

When Social Assistance Meets Market Power: A Mixed Duopoly View of Health Insurance in the United States

Ashantha Ranasinghe* Xuejuan Su†

December 2021

Abstract

We develop a mixed duopoly model with quality-differentiated products. The public firm chooses its product quality and offers it for free to eligible individuals as a form of social assistance to maximize consumer welfare, while the private firm chooses both the product quality and price to maximize profit. We first characterize the pattern of market segmentation for given product offerings, highlighting the non-monotonic relationship between market participation and individual income. We then calibrate the model to health insurance for the U.S. working-age population, with Medicaid acting as the public firm. We examine the distributional implications of policy changes, both actual and hypothetical, that lead to various degrees of public program expansion. Despite potentially significant inefficiency of the public firm, the overall effect of its expansion is welfare improving. Central to these findings is the significant market power enjoyed by the private firm that results in high profit margins if left unchecked. As more individuals become eligible for the public program, the resulting increase in competitive pressure disciplines the private firm's ability to exercise market power.

Keywords: Mixed duopoly; quality differentiation; public provision of private goods; funding of public services; distribution.

JEL codes: D21; D43; H11; H42; H44; I00; L38

*Department of Economics, University of Alberta. E-mail: ranasinghe@ualberta.ca.

†Department of Economics, University of Alberta. E-mail: xuejuan1@ualberta.ca.

1 Introduction

In the United States, many services are provided by both private and public providers. Prominent examples include health insurance, health care, education, housing, and transportation, etc. In these sectors, direct public intervention is typically motivated to ensure access to low income individuals, and more generally, by concern for distributional equality (Poterba, 1996). Accordingly, public provision can take the form of social assistance programs with eligibility requirements (e.g., Medicaid and public housing), or universally available public options (e.g., public hospitals, public schools, and public transit systems).

In these mixed markets, the co-existence of public and private providers naturally gives rise to strategic interactions between them. While private firms are driven by a conventional profit motive, the public firm acts as a government agency tasked with protecting consumer welfare. Although the public firm may not actively compete with private firms to attract customers, its presence in the market *per se* changes the competitive landscape, to which the private firms need to respond. Accordingly, the public firm can influence not only the well-being of its customers directly, but also the well-being of the private firms' customers indirectly. While such an indirect effect would be minimal if private firms are highly competitive, it becomes important when they have significant market power. This is the focus of our paper, namely the role the public firm plays in disciplining the private firms' ability to exercise market power.

We develop a mixed duopoly model where a public firm and a private firm engage in Stackelberg competition and offer quality-differentiated products, with the public firm being the first mover and the private firm the follower. We highlight the social assistance aspect of the public program by assuming the public firm chooses an income-based eligibility requirement and makes its product freely available to all eligible individuals to maximize social welfare (aggregate consumer welfare for all individuals).¹ Thus, with a

¹This assumption is innocuous and includes universal public coverage as a special case, i.e., when the eligibility requirement is non-binding for the entire population. We explicitly examine the distributional impact of such a policy experiment in Section 4.

given budget, the public firm faces a tradeoff between the quality of its product and the size of its customer base, where a more stringent eligibility requirement restricts access but allows for higher quality, and a less stringent eligibility requirement expands access but with lower quality. The private firm chooses both the quality and price of its products to attract customers. Moreover, it acts as the monopoly firm over the “residual demand” in the market, the size of which shrinks as the eligibility requirement for the public firm’s product becomes less stringent. Overall, the Stackelberg equilibrium features optimal product offerings by both firms, and consumer segmentation across these products.

We derive analytical characterizations of the pattern of market segmentation for given product offerings. We highlight the potential non-monotonic relationship between market participation and individual income, which follows directly from the social assistance nature of the public program. More specifically, for a given eligibility requirement, all eligible individuals obtain the public firm’s product for free and hence never opt out, yet some individuals just above the income threshold may find it optimal to opt out. This income-based threshold can result in a downward jump in utility for similar individuals just below and above the threshold. Thus, a change in the eligibility requirement is felt mostly by individuals whose incomes are close to the threshold. We then turn to a quantitative approach to analyze optimal product offerings, and the impact of various policies on market segmentation and social welfare.

Our mixed duopoly model is general and not sector-specific. In this paper, we calibrate the model to health insurance for the U.S. working-age population, motivated by two factors. First, among the working-age population, Medicaid plays an important role in providing social assistance via health insurance coverage to low income individuals, thus corresponding to the public firm in our model.² Second, there are significant and growing market power concerns in the private health insurance sector. For example, the American Medical Association (AMA) found that across 384 U.S. Metropolitan Statistical Areas (MSAs), 74 percent of MSA-level markets were “highly concentrated” in 2019,

²In contrast, *Medicare* is best viewed as universal public coverage instead of social assistance as over 90 percent of the retirement-age population is enrolled in Medicare.

and 56 percent of markets experienced an increase in market concentration between 2014 and 2019.³ Furthermore, the economics literature has also shown increased market concentration is the cause of significant price increases for private health insurance (Dafny 2010; Dafny *et al.* 2012; Starc 2014).⁴ These two features—an important public firm that provides social assistance and private firms that have significant market power—make the health insurance industry an ideal target for our mixed duopoly model.

As a benchmark, we evaluate the quantitative implications of our model by calibrating it to the U.S., prior to the major Medicaid expansion under the 2010 Affordable Care Act (ACA). Key parameters in the model are pinned down by matching the share of individuals that choose private, public or opt-out of health insurance, health expenditure as a percentage of GDP, and relative quality (measured by reimbursement fees) between the private and the public firm’s products. Our quantitative setting features a wide range of heterogeneity across individual income and initial health, allowing for a clear assessment of welfare across the distribution of individuals for a given policy change. Our modeling choice that the public firm provides social assistance is of consequence, as in our benchmark the individuals that opt-out have income levels that are too high to be eligible for the public program, but yet too low to warrant purchase from the private firm despite its affordability. Thus, market participation does not monotonically increase with individual income.

We use our calibrated model to evaluate the welfare implications of several policy changes that expand the size of the public program to varying degrees. Throughout, we hold the product quality offered by the public firm constant at the benchmark level, and consider changes to the size of eligible individuals in the population. This allows us to focus on the policy impact triggered by changes in the eligibility requirement, without potentially confounding effects from simultaneous changes in public quality. Moreover,

³“Competition in Health Insurance: A Comprehensive Study of U.S. Markets, 2020 Update.” <https://www.ama-assn.org/system/files/2020-10/competition-health-insurance-us-markets.pdf>, last accessed April 10, 2021.

⁴Similar patterns of increased market concentration and its effect on price increases are also found for hospitals and physician practices (Krishnan 2001; Capps *et al.* 2018; Dafny *et al.* 2019).

we assume public program expansion is financed by a proportional tax on income over a fixed tax exemption amount. The main results from our quantitative experiments are as follows.

First, we consider Medicaid expansion under the ACA, where the public program maintains its social assistance nature but increases its coverage from the bottom 12.5% of the income distribution to the bottom 20%. We find that as more individuals become eligible for the public product, the private firm responds to the increased competitive pressure by raising the average quality of its products by 8% and lowering the average markup by 4%; i.e., a notable reduction in its ability to exercise market power. Thus, Medicaid expansion indeed generates both a direct and an indirect benefit: the former accrues to individuals who are newly eligible for the public program, and the latter accrues to individuals who buy the more appealing product from the private firm. However, as middle and high income individuals bear disproportionately high tax burdens (based on a 1% proportional tax on income), the indirect benefit from the more appealing private product is not sufficient to offset the tax burden. So on net, they are worse off. Overall, Medicaid expansion (up to 50% coverage of the population) is welfare improving on aggregate, but there is considerable heterogeneity across individuals. Additionally, despite the lessening stringency of the eligibility requirement, a non-trivial share of individuals still optimally opt out. Thus, as long as Medicaid remains a social assistance program, its expansion alone does not ensure universal participation in the market.

Next, to the extent that universal participation is a desirable policy outcome, we consider a hypothetical public program that expands to cover the entire population. In this policy experiment, the public firm provides a “one size fits all” product to the population, and the eligibility requirement becomes non-binding. Within this setup, we consider alternative treatments regarding the private firm.

Motivated by universal public coverage programs observed in many developed economies (including Medicare for the U.S. retirement-age population), we first consider the case when the private firm is removed from the market *ex ante*, and hence the public firm acts

as the state monopoly. Although the quality of the public program is lower than what the private firm would offer in the benchmark case, we find that consumer welfare improves on aggregate. Moreover, the benefit accrues not only to low and middle income individuals (as in the previous case when Medicaid remains a social assistance program), but also to upper-middle income individuals. Only high income individuals (those above the 75th percentile of the income distribution) are worse off under universal public coverage, because the product quality is too low and more importantly, the additional tax burden too high.

Finally, taking the potential strong push-back against a “one size fits all” system into account, we consider alternative cases when the private firm is allowed to operate in the market. Depending on whether individuals forfeit their benefit from the public program when they buy from the private firm, the public firm offers either a take-it-or-leave-it option or a portable voucher option to all individuals. Interestingly, in the take-it-or-leave-it option, the private firm cannot offer a product to profitably compete for any customer. Thus, even though the private firm is allowed to operate in the market *ex ante*, it optimally shuts down *ex post*, resulting in the same market outcome as that under universal public coverage. Alternatively, in the voucher case, all individuals are entitled to the same valued voucher, which can be used either to obtain the public product for free or to offset part of the price for the private product. Thus, the portable voucher partially “subsidizes” the private firm and lessens the competitive pressure it faces when the public product is universally available. We find that the private firm indeed increases its market share under the voucher system, attracting customers with higher quality and a lower markup of its product relative to the benchmark. Moreover, the private firm extracts most of the rent associated with the voucher program. For individuals who switch to the private firm, despite enjoying a higher quality product, their utility improves marginally compared to the case with universal public coverage.

Related Literature. This paper is related to several strands of the literature. First, our modeling approach brings together elements from both industrial organization and

public economics. On the one hand, while quality competition between private firms has received considerable attention in industrial organization (Gabszewicz and Thisse 1979; Shaked and Sutton 1982; Motta 1993; Wauthy 1996; Lehmann-Grube 1997; Wang and Yang 2001), the lack of a public firm makes these models unsuitable for examining policies that affect the size of the public firm.⁵ On the other hand, motivated by the liberalization of industries formerly comprised of state-owned monopolies, mixed markets have been extensively studied in public economics (Cremer *et al.* 1991; Barros 1995; Anderson *et al.* 1997; Matsumura 1998; Ishida and Matsushima 2009; Lasram and Laussel 2019). However, this literature largely abstracts from quality competition between the public and private firms. To our best knowledge, the only exception is Klumpp and Su (2019). In that paper, both the public and private firm choose the quality and price of their products, so policies that directly affect the public firm's product lead to strategic responses by the private firm to adjust its product, which generates distributional impacts across the entire population. Their main interest is to characterize the distributional effects analytically. In comparison, our model highlights the social assistance aspect of the public program, namely the public firm chooses an income-based eligibility requirement instead of a price. Unlike Klumpp and Su (2019), our model can generate a non-monotonic relationship between market participation and individual income. This non-monotonicity allows us to capture a rich set of tradeoffs under a policy change but also poses significant analytical challenges. As a result, we rely more heavily on the quantitative method to characterize our main results.

Our work also relates to the macro-health literature that studies the life-cycle implications of health shocks across individual labor supply, savings and mortality. Capatina (2015) and Hosseini, Kopecky and Zhao (2020) evaluate the effects of health shocks on income inequality over the life-cycle, while in a related setting Jeske and Kitao (2009) examine the role of tax policy on demand for health insurance. Chen, Feng and Gu

⁵Gaynor *et al.* (2015) provides a detailed survey of the industrial organization literature on the health care market, focusing on the impact of competition on price, quality and treatment decisions for health care providers and health insurers.

(2020) consider a life-cycle setting where health shocks affect earnings which feeds back to health amplifying inequality. These models focus on the demand for healthcare and abstract to a firm that earns zero profit or charges a fixed markup over cost. While we neglect the rich dynamics highlighted in these models, we emphasise the strategic price-quality decisions of a private and public healthcare provider which plays a crucial role in our setting. Specifically, we show that firms optimally adjust price-quality offerings and markups in response to the policies we explore, which affects demand for healthcare and selection across private and public providers.

The remainder of the paper is organized as follows. In Section 2 we present our formal model and Section 3 outlines the model calibration to health insurance in the United States. Section 4 conducts a number of policy exercises and reports their distributional and welfare implications. Section 5 concludes with a synthesis and discussion of our findings. Proofs that are too long to be included in the main text and supplemental analysis are in the Appendix.

2 The model

Consider an economy with heterogeneous individuals characterized by income m and initial health ε , jointly distributed with a cumulative distribution function (CDF) $F(m, \varepsilon)$. An individual's utility depends on three factors: his health and two consumption goods. More specifically, an individual can improve initial health ε by utilizing a health product of quality θ , with a health production function given by $h(\theta)$ where $h(0) = 0$, $h' > 0$, and $h'' < 0$.⁶ Of the two consumption goods, one exhibits complementarity with health (denoted by x) while the other does not (denoted by n). The individual's utility function can be written in following quasi-linear form:

$$U(n, x, \theta) = n + u(x, h(\theta) + \varepsilon),$$

⁶In this model we do not distinguish health insurance from health care, where the quality θ can be interpreted as the expenditure of all medical treatment that is covered in a health insurance product.

with $u_1 > 0$, $u_2 > 0$, $u_{11} < 0$, $u_{22} < 0$, and $u_{12} > 0$. The price of the numeraire good n is normalized to 1. Without loss of generality, the price of the consumption good x can also be normalized to 1 if its measurement unit is re-scaled to $\frac{1}{p_x}$ for an arbitrary price p_x . Thus, the individual's budget is given by $n + x + p \leq m$, where p is the price of the health product relative to the consumption goods. Note that the complementarity condition ($u_{12} > 0$) implies that higher-income individuals are willing to pay more for quality θ , holding initial health ε constant. At the same time, holding income constant, individuals with good initial health are less willing to pay for quality θ , and may even opt out entirely (not buy θ).

There are two providers of the health product, a public firm and a private firm. The public firm imposes an income-based eligibility requirement \bar{m} , and offers its product for free to all eligible individuals (i.e., when $m \leq \bar{m}$). We denote the public firm's product as $(\theta_1, p_1 = 0)$. The private firm can offer up to two products, denoted as (θ_2, p_2) and (θ_3, p_3) . Together with the outside option, denoted as $(\theta_0 = 0, p_0 = 0)$, there are thus up to four product choices in the market, i.e., (θ_k, p_k) for $k = \{0, 1, 2, 3\}$.

To provide a certain quality θ , we allow for differences in efficiency across the private and public firm. More specifically, the cost of providing quality is $c_r(\theta) = \theta$ for the private firm, and $c_b(\theta) = \xi\theta$ for the public firm, where $\xi \geq 1$.⁷ The public firm is equally efficient when $\xi = 1$, but less efficient when $\xi > 1$.

2.1 The market segmentation

Facing the health products (θ_k, p_k) for $k = \{0, 1, 2, 3\}$, an individual chooses one and only one of the alternatives. Below we characterize the optimal product choice for consumers, and the resulting market segmentation across products. In particular, we separately consider the choices involving only non-public varieties $k = \{0, 2, 3\}$ first, and then the choice involving the public variety $k = 1$ and one of the non-public varieties.

⁷Note that in our model, the cost function $c(\theta)$ and the health production function $h(\theta)$ cannot be separately identified. Any cost function satisfying $c'(\theta) > 0$ and $c''(\theta) > 0$ is already captured in our setup when the health production function is adjusted to $h(c^{-1}(\theta))$ instead of $h(\theta)$.

Lemma 1. *Compare two non-public varieties (including the outside option) (θ, p) and (θ', p') . If $\theta \leq \theta'$ and $p \geq p'$, the variety (θ, p) is (at least weakly) dominated by the variety (θ', p') and attracts no customer. Thus, for both varieties to attract some customers, it has to be the case that $(\theta - \theta')(p - p') > 0$.*

The proof is straightforward and follows directly from the conditions $u_1 > 0$ and $u_2 > 0$. With the outside option $(\theta_0 = 0, p_0 = 0)$, for both private varieties to attract some customers, higher quality has to be matched with higher price. Without loss of generality, we denote $0 < \theta_2 < \theta_3$ and $0 < p_2 < p_3$, namely, (θ_2, p_2) represents the basic variety and (θ_3, p_3) represents the premium variety.

Proposition 2. *Compare two non-public varieties (including the outside option) (θ_l, p_l) and (θ_h, p_h) with $\theta_l < \theta_h$ and $p_l < p_h$. (i) If an individual with (m, ε) prefers (θ_l, p_l) over (θ_h, p_h) , then all individuals with (m', ε') satisfying $m' \leq m$ and $\varepsilon' \geq \varepsilon$ prefers (θ_l, p_l) over (θ_h, p_h) . (ii) If an individual with (m, ε) prefers (θ_h, p_h) over (θ_l, p_l) , then all individuals with (m', ε') satisfying $m' \geq m$ and $\varepsilon' \leq \varepsilon$ prefers (θ_h, p_h) over (θ_l, p_l) .*

The proof is in the Appendix. This proposition shows that among the non-public varieties, market segmentation features the conventional “single crossing” property in both the income dimension (holding ε fixed) and the initial health dimension (holding m fixed). Moreover, the two indifference contours—one separating the outside option from the basic variety, and the other separating the basic variety from the premium variety—are both upward-sloping in the (m, ε) space.

However, adding the public variety with an income-based eligibility requirement \bar{m} disrupts this market segmentation. Depending on an individual’s income m , their choice set may include the public variety or not, and accordingly the optimal product choice is

$$s^* = k \quad \text{if} \quad \begin{cases} m \leq \bar{m} : U_k^* = \max\{U_0^*, U_1^*, U_2^*, U_3^*\} \\ m > \bar{m} : U_k^* = \max\{U_0^*, U_2^*, U_3^*\} \end{cases} \quad (1)$$

Thus, among eligible individuals, the presence of the public variety can impact their optimal product choices.

Lemma 3. *For eligible individuals ($m \leq \bar{m}$), consider a non-public variety (including the outside option) (θ, p) and the public variety $(\theta' > 0, p' = 0)$. If $\theta \leq \theta'$ and $0 = p' \leq p$, the non-public variety (θ, p) is (at least weakly) dominated by the public variety (θ', p') and attracts no customer among the eligible individuals. For both varieties to attract customers among eligible individuals, it has to be the case that $0 < \theta' < \theta$ and $0 = p' < p$.*

The proof parallels that of Lemma 1. It is immediately clear that the outside option is strictly dominated by the public variety ($\theta_0 = 0 < \theta_1$ and $p_0 = 0 = p_1$), so no eligible individual will opt out. Similarly, a private variety could be dominated by the public variety and attract no customers among eligible individuals (while yet attracting ineligible ones).

Proposition 4. *For eligible individuals ($m \leq \bar{m}$), compare a private variety (θ_r, p_r) and the public variety $(\theta_b, p_b = 0)$ with $0 < \theta_b < \theta_r$ and $0 = p_b < p_r$. (i) If an individual with (m, ε) prefers (θ_b, p_b) over (θ_r, p_r) , then all individuals with (m', ε') satisfying $m' \leq m$ and $\varepsilon' \geq \varepsilon$ prefers (θ_b, p_b) over (θ_r, p_r) . (ii) If an individual with (m, ε) prefers (θ_r, p_r) over (θ_b, p_b) , then all individuals with (m', ε') satisfying $m' \geq m$ and $\varepsilon' \leq \varepsilon$ prefers (θ_r, p_r) over (θ_b, p_b) .*

The proof parallels that of Proposition 2. It shows that among eligible individuals, market segmentation again features “single crossing” among the non-dominated varieties. The presence of the public variety does not preclude the private firm from competing for eligible individuals *per se*, but it does make it much harder to attract them. On the other hand, among ineligible individuals, market segmentation is given in Proposition 2 and may feature a positive share of individuals choosing the outside option (or a dominated private variety) because the public variety is not available to them. This dichotomy is what we intend to capture, that market participation may not be monotonically increasing in income under a social assistance program.

2.2 Stackelberg competition

Taking the public firm's product offering $(\theta_1, p_1 = 0, \bar{m})$ as given, the private firm chooses its product offerings (θ_2, p_2) and (θ_3, p_3) to maximize its profit:

$$\pi_r = Q_2(p_2 - \theta_2) + Q_3(p_3 - \theta_3), \quad (2)$$

where Q_2 is the share of individuals optimally purchasing the basic variety ($s^* = 2$) according to (1), $p_2 - \theta_2$ is the per-unit profit margin of the basic variety, and similarly, Q_3 is the share of individuals optimally purchasing the premium variety ($s^* = 3$), and $p_3 - \theta_3$ is the per-unit profit margin of the premium variety.

The public firm is an agent of the government and is given a budget B to maximize consumer welfare. In particular, the public firm's budget constraint can be written as

$$\pi_b = -Q_1\xi\theta_1 \geq -B, \quad (3)$$

where Q_1 is the share of individuals optimally purchasing the public variety ($s^* = 1$) according to (1), $-\xi\theta_1$ is the per-unit loss of the public variety (as $p_1 = 0$), and $B \geq 0$ is the fixed budget the public firm receives. The budget constraint implies the public firm cannot incur more operational loss than its given budget.

Taking the private firm's best response into account, the public firm chooses its product offering $(\theta_1, p_1 = 0, \bar{m})$ to maximize consumer welfare for all individuals:

$$W = \left(\int \int V^{1-\psi} dF(m, \varepsilon) \right)^{\frac{1}{1-\psi}}, \quad (4)$$

a constant-elasticity-of-substitution (CES) social welfare function (SWF) where the substitution elasticity $1/\psi$ reflects the public firm's (or government's) redistributive concern. Note that $\psi = 0$ corresponds to a utilitarian welfare function, which morphs into a Cobb-Douglas welfare function as $\psi \rightarrow 1$, and a Rawlsian welfare function as $\psi \rightarrow \infty$.

The larger is the value of ψ , the more the public firm is concerned about distributional equality.

Note that when choosing its product offering $(\theta_1, p_1 = 0, \bar{m})$, the public firm takes into account the optimal response of the private firm for both the basic variety (θ_2, p_2) and the premium variety (θ_3, p_3) . In other words, we assume a Stackelberg model of competition, with the public firm being the leader and the private firm being the follower.⁸ Moreover, by assuming that the public firm’s objective is perfectly aligned with social welfare, we can abstract from potential principal-agent problems when the public firm acts as a government agency.

2.3 Remarks

In focusing on the strategic interactions between the public and private firm, we have made some simplifying modeling choices that abstract from other important considerations. Some of these choices are more innocuous than others. Here we briefly discuss their implications.

First, in our mixed duopoly setting, we use a single private firm as a stand-in for a small number of private firms that enjoy significant market power. Qualitatively, all our main results are not sensitive to the actual number of firms per se, as long as a few of them dominate the market. This is indeed the case for the health insurance market in the U.S., the target of our quantitative analysis in the next Section. For example, the *Horizontal Merger Guidelines* issued by the U.S. Department of Justice and Federal Trade Commission (2010) uses the Herfindahl-Hirschman Index (HHI) to measure market concentration: a market is considered “highly concentrated” if its HHI is above 2500.⁹ According to this definition, the American Medical Association (AMA) found that across

⁸Clemens and Gottlieb (2017) have shown that Medicare acts as the first mover, and private insurance companies respond as the followers.

⁹Hypothetically, a market consisting of four equal-sized firms with market shares of 25% each would have HHI equal to $25^2 + 25^2 + 25^2 + 25^2 = 2500$.

384 Metropolitan Statistical Areas (MSAs) in the U.S., 74 percent of the MSA-level markets were “highly concentrated” in 2019, with the average HHI of 3473.¹⁰

Second, to focus on market power and strategic consideration as the main source of market failure, we intentionally abstract from any asymmetric information issues. In particular, our model does not account for adverse selection in the conventional insurance market framework. The health product in our model is such that all customers, regardless of initial health type ε , pay the same price p and receive the same quality θ . This excludes any *ex ante* redistribution across health types, i.e., no adverse selection resulting from good health types cross-subsidizing bad health types. If individuals choose to opt out in our model, it is not because their θ from a given health product is smaller, but rather the same θ is valued less.¹¹ Thus, even though our model is capable of generating market segmentation similar to that under adverse selection, the mechanism is not through asymmetric information. Similarly, even though our model does not explicitly allow for moral hazard behaviors such as over-use or ineffective use of medical services, it implicitly accounts for wasteful spending in healthcare through both the public firm’s inefficiency and the private firm’s markup rates in its products.

Third, on the demand side, we assume that all individuals have exogenous income and initial health. This implies that an individual’s labor supply is perfectly inelastic and independent of health, an assumption that is unlikely to hold in the data. However, we believe that endogenizing the labor supply decision would only affect our results quantitatively but not qualitatively. More specifically, if we model individual labor supply as endogenously chosen and positively related to initial health, we would arrive at an income-health distribution where income and initial health are positively correlated, instead of independently distributed as we later assumed. At the same time, some individuals may intentionally reduce their labor supply to meet the eligibility requirement

¹⁰“Competition in Health Insurance: A Comprehensive Study of U.S. Markets, 2020 Update.” <https://www.ama-assn.org/system/files/2020-10/competition-health-insurance-us-markets.pdf>, last accessed April 10, 2021.

¹¹The private firm also does not distinguish good versus bad customers as they all generate the same profit margin when purchasing the same product.

for the public product, thus create “bunching” at the income threshold. Despite these differences, all subsequent steps—consumer segmentation and the Stackelberg equilibrium between the firms—would follow the same procedure to solve our model. We opt for the simpler approach because, without reliable data measuring initial health, it is all but impossible to pin down the correlation between income and initial health convincingly.

Finally, following Atkinson (1970), we use a CES social welfare function to model the public firm’s (or government’s) concern for distributional equality. While our approach is broadly consistent with the optimal taxation literature (Amiel *et al.* 1999; Kuziemko *et al.* 2015; Lockwood and Weinzierl 2016; Saez and Stantcheva 2016), our calibrated parameter ψ and the substitution elasticity $1/\psi$ should not be directly compared to the values typically used to examine the optimal tax structures. With our assumption of exogenous income, our model cannot capture the excess burden of taxation, a key component in the optimal taxation literature. To ensure that our policy exercises remain in the reasonable realm, we model any public budget change as being *internally funded* through a flat-rate income tax (i.e., the direct tax burden), and restrict the changes in the public budget to small or moderate levels (i.e., limited excess burden of taxation).

3 Calibration

We calibrate the mixed duopoly model to replicate key features of health insurance for the U.S. working-age population (18-64 years of age) prior to the major Medicaid expansion under the Affordable Care Act.¹² We think of the public firm in our model as providing Medicaid, and the private firm as a stand-in for private insurance companies. For private insurance products, we abstract from copay and deductible by subsuming them in the effective price.¹³

¹²The vast majority of people 65 and above are enrolled in Medicare, and those below 18 are covered by their parents plans and hence do not make meaningful choices.

¹³We allow the private firm to offer two varieties to match the private firm’s market share. If restricted to only one variety, the private firm would optimally serve only the low-health-high-income segment and generate too small a market share compared to the data.

We need to take a stance on the utility function, the distribution of income and health across the population, and assign corresponding parameter values. We assume utility is $U = n + \delta(x^\alpha(h(\theta) + \epsilon)^{1-\alpha})$ where health is $h(\theta) + \epsilon$; $h(\theta)$ transforms healthcare quality into a health outcome and is additive with innate health. We assume $h(\theta) = \theta^\gamma$ where $\gamma > 0$ measures the importance of healthcare quality on health outcomes. Consumption includes two components: one that is complementary with health, x , entering through a Cobb-Douglas aggregator, and a second that is independent of health, n . The numeraire good n establishes a lower bound on the marginal utility of income and hence disciplines the private firm's markups (or profit margin).¹⁴ Parameter δ measures the importance of consumption good x and health relative to the numeraire good, and also allows us to normalize the price of x to 1 (the same as the price of the numeraire good).

The distributions for income m and initial health ϵ are exogenous and independent (though in equilibrium there is a positive correlation between income and overall health). As utility is a product of consumption and health, which are in terms of income and health outcomes respectively, the range of values assigned to these distributions matter. To limit additional structure, we assume income is uniformly distributed between 0 and 1, and innate health ϵ is uniformly distributed between 0 and ω ; $m \sim U(0, 1]$ and $\epsilon \sim U(0, \omega]$.¹⁵

To find the Stackleberg equilibrium, we solve the nested problem with backward induction. First, for any set of product offerings, we can find the equilibrium sorting of individuals across private, public, and the outside option. Second, taking the public variety as given, and anticipating the market segmentation from individual choices, the private firm chooses its optimal varieties $(p_i^*, \theta_i^*, i = 2, 3)$ to maximize profit. Finally, taking the private firm's response and the resulting market segmentation into account, for a given budget B , the public firm optimally chooses its eligibility requirement \bar{m} and accordingly, the quality θ_1^* to maximize social welfare.

¹⁴Without the numeraire good, the private firm would charge much higher markups, and the aggregate health expenditure as a share of individual income would far exceed the data counterpart.

¹⁵ We have no strong prior for the distribution of innate health across the population and so we assume a uniform distribution. For purposes of symmetry we take the same stance on income, though we recognize the income distribution in the U.S. is right skewed. In the Appendix we show that our main results hold if income is Pareto distributed.

Table 1: MODEL FIT AND PARAMETER VALUES

Target Moments	Model	Data	Parameter
Shares rel. to GDP (%):			
Healthcare spending	11.9	12.4	$\gamma = 0.60$
Population choosing (%):			
Public, Q_1	12.5	12.6	$\psi = 1.55$
Private, $(Q_2 + Q_3)$	71.0	70.5	$\omega = 0.82$
Opt-out, Q_0	16.5	16.9	$\delta = 1.47$
Private premium to Public quality:			
θ_3^*/θ_1^*	2.0	2.0	$\xi = 1.70$
Equilibrium product offerings:	Market shares:		
$(\theta_1^*, \bar{m}) = (0.026, 0.114)$	$Q_1 = 12.5$		
$(\theta_2^*, p_2^*) = (0.0129, 0.0394)$	$Q_2 = 10.4$		
$(\theta_3^*, p_3^*) = (0.0519, 0.0912)$	$Q_3 = 60.6$		

Notes: Moments are from the National Center for Health Statistics which are based on the National Health Expenditure tables. See text for details.

In total there are 7 parameters to calibrate: one for social welfare ψ , three for the utility (δ , α , γ), the upper bound on innate health ω , public firm inefficiency ξ , and the public firm budget B . Two of these parameters are set a priori. For the Cobb-Douglas portion of the utility function, we set $\alpha = 1/2$ to reflect equal importance between consumption good x and individual health. And since income is exogenous, we choose B to target spending on Medicaid relative to GDP which is 1.1 percent.¹⁶ (This target and all targets that follow are from the National Center for Health Statistics.) The remaining parameters are jointly determined by solving the equilibrium of the model so that a set of relevant model generated outcomes match corresponding data moments of the U.S. healthcare system.

¹⁶While our framework does not explicitly model a supply side, we use aggregate income to proxy overall output/GDP. We arrive at a Medicaid to GDP ratio of 1.1 percent as follows: National Health Expenditures (NHE) accounts for 12.4 percent of GDP (we remove Medicare expenditures since individuals between 18-64 years of age are not eligible for Medicare). Medicaid accounts for about 21 percent of NHE, but this includes individuals below the age of 18 and those above 65. As there is no detailed breakdown by age groups, we refer to the Population Reference Bureau analysis of the American Community Survey which shows that over 55 percent of the population receiving Medicaid in 2015 are those below 18 and above 65. Assuming a similar proportion for 2004 and applying this correction, we arrive at a Medicaid to GDP target of 1.1 percent.

While no one parameter uniquely identifies any data moment they have some direct impact on a specific moment(s), and so we describe these in turn to motivate their relevance. Parameter γ converts healthcare quality into a health outcome and affects individual spending on healthcare, with higher values for γ implying a given quality of healthcare has a bigger impact on health outcomes. We target total spending on healthcare relative to GDP which is 12.4 percent. Parameters ω , ψ and δ impact the proportion of individuals choosing private, public or no health coverage which are 70.5, 12.6 and 16.9 percent, respectively (as well as spending on healthcare). Specifically, ω affects the extreme value of the innate health distribution with higher values reflecting a ‘healthier’ population, and affects overall demand for healthcare (private or public) and decisions to opt-out. Parameter ψ measures the government’s preference for redistribution through its impact on \bar{m} (the income threshold for Medicaid eligibility). In particular, a higher value of ψ implies a deeper concern for redistribution, which given a budget B affects the public firm’s choice of \bar{m} and θ_1^* , thus influencing the proportion of individuals that choose public, private or opt-out of healthcare. Parameter δ measures the importance of health and health related consumption x which primarily affects demand for private healthcare and whether to opt out.¹⁷

Lastly, for government inefficiency ξ , we target a relative quality ratio of 2 between the private firm’s premium variety and the public variety. While this is admittedly ad-hoc, we think this is a reasonable target because, (1) Medicaid fees have been fairly stable hovering around 70 percent of Medicare fees (e.g., 69% in 2003, 72% in 2008, Zucherman *et al.* 2004, Zuckerman *et al.* 2009), and (2) Medicare reimbursement ranges between 60-80 percent for surgical and non-surgical services as that in the private market (Clemens and Gottlieb 2017). This implies that Medicaid reimbursement is about half of that offered by the private firm. We use this measure of inefficiency as a sensitivity check when evaluating policies that expand public coverage.

¹⁷We note that targeting the proportion that choose private and public coverage pins down the proportion opting out (or any two of the three). Nonetheless, we target all three proportions since in our setting these parameters are interconnected and jointly informative of individual decisions.

Again, we stress that these parameters jointly affect the model generated equilibrium moments, and do not have an independent impact on any one target. This is in part because the private firm optimally adjusts its price-quality offerings in response to the public firm’s competitive pressure and resulting market shares. For income and health we use a grid with 80 and 10 discrete points respectively.

Table 1 reports the jointly calibrated parameter values, the model generated moments and corresponding data moments. Overall, our simplified mixed duopoly setting can reproduce the broad aggregate patterns related to market shares across private and public health coverage, the share that opts out, and the share of spending on healthcare. To generate that 12.6 percent of the population receives Medicaid requires that $\psi = 1.55$. The income threshold for Medicaid eligibility is $\bar{m} = 0.114$ implying that up to 12.5 percent of the population (ascending in income) are eligible for Medicaid, which is consistent with the income threshold for Medicaid eligibility based on the federal poverty line.¹⁸ $\gamma = 0.6$ implies diminishing returns to higher health quality, and $\delta = 1.47$ implies that health and related consumption is more important than non-health related consumption. Government inefficiency $\xi = 1.70$ implies that the cost to provide a unit of quality is 70 percent higher for the public firm than the private firm. Also reported are equilibrium private firm price and quality for both varieties, and market shares. For context, the basic and premium prices account for about 8 and 18 percent of the median income.

Figure 1 shows the distribution of individual types across health coverage and their respective shares. Everyone that is eligible for Medicaid optimally chooses the public option. For individuals above the eligibility requirement, market segmentation follows the “single crossing” property established in Proposition 1, with two upward sloping contours separating the three choices. More specifically, individuals choose the outside option if they have low-income and high-innate-health; the basic variety if their income is higher

¹⁸Medicaid eligibility varies across U.S. states but a general approximation is that people with income between 100 to 138 percent of the federal poverty level are eligible for Medicaid based solely on income. This implies an income eligibility range for Medicaid between \$12.5-17.5 thousand, or roughly 10-15 percent of the population based on the distribution of income in the U.S. around this time period. The model prediction is the mid-point of this range.

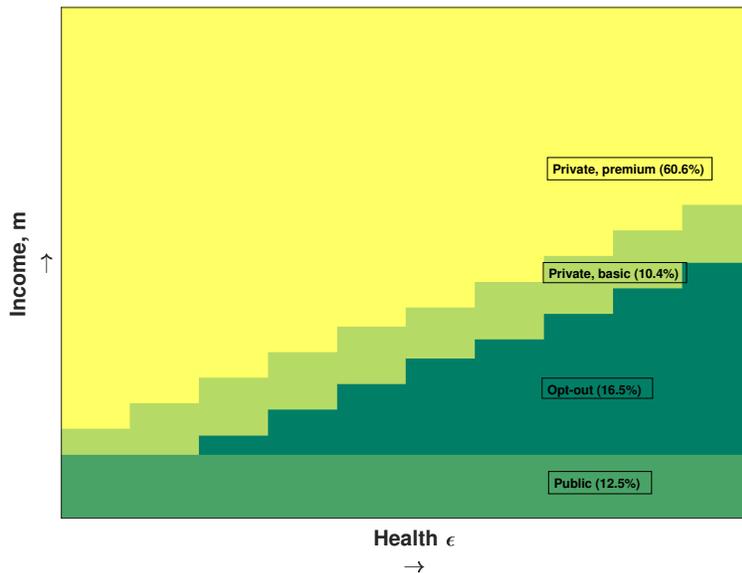


Figure 1: Healthcare coverage across individual types

and/or innate health is lower; and the premium variety if their income is even higher and/or innate health is even lower.

Other model predictions. Before considering various policy implications, we discuss some of the features implied by our model. While there are no estimates we can rely on to validate our values for γ , δ , ω , ψ and ξ , the overall predictions of the model are reasonable and so we use it as a benchmark to evaluate the quantitative impacts of the policy experiments that follow.

For the premium quality, the private firm charges a markup—price over cost—of over 175 percent. While this is on the high end for markups, it should not be interpreted literally for two reasons. First, our model does not account for fixed costs of operation, which can represent a significant portion of total costs in the data.¹⁹ Second, the underlying cost function can be scaled-up by an arbitrary factor without having any material impact on the results, in which case the markup rate will be scaled-down by the

¹⁹A recent report in Time Magazine documents that private insurer administrative and overhead costs account for over 30 percent of total healthcare expenditures in the U.S., and is about five times the cost per person than in Canada.

same factor.²⁰ Hence, it is the *change* in the markup that is informative and what we focus on. Interestingly, the private firm’s basic variety features low quality (50 percent lower than what the public firm offers) and high markup (about 300 percent), making it strictly dominated by the public variety. The basic variety is offered to target the low health individuals just above the Medicaid eligibility threshold, who can afford the premium variety but find it sub-optimal (about 10 percent of the population). These strategic markups are a central feature of the mixed-duopoly setting where the private firm squeezes rents from low income non-Medicaid eligible individuals, features that are absent in more perfectly competitive settings. Finally, the model predicts that spending on non-health related consumption (good n) accounts for 44 percent of all consumption spending, and that individuals choosing private healthcare account overwhelmingly for the majority of income.

A further implication of our mixed duopoly setting is the private firm offers a higher quality of the premium variety than what individuals would ideally prefer. To make this point clear, we consider a perfectly competitive setting with infinite private firms that earn zero profit, and each individual chooses optimal healthcare quality.²¹ Under perfect competition, an overwhelming majority of individuals choose an optimal quality of healthcare above the basic variety but below the premium variety offered in the benchmark economy. That is, in a mixed-duopoly setting the private firm sets the quality of the basic variety sufficiently low to induce selection in to the premium variety, a higher quality than what most individuals prefer. Said differently, the private firm over supplies healthcare quality than what is socially optimal in first-best sense, and at a higher cost to individuals. This particular feature has key implications for understanding some of the policy implications in what follows.

²⁰Our model cannot separately identify the cost function of providing quality $c(\theta)$ and the health production function $h(\theta)$, but rather only $h(c^{-1}(\theta))$ (see footnote 7). We normalize the private firm’s cost function to be $c_r(\theta) = \theta$. Our calibration remains unchanged if the underlying cost function is scaled up by an arbitrary factor λ , in which case the resulting markup rate would be scaled down by $1/\lambda$.

²¹A setting with free entry, where each individual i optimal chooses θ_i at a cost $p_i = \theta_i$ given the zero profit assumption.

4 Policy exercises

We now use our calibrated benchmark economy to evaluate the impact of various policies relevant to the U.S. healthcare debate. Specifically, we examine: (1) the implications of Medicaid expansion under the Affordable Care Act, (2) universal public coverage, and (3) universal public coverage with private firm competition. For each of these policies we keep public quality at the benchmark level to isolate and quantify the impact triggered by changes in the eligibility requirement. We then specifically focus on public and private firm market shares, consumer welfare, healthcare quality and markups. For convenience, we report private firm quality as a weighted average of its two varieties, and discuss each variety when relevant.

4.1 Medicaid expansion under ACA

We evaluate the implications of Medicaid expansion where the public firm plays a more prominent role in healthcare provision. Specifically, we consider a policy that raises the public firm's budget B so that 20 percent of the population is covered by Medicaid (up from 12.5 percent), in line with one of the ACA's broader objectives.²² We assume the increase in B is funded through a uniform, proportional tax on income exceeding the Medicaid income eligibility threshold.²³ We recognize this can affect labor supply and earnings, which in our setting is exogenous. Nevertheless, we use such a tax primarily as a modeling convenience to highlight the weight of the tax on welfare, and note that a lump-sum tax that kicks in above an income threshold is an alternative to fund the expansion (and have a lesser impact on the intensive margin of labor supply).

Table 2, column (2) reports the results when 20 percent of individuals are Medicaid eligible and public quality is held at the benchmark level (results are relative to the

²²While the specifics of the ACA are wide ranging and complex, and often varying across states, a main component involved the expansion of Medicaid eligibility (to lower the number of uninsured individuals) and the public healthcare budget.

²³Specifically, $\tau \times (m - \bar{m})$. We note that our central results generalize to other progressive tax schemes.

Table 2: MEDICAID EXPANSION

	(1) Benchmark	(2) Eligibility 20% & fixed θ_b
Healthcare quality, θ		
Public	1.00	1.00
Private	1.00	1.08
Population choosing (%)		
Public, Q_b	12.5	20.0
Private, Q_r	71.0	68.9
Opt-out, Q_0	16.5	11.1
Social Welfare		
Entire population	1.00	1.01
Medicaid eligible (20%)	1.00	1.04
Equilibrium product offerings: Market shares:		
$(\theta_1^*, \bar{m}) = (0.026, 0.19)$	$Q_1 = 20.0$	
$(\theta_2^*, p_2^*) = (0.0134, 0.0403)$	$Q_2 = 8.6$	
$(\theta_3^*, p_3^*) = (0.0553, 0.0946)$	$Q_3 = 60.3$	

Notes: Statistics are relative to the benchmark economy except for population shares across healthcare coverage and equilibrium product offerings (bottom panel). Private firm statistics (top panel) are a weighted average of their varieties (by market shares).

benchmark economy except for population shares). We note that this scenario is an out of equilibrium outcome where we impose that 20 percent of the lowest income individuals qualify for Medicaid and public quality is unaltered.²⁴ As intended the policy lowers the share of individuals that opt-out, from 16.5 to 11.1 percent, and lowers the private firm's market share by 2 percent (majority coming from the basic variety). Medicaid now accounts for 1.75 percent of GDP and is funded with a 0.9 percent tax on income. An implication of Medicaid expansion is that private firm average quality rises by 8 percent and markups fall by 4 percent (primarily from an increase in the premium variety, and decrease in the proportion choosing the basic variety). We elaborate on private quality and markups below.

²⁴An alternative is to solve the equilibrium of the model where B is sufficiently raised so that the public firm optimally chooses \bar{m} where 20 percent of individuals are eligible for Medicaid. The results are qualitatively similar to what is reported in column (2).

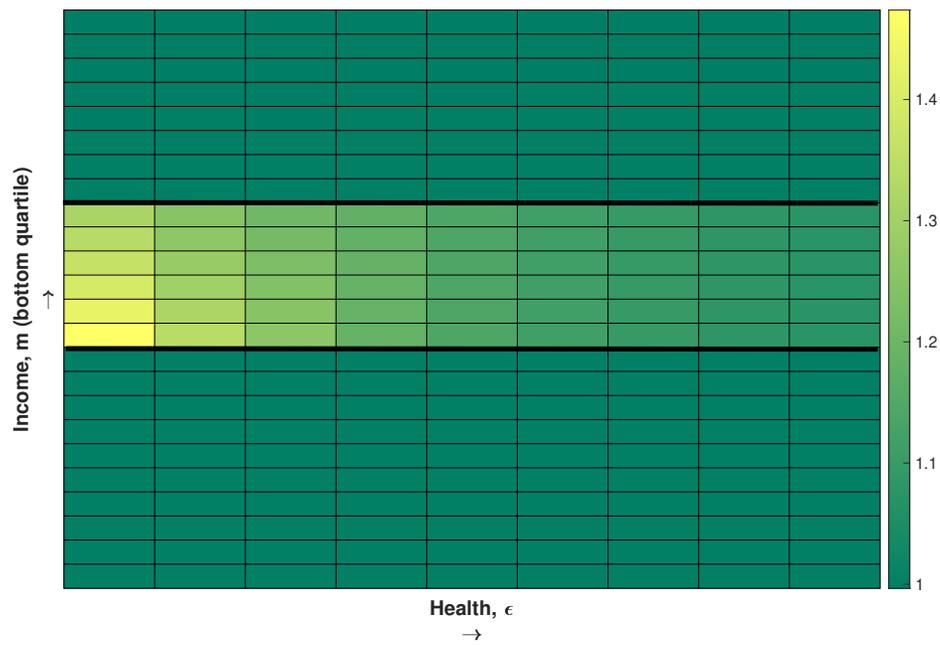


Figure 2: Utility change across individuals

Notes: The shades represent the magnitude of individual utility change: lighter shades correspond to larger increases in utility, and darker shades to smaller gains/decreases in utility. The figure shows the bottom quartile of the income distribution. The horizontal black lines show eligibility pre and post Medicaid expansion.

The impact on welfare is positive, rising by 1.2 percent (including the cost of the proportional tax), and rising by 4.1 percent among Medicaid eligible individuals. However, this masks considerable heterogeneity across individuals. Figure 2 provides a disaggregated view of utility changes for individuals in the bottom quartile of income across health, relative to the benchmark (values above/below 1 represent an increase/decrease in utility). For individuals previously eligible for Medicaid (below the lower solid black line), and noting that public quality is held fixed, there is no impact on utility. Utility increases for newly eligible Medicaid recipients (those between the two solid black lines), especially low health individuals whose utility rises between 30-50 percent. At the same time, the private firm's customers also see a notable indirect benefit. In response to the Medicaid expansion, the private firm increases the average quality of its products by 8 percent and reduces the average markup by 4 percent. More specifically, this indirect benefit is driven by two factors: the private firm raises the product quality and lowers the markup on both the basic and premium variety; and in addition, these product adjustments induce a larger *share* of its customers to choose the premium over the basic variety. Nevertheless, this indirect impact is offset by the proportional taxes on income required to fund Medicaid expansion, such that on net, utility falls in the range of 0.1-0.5 percent among non-Medicaid eligible households. (Households above the bottom income quartile experience similar utility changes as this latter group.)

We now use this setting to highlight the importance of Medicaid expansion on the competitive landscape and private firm responses. Figure 3 panel (a) plots private firm markups and product quality for different Medicaid eligibility thresholds. For each threshold, the public budget is set so that public quality is as in the benchmark, and funded through taxes as before. As Medicaid eligibility rises, a distinct pattern is that private firm quality rises and markups fall, as much as 20 and 10 percent respectively.²⁵

²⁵While quality is above the benchmark across all cases, there is a dip in quality when Medicaid eligibility is between 20 to 30 percent of the population. This is due to higher taxes, which especially affect low/middle income individuals that are just above the respective Medicaid eligibility threshold, and so the private firm optimally adjusts its quality. A similar reasoning applies for markups. Figure A.1 shows a breakdown by each private variety.

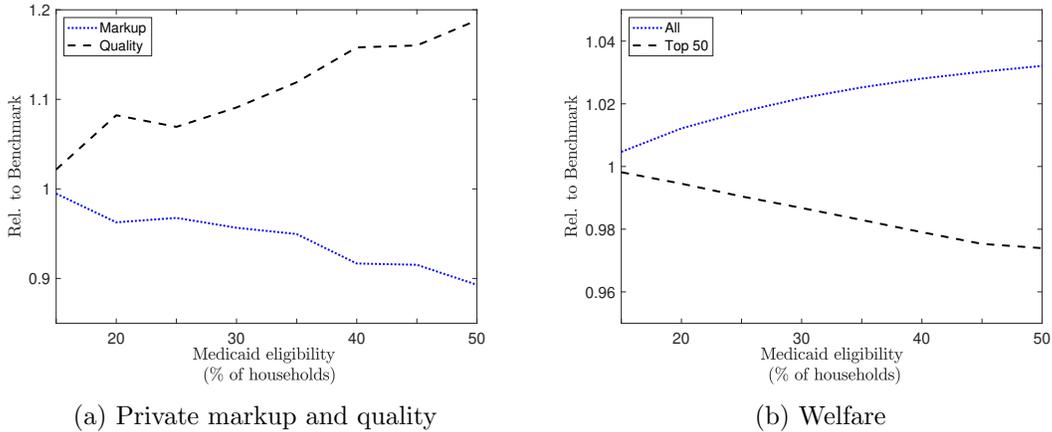


Figure 3: Markups, Welfare and Medicaid Eligibility Thresholds

Notes: Panel (a) shows weighted private firm quality and markups relative to the benchmark economy, and panel (b) shows welfare for all individuals and those above the median income, for various Medicaid eligibility thresholds (e.g., 30 on the axis implies the bottom 30 percent of individuals received Medicaid). Taxes required to fund the various Medicaid eligibility thresholds range from 0.3 to 4.3 percent. Welfare for the Top 50 is shown because we entertain Medicaid eligibility up to the median income.

Panel (b) shows the impact on welfare, which rises for the population as Medicaid expansion increases. While panel (a) shows private firm customers benefit from higher quality and lower markups, panel (b) shows the top 50 percent of individuals (and in fact all private customers) are worse off on net due to higher taxes. So in this case, the benefit of public expansion primarily goes to Medicaid recipients. For private firm customers, the benefit from increased competition is small relative to the additional tax burden, so the net effect is negative. This is consistent with the social assistance nature that Medicaid maintains under this policy.

We conclude this section by briefly addressing the sensitivity of our results to two key parameters in our model. First, when expanding Medicaid coverage we have assumed public firm inefficiency ξ remains fixed. As we show in Figure A.2 panel (a), expanding Medicaid coverage to 20 percent of individuals is welfare improving as long as public firm inefficiency does not exceed $2.5 \times \xi$; that is, the cost of providing quality does not rise by over 250 percent. Second, we recognize that the impact on welfare is sensitive to the calibrated parameter $\psi = 1.55$, which measures social welfare. Figure A.2 panel (b)

shows that for a plausible range $1 < \psi < 2$, Medicaid expansion is also welfare improving. Finally, we have assumed the social welfare function consists of only consumer welfare and excludes private firm profit (producer surplus), which decreases under Medicaid expansion. Our results, in particular that Medicaid expansion is welfare improving, remains robust even when producer surplus is included, as long as its weight is not too large.²⁶

4.2 Universal public coverage

In many developed countries such as the United Kingdom, France, and Canada, health insurance is exclusively provided by the public sector, with little to no role left for the private sector.²⁷ More recently in the United States, the state of Massachusetts adopted a similar program. To evaluate the impact of universal access to public health insurance *without* private competition, we make two modifications to our setting. First, we shut down the private firm so that the public firm is the sole provider of health insurance, i.e., “Medicare for all,” and where health insurance is offered at no charge, $p_1 = 0$. Second, we assume the public firm budget B is sufficient to keep public quality at the benchmark level. As in Section 4.1, the expansion in the public budget is funded with a proportional tax on income. We also hold government inefficiency ξ fixed for easy comparison with the benchmark.

Table 3 shows the results. Public spending on healthcare relative to GDP is about 9 percent—lower than healthcare spending shares observed across countries with universal public health coverage—and requires a 10 percent proportional tax on individual income exceeding the exemption amount. Again, we stress that we use a proportional tax to highlight the magnitude of a tax, and note that taxes in this range can impact labor supply and earnings.

²⁶In a setting where private firm profit and an aggregation of individual utility are equally weighted—where social welfare is $W + \pi_r$ based on equations 2 and 4—the increase in W exceeds the decrease in π_r . Of note, this is despite a potentially conservative stance of equally weighting private profit and individual utility.

²⁷While certain types of health coverage (e.g., dental, chiropractic etc.) are provided by private firms, we focus primarily on general practitioners and hospital care which are exclusively public.

Table 3: UNIVERSAL PUBLIC HEALTHCARE

	(1) Benchmark	(2) Universal Coverage
Healthcare quality, θ		
Public	1.00	1.00
Private	1.00	–
Government		
B/Y (%)	1.10	8.82
Tax rate (%)	0.00	9.95
Social Welfare		
Entire population	1.000	1.036
Medicaid eligible (12.5%)	1.000	1.000

Notes: Statistics are relative to the benchmark economy except for the Public budget.

Despite higher taxes, welfare rises by 3.6 percent. Again, this masks considerable heterogeneity across individuals as shown in Figure 4. Specifically, there is no impact on the bottom 12.5 percent of individuals based on income (previously Medicaid eligible) since they do not pay additional taxes and public quality remains fixed. Households above 75th percentile of the income distribution are worse off under universal healthcare, which becomes amplified as income rises (due to higher taxes) and as innate health falls (lower quality healthcare). All remaining individuals, those between 12.5-75th percentiles of the income distribution are better off under the policy, suggestive that universal healthcare is primarily favored by low, middle, and upper-middle income individuals. In particular, low-health individuals above the previous Medicaid threshold benefit the most with utility rising between 20-45 percent.

What accounts for higher welfare under universal public insurance? We rationalize this from three vantages. First, we interpret our results consistent with a private firm in the benchmark economy that primarily targets its premium variety to high income individuals, in essence providing excess quality for the majority of (middle income) individuals that buy the premium variety, and a basic variety targeted at individuals just above the Medicaid threshold (and with a high markup). This is driven by the

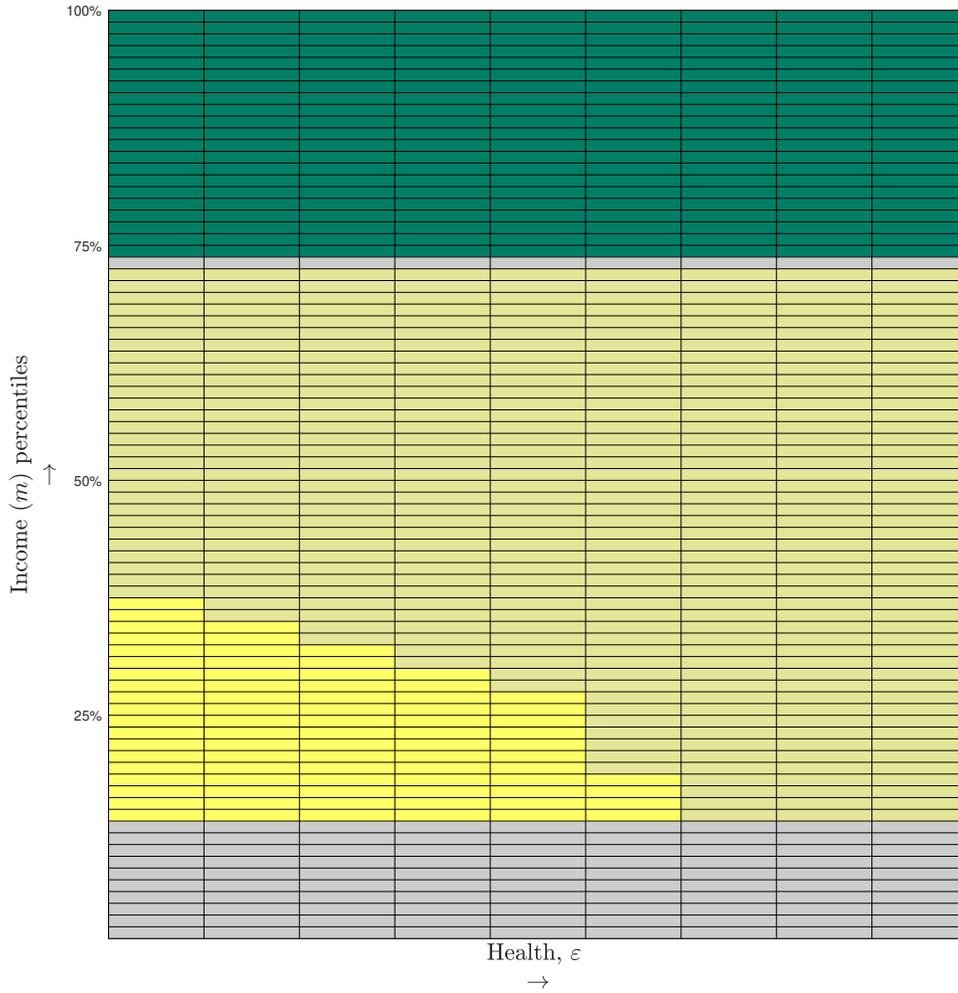


Figure 4: Universal health: Utility change across individuals

Notes: The gray shade shows individuals whose utility does not change (those in the bottom 12.5% of income and those just below the 75th percentile). The green shade shows individuals whose utility falls between 0-3 percent (those in the top quartile of income). The various yellow shades show individuals who benefit: light yellow represents individuals whose utility rise by less than 10 percent, and bright yellow represents those that see utility increases between 10 to 45 percent.

private firm’s monopoly power among non-Medicaid eligible individuals allowing for high markups and squeezing of rents. Second, we have held government inefficiency ξ constant, a strong assumption given there is a 7-fold expansion in the number of individuals the public firm now covers. Our simulations imply that universal health coverage is welfare improving if public firm inefficiency, or alternatively the cost of providing healthcare per person, does not rise by more than 85 percent relative to the benchmark setting. Finally, we have abstracted from labor supply and income effects from the substantial taxes required to fund universal healthcare. As long as income taxes do not uniformly lower individual earnings by more than 7 percent, universal healthcare can be welfare improving.

4.3 Universal public option

Finally, taking the potential strong push-back against a one-size-fits-all public program into account, we consider alternative cases when the private firm is allowed to operate in the market. Specifically, we consider cases where the public firm either acts as offering a take-it-or-leave-it option or offering portable vouchers, which depends on whether individuals forfeit the benefit from the public program when they choose to buy from the private firm.²⁸

Interestingly, in the take-it-or-leave-it option case, even though individuals are free to buy from the private firm if they deem the public firm’s product is inadequate, we find that the private firm cannot offer a product to profitably attract any individual away from the free public variety. Thus, although the private firm is allowed to operate in the market *ex ante*, it optimally shuts down *ex post* when the eligibility requirement is non-binding for the entire population. Accordingly, the market outcome remains the same as that under the universal public coverage (Section 4.2).

²⁸Outside the health context, historically, public schools and public housing projects are best viewed as take-it-or-leave-it options. More recently, vouchers have gained popularity as a way to introduce competition into the public system, and have been used in education reforms (education vouchers) and public housing reforms (housing vouchers) in the U.S.

Alternatively, in the voucher case, all individuals receive the same valued voucher from the public program. They can use the voucher either to obtain the public firm’s product for free (incurring zero out-of-pocket expenditure), or to offset part of the price if they choose to buy the private firm’s product. A notable example of such a system in practice is Germany. Compared to the take-it-or-leave-it public option, portable vouchers “subsidize” the private firm and lessen the competitive pressure it faces when the public program is universally available. As before, budget expansions are funded via a proportional tax on income exceeding the exemption amount.

Table 4: VOUCHER

	(1) Benchmark	(2) $B = B_b \times 8$
Healthcare quality, θ		
Public	1.00	1.00
Private	1.00	1.22
Population choosing (%)		
Public, Q_b	12.5	33.5
Private, Q_r	71.0	66.5
Opt-out, Q_0	16.5	0.0
Government		
B/Y (%)	1.10	8.80
Tax rate (%)	0.00	9.93
Social Welfare		
Entire population	1.000	1.036
Medicaid eligible (12.5%)	1.000	1.000
Equilibrium product offerings: Market shares:		
$(\theta_1^*, \bar{m}) = (0.026, 1)$	$Q_1 = 33.5$	
$(\theta_2^*, p_2^*) = (0.0385, 0.0421)$	$Q_2 = 4.5$	
$(\theta_3^*, p_3^*) = (0.0575, 0.0629)$	$Q_3 = 62$	

Notes: Statistics are relative to the benchmark economy except for market shares and the Public budget.

Table 4 shows the results where the public budget (at 8 times the benchmark budget B_b) is sufficient to keep public quality at the benchmark level, and where the voucher is equal to about 5 percent of the median income. Funding this expansion requires a 10 percent tax on income exceeding the exemption amount, and the public budget

accounts for 9 percent of GDP. Compared to the universal public coverage outcome (Table 3), we find that a portable voucher increases the private firm's market share from zero to 66.5 percent. However, to attract customers away from the public program, the private firm now offers both the basic and the premium varieties at quality levels notably higher than the public variety, and with much lower markups (relative to the benchmark). Interestingly, despite higher-quality varieties offered by the private firm, utility for individuals increase only marginally. The private firm is able to extract most of the rent associated with the portable voucher by setting the price just below the indifference point for its customers, but its ability for such rent extraction is limited by the voucher value. Overall, the voucher system generates similar impact on consumer welfare as the universal public coverage, but has drastically different implications for the private firm.

5 Conclusion

Tasked with providing social assistance, the presence of a public firm *per se* changes the competitive landscape of the market. This influence is particularly important in sectors where private firms enjoy significant power. In this paper, we develop a mixed-duopoly model with quality-differentiated products to explicitly examine the private firm's strategic responses to policies that expand the public program. With social assistance, and since market participation does not increase monotonically with individual income, the distributional impact of a policy change also exhibits rich non-monotone patterns. When the model is calibrated to health insurance for the U.S. working-age population, we show that Medicaid expansion increases the competitive pressure on the private firm, who responds by increasing its product quality and lowers its markup. Thus, Medicaid expansion is generally welfare improving, even though the public firm may be significantly less efficient than the private firm, and additional tax revenue is needed to fund the program expansion.

Moreover, we show that the specific form of the public program expansion matters. On the one hand, when Medicaid remains a form of social assistance (its expansion covering less than 50% of the population), the benefit mostly accrues to low-income individuals who become eligible for the public program. On the other hand, when the public program expands to offer universal coverage to the entire population, its benefits extend to low, middle, and even upper-middle individuals. In addition, as more individuals become eligible for the public program, the private firm is less able to exercise market power, so it earns lower profit or optimally chooses to shut down. Accordingly, the strongest push-back against a policy proposal to expand the public program tends to arise from special interest groups representing the private sector in the relevant market as well as high-income individuals. Incorporating portable vouchers into the universally available public option is an effective way to induce the private firm to continue operating in the market, and allow high-income individuals to obtain products with higher-quality than that offered by the public program.

It is worth pointing out that our mixed-duopoly model is general and not sector-specific. Even though it is calibrated to the health insurance sector, the main results may be applicable to other sectors. For example, in the education sector, historically public schools can be viewed as either offering universal public coverage (all parents can send their children to public schools for free) or a take-it-or-leave-it option (parents forfeit the benefit from the public school system if they send their children to private schools). Recently, several states have pursued education reforms by introducing school vouchers, allowing parents to keep the benefit from the public school system even when they choose private or charter schools. Similarly, public housing authorities have historically relied on public housing projects to provide social assistance (eligible households can get an apartment with reduced or zero rent). Recently a number of jurisdictions have switched to housing vouchers, allowing households to rent apartments in the private market if they deem the quality offered in the public program is inadequate. With minor adjustments, our model can be useful for understanding the distributional implications of such reforms.

References

- [1] Amiel, Yoram, John Creedy, and Stan Hurn (1999). “Measuring attitudes towards inequality.” *Scandinavian Journal of Economics* 101(1), 83–96.
- [2] Anderson, Simon P., André de Palma, and Jacques-François Thisse (1997). “Privatization and Efficiency in a Differentiated Industry.” *European Economic Review* 41(9), 1635–1654.
- [3] Atkinson, Anthony B (1970). “On the Measurement of Inequality.” *Journal of Economic Theory* 2(3), 244–263.
- [4] Barros, Fátima (1995). “Incentive Schemes as Strategic Variables: An Application to a Mixed Duopoly.” *International Journal of Industrial Organization* 13(3), 373–386.
- [5] Capatina, Elena (2015). “Life-cycle Effects of Health Risk.” *Journal of Monetary Economics* 74, 67–88.
- [6] Capps, Cory, David Dranove, and Christopher Ody (2018). “The Effect of Hospital Acquisitions of Physician Practices on Prices and Spending.” *Journal of Health Economics* 59, 139–152.
- [7] Chen, Chaoran, Zhigang Feng, and Jiaying Gu (2020). “Health, Health Insurance, and Inequality.” *mimeo*.
- [8] Clemens, Jeffrey, and Joshua D. Gottlieb (2017). “In the Shadow of a Giant: Medicare’s Influence on Private Physician Payments.” *Journal of Political Economy* 125(1), 1–39.
- [9] Cremer, Helmuth, Maurice Marchand, and Jacques-François Thisse (1991). “Mixed Oligopoly With Differentiated Products.” *International Journal of Industrial Organization* 9(1), 43–53.
- [10] Dafny, Leemore S (2010). “Are Health Insurance Markets Competitive?” *American Economic Review* 100(4), 1399–1431.

- [11] Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan (2012). “Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry.” *American Economic Review* 102(2), 1161–1185.
- [12] Dafny, Leemore, Kate Ho, and Robin S. Lee (2019). “The Price Effects of Cross-Market Mergers: Theory and Evidence from the Hospital Industry.” *RAND Journal of Economics* 50(2), 286–325.
- [13] Gabszewicz, J. Jaskold and Jacques-François Thisse (1979). “Price Competition, Quality and Income Disparities.” *Journal of Economic Theory* 20, 340–359.
- [14] Gaynor, Martin, Kate Ho, and Robert J. Town (2015). “The Industrial Organization of Health-Care Markets.” *Journal of Economic Literature* 53(2), 235–284.
- [15] Hosseini, Roozbeh, Karen Kopecky, and Kai Zhao (2020). “How Important Is Health Inequality for Lifetime Earnings Inequality?.” *mimeo*.
- [16] Ishida, Junichiro and Noriaki Matsushima (2009). “Should Civil Servants Be Restricted in Wage Bargaining? A Mixed-Duopoly Approach.” *Journal of Public Economics* 93(3), 634–646.
- [17] Jeske, Karsten and Sagiri Kitao (2009). “U.S. tax policy and health insurance demand: Can a regressive policy improve welfare?” *Journal of Monetary Economics* 56, 210–221.
- [18] Krishnan, Ranjani (2001). “Market Restructuring and Pricing in the Hospital Industry.” *Journal of Health Economics* 20(2), 213–237.
- [19] Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva (2015). “How elastic are preferences for redistribution? Evidence from randomized survey experiments.” *American Economic Review* 105(4), 1478–1508.

- [20] Lasram, Hejer and Didier Laussel (2019). “The Determination of Public Tuition Fees in a Mixed Education System: A Majority Voting Model.” *Journal of Public Economic Theory* 21(6), 1056–1073.
- [21] Lehmann-Grube, Ulrich (1997). “Strategic Choice of Quality When Quality is Costly: The Persistence of the High-Quality Advantage.” *RAND Journal of Economics* 28(2), 372–384.
- [22] Lockwood, Benjamin B. and Matthew Weinzierl (2016). “Positive and normative judgments implicit in US tax policy, and the costs of unequal growth and recessions.” *Journal of Monetary Economics* 77, 30–47.
- [23] Matsumura, Toshihiro (1998). “Partial Privatization in Mixed Duopoly.” *Journal of Public Economics* 70(3), 473–483.
- [24] Motta, Massimo (1993). “Endogenous Quality Choice: Price vs. Quantity Competition.” *Journal of Industrial Economics* 41(2), 113–131.
- [25] Poterba, James M. (1996). “Government Intervention in the Markets for Education and Health Care: How and Why?” In: Victor R. Fuchs, ed., *Individual and Social Responsibility: Child Care, Education, Medical Care, and Long-Term Care in America*. Chicago, IL: University of Chicago Press.
- [26] Saez, Emmanuel and Stefanie Stantcheva (2016). “Generalized social marginal welfare weights for optimal tax theory.” *American Economic Review* 106(1), 24–45.
- [27] Shaked, Avner and John Sutton (1982). “Relaxing Price Competition Through Product Differentiation.” *Review of Economic Studies* 49(1), 3–13.
- [28] Starc, Amanda (2014). “Insurer Pricing and Consumer Welfare: Evidence from Medigap.” *RAND Journal of Economics* 45(1), 198–220.

- [29] Wang, X. Henry and Bill Z. Yang (2001). “Mixed-Strategy Equilibria in a Quality Differentiation Model.” *International Journal of Industrial Organization* 19(1), 213–226.
- [30] Wauthy, Xavier (1996). “Quality Choice in Models of Vertical Differentiation.” *Journal of Industrial Economics* 44(3), 345–353.
- [31] Zuckerman, Stephen, Joshua McFeeters, Peter Cunningham, and Len Nichols (2004). “Trends: Changes in Medicaid Physician Fees, 1998–2003: Implications for Physician Participation.” *Health Affairs* (Jan–Jun Suppl Web Exclusive): W4-374–W4-384.
- [32] Zuckerman, Stephen, Aimee F. Williams, and Karen E. Stockley (2009). “Trends in Medicaid Physician Fees, 2003–2008.” *Health Affairs* 28(3), 510–519.

A Appendix

A.1 Proofs

Proposition 2.

Proof. For an individual with initial endowment (m, ε) facing a product (θ, p) , if the individual can afford it $m \geq p$, his optimal allocation between the two consumption goods after purchasing this product is

$$(n^*, x^*) = \begin{cases} (0, m - p) & \text{if } u_1(m - p, h(\theta) + \varepsilon) > 1, \\ (0, m - p) & \text{if } u_1(m - p, h(\theta) + \varepsilon) = 1, \\ (m - p - \bar{x}, \bar{x}) & \text{if } u_1(m - p, h(\theta) + \varepsilon) < 1 \ \& \ u_1(\bar{x}, h(\theta) + \varepsilon) = 1 \end{cases} \quad (5)$$

With the optimal allocation, his highest utility while purchasing this product is denoted $U(n^*, x^*, \theta)$. Note that (n^*, x^*) is continuous but not continuously differentiable with respect to (m, ε) , and so is $U(n^*, x^*, \theta)$.

(i) Facing two non-public varieties (θ_l, p_l) and (θ_h, p_h) with $\theta_l < \theta_h$ and $p_l < p_h$, if this individual prefers (θ_l, p_l) over (θ_h, p_h) , we know $U_l^* \geq U_h^*$. Moreover, since $u_{11} < 0$ and $u_{12} > 0$, we know $u_1(m - p_l, h(\theta_l) + \varepsilon) < u_1(m - p_h, h(\theta_h) + \varepsilon)$. Thus, there are three different scenarios to consider: (a) $1 \leq u_1(m - p_l, h(\theta_l) + \varepsilon) < u_1(m - p_h, h(\theta_h) + \varepsilon)$; (b) $u_1(m - p_l, h(\theta_l) + \varepsilon) < 1 \leq u_1(m - p_h, h(\theta_h) + \varepsilon)$; and (c) $u_1(m - p_l, h(\theta_l) + \varepsilon) < u_1(m - p_h, h(\theta_h) + \varepsilon) < 1$. For each of the three scenarios, we consider another individual with (m', ε') such that $m' = m - dm$ and $\varepsilon' = \varepsilon + d\varepsilon$, with $dm \geq 0$ and $d\varepsilon \geq 0$ representing non-negative changes.

(a) When $1 \leq u_1(m - p_l, h(\theta_l) + \varepsilon) < u_1(m - p_h, h(\theta_h) + \varepsilon)$, we have $1 \leq u_1(m' - p_l, h(\theta_l) + \varepsilon') < u_1(m' - p_h, h(\theta_h) + \varepsilon')$ as well. Denote the highest utility for the individual with (m', ε') as $U_l'^*$ and $U_h'^*$. For sufficiently small dm and $d\varepsilon$, $U_l'^* - U_l^* \approx -u_1(m - p_l, h(\theta_l) + \varepsilon)dm + u_2(m - p_l, h(\theta_l) + \varepsilon)d\varepsilon$ and $U_h'^* - U_h^* \approx -u_1(m - p_h, h(\theta_h) + \varepsilon)dm + u_2(m - p_h, h(\theta_h) + \varepsilon)d\varepsilon$, so we have $U_l'^* - U_h'^* \approx U_l^* - U_h^* + (u_1(m - p_h, h(\theta_h) + \varepsilon) - u_1(m - p_l, h(\theta_l) + \varepsilon))dm + (u_2(m - p_l, h(\theta_l) + \varepsilon) - u_2(m - p_h, h(\theta_h) + \varepsilon))d\varepsilon > 0$.

(b) When $u_1(m - p_l, h(\theta_l) + \varepsilon) < 1 \leq u_1(m - p_h, h(\theta_h) + \varepsilon)$, we have $u_1(m' - p_l, h(\theta_l) + \varepsilon') < 1 \leq u_1(m' - p_h, h(\theta_h) + \varepsilon')$ for sufficiently small dm and $d\varepsilon$. Now $U_l'^* - U_l^* \approx -dm + u_2(m - p_l, h(\theta_l) + \varepsilon)d\varepsilon$ and $U_h'^* - U_h^* \approx -u_1(m - p_h, h(\theta_h) + \varepsilon)dm + u_2(m - p_h, h(\theta_h) + \varepsilon)d\varepsilon$, so we have $U_l'^* - U_h'^* \approx U_l^* - U_h^* + (u_1(m - p_h, h(\theta_h) + \varepsilon) - 1)dm + (u_2(m - p_l, h(\theta_l) + \varepsilon) - u_2(m - p_h, h(\theta_h) + \varepsilon))d\varepsilon > 0$.

(c) When $u_1(m - p_l, h(\theta_l) + \varepsilon) < u_1(m - p_h, h(\theta_h) + \varepsilon) < 1$, we have $u_1(m' - p_l, h(\theta_l) + \varepsilon') < u_1(m' - p_h, h(\theta_h) + \varepsilon') < 1$ for sufficiently small dm and $d\varepsilon$. Now $U_l'^* - U_l^* \approx -dm + u_2(m - p_l, h(\theta_l) + \varepsilon)d\varepsilon$ and $U_h'^* - U_h^* \approx -dm + u_2(m - p_h, h(\theta_h) + \varepsilon)d\varepsilon$, so we have $U_l'^* - U_h'^* \approx U_l^* - U_h^* + (u_2(m - p_l, h(\theta_l) + \varepsilon) - u_2(m - p_h, h(\theta_h) + \varepsilon))d\varepsilon > 0$.

Overall, for arbitrarily large dm and $d\varepsilon$, we can break them into the sum of a sequence of sufficiently small changes $dm = \sum_{k=1}^K dm^k$ and $d\varepsilon = \sum_{k=1}^K d\varepsilon^k$, and apply the reasoning above recursively. This establishes the result for any individual with (m', ε') satisfying $m' \leq m$ and $\varepsilon' \geq \varepsilon$.

The Proof for (ii) mirrors the steps for (i), with the three different scenarios being (a) $1 < u_1(m - p_l, h(\theta_l) + \varepsilon) < u_1(m - p_h, h(\theta_h) + \varepsilon)$; (b) $u_1(m - p_l, h(\theta_l) + \varepsilon) \leq 1 < u_1(m - p_h, h(\theta_h) + \varepsilon)$; and (c) $u_1(m - p_l, h(\theta_l) + \varepsilon) < u_1(m - p_h, h(\theta_h) + \varepsilon) \leq 1$. \square

A.2 Figures

Section 4.1, Figure 3 shows the impact of Medicaid expansion on weighted private firm quality and markups. Figure A.1 shows the impact from this same scenario for the private firm basic and premium varieties. For both varieties, markups fall and quality rises with Medicaid expansion with the largest impacts occurring for the basic variety.

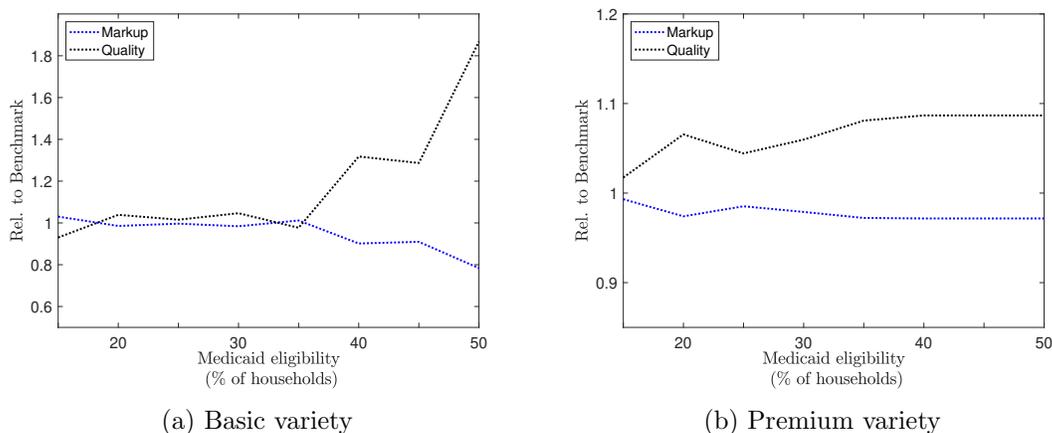


Figure A.1: Private Variety Markups and Quality across Medicaid Eligibility Thresholds

Notes: This is a disaggregated view, by private firm varieties, of Figure 3 panel (a). See Figure 3 notes for details.

Also in Section 4.1 we discuss the sensitivity of our results to two parameters in our model. Figure A.2 panel (a) shows welfare for all and the top 50 percent of individuals by income for different values of inefficiency ξ . We maintain that 20 percent of individuals are eligible for Medicaid and adjust the public budget (and taxes) so that public quality is as in the benchmark economy. As inefficiency rises, welfare falls. Expanding Medicaid coverage to 20 percent of individuals is welfare improving if public inefficiency does not exceed $2.5 \times \xi$. We also note that welfare for the bottom 20 percent of individuals (Medicaid eligible) does not change with ξ since public quality is kept fixed.

Figure A.2 panel (b) shows the sensitivity of welfare when 20 percent of individuals are eligible for Medicaid for values of ψ between 1 and 2. Specifically, panel (b) plots relative welfare—when 20 and 12.5 percent of individuals are eligible for welfare—for all and the bottom 20 percent of individuals for different values of ψ . Welfare for the benchmark economy, where 12.5 percent are Medicaid eligible, is recalculated for each value of ψ . For the range we consider, Medicaid expansion is welfare improving, though higher values of ψ imply smaller gains for those receiving Medicaid. We note that welfare from Medicaid expansion approaches welfare in the benchmark economy for values of $\psi > 2.8$.

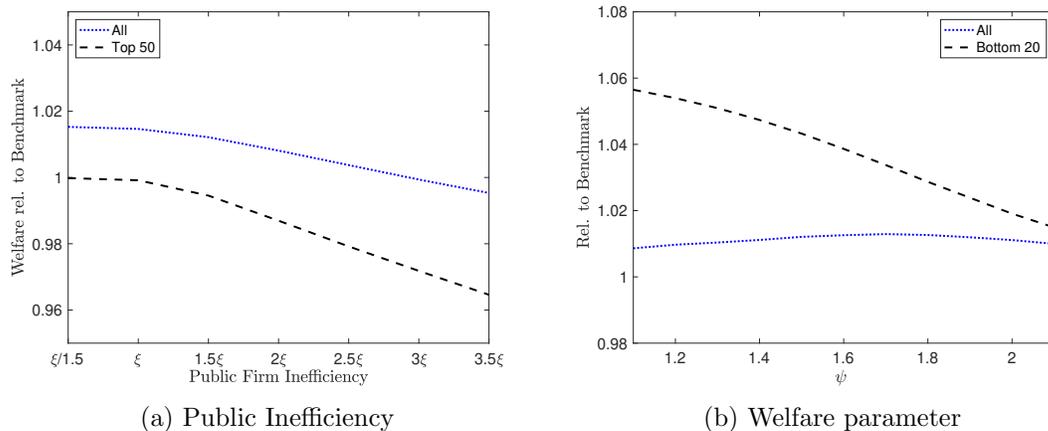


Figure A.2: Welfare gains sensitivity

Notes: Panel (a) shows welfare for varying levels of public firm inefficiency relative to benchmark inefficiency ξ , and panel (b) shows the sensitivity of welfare increases from Medicaid expansion for values of ψ . See text for details.

A.3 Alternate income distribution

In calibrating the model we assumed individual income is uniformly distributed. We show below that our main results are not driven by this assumption. We now let income be described by a Pareto distribution with shape parameter κ . Since income is exogenous, we choose this additional parameter κ so that the top decile of the income distribution accounts for 27 percent of all income (based on the Panel Study of Income Dynamics and taken from Chen, Feng and Gu, 2020). We then follow our calibration strategy in Section 3, except that we keep government inefficiency ξ as before.²⁹

²⁹The model cannot hit a private premium to public quality ratio of 2 while matching the other targets. Since our focus is to compare results from public firm expansion under a uniform and Pareto income distribution, we prioritize matching market shares, especially for the public firm.

Table A.1: MODEL FIT AND PARAMETER VALUES

Target Moments	Model	Data	Parameter
Income share, top decile	0.27	0.27	$\kappa = 1.9$
Healthcare spending (share of GDP)	15.4	12.4	$\gamma = 0.60$
Population choosing (%):			
Public, Q_1	12.3	12.6	$\psi = 1.60$
Private, $(Q_2 + Q_3)$	70.4	70.5	$\omega = 0.31$
Opt-out, Q_0	17.4	16.9	$\delta = 1.47$
Equilibrium product offerings:	Market shares:		
$(\theta_1^*, \bar{m}) = (0.018, 0, 0.0381)$	$Q_1 = 12.3$		
$(\theta_2^*, p_2^*) = (0.010, 0.033)$	$Q_2 = 16.3$		
$(\theta_3^*, p_3^*) = (0.049, 0.085)$	$Q_3 = 51.1$		
Private premium to Public quality:			
θ_3^*/θ_1^*	2.76		

Notes: Moments are from the National Center for Health Statistics which are based on the National Health Expenditure tables. See text for details.

Table A.1 shows the model fit. Since income is Pareto distributed the Medicaid eligibility threshold is much lower, and to hit the target moments much of the adjustment falls on ω . Table A.2 shows the results from our policy exercises, where in each case public quality is held fixed at the benchmark. Column (2) shows when Medicaid is expanded to cover 20 percent of the population, column (3) shows when there is universal coverage, and column (4) shows when there is universal coverage with a voucher option; the results in these columns should be compared to Tables 2, 3 and 4, respectively.

We highlight several points. First, public firm expansion and increased competition is consistent with lower private firm markups. Second, while private quality rises in the voucher system as before, we find that private quality falls under Medicaid expansion, and this is a pattern that holds for other thresholds of Medicaid eligibility as well. Third, the impacts on welfare from public firm expansion remains positive, and is quantitatively larger.

Table A.2: RESULTS WITH PARETO DISTRIBUTION

	(1) Benchmark	(2) Medicaid Exp. 20%	(3) Universal Coverage	(4) Voucher
Healthcare quality, θ				
Public	1.00	1.00	1.00	1.00
Private	1.00	0.97	–	1.32
Private markup	1.00	0.99	–	0.55
Population choosing (%)				
Public, Q_b	12.2	20.3	100.0	36.0
Private, Q_r	70.4	69.4	0.0	64.0
Opt-out, Q_0	17.4	10.3	0.0	0.0
Government				
B/Y (%)	1.10	1.82	8.97	8.92
Tax rate (%)	0.00	0.78	11.31	11.23
Social Welfare				
Entire population	1.000	1.027	1.083	1.083
Bottom 12.5%	1.000	1.000	1.000	1.000
Bottom 20%	1.000	1.064	1.065	1.064

Notes: Statistics are relative to the benchmark economy except for population shares across healthcare coverage. Private firm statistics are a weighted average of their varieties (by market shares).