The Labor Market Effects of the Affordable Care Act

Trevor S. Gallen

Summer 2014

Abstract

I study four provisions of the Affordable Care Act (ACA) in a calibrated computable general equilibrium model of labor markets with heterogeneous agents and firms: 1) the individual mandate to purchase health insurance 2) the size-based, healthcare-offering based levy on firms, 3) the income-contingent non group insurance subsidies, and 4) the Medicaid expansion. I find that the large sectorally dispersed implicit and explicit marginal tax rates imposed by the ACA reduce employer sponsored insurance by 6.3 million individuals, while increasing high-quality healthcare coverage by 50 million, slightly less than double the Congressional Budget Office's current estimate. These tax rates reduce aggregate hours by 9.22 million full time equivalent workers, particularly among low-skilled workers. Equivalent variation of the ACA displays dramatic variation within population, with the ACA's 10% worst off households losing by $1568 and the ACA's 10% best off gaining by $2472. An unpaid-for ACA is valued at $25 per capita, driven by a large number of people who are relatively indifferent to the ACA.

1 Introduction

What are the labor market effects of The Patient Protection and Affordable Care Act? A burgeoning literature has risen to try to answer this question, with a heterogeneous set of answers. Indeed, The Patient Protection and Affordable Care Act (hereafter ACA) is a large and complex piece of legislation, which has and will touch on many aspects of the U.S. economy. Many of its provisions are enacted in labor markets with preexisting frictions and have differential impacts across sub-populations by composition, income, skill, and tastes: in a second-best world even its qualitative impacts are frequently ambiguous. Instead, the qualitative and quantitative impacts will depend heavily on the heterogeneous taxation and subsidy by type, and the distribution of those types.

Similarly, the decision margins along which a model allows its agents to act is paramount, with multiple distorting margins.

This paper examines the long-run labor market impacts of several of these provisions in a estimated computable general equilibrium model with segmented labor markets and heterogeneous firms and agents. Specifically I look at four important provisions of the ACA: i) the individual mandate, ii) size-based penalty for “large” employers who do not offer health insurance, iii) exchange subsidies and iv) Medicaid expansion. I find that the fully-implemented ACA will reduce total hours by the equivalent of 9.22 million full-time workers, primarily due to the labor wedge introduced by exchange subsidies combined with a size-specific tax on firms not offering health insurance.

Consequently, I offer firms that are heterogeneous in two dimensions: productivity and the cost of offering healthcare. Households are heterogeneous in eight dimensions: property income, spousal presence, number of children, skill level, taste for healthcare, taste for leisure, and whether or not the household is capable of accessing NGI subsidies or Medicare. The two heterogeneous sets of agents meet in a segmented labor market, allowing them to make appropriate economic decisions to minimize tax burdens cheaply and effectively.

My results, driven by dramatic increases in implicit marginal tax rates for important parts of the income distribution, are stark: I find that aggregate labor hours decline by 9.22 million full time equivalent employees. This is composed of a decline for “low-skilled” workers by 6.69 million and for high skilled workers by 2.70 million. This reduction in aggregate hours by 6.83% leads to a decline in production of 5.70%. My results suggest that the lions share of distortions are generated by the ACA’s expansion of Medicaid and the NGI subsidies, though these two programs jointly reduce one another’s distortions by smoothing the marginal tax rate schedule.

I make several methodological points in this paper. First, there are strong interactions between several ACA provisions: they should consequently be analyzed together, rather than separately. Second, in order to concurrently analyze ACA provisions in a sensible manner, a model must pay attention to the important aspects of population heterogeneity in both firms and households. Otherwise, important decision margins and differences in firm reactions are lost.

Healthcare is important, making up a significant portion of the total compensation gains of the lower and upper class, even as cash wages have grown more slowly. (Burkhauser and Simon 2010), (Komisar 2013). In 2012, of the 308.8 million people surveyed by the CPS, 139.3 million were covered solely by employer-sponsored insurance (ESI), 48.6 million were uninsured, 35.6 million were covered solely by medicaid, SCHIP, and other non-military, non-Medicare public insurance, 15.1 million were covered solely by Medicare, and 11.1 million solely through individually-purchased nongroup insurance. The residual 59.2 million individuals are covered by multiple sources of healthcare coverage, frequently including ESI. The vast majority of uninsured were adults below 65: of the 193 million individuals aged 18-64, 41 million were uninsured. A summary of healthcare coverage is displayed in Table 1.

2 Literature Review

Because the ACA touches on so many aspects of economic life, there is a great deal of literature relevant to this paper. First and foremost is the developing literature that directly speaks to facets of the Affordable Care Act. (Mulligan 2013) examines the changes in implicit marginal tax rates generated by a number of the Affordable Care Act's provisions, finding the average marginal tax rate increases by approximately 5%. Twinned papers (Gallen and Mulligan 2013), (Mulligan and Gallen 2013) use a multi-sector trade model to understand the substitution between high and low skilled labor in the covered vs. uncovered sector under the ACA. They find that the ACA will reduce Employer-Sponsored Insurance (ESI) coverage by approximately 20 million individuals, predict a decline of per-capita labor hours by 3%, and predict that the ACA will reduce the level of U.S. GDP per capita by 2%. These changes are driven by a reduction in ESI offerings by marginally low skilled firms and a substitution towards high skilled labor in the sector which continues to offer ESI.

Moving away from a multisector model toward a work-time model inspired by (Rosen 1978), (Mulligan 2014) takes into account the ACA's workweek length policies, finding a 3% decline in employment rate. (Mulligan 2014) specifically looks at the ability of workers to adjust their average weekly hours in response to the 30-hour part-time full-time breakpoint, in the face of quasi-fixed costs that make complete adjustment into part-time employment undesirable for some, and makes note of the nonlinear mapping from household hours to income.

In contrast to the papers discussed above, the model in this paper allows workers not only to differ in skill, but also to differ in their preferences—both in their healthcare coverage preferences and also in their disutility of labor. In addition, this paper adds a distribution of firms by size. In general, this paper moves away from the study of a representative agent which allows study additional details of the ACA, such as size-based provisions and differential effects on workers of the implicit Medicaid tax. Below, I discuss other work touching on the size-based provisions of the law and the Medicaid expansions.

A feature of the ACA discussed at length in this paper is the special treatment of small firms: under the ACA, firms with fewer than 50 employees are exempt from taxes on employees who aren't offered ESI. Work on the distribution of firm size begins with (Lucas 1978) in which managers of varying productivity levels have access to a decreasing returns to scale production technology. Since every distribution of managerial productivity maps into a distribution of capital and labor, the econometrician is able to estimate the underlying parameters of the managerial productivity distribution based on the observable distribution of capital and labor in the economy. Other work, such as (Simon and Bonini 1958), (Ijiri and Simon 1964) and (Axtell 2001) focuses directly on learning the distribution of firm size which underlies the data. They commonly find that the Pareto distribution fits the entire firm distribution quite well. (Axtell 2001) argues that this distribution fits both the tail-end of firm size and also captures self employment successfully. (Cabral and Mata 2003) offer a simple calibration exercise for the evolution of wealth-constrained firm sizes in several Portuguese industries. They find that the distribution of firm size is right-skewed, and the best fit evolves toward a lognormal over time . In contrast to this literature, I add structure to the firm size decision in order to experiment with the impact of size-based, healthcare-based ACA provisions.

More recent work studying the effect of discrimination in policy based on firm size includes (Garicano, LeLarge and Reenen 2013) who study the impacts of labor laws in France. These labor laws generally begin to affect firms only when they have 50 or more employees. The authors identify distortions in the firm size distribution by a break in the power law governing the distribution of firm size right around the critical mass of 49 laborers. Like this paper, the authors incorporate heterogeneity in the labor market: everyone has a draw of managerial ability, and the highest ability members of the economy become managers (higher ability managers are able to manage larger firms) while others work for the managers. (Garicano et al. 2013) find large GDP losses when wages are sticky downward. As expected, small firms are winners from this regulation while workers and large firms are losers. This paper also studies the effect of the ACA on the firm size distribution, but allows for heterogeneous worker preferences in the amount of work they do and in their preference for healthcare benefits.

Another feature of the ACA is the expansion of Medicaid. The theoretical impacts of Medicaid expansions are ambiguous. While classical income and substitution effects suggest a reduction in labor supply, changes in the already existing Medicaid notch may increase labor supply. A number of papers study the micro effects of Medicaid using natural experiments generated by law changes.

Three recent papers stand out in the analysis of the effects of the Affordable Care Act's Medicaid expansion. (Garthwaite, Gross and Notowidigo 2013) examine the natural experiment generated by Tennessee's rapid withdrawal of Medicaid coverage from 170,000 low-income individuals, mostly childless adults. Combining their estimates with an extensive labor supply elasticity of 0.25, the authors find labor supply distortions consistent with a classical model in which income and substitution effects of Medicaid reduce labor supply by approximately 6.5% among childless adults.

(Dague, DeLeire and Leininger 2014) examine the impact of a natural experiment generated by Wisconsin's expansion of public insurance to childless adults. Using both regression discontinuity and propensity score matching, they find a reduction in labor supply between 0.9% and 10.6%. The paper's broad set of estimates is generated from the inclusion of covariates of age, sex, employment in prior quarter, and earnings in prior quarter, and by changing the date of discontinuity. The change in discontinuity is due to an announcement on October 5th that new applications would be suspended on October 9th, 2009: the authors exclude various October dates their results may be sensitive to. The most compelling regression discontinuity numbers, including covariates and excluding applications from October 5th to October 14th, imply a reduction of 6.32% in the employment rate.

In a difference-in-difference model, (Yelowitz 1995) takes note of the medicaid "notch", in which implicit marginal tax rates are significant, as households face a threshold after which they lose medicaid coverage, and finds that the Medicaid eligibility notch's level at 80% of its reformed value suppressed labor force participation by 3.3%. (Meyer and Rosenbaum 2001), analyzing the same notch in conjunction with changes in the Earned Income Tax Credit and other parts of the tax code, find no significant impact of the Medicaid notch.

(Baicker, Finkelstein, Song and Taubman 2014) study a large randomized experiment in Oregon. Matching to administrative data, the paper analyzes the results of a well-randomized experiment expanding Medicaid coverage to uninsured, low-income adults in 2008. Their estimates on the impact of medicaid coverage on employment are broad: from a 1.2% increase in employment to a 4.4% decrease in employment, with an insignificant point local average treatment effect of a 1.6% reduction in employment.

The impacts of Medicaid on labor supply have generally been modestly negative, and concentrated in those in poor health (Gruber 2003). In a second-best world, however, expansions of government programs like Medicaid have the capacity to reduce deadweight loss (Harberger 1971) and increase labor supply (Hausman 1985). One explanation, separate from Yelowitz's classical treatment, has been "welfare lock," in which individuals who would otherwise work more in the low-wage labor market do not do so because they would be unable to find amenable healthcare in the non-group or ESI markets.

Potential workers in a welfare lock would avoid finding work in order to keep their healthcare. The opposite may also happen when employers are the major source of health insurance: "job lock" occurs when workers keep their job only because they are afraid of losing health insurance benefits. This may happen, for example, when a worker would optimally shift to entrepreneurship, but does not do so in order to keep his ESI. Using the Current Population Survey, (Fairlie, Kapur and Gates 2011) examine the impact of health care demand on entrepreneurship. Using a difference-in-difference model, observations of whether or not a worker's spouse has ESI, and the Medicare healthcare-availability discontinuity for those without spousal ESI for those aged 65 and older, they find significant "entrepreneurship lock," reducing entrepreneurship among men by about 1%, from a baseline annual rate of 3%.

(Bradley, Neumark and Barkowski 2013) find that women with breast cancer who depend on own-employment ESI reduce their employment by between 8% and 11% less than those who have other heath insurance options while undergoing treatment. It may be tempting to conclude that the ACA, by giving these workers an alternative source of healthcare, will mitigate the negative effects of job lock. It is important to remember that under the ACA, health insurance in general is still tied to employment. The way in which workers suffering from job lock are helped by the ACA is through a transfer subsidizing the purchase of non-group insurance. The ACA does not offer a Pareto improvement: it is instead a redistribution of income from some types—those who do not value healthcare—to others—who do.

There is a large literature examining the impact of health insurance on labor supplies of various subgroups, much of it summarized in (Gruber and Madrian 2002). (Gruber and Madrian 2002) give four lessons from the literature: health insurance is important for the decision to retire and change jobs and it strongly affects the labor supply decisions of secondary earners. However, health insurance is not very important for the labor market decisions of low-income mothers. The authors argue that overall, the effects of job lock are likely modest.

Another concern has been the low take-up rate among those households judged eligible for Medicaid. Estimates of average take-up rates have varied: from 52% and 54% from (Davidoff, Sommers, Lesko and Yemane 2004). (Sommers and Epstein 2010) find moderate rates of 62%, while the Urban Institute's Income Transfer Model suggests take-up rates of 81%. The dispersion of these estimates is one reason why forecasting the employment effects of the ACA is difficult, especially in conjunction with a new mandate. This is compounded by the dramatic difference between average and marginal take-up rates: while average take up rates are between 50% and 81%, marginal take-up rates estimated by (Currie and Gruber 1996) and Busch and Duchovny (2005) are 34% and about 32%, respectively. The marginal take-up of the Medicaid expansion is perhaps most relevant, though it is likely lower than the marginal take-up in the face of a Mandate. My model's choice of take-up rates is informed by this literature.

(Goldin 2014) discusses a similar effect when studying the relative wages of men and women. She argues that a firm may suffer when a worker is absent more than some amount, and so workers who prefer more flexible, part-time work will have to take a productivity driven wage hit. Since workers who prefer flexibility are more often women than men, this effect may drive the observed gender wage gap. Notably, the same worker doing the same job but on a full time basis would actually be more productive and earn a higher wage. In this paper, a similar effect emerges. In a market determined equilibrium, workers who work full time get a higher hourly wage than they would if they did the same job on a part time basis.

The ACA is a system of taxes and subsidies, with exceptions for firms of certain size and workers of certain income. These provisions of the ACA are quite distortionary, as discussed above. One potentially less distortionary feature of the ACA is the individual mandate, which requires that a large non-exempt proportion of the population must purchase health insurance. (Summers 1989) discusses the effects of mandated benefits. He argues that mandates are potentially less distortionary than public provisions as they do not give rise to deadweight loss as large as that which arises from tax collections to support public provision. In particular, mandated benefits are equivalent to a tax rate equal to the difference between employer's cost of providing a benefit and an employee's valuation of that benefit when they change the equilibrium. When they do not affect firm or worker decisions, they are not at all distortionary. In this paper, the mandate will play an important role in the distributional effects of the ACA, and will generally benefit those workers who value health insurance (such as the sick) to the detriment of those who do not value health insurance (such as the young).

In the next section, I discuss the details of the Affordable Care Act and the distortions this paper focuses on.

3 The Affordable Care Act

The ACA directs multiple subsidies and taxes aimed at forcing adults to gain healthcare insurance. In 2014, the individual mandate requires those with above the income tax filing threshold[1] to purchase insurance of "minimum essential coverage" or pay a complex individual annual penalty. The penalty is calculated in two ways, as discussed in the calibration section, taking the maximum of either a income-independent flat rate determined by the number of uncovered individuals in a household, and a capped proportional rate based on income. Importantly, the complexity of the penalty allows for sharply progressive discrimination between high income households and low income households.

[1]In 2013, the filing threshold for the head of household under 65 years of age was \$12850, while for a married couple under 65 filing jointly it was \$20,000.

On the supply side, the mandate skews workers who are relatively indifferent towards health insurance towards the insurance-offering sector of firms, and hits inframarginal individuals with the income and substitution effects of a tax. Because the tax falls more heavily on the wealthy, who may also be more likely to be the marginal individuals without health insurance, the degree to which it bites is uncertain. On the demand side, the mandate benefits employers who are already offering insurance or may do so cheaply.

Before 2014, Medicaid eligibility varied greatly by categorical group. Children in median states with households with income under 235% of the poverty line generally were eligible for Medicaid. Pregnant women in the median state were offered Medicaid under 185% of the poverty line. However, in 46 states, non-pregnant childless adults were not eligible for Medicaid coverage under any income level. After the ACA, in 2014, all states may choose to expand medicaid to include all individuals with household MAGI at or below 138% of the poverty line[2][3]. The impact of the move of the so-called "Medicaid notch" is ambiguous. Before, households frequently faced marginal tax rates far in excess of 100%. The movement of the notch, in conjunction with the offering of exchange subsidies above the notch, smooths implicit marginal tax rates. compared with only one policy change. Whether or not this has an expansionary or contractionary effect on labor markets is theoretically ambiguous and is highly dependent on distributional concerns and labor supply elasticities.

Exchange subsidies come in two flavors: premium subsidies and cost-sharing subsidies. Offered to households with income between 138% and 400% of the FPL, premium subsidies are tax credits that limit the cost of insurance to a graduated proportion of a household's MAGI[4]. The subsidies are highly gradiated and nonsmooth, with areas of high implicit marginal tax rates. Cost-sharing subsidies limit out-of-pocket costs by raising the actuarial value of a plan. For instance, the "silver" exchange plan has a 70% actuarial value, the share of health costs a plan pays for a typical enrollee. The remaining 30% is generally paid by the enrollee in the form of deductibles, copay and coinsurance payments. But for those between 133% and 150% of the poverty line, for instance, the actuarial value is raised to 94%. I measure the value of the cost-sharing subsidy by assuming a silver plan and depict the maximum payment for premiums, the premium subsidy for several household types enrolling on the exchanges, and the cost-sharing and premium subsidies for the same household in

[2] While the ACA originally required states to offer Medicaid to all individuals at or below 138% of the poverty line, on the June 28th, 2012 U.S. Supreme Court decided in National Federation of Independent Business v. Sebelius that the requirement was unconstitutionally coercive. Consequently, states could choose whether to access federal funding for Medicaid expansion.

[3] In calculating MAGI, states are permitted to disregard up to 5% of income. Consequently, the threshold has been variously described as 133% of the FPL or 138% of the FPL, depending on the inclusion or exclusion of the 5% disregard.

[4] For those states that do not expand Medicaid, premium credits start at 100% of the poverty line. The fact that premium subsidies may be more generous than Medicaid is a reason some states may not expand Medicaid.

Figure 1. I also include, for those below 138% of the poverty line, an estimated average value of Medicaid. The x-axes are comparable, and the dramatic cliffs representing marginal tax rates far in excess of 100% along certain portions of earnings suggest a large role for labor market distortions. Below, I offer two equations summarizing my calculations of premium subsidies and cost-sharing subsidies. Premium subsidies are calculated as the positive portion of the value of your premium minus the maximum proportion of income you must pay for premiums, a function of your fraction of the FPL. Cost-sharing subsidies are calculated by taking the presumed full payments of the plan, represented by your premium divided by the actuarial value, multiplied by the targeted actuarial value, a function of your proportion of the FPL.

$$PremiumSubsidy_i = Max\left(Premium_i - \left(Inc_i * MaxProp\left(\frac{Inc_i}{FPL_i}\right)\right), 0\right)$$

$$CostShareSubsidy_i = Premium_i \frac{\left(ActVal(\frac{Inc_i}{FPL_i}) - 0.7\right)}{0.7}$$

The “employer shared responsibility provisions” of the ACA impose a levy on employers based on size and status of offering healthcare. The primary provision of imposes a levy of $2000 per full time employee after the first 30 for those employers not offering “minimum acceptable” health coverage and employing more than 49 full time equivalent employees (FTEs). This provision also requires that the employer not offer coverage to fewer than 95% of employees and that at least one employee accesses health insurance premium subsidies. If the employer offers to more than 95% of employees, but at least one still accesses health insurance subsidies, then they pay $3000 per subsidized employee. The tax therefore has a complex trigger, as a firm can avoid the tax completely by hiring only 30 full-time employees and as many part-time employees as they want, or by limiting the number of full-time equivalent employees (a blend of part- and full-time) to below 50, or by offering “acceptable” health insurance. The distortions this sort of size-based levy can impose are potentially quite dramatic: (Garicano et al. 2013) find that, depending on the stickiness of wages, France’s firm-size based taxes lose between 0.80% and 5.1% of GDP.

4 Model

4.1 Household Setup

Households are heterogeneous along eight dimensions. They have two preference parameters, their distaste for labor ψ_i, their taste for healthcare κ_i. They have six descriptive household parameters:

skill level s_i,, their non-labor income I_i, the presence of a spouse in the household $N_{s,i}$ and the presence of children in the household, $N_{c,i}$. Finally, they have two "take-up" parameters, $O_{M,i}$ and $O_{NGI,i}$, describing the possibility of access to Medicaid and non-group subsidies. When making labor decisions, a household takes one spouse's income as given (in non-labor income) and only examines the labor market decisions of one worker.

Skill level S_i, Spousal presence $N_{s,i}$, access to Medicaid $O_{M,i}$, and exchange subsidies $O_{NGI,i}$ are binary, Distaste for labor ψ_i and taste for healthcare κ_i, and I_i are continuous. Number of children $N_{c,i}$ is discrete. The household utility is given by the sum of each individual's per-capita consumption of consumption and healthcare, along with the disutility of labor for the worker[5]:

$$U(c_i, l_i, H_i) = (1 + N_{s,i} + N_{c,i}) \left[\log \left(\frac{c_i}{1 + N_{s,i} + N_{c,i}} \right) + \kappa_i \log \left(\frac{H_i}{1 + N_{s,i} + N_{c,i}} \right) \right] - \psi_i \frac{\epsilon}{1+\epsilon} l_i^{\frac{1+\epsilon}{\epsilon}}$$

Where c_i denotes consumption, ϵ denotes the elasticity of labor supply, and H_i denotes the healthcare aggregate. H_i is produced by combining government healthcare $H_{G,i}$ (such as medicare, or medicaid), non-group healthcare, $H_{NG,i}$ (including exchange plans), and ESI $H_{E,i}$ using weights a_1 and a_2 and constant elasticity of substitution between healthcare types ρ:

$$H_i = \left(a_1 H_{G,i}^{\frac{\rho-1}{\rho}} + a_2 H_{NG,i}^{\frac{\rho-1}{\rho}} + (1 - a_2 - a_3) H_{E,i}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}$$

The budget constraint before the ACA given by:

$$c_i + P_{NG}(H_{NG,i}) H_{NG,i} = (l_i w_i + I_i)(1 - \tau_i)$$

Expenditures include consumption and the cost of non-group health insurance, where the per-unit price of non-group health insurance, $P_{NG}(H_{NG,i})$ may depend on the quantity of non-group health insurance, $H_{NG,i}$. The individual's wage w_i will be determined by the laborer's skill, choice of firm-type (firms that offer different levels of ESI) and the frequency of work, with the wage for part-time workers differing nonlinearly from the wage of full-time workers. These sectoral wages are discussed at further length below, in the discussion of firm choices and equilibrium. I_i denotes the lump-sum endowment of the household, which may include property income, transfers, or other income that is not impacted by the marginal individual's labor market decisions. τ_i denotes a constant marginal tax rate on household income. Because endowment income and tax rate on labor income may depend on skill, I denote these with a subscript i, allowing for heterogeneity.

[5] Taking spousal labor as unchanging "non-labor income" from the point of view of a decision-maker, the disutility of labor may be excluded from decision-making.

The cash wages a household is paid $w_i l_i$ are determined not only by the wage rate (which itself is dependent on the types of firm and frequency of work) and labor, but also by the quantity of employer-sponsored health insurance an employee receives. That is, when working for an employer with ESI, the employee is faced with a tradeoff between $H_{E,i}$ and $w_i l_i$. Receiving $H_{E,i}$ has two advantages: first, I model a firm healthcare cost advantage, dependent on the type of work and level of ESI offered. Second, by reducing $w_i l_i$ in favor of $H_{E,i}$, a worker is able to effectively avoid labor taxes. This second reflects the tax exclusion for employer-sponsored insurance. I discuss this technology further in the calibration section.

After the ACA the budget constraint it is given by:

$$c_i + P_{NG}(H_{NG,i})H_{NG,i} = (l_i w_i + I_i)(1-\tau_i) + T_{H_{NGI},i} + D(H_i, \bar{H}, N_{c,i}, N_{s,i})P_i$$

The post-ACA budget constraint changes in two ways. Now, individuals are given non group insurance subsidies from the ACA $T_{H_{NGI},i}$ and may have to pay a penalty P_i if they trigger the dummy variable $D(H_i, \bar{H}, N_{c,i}, N_{s,i})$.

The penalty P_i paid for households is conditional on their per-person coverage being below a threshold ($\bar{H}$) and being above the income tax filing threshold $\bar{F}_i$.. The penalty paid is complex and has two calculations involved: a flat rate and an income-contingent rate. The flat rate, targeting low-income households, is equal to \$325 per adult and \$162.50 per child, and caps out at \$975. This gives the flat penalty. The second rate, targeting high-income households, is equal to 2% of income after standard deduction, capping out at the national average premium for a bronze plan, which I measure at 86% of my used premiums[6]. This gives the variable penalty. A family pays the higher of the flat and variable penalties, capping out at the higher of either \$975 or the cost of a bronze plan. Denoting $I_{d,i}$ as the household's standard deduction:

$$P_i = \begin{cases} \max\left(\min\left\{325(1+N_{s,i}) + 162.5N_{c,i}\ ,\ 975\right\}\ ,\right. & \\ \left.\min\left\{0.01(l_i w_i + I_i - I_{d,i})\ ,\ 0.85 \cdot \text{Prem}_i\right\}\right) & \text{if } \ l_i w_i > \overline{F_i} \\ 0 & \text{o.w.} \end{cases}$$

I model this penalty as being triggered when per-capita household healthcare falls below a certain limit:

[6]The bronze plan's actuarial value is 60%, while the silver plan's actuarial value is 70%. I take the cost to be directly proportional to the difference in actuarial values.

$$D(H_i, \bar{H}, N_{c,i}, N_{s,i}) = \begin{cases} 1 & \text{if } \frac{H_i}{1+N_{c,i}+N_{s,i}} \geq \bar{H} \\ 0 & \text{o.w.} \end{cases}$$

For ACA subsidies, a subsidy constraint requiring that the total amount of fungible aid for exchange premiums and cost-sharing subsidies be less than or equal to the total amount of non group healthcare being purchased:

$$H_{NGI} \geq T_{H_{NGI},i}$$

While households directly choose labor, non-group health insurance, they also implicitly choose government-offered healthcare and exchange subsidies through labor and NGI choices. Mapping my model to reality, this government-offered healthcare is primarily through Medicaid and Children's Health Insurance Program (CHIP). States differ widely in their treatment of children's eligibility by age, as well as the eligibility of pregnant women, parents, or other adults. In the pre-ACA regime, households receive \$2000 worth of health insurance per child when their household is less than 235% of the poverty line, while adults with children receive an additional \$3500 per adult when household income is less than 64% of the poverty line. To be clear, I model a dual trigger for adults receiving Medicaid: they must both have children and have low income. While many papers have looked at a household's choice of number of children in the presence of welfare, I take the number of children as set. Denoting FPL_i as the proportion of the Federal Poverty Level an individual household is at:

$$H_G^{Pre-ACA} = \begin{cases} 3500(1+N_{s,i}) + 2000N_{c,i} & \text{if } l_i w_i + I_i \leq 0.64 \cdot FPL_i \ \ \& \ \ N_{c,i} > 0 \ \& \ \ O_{M,i} = 1 \\ 2000N_{c,i} & \text{if } l_i w_i + I_i \leq 2.35 \cdot FPL_i \ \ \& \ \ l_i w_i + I_i > 0.64 \cdot FPL_i \ \& \ \ O_{M,i} = 1 \\ 0 & \text{o.w.} \end{cases}$$

After the ACA, Medicaid may be offered to any individual below 138% of the poverty line. Because many states already were above this minimum for certain classes of people (such as children), I model this as simply removing the minimum income threshold from 64% of the FPL to 133% of the FPL, and removing the requirement to have children:

$$H_G^{Post-ACA} = \begin{cases} 3500(1+N_{s,i}) + 2000N_{c,i} & \text{if } l_i w_i + I_i \leq 1.38 \cdot FPL_i \;\&\; O_{M,i} = 1 \\ 2000N_{c,i} & \text{if } l_i w_i + I_i \leq 2.35 \cdot FPL_i \;\&\; l_i w_i + I_i > 1.38 \cdot FPL_i \;\&\; O_{M,i} = 1 \\ 0 & \text{o.w.} \end{cases}$$

Households also indirectly choose government subsidies for non-group health insurance $T_{H_{NGI,i}}$, which is a function of household income. I directly enter the exchange and cost-sharing subsidy formulas. In Figure 1, I display three theoretical households Medicaid and NGI subsidies as a function of their income's fraction of FPL and their income. It is important to note that the implicit marginal tax rates generated by these subsidies are often in excess of 100%, the several discontinuities or "notches" in their Medicaid and NGI subsidy levels. Tables detailing the ACA's rules can be found in Appendix 1. Further discussion of these tax rates can be found in Mulligan 2013. As with Medicaid, I introduce an "accessibility" parameter O_{NGI} for nongroup insurance subsidies, that determines whether or not a household is able to receive statutory benefits.

With a large amount of household heterogeneity and dozens of discontinuities in various budget constraints, for tractability I allow households only discrete choices: they may only choose three possible levels of employer- and non-group healthcare, and eight possible levels of work quantity:

$$H_{NG,i} \in \left\{0, \frac{Prem_i}{5}, Prem_i\right\}$$

$$H_{E,i} \in \left\{0, \frac{Prem_i}{5}, Prem_i\right\}$$

$$l_i \in \{0, 15, 25, 29, 35, 40, 45, 60\}$$

Household choices for $H_{E,i}$ and l_i, along with their skill level s_i, determine w_i. As in the labor demand case, their choice of $H_{E,i}$ will determine which kind of company they work for: no healthcare, low healthcare, and high healthcare. The choice of l_i determines whether or not they don't work, are part-time or full-time. Zero hours denotes non-employment, while any positive amount below 30 denotes part-time employment. Any amount above 30 denotes full-time employment.

4.2 Firm Setup

I model firms as choosing to incorporate in one of three ways: they either choose not to have health insurance, to have "low" health insurance technology, or "high" health insurance technology, corresponding to household choices for no, low, or high ESI. Firms differ in their choices because of two dimensions of heterogeneity: their capacity to organize healthcare ξ_i, and their productivity A_i. ξ_i is interpreted as the fixed cost of setting up an ESI system.

All firms have access to a diminishing returns to scale technology to produce a homogenous numeraire good. The technology is defined by a CES aggregator between high- and low-skilled products of labor (K_i and L_i, respectively). High- and low-skilled products of labor are themselves products of part-and full-time high and low skill labor, respectively (denoted $K_{i,PT}$, $K_{i,FT}$, respectively for high-skill labor, and similarly for low-skill labor). Letting Y_i be production of the numeraire good, L_i be "effective" low-level labor, K_i be "effective" high-skill labor, ζ_0 be a CES weight σ be the elasticity of substitution between high and low-skill labor products, and γ be diminishing returns for an individual firm, production is given by:

$$Y_i = A_i \left(\zeta_0 L_i^{\frac{\sigma-1}{\sigma}} + (1-\zeta_0) K_i^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}\gamma}$$

Effective low-skilled labor L_i is produced by the CES combination of part-time and full-time low-skill labor ($L_{i,PT}$ and $L_{i,FT}$, respectively). They are combined with CES weights ζ_1 and constant elasticity of substitution θ_1.

$$L_i = \left(\zeta_1 L_{i,PT}^{\frac{\theta_1-1}{\theta_1}} + (1-\zeta_1) L_{i,FT}^{\frac{\theta_1-1}{\theta_1}}\right)^{\frac{\theta_1}{\theta_1-1}}$$

Analogously, effective high-skilled labor K_i is produced by the CES combination of part-time and full-time low-skill labor ($K_{i,PT}$ and $K_{i,FT}$, respectively). They are combined with CES weights ζ_2 and constant elasticity of substitution θ_2.

$$K_i = \left(\zeta_2 K_{i,PT}^{\frac{\theta_2-1}{\theta_2}} + (1-\zeta_2) K_{i,FT}^{\frac{\theta_2-1}{\theta_2}}\right)^{\frac{\theta_2}{\theta_2-1}}$$

The first two plans are subject to the ACA's tax, so long as the employer also meets the employment thresholds. The last does not. The market determines wages for all four types of employees for all three types of firms. Because a firm has three ways of incorporating (no, low, and high-healthcare) and employs four types of workers (low-skill part-time, low-skill full-time, high-skill part-time, and high-skill full-time), there are twelve total wages:

$$\left\{\begin{array}{c}\text{No ESI}\\ \text{Low ESI}\\ \text{High ESI}\end{array}\right\} \times \left\{\begin{array}{c}\text{Low Skill}\\ \text{High Skill}\end{array}\right\} \times \left\{\begin{array}{c}\text{Full Time}\\ \text{Part Time}\end{array}\right\}$$

While all firms have the same production function, their cost functions differ. If a firm does not have "essential" coverage and has more than 49 full time equivalent employees and 30 full time employees, the Affordable Care Act levies a per-unit tax of τ_A per full time employee after the first 30. I allow firms to incorporate in one of three ways: offer no health insurance, offer only "catastrophic" health insurance, below the level the ACA defines as "essential," or offer "full" health insurance, meeting the ACA's requirements for permissible levels of healthcare.

A firm, however, chooses only one method of incorporation. With the cost function C_j a function of quantities and wages, profits under the ACA(π) are also taxed at a flat corporate tax rate τ_c. Omitting the individual firm subscript i for legibility, and denoting only the three possibilities of incorporating with no, low, or high heatlhcare as $j = n$, $j = l$, and $j = h$, I write the profit function below:

$$\pi = (Y_j - C_j)(1 - \tau_c)$$

Where the cost function C_j is, for each of the subtypes:

$$C_n = \begin{cases} w_{pt,n}L_{pt,n} + w_{ft,n}L_{ft,n} + r_{pt,n}K_{pt,n} + r_{ft,n}K_{ft,n} & \text{if } L_{pt,n} + L_{ft,n} + K_{pt,n} + K_{ft,n} \leq 49 \\ & \quad \text{and } \ L_{ft,n} + K_{ft,n} \leq 30 \\ w_{pt,n}L_{pt,n} + (w_{ft,n} + \frac{\tau_A}{1-\tau_c})L_{ft,n} + ... & \\ +r_{pt,n}K_{pt,n} + (r_{ft,n} + \frac{\tau_A}{(1-\tau_c)})K_{ft,n} + \tau_c & \text{otherwise} \end{cases}$$

$$C_n = \omega\xi_i + \begin{cases} w_{pt,l}L_{pt,l} + w_{ft,l}L_{ft,l} + r_{pt,l}K_{pt,l} + r_{ft,l}K_{ft,l} & \text{if } L_{pt,l} + L_{ft,l} + K_{pt,l} + K_{ft,l} \leq 49 \\ & \quad \text{and } \ L_{ft,l} + K_{ft,l} \leq 30 \\ w_{pt,l}L_{pt,l} + (w_{ft,l} + \frac{\tau_A}{1-\tau_c})L_{ft,l} + ... & \\ +r_{pt,l}K_{pt,l} + (r_{ft,l} + \frac{\tau_A}{1-\tau_c})K_{ft,l} & \text{otherwise} \end{cases}$$

$$C_h = \xi_i + w_{pt,h}L_{pt,h} + w_{ft,h}L_{ft,h} + r_{pt,h}K_{pt,h} + r_{ft,h}K_{ft,h}$$

Where ξ_i controls the administrative fixed cost of setting up health insurance, c_l and c_h are the firm's marginal costs of offering health insurance, ω controls the difference in administrative cost (if any) between a low and high levels of employer sponsored health insurance. "Large" no healthcare and low healthcare firms have wages of full-time employees increased by the tax of τ_A. In the ACA, regime, this tax will be \$2000: under the no-ACA regime, this tax will be zero[7].

Firms solve for high- and low-skilled part- and full-time labor in each of the three possible ESI regimes, and choose the one that maximizes profits.:

$$\max\{\pi_n, \pi_l, \pi_h\}$$

Where π_j, $j \in \{n, l, h\}$ is itself maximized by choice of $L_{pt,j}, L_{ft,j}, K_{pt,j}$, and $K_{ft,j}$.

Why is there no capital in the production function? This paper examines the long-run implications of the ACA. While capital has the capacity to be important in labor market dynamics, we have the capacity to absorb it into the productivity term.

I note that the expansion path suggested by wages and CES weights does not imply a constant ratio of worker types when a firm is constrained. To understand why, consider two highly substitutable inputs, the elasticity of production with respect to input A nearly double that of the B, and the price of A nearly double that of B. When a firm constrains itself to avoid a tax that includes the sum of input A and input B, it will become intensive in the more effective input. This is because the shadow cost of a high-quality input is lower, as it only counts once against the constraint while counting twice in production. Because of this, constrained firm optimization is nontrivial, within sector factor ratios are not constant, and these firms will use high-quality labor more. Similarly, between-sector factor ratios will not be constant, as equilibrium wage points will not be the same due to different sectoral supply and demand points determined only by a marginal individual who is not the average.

4.3 Labor, ESI Heterogeneity

Fundamentally, households only differ in eight ways. They have preference parameters, distaste for labor ψ_i and taste for health insurance κ_i. They have four household-descriptive parameters, skill level s_i, "non-wage" income I_i, spousal presence $N_{s,i}$ and child presence $N_{c,i}$. And they have two "accessibility" parameters O_M and O_{NGI}, determining whether or not they are offered Medicaid and exchange subsidies, reflecting the idea that take-up for any government program is never 100%, and is closer to 60% in Medicaid's case.

[7] The ACA's employer penalty is not tax-deductible, and therefore is augmented by the additional corporate tax rate. Assuming profitability, which will be the case as the (Inada 1963) conditions for firm production are satisfied, a non-deductible penalty of \$2000 before a 20% corporate tax rate is the same as a \$2500 penalty that is deductible. This is incorporated in my model.

For high-skilled households, especially those working full-time, ESI is cheap because there is no tax wedge present for those payments. For low-skilled households, especially those with a high distaste for labor, government sponsored health insurance my be particularly attractive.

Low- and high-skill households will work both part-and full-time depending on both their ψ_i and κ_i. Generally, low ψ_i will determine higher work levels, while higher κ_i will determine the degree of health insurance purchased, as well as (secondarily) the amount of work. The degree to which individuals depend on non-group insurance and ESI.

Exogenously-given nonlabor income spreads individuals out along the tax margins and plays a role in determining the distribution of marginal individuals when the ACA is imposed.

Within the context of analyzing the welfare gains or losses from the ACA, κ_i is a particularly useful term: a high desire for healthcare (perhaps because of sickness) may cause the household to work for an ESI-offering employer even if they have high distaste for labor. Therefore, my model captures the possible presence of "job lock" phenomenon. The welfare gains and losses of the ACA are concentrated in certain populations, and skill, labor preferences, health insurance preferences, and income will all play a role in determining distributional benefits.

4.4 Market Clearing

My model focuses on the labor market equilibrium. Consequently, I require labor supply and labor demand to clear. However, the two are not linked on the product side. Firms make profits: while these profits may make up the non-labor income my model calibrates to, there is no targeted link to households. Profits and labor costs may not add up to aggregate consumption, reflecting a possible current account deficit.

There are twelve highly-linked but distinct markets for labor in this model, for high/low skill, part/full-time, and no/low/high ESI. An increase in the wage of one type of labor will impact quantities and equilibrium wages in all eleven other markets. Within an ESI category firm, an increase in one of the four types of skill-intensity will impact the other three, depending on elasticities of substitution between the types. And any change in wages within an ESI type will change quantities of skill-intensity, so long as there is a continuous frontier of marginal households willing to switch employment. Other than market clearing, there is no matching problem in this model: while part-time workers differ in how much part-time labor they provide, a firm only hires the "standard" 25-hour employee. While matching on intensity of employee hours is a possible way around the ACA (for instance, a constrained firm may only hire 45-hour full-time employees and expand production) my model does not allow for the precision picking and choosing by firms of exactly the work frequency of their employees.

Market clearing conditions simply state the sum of all labor by individuals $L_{i,J}$ supplied for a type J firm equals the total labor demanded.

$$\sum_{i\in\{\Omega_J\}} L_{i,J}^D = \sum_{i\in\{\Omega_J\}} L_{i,J}^S \quad J \in \left\{ \left\{ \begin{array}{c} \text{No ESI} \\ \text{Low ESI} \\ \text{High ESI} \end{array} \right\} \times \left\{ \begin{array}{c} \text{Low Skill} \\ \text{High Skill} \end{array} \right\} \times \left\{ \begin{array}{c} \text{Full Time} \\ \text{Part Time} \end{array} \right\} \right\}$$

Where Ω_J denotes the set of individuals/firms working for or organizing a type J firm.

5 Calibration

There are many distributions and parameters to calibrate. In calibrating this model, I minimize the degree of "casual" calibration, avoiding prior studies for fundamental parameters, and opting to directly calibrate them within the context of the model when possible, in the spirit of (Cooley 1997). Fortunately, this model's complexity frequently allows a close mapping of theory to data. Consequently, many parameters will be set literally. For instance, from the perspective of an individual, I will be able to generate the distribution of "non-labor income"[8] by spousal presence by number of children. I assume a structural form of log-normality, allowing household income to vary by spouse and child presence. This is crucial to estimating implicit marginal tax rates effectively, as the federal poverty level is so heavily impacted by the presence of others in the household.

I allow for this sort of correlation between parameters only when necessary: with so many dimensions of heterogeneity, a guiding point will be to assume independence unless it seems to be a first-order error, as in the case of income heterogeneity by household size. For instance, I will be assuming that the preferences of households do not differ by skill level, and average differences in choices by skill or income level are generated by differences in productivity, rather than tastes, adhering as closely as possible to a production-based modeling approach that (Becker 1978) puts forth while being able to match the data.

I will be calibrating the distributions of both observable and unobservable covariates. For those that are directly observable, such as non-labor income, number of children, labor market quantities, and wages, I will set directly, typically via the Current Population Survey (CPS) data. For other

[8]Again, including an inelastically supplying partner's income

covariates and parameters not directly observable, such as preference parameters, I use a two-stage method of simulated moments procedure. This paper uses the Simulated Method of Moments (SMM) methodology detailed in (McFadden 1989) and (McFadden and Rudd 1994) and exemplified for welfare program participation by (Keane and Moffitt 1996) .

My two-stage estimator is used for tractability, at a loss of efficiency. Specifically, I set the simulated wage at the wage target. Given this wage and other parameters, I can separately minimize a quadratic form for both the demand and supply side, where the quadratic form summarizes the weighted squared errors between simulated outcomes and data. At the minimum point, both demand and supply will be near one another conditional on the same wage. The second stage takes all parameters as given save the wage, and clears markets between supply and demand. This eliminates the requirement that I find the market-clearing wage at every parameter evaluation.

This method is still a simulated method of moments: it is lexicographic weighting of wage squared errors over all others, which can is arbitrarily close to a very high weighting of wage squared errors to others. While my weighting matrix is unusual for tractability's sake, it is consistent, though not efficient.

5.1 Definitions and Data

By necessity, this model is an abstraction, and it prominently uses "high-skill" and "low skill" individuals: bringing theory and empirics together requires definitions of these terms. Using the 2013 Current Population Survey (CPS), I examine all heads of household and spouses aged 18-64, representing approximately 193 million individuals, of the 311 million in the CPS, 138 million of whom have a nonzero wage. The "high-skill" and "low-skill" categories are defined by whether or not an individual has a college education. The average wage (conditional on work) for a high-skill individual is approximately $71,000, while the same wage for low-skill individuals is nearly $29,000. This method suggests that there are approximately 136 million low-skill individuals, 91 million of whom work. Similarly, there are about 58 million high-skill individuals, 47 million of whom work.

For family definitions, I count spouses if they live in the same household (separations are not counted). Non-labor income (from the perspective of a spouse or head of household) is income not earned by that individual: the income of a spouse is counted as non-labor income, taken as a given. The only other people included in the construction of a household are children. Federal poverty lines and tax filing exemptions are recalculated using the 2013 definitions using these new households. CPS data only has whether or not an individual is covered by ESI, rather than if they were offered ESI. Consequently, my quantity moments are informed by separate Survey of Income and Program

Participation (SIPP) data, which asks whether or not ESI was offered. I define an employee as having "high" coverage if they receive ESI and the self-reported employer contribution was greater than or equal to \$2000, while they have "low" coverage if they receive ESI with contribution less than \$2000. This is potentially a weakness: workers and firms do not care about earmarked funds, but rather total compensation, and who pays for it.

5.2 Joint Targets

Both the supply-side estimation and demand-side estimation must agree on both wages and quantities: markets must clear. Consequently, for the twelve markets there are twenty-four moments, with twelve wages and twelve quantities to match. While quantities, displayed in Table 2 may be measured straightforwardly, wages may not. Within the data, unobserved heterogeneity is likely to see those workers with a college degree working with ESI to receive higher cash wages than those workers with a college degree not working for ESI. This issue of sample selection causes the conditional averages to be mis-measured to such a degree that the compensating differential for not getting health insurance is negative. To better measure the wage in the face of sample selection, I use a simple model of sample selection to correct for this.

Fundamentally, the problem is that we want the mean effect of ESI on wages conditional on skill. However, there is unobserved heterogeneity in skill even holding constant education. A panel estimation method is under construction. In the meantime I am informed by the yearly \$1488 estimation of (MaCurdy and Rapoport 2003) for low-quality healthcare, and the compensating differential figure of \$5350 in (Kolstad and Kowalski 2014) as wage differentials to target. This yields the median wages in Table 2. I also display the comparable results of my simulated baseline calibration in the next two rows.

Assessing the fit on these moments, the average absolute deviation of simulated to targeted quantities is 1.08 million workers, of a targeted population of 193 million (potential) workers. The worst deviation is 1.76 million workers, for high-skill part-time high-ESI workers, and high-skill part-time no-ESI workers. The phenomenon of high-skilled workers working part-time is the most significant problem of my fit, in particular because the population is so small in the data (8% of all high-skill workers work part-time). On the wage-side, the average absolute deviation is \$2283 dollars, of an average unweighted wage of \$40,000. Most of this comes from missing the same two high-skilled, part-time targets: no ESI at \$6,100 and high ESI at \$8,000. Again, these large errors are partially generated by minimizing the squared distance between levels, rather than the percentage errors.

5.3 Household Parameters

There are several household parameters that may be pinned down by direct estimation. Because household size is so important for measuring a family's income as a percent of FPL, I estimate both skill frequency and presence of a spouse using the bernoulli distribution (by skill), and the number of children as a poisson distribution (by skill by spousal presence). The parameters for skill frequency, frequency of spousal presence crossed with skill, and frequency of children crossed with skill and spousal presence are found in Table 3.

$$S \sim \text{Bernoulli}(\beta_1)$$

$$N_s \sim \text{Bernoulli}(\beta_{2,j}) \quad j \in \{\text{low, high}\}$$

$$N_c \sim \text{Poisson}(\beta_{3,j}) \qquad j \in \{\text{low, high}\} \times \{\text{no spouse, spouse}\}$$

The distribution of "non labor" income by skill and spousal presence is straightforward, imposing log normality on the distribution. The estimated non-labor income moments for each subtype are displayed in Table 4 using transformed non labor income's sufficient statistics.

$$I_j \sim \log \mathcal{N}\left(\beta_{4,j}, \beta_{5,j}\right), \quad j \in \{\text{low,high}\} \times \{\text{no spouse,spouse}\}$$

On the household side, I directly set σ, the elasticity of substitution between types of healthcare, and ϵ, the Frisch elasticity of labor supply. For σ, I choose $\sigma = 9$, indicating a high degree of substitution between the different types of healthcare. The elasticity of labor supply is a more controversial issue, and I choose a conservative value for the Frisch elasticity of labor supply of aggregate hours in a microeconomic model: 0.5. This can be compared to the (Chetty, Guren, Manoli and Weber 2011) average micro-model value for aggregate hours, including both extensive and intensive elasticities.

It is left to set the relative weights of healthcare types in the utility function, and the distributional parameters for ψ_i and κ_i. I do so in a joint maximization problem described above. I report the a's in Table 5, and display the probability density function of the joint distribution in Figure 2.

5.4 Firm Parameters

I set the two firm elasticities of substitution for the part-time and full-time high- and low-skill labor, as well as the elasticity of substitution for low-skill and high-skill effective labor at 2. However, given wages and quantities, these parameters are estimable from aggregate firm behavior using time-series evidence. First order conditions for the firm imply that:

$$\frac{L}{K} = \left(\frac{w_{pt}}{w_{ft}} \frac{\zeta_1}{(1-\zeta_1)}\right)^{\theta}$$

Conditional on measuring L, K, and wages w_{pt} and w_{ft} for several time periods, we can estimate ζ_1 and θ as the values that minimize the sum of squared errors. Specifically, given a generic profit function of the sort:

$$\pi = \left(\delta K^{-\rho} + (1-\delta)L^{-\rho}\right)^{-\frac{1}{\rho}} - wL - rK$$

One can take log transformations of the first order conditions to get:

$$\log\left(\frac{L}{K}\right) = \left(\frac{\delta}{1-\delta}\right) + \sigma \log\left(\frac{w}{r}\right)$$

And recover the coefficients for δ and σ from a linear regression, assuming exogenous price levels. While this method is currently under construction, I use the simple assumptions above, of elasticities of substitution equal to 2. The identification of this method are detailed in Appendix 2.

Households choose health insurance given total compensation with a variable $P(H_{E,i})$. I set to offer a discount: a worker may lose \$1 in compensation while accruing \$1.05 worth of ESI. I choose a 1% discount for part-time low-skill workers, a 5% discount for full-time low-skill workers and part-time high skill workers, and a 10% discount for full-time high-skill workers. Conditional on these parameters, using the labor demand procedure outlined above, I estimate an equally-weighted quadratic form of γ, ζ_0, ζ_1, and ζ_2 along with a kernel density estimation of the joint distribution between A_i and ξ_i. The γ and ζ's are displayed in Table 5. The joint distribution of firm parameters A_i and ξ_i is displayed in Figure 3. Secondary firm moments, focussing on the distribution of firm size the distribution of A and the proportion of employers offering ESI, along with fits, can be found in Table 6. The first set of moments primarily informs the distribution of A, which is the primary determinant of size. The second set of moments informs both the distribution of A and the distribution of ξ, as both parameters play a role in determining whether or not a firm offers ESI.

5.5 Identification

The joint distributions of productivity and healthcare costs for firms, and disutility of labor and utility from healthcare for households, are indirectly inferred via simulation. The process, the fundamental idea of which is outlined in detail in (Gouriéroux and Monfort 1997), is simple. I have a set of unknown structural parameters Θ, that I need to solve my model. My theoretical model, $G(\Theta)$ can produce a set of measurements $y_{Simul} = G(\Theta)$ that can be compared directly to the data, y_{Data}. An estimator for Θ, $\hat{\Theta}$, may be produced by minimizing the sum of squared deviations between y_{Data} and y_{Simul}.

Fortunately, it is relatively easy to understand how my joint distributions are identified from data. As discussed in the results section, Figure 4 is the policy function for firms, an endogenous set of policies of the calibrated model. Figure 5 is a modified version of Figure 4, which visually displays the policy function of firms at equilibrium. In Figure 5, I superimpose two areas, demarcated by a blue and black line. The blue triangular line denotes all firms smaller than 25 employees that offer some form of health insurance to their employees, lining up with the censored data moment in Table 6. The black triangular hull denotes all those firms whose size is between 10 and 25 employees, respectively. Both moments, combined with the simulation, give information about how many firms should be represented in a given area. The kernel weights then attempt to fit these areas as best they can.

Of course, there are many ways to get the right amount of area in a given box: one can have all focused one type of person, or uniformly diffuse within all possible types. Key to my identification is the bandwidth of the kernels I use to generate the population: I force smoothness on the distribution of preferences, rather than assuming the absence of a continuous margin of household types. I view the choice of a smooth distribution of preferences as highly desirable in this case: for instance, there is little reason to think that there is a reasonably large quantum for productivities or healthcare costs for firms.

6 Results

6.1 Mechanisms

With twelve linked markets and a highly heterogeneous population of households and firms, many comparative statics are theoretically ambiguous. For example, while the large implicit marginal tax rates from the exchange subsidies reduce work via both income and substitution effects, the curtailment of the medicaid notch removes implicit marginal tax rates of above 100% for some subsection of the population. As I have noted, the distribution of preferences will determine the empirical response. Additionally, general equilibrium effects with changing wages make signing effects difficult. For example, the exchange subsidy's implicit marginal tax is likely to push both low-income and some high-income workers to work part-time. On the demand side, this tends to raise the marginal productivity of full-time workers of both types, increasing their wage.

On the supply side, the medicaid and exchange subsidies raise implicit marginal tax rates even while increasing income, yielding all types of workers to work less, yielding less work on the intensive margin in both groups, and leading more full-time workers to work part-time. This effect is more pronounced on low-skill individuals who are the primary targets of the ACA's income thresholds: high skill workers largely fall outside the ACA's main implicit marginal tax rate increases. The mandate increases individual desire for healthcare: this causes all types of insurance, including medicaid, non group insurance, and ESI to become more desirable. Theoretically, the mandate's negative income effect could cause individuals to work more. However, low-skill individuals, (who are impacted more by the mandate because high-skill individuals typically already have insurance) are more likely to find insurance from Medicaid or NGI with exchange subsidies than through working. High-skill individuals are slightly more likely to work. These effects of the mandate and exchange subsidies have an interesting sectoral push-pull effect: low-skill workers are more likely to desire no- and low-healthcare employers, so they gain access to the subsidy. High-skill workers, unable to access exchange subsidies, are more likely to desire high-healthcare employers to avoid the mandate's penalty.

Unfortunately, the comparative statics of labor demand are far more ambiguous. Specifically, the cross-price elasticity of demand for a given input is dependent on both the diminishing returns to scale and the elasticity of substitution. As the elasticity of substitution grows higher, an increase in the price of a substitutable input A increases total demand for input B. However, as diminishing returns become less of a factor (firms move closer to constant returns to scale), an increase in the price of a substitutable input A can decrease the total demand for input B. To understand the intuition when dealing with a constant elasticity production function, note that while the elasticity of substitution completely controls the relative use of two factors for all non constrained firms, an increase in relative use can be more than overcome by a diminishing of scale. This is shown in Appendix 3. This fact allows for cross-price elasticities to be ambiguous, even as own price elasticities are always negative. Within my calibration, I typically have moderately positive cross-price elasticities, on the order of 0.28. For my paper, this suggests that a change in compensating differentials for one group of workers has two offsetting effects. For example, an increase in the wage for low-skilled workers will cause firms to shrink and use more high-skilled labor: the latter effect is larger, leading to a rise in high-skilled workers.

On the demand side, large firms with low health insurance at the margin are pushed to offer high health insurance to avoid the tax. While for large firms high health insurance becomes more prevalent, small firms that were previously offering high levels of insurance are more likely to drop their high ESI in order to give employees access to exchange subsidies without accruing the penalty for doing so. This “trade” of space, as high insurance expands for larger firms and low insurance expands for small firms, increases inefficiency: the small firms that were offering high health insurance and are now offering low health insurance necessarily had lower costs of doing so than the large firms that were offering low health insurance and are now offering high health insurance. Additionally, a large number of firms near but above the 49-employee cutoff will contract their employment in order to cheaply avoid the penalty while not changing behavior. These firms will typically be far from the margin of offering health insurance, but at the size margin.

6.2 Labor and Healthcare Counterfactuals

The simulated labor market and health insurance impacts of the ACA are substantial. Table 7 offers the central results of this paper, displaying the main labor market, production, and healthcare impacts of the fully-implemented ACA as compared to the simulated baseline economy in Column 3. Table 7 additionally displays two ACA counterfactuals that offer an insight into the importance of modeling the ACA jointly. The fourth column models what would happen had the ACA only expanded medicaid to 133% of the poverty line, while not mandating coverage, offering subsidies, or taxing noncomplying businesses. The fifth column looks at what would happen if the ACA only included a mandate and penalty to purchase HI, without Medicaid expansion or NGI subsidies.

The full implementation of the ACA reduces aggregate labor hours by 6.83%, or 9.22 million full time equivalent workers. This net decrease in labor hours reflects a decline of both high-skilled labor and low-skilled labor, for the reasons described above: cross wage elasticities are negative in this decreasing returns model, in spite of the two being substitutes. This leads high-skill workers to decrease aggregate labor market hours by approximately 2.7 million employees, even as low-skilled workers decrease by an equivalent of 6.7 million FTEs. The decrease in labor market hours and a less-efficient composition of workers leads production to decline by 5.7% and overall labor market income to decline by 7.09%.

Because of the large advantage to the average firm-worker pair by switching to non-ESI coverage, I find the total number of people (including children and spouses) receiving a “minimum acceptable” level of ESI declines by 6.3 million workers, or 2.72%. While the number of individuals on ESI declines, the ACA succeeds in increasing the total number of individuals with acceptable health insurance. The number of individuals with minimum acceptable coverage increases by 17.95%. This is unsurprising: while the ACA discourages ESI, it does so by making Medicaid and NGI more desirable. The pre-ACA and post-ACA household policy functions are depicted in Figure 6.

I measure the equivalent variation of an "unfunded" ACA in the bottom three rows of Table 7. The EV is unfunded in that I do not impose an increase in taxes to fund exchange subsidies and medicaid expansion. This may be considered a measure of the ACA funded entirely by reducing purely wasteful spending. On average, this free spending benefits each household by the utility equivalent of $25. This meager average hides great dispersion, with some households suffering and others gaining dramatically. Typically, households facing a decline are poor, and whose take-up of exchange subsidies is low. The average of households that are made worse off by the ACA is a decline of $1549 worth of utility, while the average gain of households made better off is $1204. The conditional differences in EV is instructive on the importance of distributional concerns of the ACA's impacts: it will benefit some types of households and harm others.

Perhaps most interestingly from a public finance point of view, the labor market impacts of the ACA would have been far more severe had the ACA only changed Medicaid by increasing the baseline FPL and including all childless adults. The presence of a Medicaid notch in a highly-represented portion of the income distribution would have yielded a decrease in total hours by 7.47%, or more than 10 million FTEs. The ACA avoids this dramatic drop by offering ACA substantial exchange subsidies after a household leaves Medicaid. This counterfactual may lend insight into why the exchange subsidies may have been politically necessary: medicaid expansion alone would have dramatic labor market consequences.

The mandate requiring individuals obtain minimum acceptable coverage has garnered a great deal of media and legal attention. However, its expected impacts on both the labor market and the health insurance market are minimal. The individual mandate, implemented alone, would decrease total labor hours by 81 thousand FTEs, and increase individuals receiving high health care by more than 15 million workers, less than a third of the fully-implemented ACA. While its outcome of increasing population coverage is lower, it also does so at a lower cost. This is because the penalty is modest for those generally far from the margin, such as low-income families, and higher for families nearer to the margin, such as wealthier families, due to healthcare being a normal good.

6.3 Wages and Quantities

Movements in wages and quantities are driven by the effect of the subsidy, as part-time labor becomes more attractive, increasing supply more than demand is decreased, and full-time labor becomes less attractive, decreasing supply more than demand is increased.

While Table 2 displays the simple impacts of the Affordable Care Act on full compensation paid out by a firm, it is difficult to directly draw out the causes of wage changes. This theoretical ambiguity is clear: a shift in either supply or demand for any input changes the equilibrium wage for all other inputs, and supply and demand for various jobs and frequencies of work are shifting in many directions.

Table 8 displays the impact of the full ACA on wages and quantities. Wages show a clear pattern: all part-time wages decline, while all full-time wages increase. Because full-time labor experiences a larger tax, and part-time labor is subsidized, workers shift to part-time labor. This increases the quantity of part-time laborers, driving down the wage. The less efficient makeup of firms as fewer full-time labors work also reduces all wages. Similarly, marginal full-time laborers must be paid more in order not to move to part-time work. While firms are less efficient, these workers are more in demand in production.

Quantites tell a similar story. All part-time quantities rise, save for high-skill high-healthcare labor. All full-time quantities fall, save for high-skill, no-healthcare labor. The reasons are clear: the ACA's impacts focus on low-skill laborers, moving them into both no-healthcare and part-time positions. The effect is strong in the union of those two categories, as part-time no-healthcare low-skill laborers increase by 41%. Consequently, more high-skill labor is required in the no-healthcare sector. As high-skill workers typically want to work full time, the effect is strongest here, increasing full-time laborers in this sector.

The decline of part-time, high-skill, high-healthcare laborers is similar: these are some of the few high-skill workers with the capacity and willingness to move into the subsidized sector, as they have not yet been excluded by high incomes: they are able to receive moderate subsidies if they switch sectors, depending on their family income and household size.

6.4 Market Inefficiencies

Note: This subsection is under construction and is incomplete!

There are several abstract sources of inefficiency in the ACA. First, implicit marginal tax rates produce wedges between the marginal rate of substitution between labor and leisure. Second, those same tax rates, by being levied on different sectors at different rates, causes a factor-substitution effect on taxed firms, diverting firm expansion paths from the optimum and increasing costs, holding production constant. Third, sector-specific exemptions on both firm and worker sides cause a sector-substitution effect, as large firms with higher costs offer healthcare and smaller firms with lower costs cease offering healthcare Finally, the size-based exemption causes a loss in production due to constrained firms decreasing labor. Getting a grasp on the magnitudes of all of these inefficiencies is a first-order problem in researching ACA.

There are several ways to measure loss due to the labor-leisure costs of the ACA. Because of my setup's unique ability to measure utility and marginal utility of consumption, I take my complete post-ACA cash equivalent of utiles and the same measure for alternative post-ACA, shutting down the labor-leisure costs alone. This entails forcing all households to offer the same aggregate labor, yielding a completely inelastic aggregate supply curve, while still allowing households to move their labor across sectors and firms to adjust their sizes and offerings. The difference between the full ACA and this alternative constant labor supply ACA is attributed to the labor-leisure costs of the ACA.

The pure factor substitution costs of the ACA is generated in a similar way to the twinned (Gallen and Mulligan 2013) and (Mulligan and Gallen 2013) papers, with two changes. I force firms to hold their production ratios constant by changing production to a factor ratio equivalent Leontief. These firms will experience the same size-based taxes, sectoral substitutions from workers, and labor-leisure possibilities as the full-ACA. Their factor ratios will remain the same, with wages and total idiosyncratic firm production moving to bring wages into alignment.

To measure the sectoral substitution effect, I shut down the ability of both firms and households to switch their pre-ACA healthcare, and time frequency sectors. To be specific, I do not allow any household members or firms to leave a sector, though I do allow them to act on their other margins of choice. For instance, I allow any household members to work zero hours, though for all other intents and purposes they remain in their current sector. For instance, a “full-time, high-healthcare” individual may decide to work zero hours, and firms may reduce production to zero. But they accrue the penalties of being in specific sectors, such as inability to access healthcare subsidies.

To measure the size-based costs, I remove the size-based provision and tax all firms a representative amount. For instance, if previously a tax on 3% of pre-ACA firms yielded one hundred million dollars, I tax all firms an amount necessary to generate a revenue of one hundred million dollars. This measures the revenue-constant cost of having size-based sectoral exemptions.

Results are forthcoming.

Note: The above subsection is under construction and is incomplete!

6.5 Firm Size Distribution

Note: The subsection below will be revised and incorporated into the “Market Inefficiencies” section above.

The ACA’s size-based, healthcare-based tax distorts the optimal size distribution of firms. The policy function for firms by log TFP and log healthcare implementation costs before and after the ACA are depicted in Figure 4. The top panel shows that before the ACA, large (productive) firms and firms with low healthcare costs offered “high” healthcare. Moderately-sized firms and firms with mild healthcare costs offered “low” healthcare. Small firms and firms with high healthcare costs offered no healthcare. The graph depicts five areas that are of interest when considering ACA implementation. Three types of firms escape the theoretical penalty: light blue firms, because they offer “high” healthcare, yellow healthcare firms that offer “low” healthcare but are too small to tax, and red firms that offer no healthcare but are too small to tax. The two types of firms that are targeted by the tax (holding behavior constant) are in green and blue: large firms that are offering low and no healthcare, respectively.

The ACA has several implications, depicted in the bottom panel in Figure 4. First, small firms that were at the margin of offering high healthcare and low healthcare are pushed into offering low healthcare: their employees are able to access exchange subsidies while they incur no tax penalty. These firms can be compared to the large firms at the margin of offering high- and low-healthcare (green firms), which are pushed to offer ESI. Because of the ACA’s size-based, healthcare-based tax, large firms with a higher cost of offering ESI may offer it to avoid the tax, while small firms with a lower cost of offering ESI will not. Firms near 50 employees (in the top panel, differentiated by the yellow-green and red-blue transition) now find it worth their while to constrain employment at 49 employees. These firms do not offer health insurance, but avoid the tax by shrinking employment and changing composition of employees. With regards to firm behavior, the ACA's primary distortions come from these five margins: 1) the policy function switch for small high-healthcare offering firms to offering low healthcare, 2) small low-healthcare firms move to offering no healthcare 3) large low healthcare firms to high healthcare, 4) large no-healthcare firms to offer low-healthcare, and 5) shrinking of moderately-sized firms near the cutoff but not at the margin of offering health insurance or not.

Examining the distortions to the firm size distribution, Figure 7 shows the conditional marginal distribution of firm size before and after the ACA. The ACA causes a large number of firms that were at the margin between offering health insurance and not to shrink, a typical result of a size-based tax. However, not all firms do so: some were already offering high health insurance or found themselves at the margin of offering high health insurance or low health insurance. Consequently, the normally stark predictions of a size-based are tempered by the healthcare-based trigger. Nevertheless, the predictions remain stark, reducing the number of firms between 50 and 70 employees by nearly 60%.

The efficiency cost of the swapping of high and low offering insurance between large and small firms is substantial: the average administrative cost of health insurance per head for the small firms previously offering high health insurance, now offering low health insurance is about half that of large firms previously offering low health insurance, now offering high.

I measure the efficiency cost of the size-based tax by calculating the cost of matching post-ACA production using pre-ACA firm structure (i.e. unconstrained firms paying no taxes). This calculates the loss in production stemming solely from inefficient use of resources. Approximately 6 billion of the ACA's 170 billion loss in production comes from the deadweight loss of inefficient firm structure due to size-based, healthcare-based taxation.

6.6 Equivalent Variation

The Affordable Care Act has large redistributive effects. Because of the large degree of heterogeneity my model allows for, as well as general equilibrium effects, it is worthwhile to analyze equivalent variation along several margins. First and foremost, because the ACA's income-contingent subsidies impact people differentially by their location on the FPL. It is reasonable to expect that the

ACA impacts households lower on the FPL by a greater amount than those higher on the FPL. However, the inaccessibility of Medicaid and exchange and cost-sharing subsidies can dramatically impact poor households utility levels, increasingly so as they pass the filing threshold and enter a region in which penalties are levied more heavily by income level.

The distribution of EV is depicted in Figure 8. A linear fit of EV by skill type and household FPL breaks up Figure 8's differences between high and low-skill labor, as depicted in Figure 9. High skill laborers tend to benefit more from the ACA as they grow richer. As we have seen the wage of full-time laborers, who make up a large portion of the highly skilled, rises in all sectors. Additionally, these individuals often already have healthcare, decreasing the negative impact of the mandate. The benefit to wealthier households occurs because they also work more hours, causing the increased wage to have a larger EV for them than households that work fewer hours.

On the other hand, low-skill households largely have a mildly positive benefit, without regard to FPL. The poorest households are generally only worse off, as they were already receiving Medicaid, or find it inaccessible, lowering their utility in conjunction with the Mandate. Wealthier households find the exchange subsidies more accessible, and therefore are more likely to benefit from the ACA, though they also get lower rewards. These offsetting effects cause there to be very little overall relationship between income and EV for low-skill households, though the region of increasing penalties does indicate a high localized level of conditionality.

A lowess fit of EV by preference type by low and high skill type can be found in Figure 10 and 11, respectively. The two tell similar stories, as can be seen by comparing Figure 10 to the policy function in Figure 6. Those benefiting most from the ACA are households that strongly value health insurance, particularly at already-working portions of the distaste for labor distribution.

7 Conclusion

This model offers a close examination of several interacting ACA provisions. By including many dimensions of heterogeneity for both firms and households in a maximizing framework, it is able to careful examine many of the Affordable Care Act's more arcane provisions, such as the individual mandate's penalty, with its four measurements determining a household's tax burden for not purchasing or obtaining adequate insurance. Because the poor are likely to be exempted or have a low penalty even without Medicaid expansions or non-group subsidies, and the wealthy are likely to already have health insurance or be close to the margin of buying it, the individual mandate's effects are far smaller than one might initially expect, the frequently assumed flat and uncapped 2% in 2015.

Similarly for the size-based firm penalties, as most large firms already offer health insurance, and the bulk of firms that do not offer health insurance are already below 50 employees. Taking a representative firm, without close attention to the joint distribution of firm size and healthcare offering, will find a dramatically higher penalty for that firm: even before the ACA, slightly less than two-thirds of employees not offered ESI are in firms below 50 employees. With only 8.5% of employees exposed to the possibility of being taxed before maximizing behavior, the penalty is much smaller than might be thought in a causal calibration.

Instead, I find the ACA's largest impacts are generated by the severe and kinked implicit marginal tax hikes by income level. Nevertheless, many individuals, especially those working individuals who are motivated primarily by a desire for health insurance, are made better off by the ACA. However a sizable number of individuals, such as those who do not desire health insurance or have difficulty accessing NGI subsidies or Medicaid, and make enough money to face a sizable penalty, are made worse off.

In this long-run model, the full ACA is successful in its goal of increasing health insurance among the whole population by nearly 18%, nearly double the CBO's long-run estimate. However, this extra coverage comes at a cost: the full ACA has a severe impact on production, reducing GDP by more than $600 billion and reducing total labor hours by nine million workers.

References

Axtell, Robert L., "Zipf Distribution of U.S. Firm Sizes," *Science*, 2001, *293*, 1818–1820.

Baicker, Katherine, Amy Finkelstein, Jae Song, and Sarah Taubman, "The Impact of Medicaid on Labor Force Activity and Program Participation: Evidence from the Oregon Health Insurance Experiment," *NBER Working Paper*, 2014, (19547), 1–13.

Becker, Gary S., *The Economic Approach to Human Behavior*, University of Chicago Press, 1978.

Bradley, Cathy J., David Neumark, and Scott Barkowski, "Does employer-provided health insurance constrain labor supply adjustments to health shocks? New evidence on women diagnosed with breast cancer," *Journal of Health Economics*, 2013, *32* (5), 833 – 849.

Burkhauser, Richard V. and Kosali Simon, "Measuring the Impact of Health Insurance on Levels and Trends in Inequality and How Health Reform Could Affect Them," 2010, (19547).

Cabral, Luís M. B. and José Mata, "On the evolution of firm size distribution: facts and theory," *American Economic Review*, 2003, *93* (4), 1075–1090.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of the Evidence on the Intensive and Extensive Margins," *American Economic Review: Papers and Proceedings*, 2011, *101* (3), 471–475.

Cooley, Thomas F., "Calibrated Models," *Oxford Review of Economic Policy*, 1997, *12* (3).

Currie, Janet and Jonathan Gruber, "Saving Babies: The Efficacy and Cost of Recent Changes in the Medicaid Eligibility of Pregnant Women," *Journal of Political Economy*, 1996, *104* (6), 1263–1296.

Dague, Laura, Thomas DeLeire, and Lindsey Leininger, "The Effect of Public Insurance Coverage for Childless Adults on Labor Supply," *Working Paper*, 2014, pp. 1–32.

Davidoff, Amy., Anna S. Sommers, Jennifer Lesko, and Alshadye Yemane, "Medicaid and State-Funded Coverage for Adults: Estimates of Eligibility and Enrollment," *Kaiser Commission on Medicaid and the Uninsured*, 2004.

Fairlie, Robert W., Kanika Kapur, and Susan Gates, "Is employer-based health insurance a barrier to entrepreneurship?," *American Economic Review*, 2011, *30* (1), 146–162.

Gallen, Trevor S. and Casey B. Mulligan, "Wedges, Labor Market Behavior, and Health Insurance Coverage under the Affordable Care Act," 2013, (19770), 1–60.

Garicano, Luis, Claire LeLarge, and John Van Reenen, "Firm Size Distortions and the Productivity Distribution: Evidence from France," *NBER Working Paper*, 2013, (18841), 1–59.

Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigo, "Public Health Insurance, Labor Supply, and Employment Lock," *NBER Working Paper*, 2013, pp. 1–67.

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

Gouriéroux, Christian and Alain Monfort, *Simulation-Based Econometric Methods*, Oxford University Press, 1997.

Gruber, Jonathan, *Medicaid*, Vol. Means Tested Transfer Programs, University of Chicago Press,

___ **and Brigitte C. Madrian**, "Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature," *NBER Working Paper*, 2002, (8817), 1–58.

Harberger, Arnold C., "Three Basic Postulates for Applied Welfare Economics: An Interpretive Essay," *Journal of Economic Literature*, 1971, *9* (3).

Hausman, Jerry A., *Taxes and Labor Supply*, Vol. 1, Handbook of Public Economics, North Holland,

Ijiri, Yuji and Herbert A. Simon, "Business firm growth and size," *The American Economic Review*, 1964, *54* (2), 77–89.

Inada, Ken-Ichi, "On a Two-Sector Model of Economic Growth: Comments and a Generalization," *The Review of Economic Studies*, 1963, *30* (2).

Keane, Michael and Robert Moffitt, "A structural model of multiple welfare program participation and labor supply," 1996, (1080-96).

Kolstad, Jonathan T. and Amanda Kowalski, "Mandate-Based Health Reform and the Labor Market: Evidence from the Massachusetts Reform," *NBER Working Paper*, 2014, (17933), 1–9.

Komisar, Harriet, "The Effects of Rising Health Care Costs on Middle-Class Economic Security," *AARP Public Policy Institute*, 2013.

Lucas, Robert E. Jr., "On the Size Distribution of Business Firms," *Bell Journal of Economics*, 1978, *9* (2), 508–523.

MaCurdy, Thomas and Dana Rapoport, "Understanding Health Insurance Coverage of Low-Skilled Workers in the US and in California: Trading Wages for Insurance," *Working Paper*, 2003, (17933), 1–102.

McFadden, Daniel, "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration," *Econometrica*, 1989, *57* (5), 995–1026.

——— **and Paul A. Rudd**, "Estimation by Simulation," *Review of Economics and Statistics*, 1994, *57* (4), 591–608.

Meyer, Bruce D. and Dan T. Rosenbaum, "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers," *Quarterly Journal of Economics*, 2001, *116* (3), 1063–1114.

Mulligan, Casey B., "Average Marginal Labor Income Tax Rates under the Affordable Care Act," *NBER Working Paper*, 2013, (19365), 1–65.

———, "The Economics of Work Schedules under the New Hours and Employment Taxes," *NBER Working Paper*, 2014, (19936), 1–36.

——— **and Trevor S. Gallen**, "Wedges, Wages, and Productivity under the Affordable Care Act," *NBER Working Paper*, 2013, (19771), 1–62.

Rosen, Sherwin, *The Supply of Work Schedules and Unemployment*, Vol. No. 28, Work Time and Employment: Special Report, National Commission for Manpower Policy,

Simon, Herbert A. and Charles P. Bonini, "The Size Distribution of Business Firms," *The American Economic Review*, 1958, *48* (4), 607–617.

Sommers, Benjamin D. and Arnold M. Epstein, "Medicaid Expansion: The Soft Underbelly of Health Care Reform?," *New England Journal of Medicine*, 2010, *363* (22), 2085–2087.

Summers, Lawrence, "Some Simple Economics of Mandated Benefits," *American Economic Review: Papers and Proceedings*, 1989, *79* (2), 177–183.

Yelowitz, Aaron S., "The Medicaid Notch, Labor Supply, and Welfare Participation: Evidence from Eligibility Expansions," *Quarterly Journal of Economics*, 1995, *110* (4), 909–939.

8 Tables

Table 1: Sources of U.S. Healthcare Coverage by Type in 2012 (millions)

Source	All Individuals	18-64	18-64, <300% FPL	18-64, ≥300% FPL	18-64, working
Employer	139.26	103.25	28.75	74.50	87.24
Joint Employer, Public Nonmilitary	16.72	3.70	2.05	1.65	1.95
Medicaid/SCHIP/Public	35.56	14.02	12.85	1.17	5.46
Medicare	15.06	3.09	2.55	0.54	0.26
Military	3.87	2.97	1.58	1.40	1.66
Other Mixed	34.19	12.79	6.04	6.76	7.11
Privately Purchased	15.55	12.43	7.01	5.41	7.75
Uninsured	48.61	40.96	31.80	9.16	25.17
Total	308.83	193.21	92.63	100.58	136.61

Table 1: This table depicts the sources of healthcare coverage, if any for the U.S. population in 2012. The first column displays counts of coverage for all individuals by source. In a given cell, the top number depicts the number of people who solely get coverage from that source. The bottom number in parentheses depicts any amount of coverage from that source. The second column replicates the first column, but limits to only individuals aged 18-64. The third and fourth column replicates the second, but breaks down coverage by household income level. The fifth replicates the second, but only for people with positive labor income in the year before.

Table 2: Targets and Simulated Results for Wages and Quantities

HI	Skill	Freq.	Quantity Target (millions)	Simulated Quantity (millions)	Wage Target (thousands)	Simulated Wage (thousands)
No	Lo	PT	12.62	14.06	21.9	22.4
No	Lo	FT	37.53	37.94	43.1	42.7
No	Hi	PT	3.19	4.95	32.6	38.7
No	Hi	FT	12.22	10.53	78.3	78.3
Lo	Lo	PT	0.47	0.89	16.6	18.9
Lo	Lo	FT	6.56	4.88	40.0	39.2
Lo	Hi	PT	0.28	1.05	29.2	34.8
Lo	Hi	FT	3.65	4.73	76.0	74.4
No	Lo	PT	0.98	1.69	15.2	16.7
No	Lo	FT	33.12	33.78	37.0	37.0
No	Hi	PT	0.73	2.49	25.4	33.4
No	Hi	FT	26.9	27.52	72.4	73.0

Table 2: This table depicts both the wage and quantity targets, generated from the CPS, as well as their simulated counterparts from the baseline calibration. Wages are measured in thousands of dollars, while quantities are measured in millions. Wages are expressed in annual equivalents for working either 25 or 40 hours per week 50 weeks a year for part-time and full-time workers, respectively.

Table 3: Skill, Spouses, and Child Parameters

Parameter	Type	Value
$\beta_{1,j}$	Skill frequency	0.340
$\beta_{2,low}$	Spousal Presence	0.550
$\beta_{2,high}$	Spousal Presence	0.646
$\beta_{3,low,no}$	Child Frequency	0.307
$\beta_{3,low,yes}$	Child Frequency	0.509
$\beta_{3,high,no}$	Child Frequency	0.169
$\beta_{3,high,yes}$	Child Frequency	0.306

Table 3: This table depicts the estimated Bernoulii and Poisson distributions for skill, spousal presence and number of children. The first panel describes 34% of the 18-64 workforce as skilled. The second panel notes that 55% of the low-skilled group have a spouse, while 65% of the high-skilled group have a spouse. Similarly, of those who are high-skilled with a spouse, the average number of children in the household is 0.31, while the average number without a spouse is 0.17.

Table 4: Estimated Nonlabor Income Parameters

Skill	Spouse	Mean Parameter $\beta_{4,j}$	Standard Deviation $\beta_{5,j}$	Median Nonlabor Income (thousands)
Low Skill	No Spouse	8.197	2.557	3.6
Low Skill	Spouse	10.300	1.462	30.0
High Skill	No Spouse	7.470	2.863	1.8
High Skill	Spouse	10.555	1.584	38.3

Table 4: This table depicts the estimated log-normal non-labor income parameters by spousal presence by skill level. Recall that "non-labor income" includes a spouse and working child's income.

Table 5: Jointly Estimated Parameters

Description	Symbol	Value
Household CES weight on ESI	a_1	0.023
Household CES weight on NGI	a_2	0.015
Firm Scale of Control	γ	0.669
Firm CES weight on Part-Time Low-Skilled Labor	ζ_1	0.497
Firm CES weight on Part-Time High-Skilled Labor	ζ_2	0.234
Firm CES weight on Effective Low-Skilled Labor	ζ_2	0.206

Table 5: This table displays the results of jointly-estimated non-preference or firm-type parameters.

Table 6: Firm Targets and Simulated Data

Description	Target (thousands)	Simulated Moment (thousands)
Number of firms with 0-4 employees	3493	3144
Number of firms with 5-9 employees	996	847
Number of firms with 10-14 employees	398	332
Number of firms with 15-19 employees	205	195
Number of firms with 20-24 employees	124	136
Number of firms with 25-29 employees	83	89
Number of firms with 30-34 employees	60	65
Number of firms with 35-39 employees	45	41
Number of firms with 40-44 employees	35	25
Number of firms with 45-49 employees	27	19
Number of firms with 49-74 employees	79	61
Number of firms with 75-99 employees	38	39
Number of firms with 100-149 employees	36	68
Number of firms with 150-199 employees	61	31
Number of firms with 200+ employees	47	13
Description	Target (percent)	Simulated Moment (percent)
Percent of 0-24 employee firms offering ESI	0.45	0.55
Percent of 25-99 employee firms offering ESI	0.75	0.60
Percent of 100-499 employee firms offering ESI	0.84	0.72
Percent of 500-1000 employee firms offering ESI	0.88	0.85
Percent of 500+ employee firms offering ESI	0.90	0.89

Table 6: This table displays non-wage non-quantity firm size targets and simulated data.

Table 7: Counterfactual ACA Implementations

Outcome	Baseline	Full ACA ACA	Only Medicaid Expansion	Only Mandate
Total Labor Hours (millions FTE)	135.10	-6.83%	-7.47%	0.06%
Total Hours: High (millions FTE)	48.85	-5.54%	-3.66%	0.05%
Total Hours: Low (millions FTE)	88.18	-7.59%	-9.65%	0.06%
Number with Zero Hours (millions)	63.19	-0.87%	5.65%	0.37%
Aggregate Production (billions)	10566.80	-5.70%	-5.15%	0.06%
Labor Income (billions)	6963.87	-7.09%	-5.06%	0.09%
Labor Income: High (billions)	3487.69	-5.69%	-3.12%	0.07%
Labor Income: Low (billions)	3476.18	-8.49%	-7.02%	0.11%
Total People (including children) Receiving High ESI	232.71	-2.72%	-1.45%	1.26%
Total Receiving High Healthcare	280.63	17.95%	9.65%	5.44%
Equivalent Variation ($)	NaN	25	506	-0.01
Equivalent Variation | EV>0	NaN	1549	1954	NaN
Equivalent Variation | EV<0	NaN	-1204	-965	-264

Table 7: This table depicts the baseline simulated scenario in the second column, along with three ACA counterfactuals: full ACA, only Medicaid, and only mandate. In the second column, values are given in the relevant level: counts are measured in millions, and all dollar values in billions, save for equivalent variation (EV) which measures household EV in dollars. The second column depicts these levels for the baseline economy. The third column displays percent deviations from the baseline level: for instance, the full implementation ACA predicts a 1.87% decline in labor hours, or a decline of 2.62 million FTE workers. For counterfactuals, EV is measured in average dollars per household. It is important to note that EV is unpaid for, purely measuring benefits and distortions from ACA subsidies and taxes, but not including any additional taxation to pay for the ACA deficit.

Table 8: Counterfactual Wages and Quantities

HI	Skill	Pt/Ft	Wage, no ACA	Wage, ACA	Quantity, no ACA	Quantity, ACA
No	Lo	PT	21882	17752	14.11	19.91
No	Lo	FT	43099	44385	37.81	32.23
No	Hi	PT	32637	30849	4.94	5.16
No	Hi	FT	78279	78964	10.43	11.46
Lo	Lo	PT	16619	9858	0.89	3.13
Lo	Lo	FT	40018	42722	4.92	4.24
Lo	Hi	PT	29239	23162	1.05	1.98
Lo	Hi	FT	75994	77360	4.82	2.37
Hi	Lo	PT	15156	13730	1.75	1.96
Hi	Lo	FT	37050	39435	33.79	28.31
Hi	Hi	PT	25448	25189	2.50	2.46
Hi	Hi	FT	72397	73562	27.57	25.63

Table 8: This table depicts the no-ACA and full-ACA wage and quantity counterfactuals. It lends insight into the economic causes of the ACA's sectoral impacts. Specifically, all ACA full-time wages increase from the baseline, while all ACA part-time wages fall, because of the strong shifts in labor supply toward part-time labor. Similarly, nearly all part-time quantities rise, while nearly all full-time quantities fall, reflecting the same phenomenon.

Table 9: Representative Wages

Description	Symbol	Value
Low-Skill, Part-Time	21009	16534
Low-Skill, Full-Time	40768	42702
High-Skill, Part-Time	30377	28082
High-Skill, Full-Time	75070	76234

Table 9: This table depicts the quantity-weighted average of full compensation (with the utility value of ESI converted into dollars) across skill and time frequency, collapsing healthcare status. It shows that the part-time low-killed and high-skilled wages decrease, as might be expected: the subsidies on part-time work (via a reduction of income) lead to an increase in supply, reducing wages. As a compensating differential for not receiving available government subsidies, full-time wages of both types increase.

9 Figures

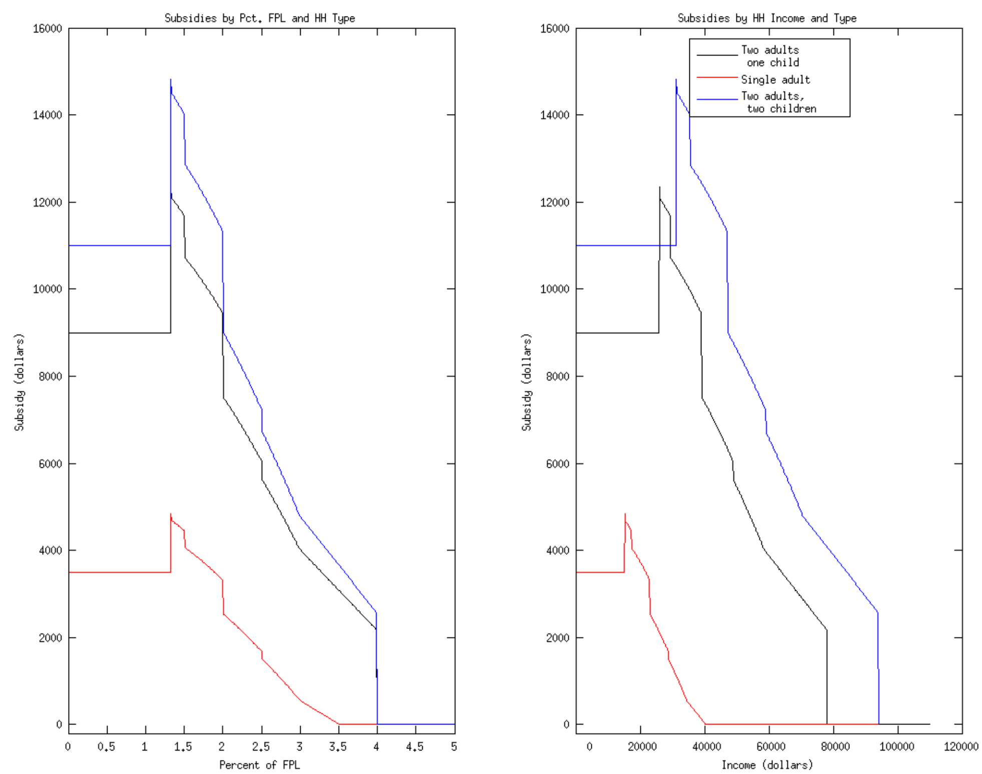

Figure 1: The figure displays three examples of post-ACA Medicaid and NGI subsidies by household type. There are three types of households: two adults, one child, two adults, two children, or a single adult. The three types have different federal poverty lines, premiums, and Medicaid offerings. The flat lines up to 138% of the poverty line is generated by \$3500/adult and \$2000/child. After 138% of the poverty line, I calculate the sum of NGI subsidies for various households by FPL level in the left panel, and income in the right panel. The sharp drops in exchange subsidies represents sometimes large losses, up to nearly \$4000 for earning an extra dollar at the wrong point. Depending on their utility valuation of Medicaid spending, it easy to see how individuals would group between 138% of the poverty line and 200% of the poverty line, though this means different levels of income for different family units. While the ACA's theoretical penalties run out by 400% of the FPL, it is possible for them to be reduced before that, as in the case of the single individual with low premiums compared to reasonable income levels.

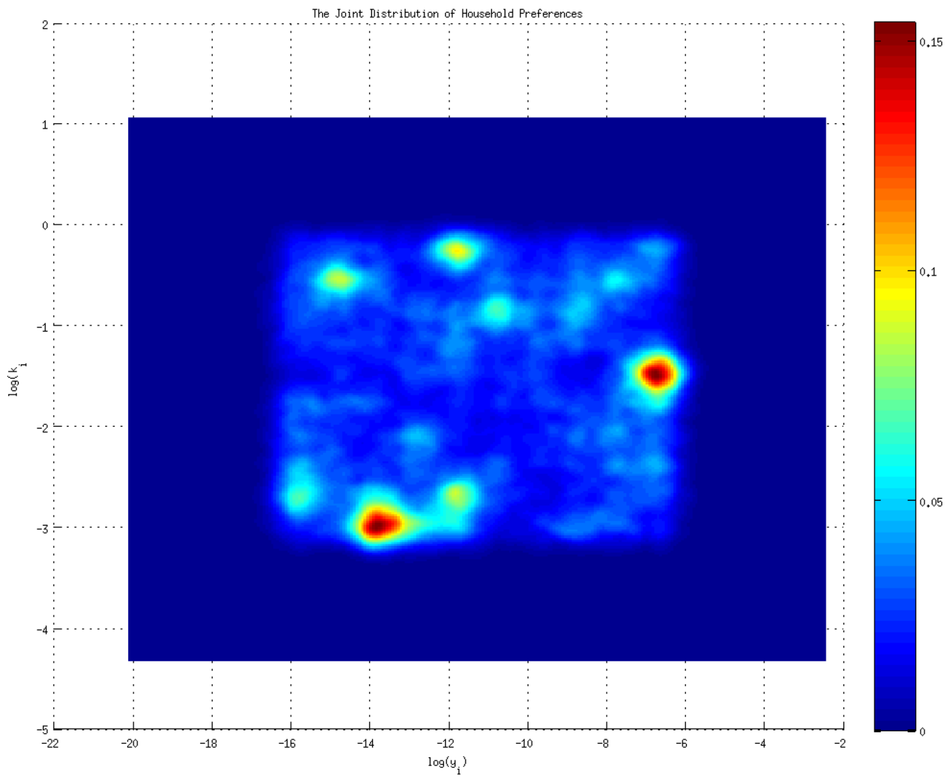

Figure 2: This figure displays the bivariate household preference distribution by log of household preferences. It is a summary of the labor preference fit, as a table of coefficients would be unwieldy.

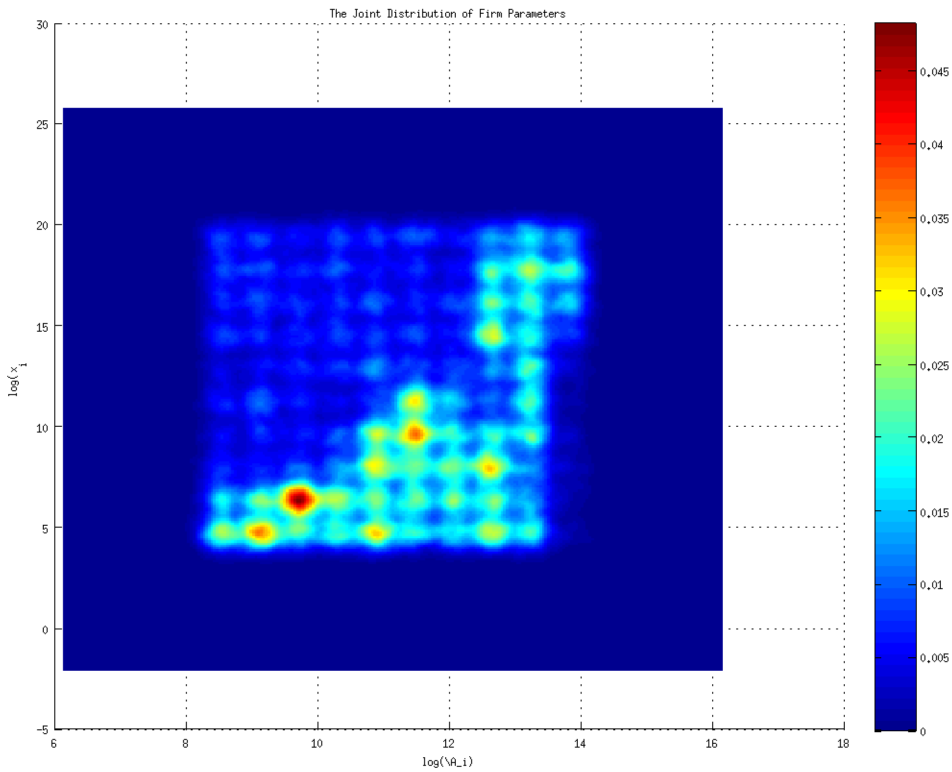

Figure 3: Like Figure 2, this figure displays the bivariate firm type distribution by log of firm parameters. It is a summary of the firm type.

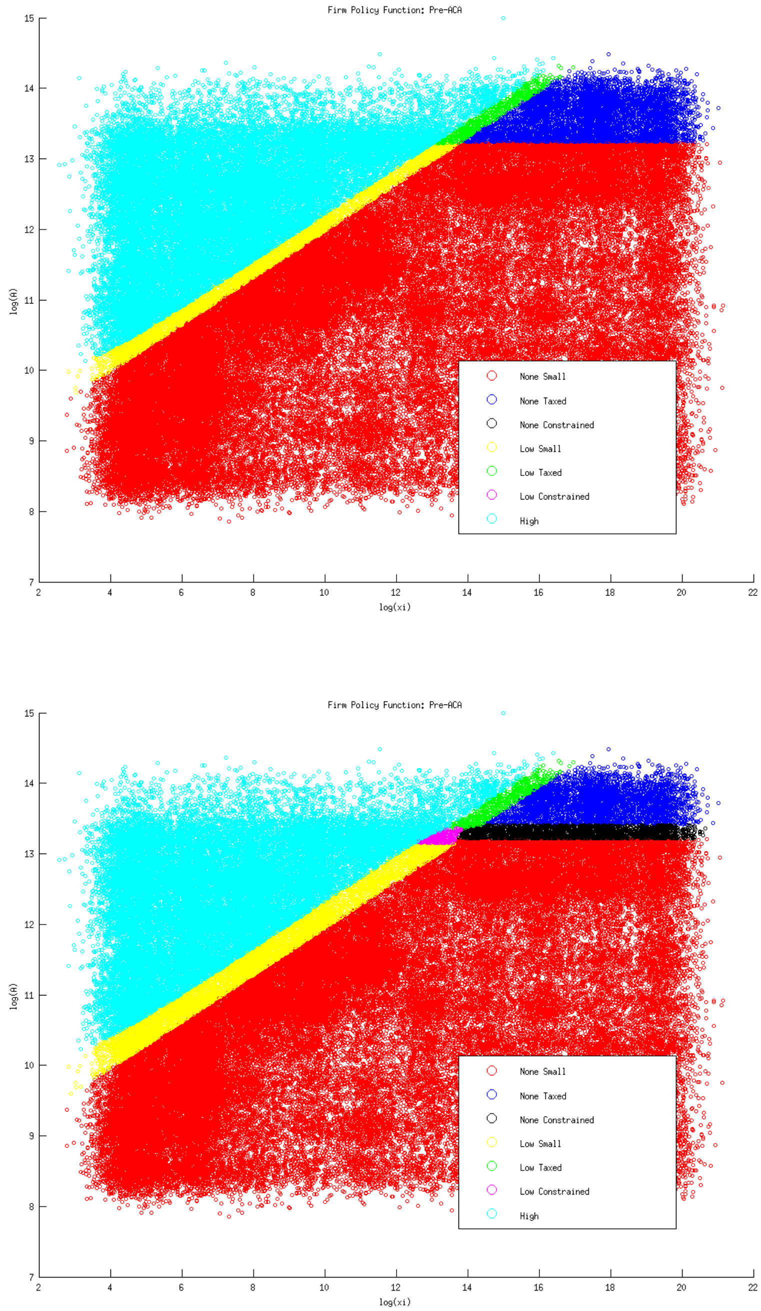

Figure 4: This Figure depicts the pre- and post-ACA policy functions for firms by log productivity A_i and log healthcare administrative costs κ_i. The top figure depicts the intuitive structure of healthcare offerings and size. High-healthcare, low-healthcare, and no-healthcare firms are arrayed from left to right, with large firms and firms with low healthcare costs offer "high" amounts of healthcare in light blue. Medium-productivity firms with mild healthcare costs offer "low" amounts of healthcare in yellow and green. Yellow firms are smaller than 50, while green firms are larger than 50 employees. While unimportant before the ACA, the difference depicts the "targeted" and "non-targeted" firms. Small firms and firms with very high healthcare costs do not offer healthcare and are depicted in red (small firms) and blue (large, targeted firms). The shift due to the ACA can be seen by comparing the bottom panel to the top panel.

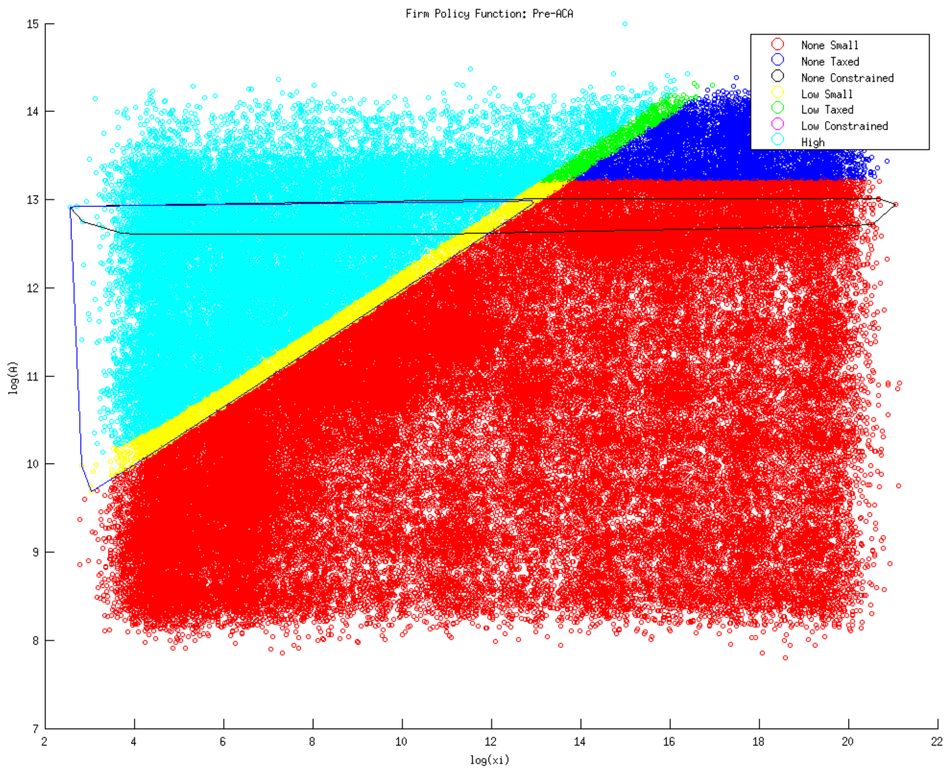

Figure 5: This figure offers an insight into the identification scheme of firm moments by depicting two of censored moments of Table 6 as a blue and black polygon. Specifically, it the black polygon depicts the size distribution of firms, conditional on being larger 40 employees and smaller than 99 empoyees. The blue polygon depicts all firms between 25 and 99 employees that offer health insurance. An estimation method based on only these two polygons attempts to get an area within these boxes commensurate with the targets in Table 6. The intersection of these two polygons gives information on the desired level of three separate areas.

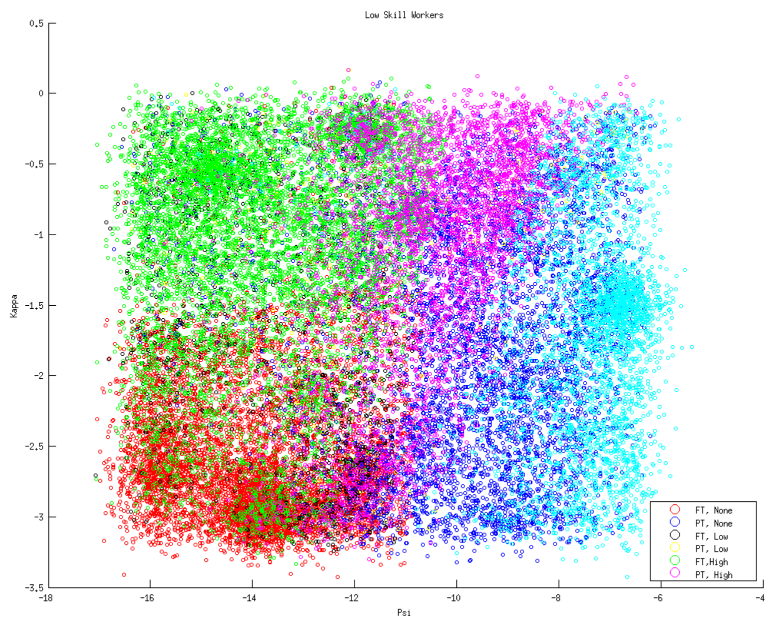

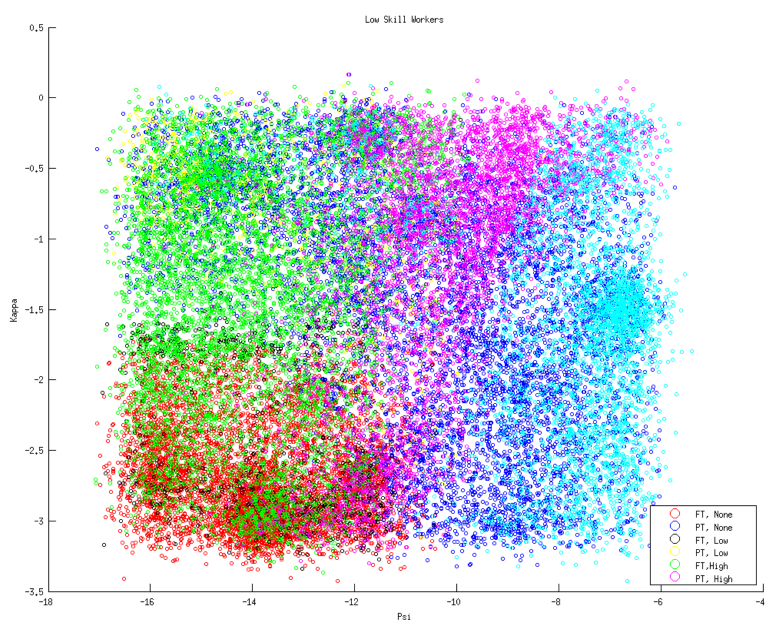

Figure 6: This Figure depicts the pre- and post-ACA policy functions for households by log distaste for labor ψ_i and log taste for healthcare κ_i. The top figure depicts the intuitive structure of worker sectoral choice. Because households vary by more than two preference parameters, especially non-labor income and family size, the policy function is more disperse. Nevertheless, it follows expected patterns: workers in green who value health insurance and do not get great disutility from labor work full-time in the high-healthcare sector. Travelling downwards along κ workers then work in the low healthcare sector full time, and finally the no healthcare sector full time. Moving right along ψ, workers who hate labor more typically work part-time, or eventually do not work at all (light blue).

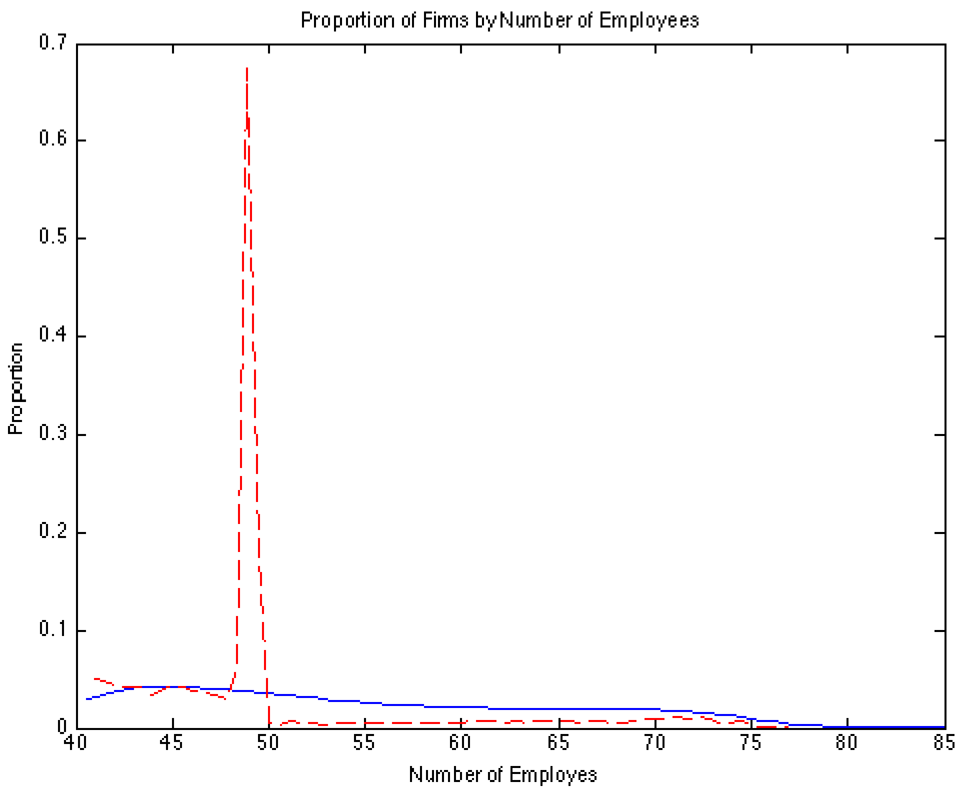

Figure 7: This figure depicts the firm size distribution, conditional on being larger than 40 employees and smaller than 80 employees, before and after the ACA¿ As might be expected from a size-based provision, firms bunch up at the constrained line of 49 employees. Because there is a dual trigger, and some firms have low costs of health insurance, not all firms switch to 49 employees: holding wages and prices constant, some firms would remain where they were with regards to employment. This graph makes clear a point: size-based provisions move the distribution away from the optimal distribution of worker allocation by firm, causing a loss in average productivity.

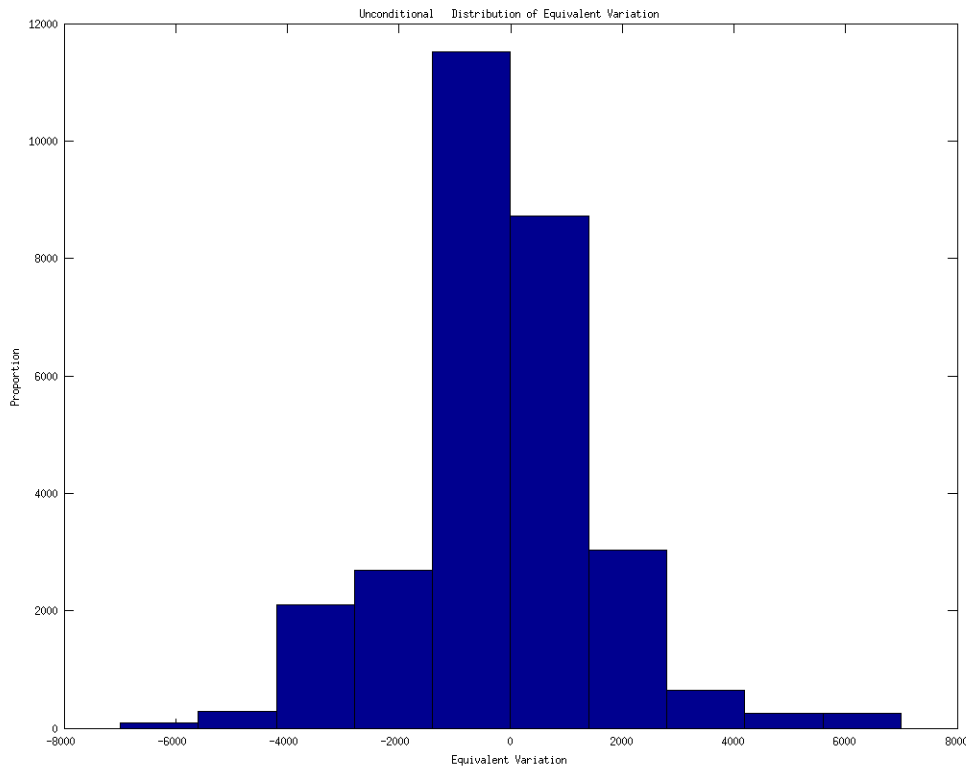

Figure 8: This figure depicts the distribution of EV, conditional on it being greater than -$7000 and less than $7000, for scalability. It makes two points: first, it is relatively symmetric with a mass near zero, denoting a large number of households relatively indifferent between the ACA and no ACA. Second, it has dispersion, noting that the ACA has distinct winners and losers.

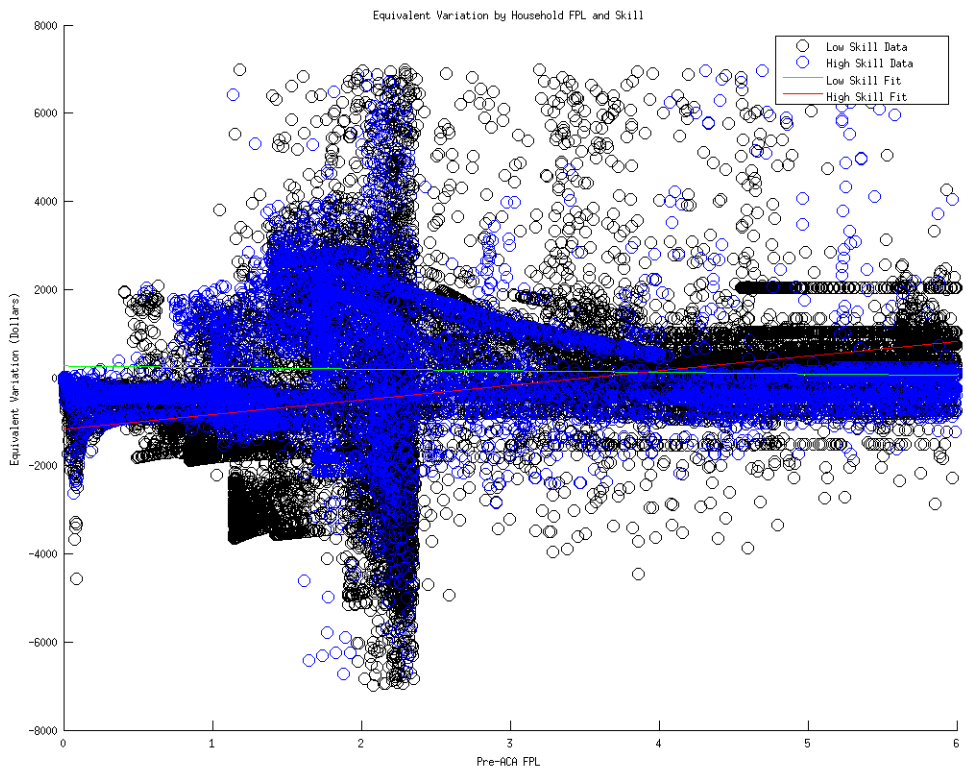

Figure 9: This figure depicts the distribution of EV over FPL and linear regression lines of EV by FPL, conditional on EV being greater than -\$7000 and less than \$7000. It makes several points: first, even in the face of not having access to Medicaid, poor households rarely suffer extremely, due to the low penalty. The households that are likely to suffer the most and to benefit the most are typically near the sharp drops in NGI subsidies at $2.5 \cdot FPL$, where the mandate starts to become punitive but exchange subsidies decline.

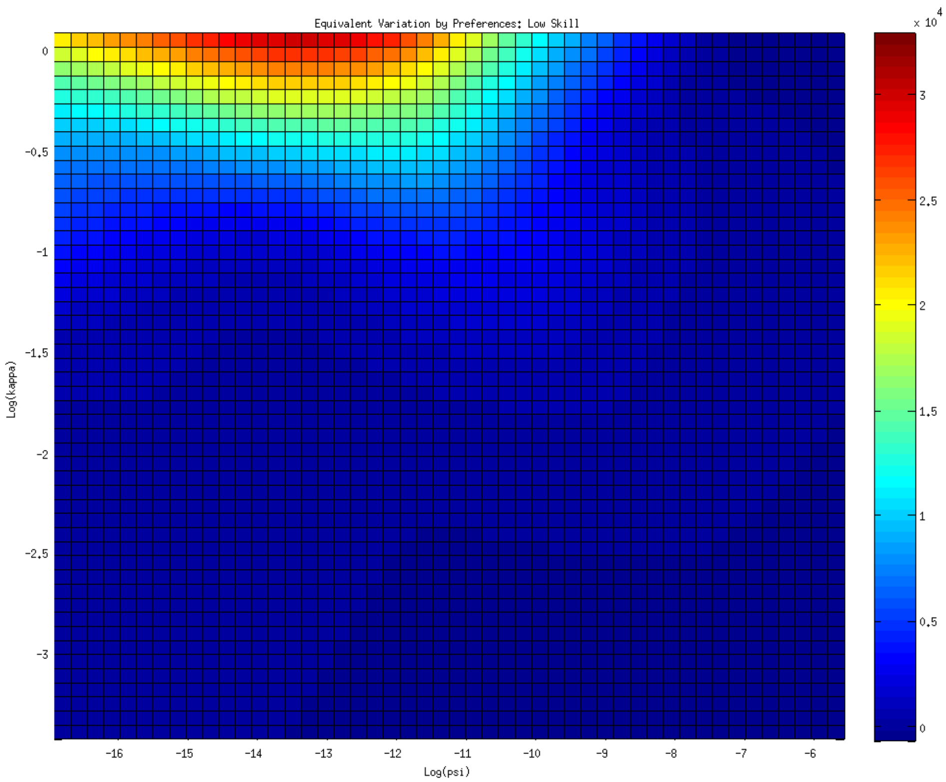

Figure 10: This figure depicts a lowess fit of equivalent variation in dollars by preference for low-skilled individuals. Those individuals who strongly disliked labor and desired healthcare were not those most benefited: they generally already were lowering labor hours enough to gain access to Medicaid. Those individuals subject to “job lock” disliked labor only moderate amounts: the primary beneficiaries were instead those individuals who were working (and therefore had only moderate distaste for labor) but had a high utility from healthcare. The primary losers from the ACA were those that did not have a high relative taste for healthcare.

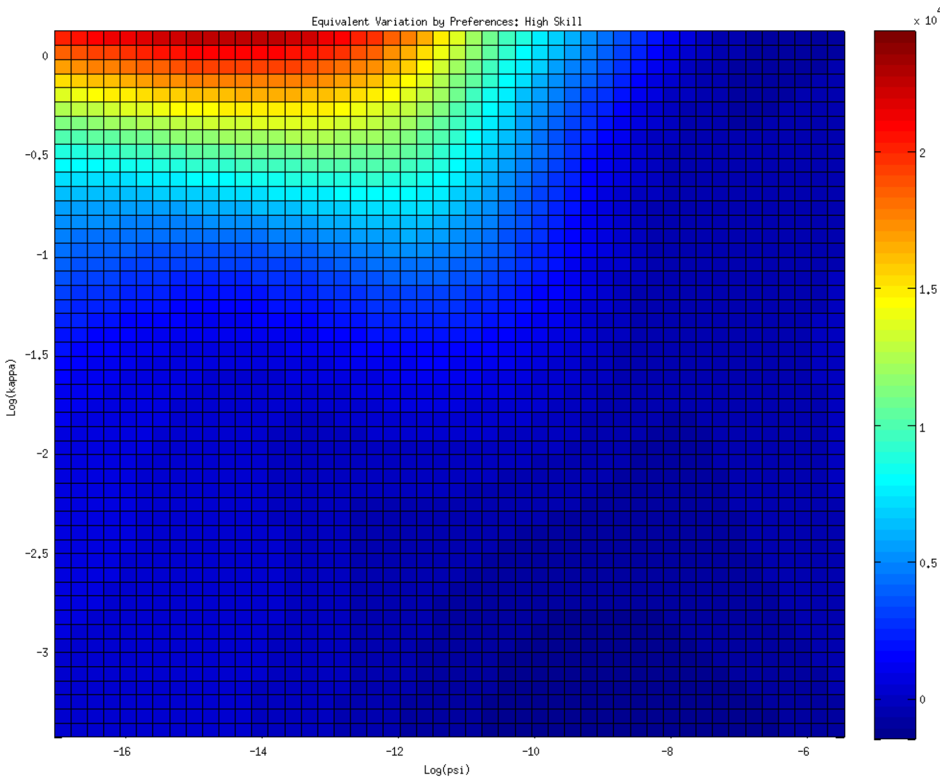

Figure 11: This figure replicates Figure 10 for high-skilled individuals. While displaying the same pattern, EV is slightly less concentrated along the healthcare preference κ_i.