graphical analysis.

Panel B shows this divergence another way. While the generous uniform policy grows with the preferred spending of the rich, both the optimal uniform and top-up policies cannot keep up with this growth. They fall relative to the distribution of privately optimal health spending, from the 90th and 85th percentiles in 1968 down to the 69th and 59th percentiles in 2045. This distributional fall also suggests the growing rationale for a top-up system. In 1968, just the richest 10% of people want health care more generous than the optimal uniform benefit. By 2015 this share grows to 22% and by 2045 to 31%. This suggests that the political economy of a uniform system without a top-up would be increasingly difficult to sustain (Baicker et al. 2013).

Underlying these simulation findings is the changing tradeoff between efficiency and (egalitarian) equity being made by the model. In the years following 1965 when health care was cheap, the cost of providing generous uniform insurance for all was low. But over time, the efficiency cost rises for two reasons. The first is the rising inequality in privately desired health spending highlighted above. A single uniform policy is increasingly divergent from what individuals would have chosen on their own.

The second reason is that the rising cost of public health insurance pushes up tax rates and therefore the marginal deadweight loss of taxes. Generous medical care becomes increasingly unaffordable even for the government. Figure 6 plots the top marginal tax rate under the different policies. This top rate rises from about 30% in 1967 to 66% in 2045 under generous uniform insurance. Shifting to the less generous optimal uniform or top-up benefits allows for a more moderate top marginal rate of about 50% in 2045.

Distributional Incidence and an Alternate Top-Up Policy

How do the uniform and top-up policies affect overall welfare and the welfare of different income groups? In the simulations above with fixed cash benefits, less generous uniform health insurance benefits (along with lower taxes) tend differentially to benefit the rich who pay higher

taxes. Moreover, shifting to a top-up system further benefits the rich because it removes a constraint on high-income people's desire to top-up the public benefit. It also allows the government to reduce the generosity of the public benefit, which is reflected in the lower *M* under top-up in all years in Figure 5A. In this sense, shifting from generous uniform to either optimal uniform or top-up is at baseline a regressive change.

Governments may choose to devote some of the reduced spending on health care to increased cash benefits, which would be a simple way to make the move away from a generous uniform medical benefit more progressive. We consider a policy exercise in which, starting from either the optimal uniform or top-up policies above, we increase cash transfers (*R*) until all individuals in the bottom half of the income distribution have utility at least equal to what they received under generous uniform benefits. Initial results (to be reported) appear to support the view that there exists a Pareto improvement in the sense of improving welfare across the income distribution.

5 Discussion and Conclusion

Means-tested in-kind transfers of housing, food, and health care are the predominant form of income redistribution to low-income households (Currie and Gavhari, 2008). Medicare is a prominent example of a uniform in-kind benefit provided to both high- and low-income populations. In this paper, we develop a model that allows us to gauge the tradeoffs involved in this uniform benefit design.

Using a stylized model implementing heterogeneous agents, we show how tax distortions, income inequality, egalitarian preferences, and technology growth affect the efficient structure of the program. Our results suggest that in 1965 when Medicare was first created, its uniform generous structure was well suited to the economic and technological environment. But by 2019 it has become much less efficient relative to a "top-up" health insurance program with basic coverage supplemented by private health insurance. Our results are consistent with many other developed countries, where the prevailing approach is to offer a basic universal public insurance plan with the opportunity for wealthier enrollees to pay extra for amenities and additional services.

Our model also helps to explain the puzzle of why the share of GDP devoted to health care grew so much faster in the U.S. than in some other countries. The combination of a U.S. system with an in-kind benefit driven by the preferences of higher-income taxpayers coupled with low

excess burdens from taxes provided an ideal environment for rapid growth since 1965. By contrast, many European countries had already elected to provide universal coverage, and given their high existing marginal tax rates, therefore had limited scope for expanding benefits (Chandra and Skinner, 2012). By contrast, Medicare covered only a fraction of the population, and with low marginal tax rates, the fiscal impact of expanding generous Medicare coverage was limited until recently. Expanding Medicare to a greater share of the population would only hasten the anticipated growth in tax rates necessary to fund it.

Our model says nothing about whether a Medicare-for-All program will increase or reduce total (public plus private) spending on health care (see Katz et al., 2019). Even if total spending remained the same, there would be a dramatic shift in how health care is paid for, with tax-financed public spending replacing insurance premium payments by individuals and employers.

We explore the implications of an alternative "basic" form of public insurance that provides more restricted benefits and allows higher income households to top-up their coverage with privately-financed plans. Under such a plan, lower income households would consume less health care than their higher-income counterparts, and perhaps less care than they do now. This naturally raises concerns about equity. Many who support a uniform benefit structure point to the "right to health care" as a foundational rationale. It is worth noting, however, that a uniform benefit is likely to devote a substantially higher share of income to health care (rather than, for example, food, education, or housing) than the typical lower-income household would choose, and is likely to result in fiscal pressures that not only raise taxes but also crowd out spending on other public goods and transfer programs (Baicker, 2000). Ultimately, the distributional implications of the plan redesign hinge on the use to which the reduced public health insurance expenditures are put – such as into increased public spending or reduced marginal tax rates.

This "basic benefit" model is similar to the voucher-like premium support suggested by Emanuel and Fuchs (2005), Aaron and Reischauer (1995), and Ryan (2012). Reinhardt (2012) suggested a similar three-tier system, with high-income households eligible for "the sky's the limit" insurance, largely unsubsidized by the government, middle-income households receiving a "reference pricing" insurance policy that reimburses only for the lowest-cost regional provider, and low-income enrollees receiving Medicaid-like public care under strict cost control rules.

It is important to note that the benefit limitation proposed here is quite different from a high-deductible health plan. Rather than increasing cost-sharing on covered services, our basic benefit

limits which services are covered based on the health benefits they produce. For example, the basic benefit might cover all services with cost-effectiveness ratios under $300,000 per life year, and might only cover the cost of such services at efficient providers.[19] In this way, it is consistent with value-based insurance design (Chernew et al., 2007) and with reference pricing, but not with catastrophic-only coverage.

Of course, identifying ex-ante which care is of sufficiently high value is no simple task, but public and private insurance plans are already experimenting with mechanisms for tailoring coverage in this way (Chernew et al., 2007; NICE; Oregon priority list). Importantly, individuals wanting more extensive coverage of treatments and/or providers could purchase top-up coverage using private funds, rather than stepping outside of the insurance plan to pay entirely out-of-pocket, as is the case for, e.g., concierge medicine or treatments not covered under the National Health Service in England.

A similar rationale applies as well to other programs such as education. Universal higher education provided by the government and financed largely by taxpayers has been proposed by several 2020 presidential candidates and is provided in many European countries. In practice, because of budgetary pressure, many such "universal" higher education systems have converged to provide lower levels of per-student expenditures, with wealthier students opting out into private universities.

Of course, political pressures may play out differently with different plans and different eligible populations. Following the birth of the National Health Service in Britain in 1948, cutbacks were quickly implemented in response to overly optimistic budget projections (Lane, 2012; pp. 191-93). A similar evolution could occur under a "Medicare for All" plan, in which budgetary pressures and monopsony power could lead to an erosion in reimbursements and a

[19] The importance of the distinction between higher cost-sharing and limited service coverage is shown in the different predictions of economic theory for these two dimensions of insurance. As our analysis and that of Hall and Jones (2007) show, health care is a normal good, with the rich demanding and being able to afford more of it. By contrast, conditional on an amount of spending when sick, health insurance is an inferior good, given the usual assumption of utility with decreasing absolute risk aversion. Intuitively, a given size risk is more costly for a poor person for whom this risk constitutes a larger portion of income. Our model abstracts from the cost-sharing dimension and considers only the amount of health care purchased. A more general theory of health insurance design might include higher copayments or deductibles, perhaps for low-value care, but this is a topic beyond the scope of this paper, which is focused primarily on choosing among different insurance plans rather than health-seeking behavior conditional on coverage.

curtailment of benefits, creating a national program that more closely resembles “Medicaid for All” but with people free to opt out into private plans.

As new technologies arrive with ever-larger price-tags, pressure will continue to mount on public budgets; and equality of access to care, rather than guaranteed access to a minimum level of care, will become increasingly costly. It is vital that policymakers consider how alternative program designs affect the overall wellbeing of households across the income distribution as they debate Medicare’s future.

References

Acemoglu, D, Finkelstein A, Notowidigdo, MJ, Income and Health Spending: Evidence from Oil Price Shocks. The Review of Economics and Statistics Vol. XCV October 2013.

Annual Report of the Boards of Trustees of the Federal Hospital Insurance and Federal Supplementary Medical Insurance Trust Funds. Washington: Government Printing Office; 2011.

Baicker K, Chandra A. Medicare Spending, the Physician Workforce, and Beneficiaries' Quality of Care. Health Affairs. 2004;23(3):W4184-97.

Baicker K, Chandra A, Skinner J. Saving Money or Just Saving Lives? Improving the Productivity of the U.S. Health Care Spending. Ann Rev Econ. 2012;4:33-56.

Baicker K, Skinner JS. Health Care Spending Growth and the Future of U.S. Tax Rates. In: Brown J, editor. Tax Policy and the Economy. Chicago: University of Chicago Press; 2011. p. 39-67.

Blümel, Miriam, and Reinhard Busse, 2019 (accessed Sept. 3rd). *The German Health Care System.* The Commonwealth Fund. https://international.commonwealthfund.org/countries/germany/

Bradley, Elizabeth, and Lauren Taylor. The American health care paradox: Why spending more is getting us less. Public Affairs, 2013.

Browning EK, Johnson WR. The Trade-Off between Equality and Efficiency The Journal of Political Economy, 1984;92(2):175-203.

Chandra A, Skinner JS. Technology Growth and Expenditure Growth in Health Care. JEL. 2012;50(3):645-80.

Chernew, Michael E., Allison B. Rosen, and A. Mark Fendrick. "Value-Based Insurance Design: By abandoning the archaic principle that all services must cost the same for all patients, we can move to a high-value health system." Health Affairs 26, no. Suppl2 (2007): w195-w203.

Chetty, R. A New Method of Estimating Risk Aversion. American Economic Review 96(5): 1821-1834, Dec. 2006

Chetty, R. Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply, Econometrica 2012;80(3): 969-1018.

Clements B, Coady D, Gupta S. The Economics of Public Health Care Reform in Advanced and Emerging Economies. Washington: I. M. Fund; 2012.

Congressional Budget Office. The Long-Term Outlook for Health Care Spending. The Congress of the United States; November 2007.

Congressional Budget Office (CBO) May 2019 Budget Baseline. https://www.cbo.gov/system/files?file=2019-05/51302-2019-05-medicare_0.pdf

Currie J, Gahvari F. Transfers in Cash and In-Kind: Theory Meets the Data. JEL. 2008;46(2):333- 83.

Cutler D. Your Money or Your Life: Strong Medicine for America's Healthcare System. New York: Oxford University Press; 2004.

Cutler DM, Reber SJ. Paying for health insurance: the trade-off between competition and adverse selection. QJE. 1998;113(2):433-66.

Davis K and R Reynolds, The Impact of Medicare and Medicaid on Access to Medical Care Chapter Author: Karen Davis, Roger Reynolds , in R. Rossett (ed.) The Role of Health Insurance in the Health Services Sector. Cambridge MA: NBER, 1976.

Doyle, J., Graves, J.A., Gruber, J., Kleiner, S., Do High-Cost Hospitals Deliver Better Care? Evidence from Ambulance Referral Patterns, Journal of Political Economy, 2015.

Economic Report of the President, 2019.

Einav, Liran, and Amy Finkelstein. 2011. Selection in Insurance Markets: Theory and Empirics in Pictures. Journal of Economic Perspectives, 25(1): 115-38.

Emanuel EJ, Fuchs V. Health Care Vouchers – A Proposal for Universal Coverage. N Engl J Med. 2005;352(12):1255-60.

Ehrlich, Issac, and Yong Yin, 2013. "Equilibrium Health Spending and Population Aging in a Model of Endogenous Growth: Will the GDP Share of Health Spending Keep Rising?," Journal of Human Capital, 7(4): 411-447

Fendrick AM, Chernew ME. Value Based Insurance Design: Maintaining a Focus on Health in an Era of Cost Containment. Am J Manag Care. 2009;15(6):338-43.

Fonseca R, Michaud PC, Galama T, Kapteyn A, On the Rise of Health Spending and Longevity, UQAM Working Paper, Montreal, August 2013.

Frakt, Austin and Elsa Pearson, 2019. "A Question Rarely Asked: What Would Medicare for All Cover? New York Times (July 29).

Gahvari, Firouz, On the Marginal Cost of Public Funds and the Optimal Provision of Public Goods. Journal of Public Economics 2006;90(6-7):1251-1262.

Getzen TE. Population Aging and the Growth of Health Expenditures. J Gerontol. 1992;17:S98-104.

Gibson TB, Wang S, Kelly E, Brown C, Turner C, Frech-Tamas F, et al. A Value-Based Insurance Design Program at a Large Company Boosted Medication Adherence for Employees with Chronic Illnesses. Health Aff (Millwood). 2011;30(1):109-17.

Gottlieb, Joshua D., Adam Hale Shapiro, and Abe Dunn, 2018. "The complexity of billing and paying for physician care." Health Affairs 37, no. 4: 619-626.

Glover, Lucinda, 2019 (accessed Sept. 3rd). *The Australian Health Care System.* The Commonwealth Fund. https://international.commonwealthfund.org/countries/australia/

Gruber, J. and E. Saez, The Elasticity of Taxable Income: Evidence and Implications, Journal of Public Economics 2002; 84(1):1-12.

Gruber, J. and K. Simon, Crowd-out 10 years later: Have Recent Public Insurance Expansions Crowded Out Private Health Insurance? Journal of Health Economics 2008; 27:201–217.

Hall RE, Jones CI. The Value of Life and the Rise in Health Spending. QJE. 2007;122(1):39-72.

Howard DH, Bach PB, Berndt, ER, Conti, RM, Pricing in the Market for Anticancer Drugs, J Economic Perspectives, 29(1), Winter 2015, 139-62.

Erlich, Isaac, and Yong Yin, "Equilibrium Health Spending and Population Aging in a Model of Endogenous Growth: Will the GDP Share of Health Spending Keep Rising?" Journal of Human Capital 2013 7:4, 411-447

Jacobs, Bas. The marginal cost of Public Funds is One. Working paper, Tilburg University. http://papers.ssrn.com/sol3/papers.cfm?abstract id=1711121.

Kaplow L. The Theory of Taxation and Public Economics. Princeton: Princeton University Press; 2008.

Katz, Josh, Kevin Quealy, and Margot Sanger-Katz, 2019. "Would 'Medicare for All' Save Billions or Cost Billions?" New York Times (April 10).

Keehan SP, Sisko AM, Truffer CJ, Poisal JA, Cuckler GA, Madison AJ, et al. National Health Spending Projections Through 2020: Economic Recovery and Reform Drive Faster Spending Growth. Health Aff (Millwood). 2011;30(8):1594-605.

Kliff, Sarah, 2019. "Bernie Sanders's Medicare-for-All plan, Explained," Vox. https://www.vox.com/2019/4/10/18304448/bernie-sanders-medicare-for-all

Kliff, Sarah, and Dylan Scott, 2019. "We Read 9 Democratic Plans for Expanding Health Care. Here's How they Work." VOX, June 21, 2019.

Lane, Joan, 2012. A Social History of Medicine: Health, Healing and Disease in England, 1750-1950. Routledge

MedPAC. Health Care Spending and the Medicare Program; June 2012.

MedPAC. Benefit Design and Cost-Sharing in Medicare Advantage Plans. Report to Congress; December 2002.

Murphy KM, Topel RH. The Value of Health and Longevity. JPE. 2006;114(5):871-904.

Nordhaus W. The Health of Nations: The Contributions of Improved Health to Living Standards. NBER Working Papers 8828; 2002.

OECD, Looking to 2060: Long-term global growth prospects, http://www.oecd.org/eco/outlook/2060%20polic 2012.

Reinhardt U. A Fork in the Road for Health Care. The New York Times. 2012 May 25.

Rothschild M, Stiglitz J. Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. QJE. 1976;90(4):629-49.

Ryan P. The Path to Prosperity: A Blueprint for American Renewal. Washington: Fiscal Year 2013 Budget Resolution; 2012.

Papanicolas, Irene, Liana R. Woskie, Duncan Orlander, E. John Orav, and Ashish K. Jha, 2019. "The Relationship Between Health Spending And Social Spending In High-Income Countries: How Does The US Compare?" Health Affairs: 10-1377.

Pashchenko, Svetlana, and Ponpoje Porapakkarm, 2019. "Reducing medical spending of the publicly insured: the case for a cash-out option." American Economic Journal: Economic Policy 11, no. 3: 390-426.

Piketty, T, Saez, E. ”Inequality in the long run.” Science 344, no. 6186 (2014): 838-843.

Schroeck, Florian Rudolf, Bruce L. Jacobs, Sam B. Bhayani, Paul L. Nguyen, David Penson, and Jim Hu, 2017. "Cost of new technologies in prostate cancer treatment: systematic review of costs and cost effectiveness of robotic-assisted laparoscopic prostatectomy, intensity-modulated radiotherapy, and proton beam therapy." European Urology 72, no. 5: 712-735.

Skinner JS. Causes and Consequences of Geographic Variation in Health Care. In: Pauly MV, McGuire TG, Barros PP, editors. Handbook of Health Economics. New York: North Holland; 2012. p. 45-94.

Skinner, J, and Staiger, D., 2015. “Technology Diffusion and Productivity Growth in Health Care,” Review of Economics and Statistics.

Thorlby, Ruth, and Sandeepa Arora, 2019 (accessed September 3rd). The English Health Care System. The Commonwealth Fund. https://international.commonwealthfund.org/countries/england/

Van de Ven, WPMM, Schut FT. Universal Mandatory Health Insurance In The Netherlands: A Model For The United States? Health Affairs 2008;27(3):771-781.

Vrangbaek, Karsten, 2019 (accessed). The Danish Health Care System. The Commonwealth Fund https://international.commonwealthfund.org/countries/denmark/

Figures

Figure 1: Graphical Model of Private and Public Medical Spending

Panel A: Private Medical Spending Choices

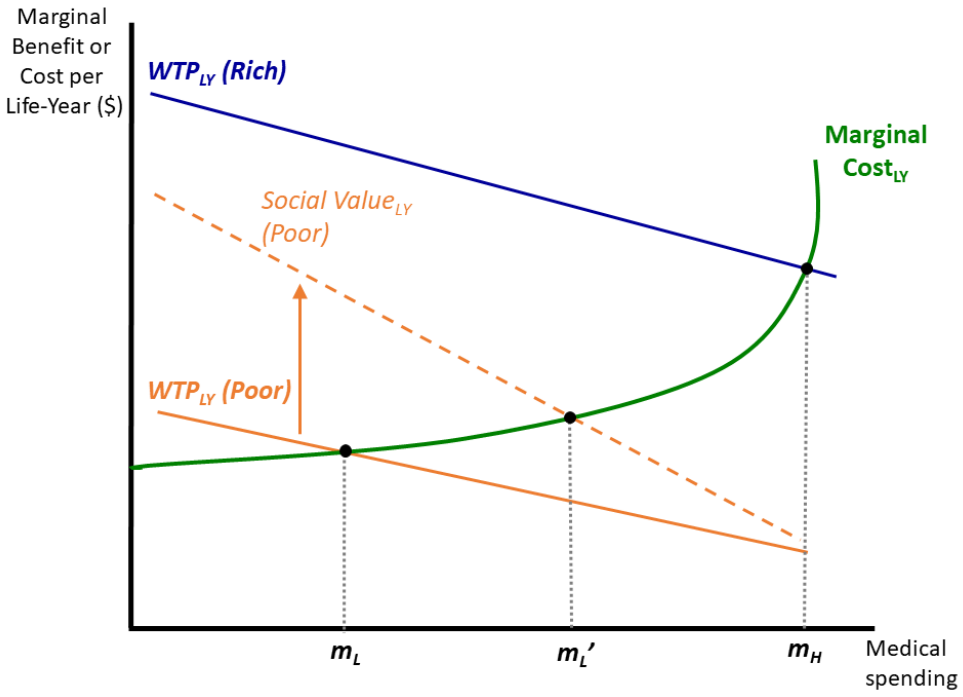

Panel B: Uniform Benefits

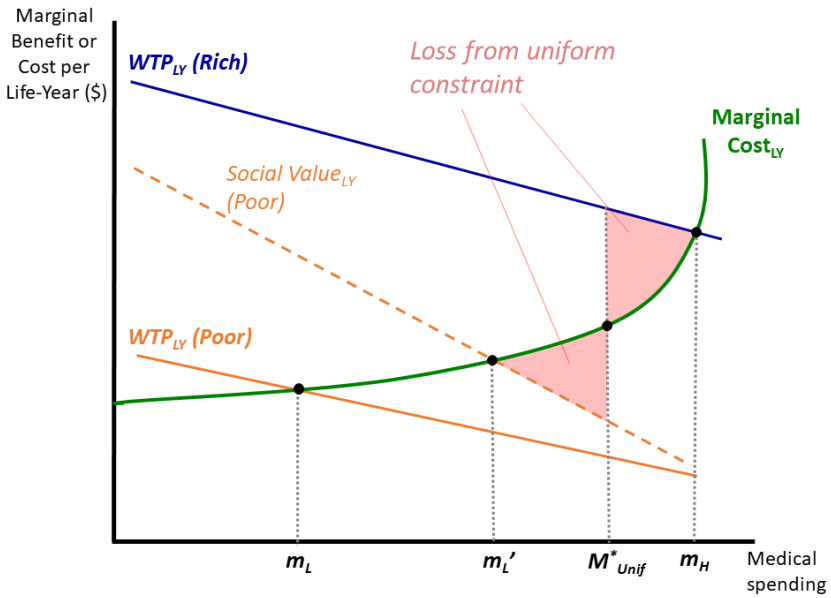

Figure 2: Rising Costs of Uniform Health Insurance Benefits

Panel A: Baseline (Medicare in 1965)

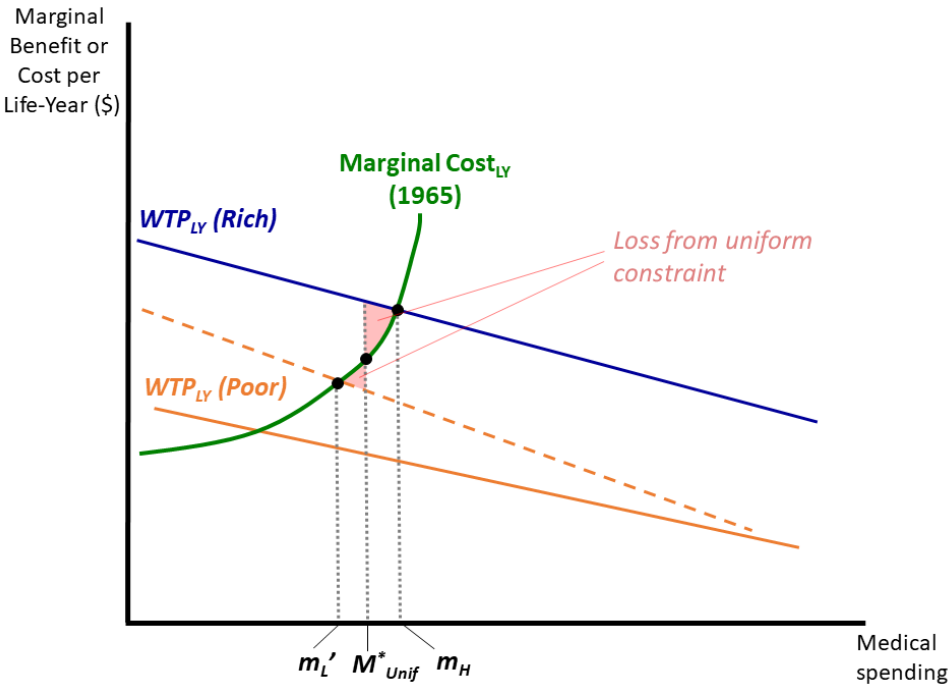

Panel B: Improving Medical Technologies

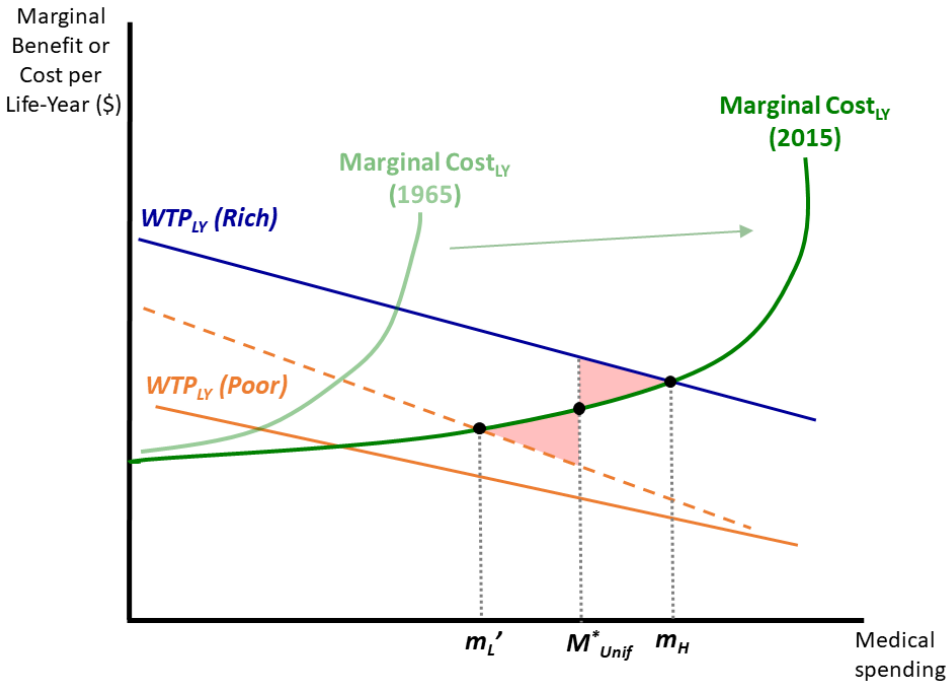

Panel C: Rising Income Inequality

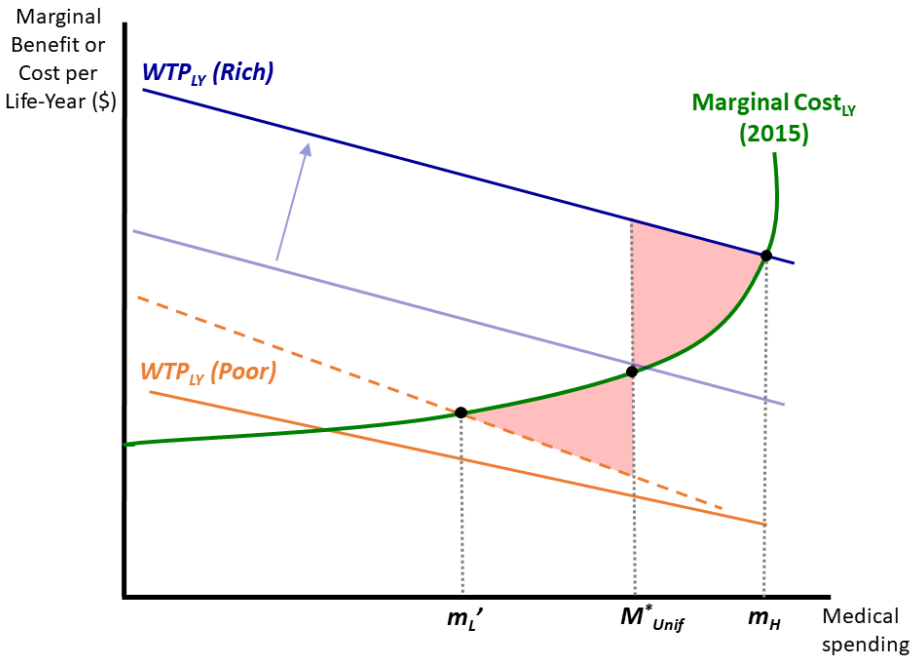

Panel D: Rising Deadweight Loss of Taxation

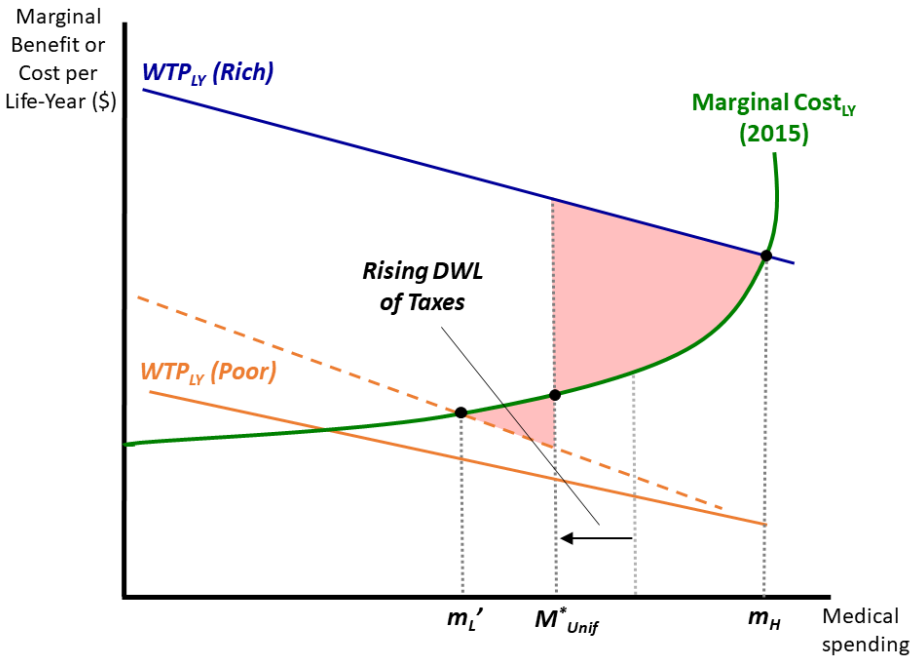

Figure 3: Top-Up vs. Uniform Benefits

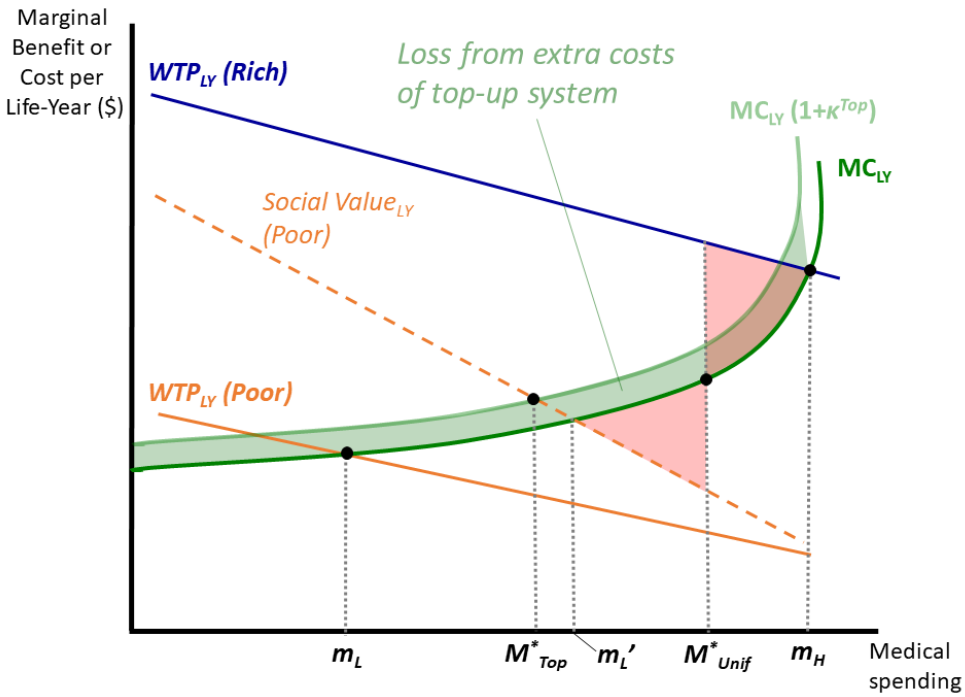

Figure 4: Simulations: Widening Inequality in Private Medical Spending

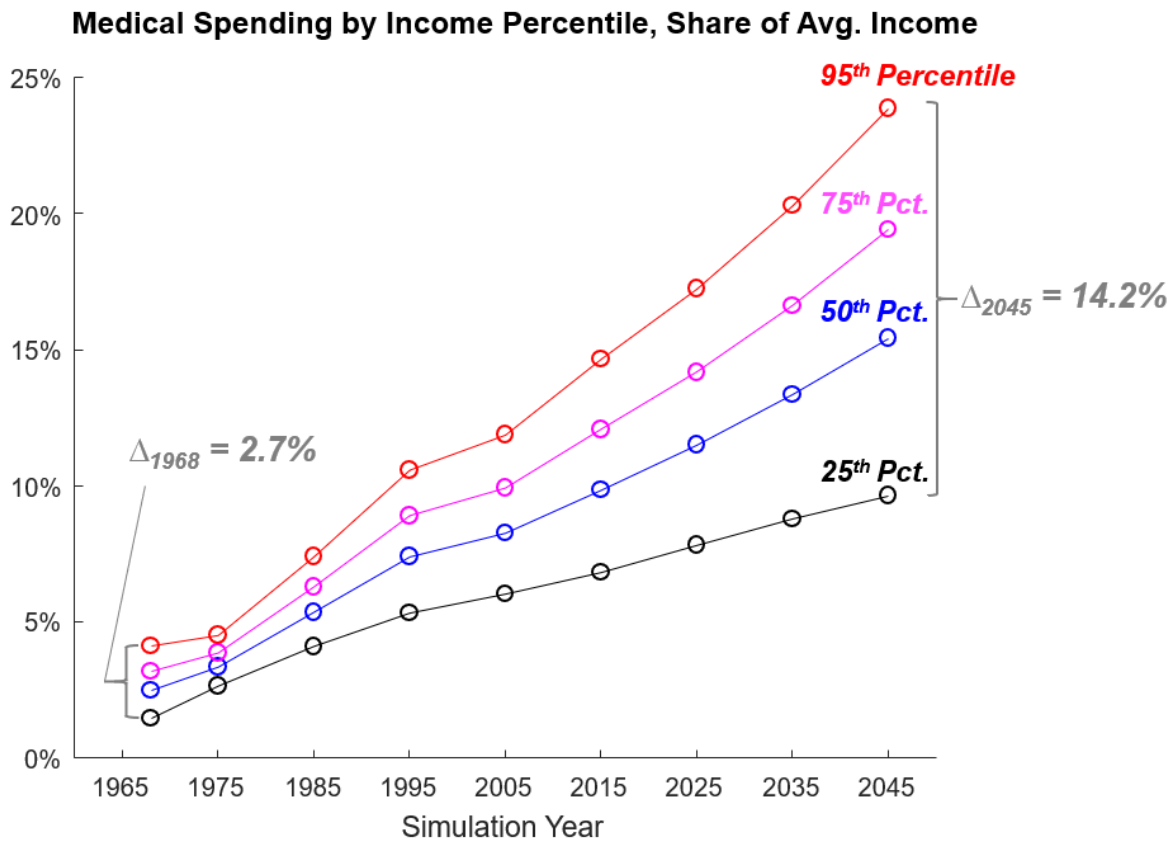

Figure 5: Simulations: Public Medical Benefit Level

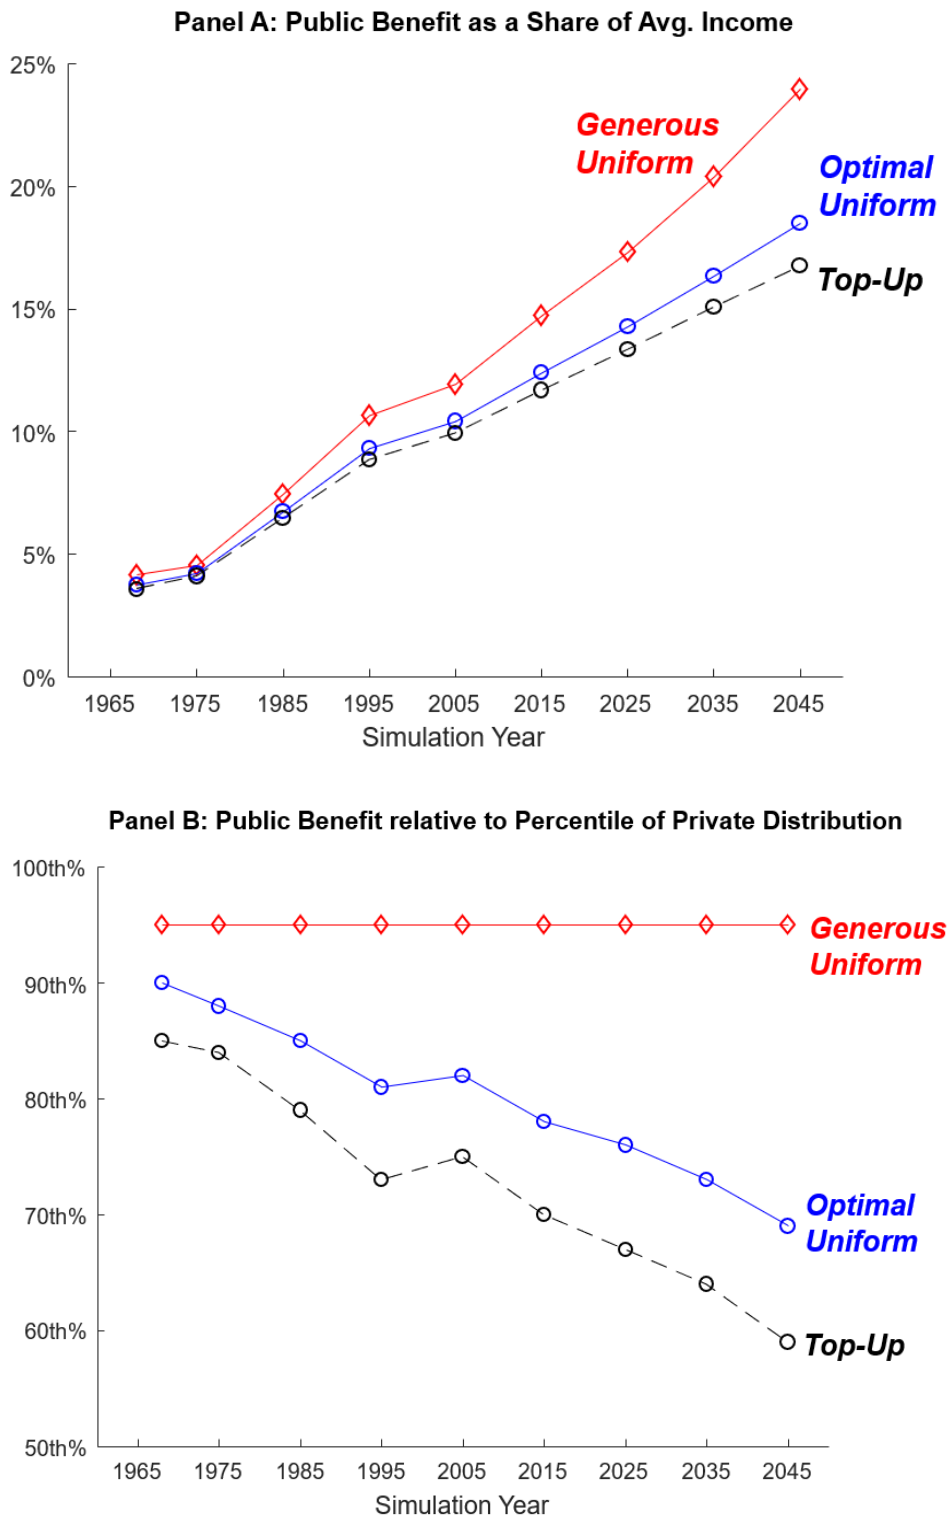

Figure 6: Simulations: Top Marginal Tax Rates

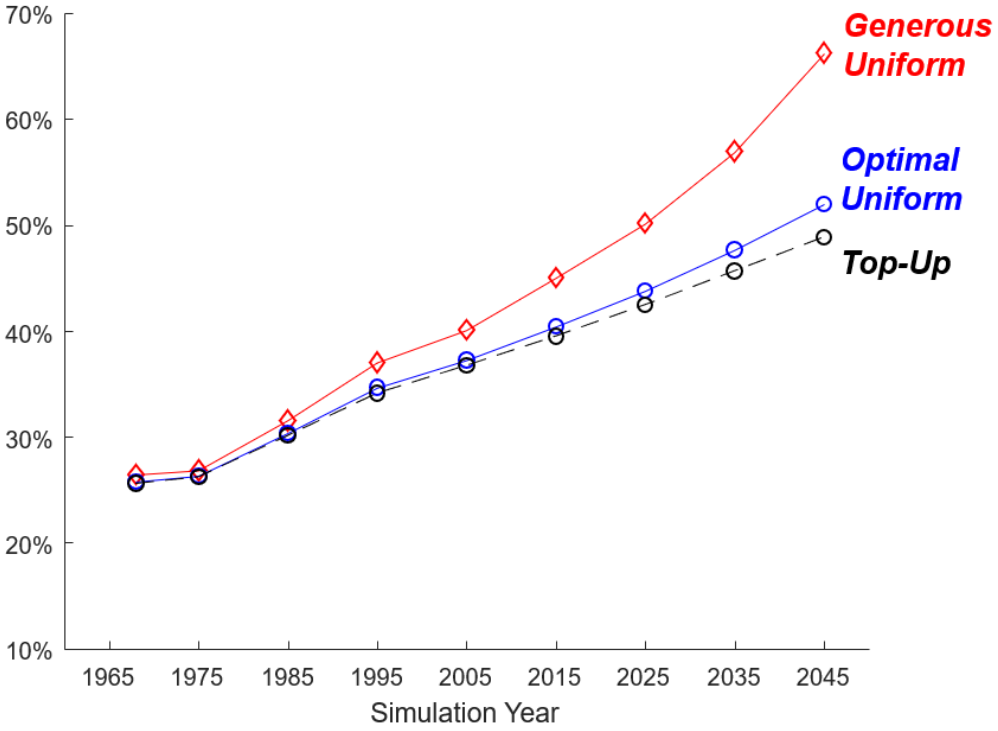

Appendix A: Uniform vs. Top-Up Benefits in a Formal Model

To be added.

Appendix B: Simulation Model Details

To be added.