

2012

The Role of Health in Labor and Marriage Markets: Dissertation Summary

Vincent Pohl
Yale University

The Role of Health in Labor and Marriage Markets

Vincent Pohl

Health plays an important role in labor markets for several reasons. First, physical health is a component of human capital. Better health leads to more available time for labor market activity, but better health may also increase wages, as healthier individuals are more productive. On the other hand, labor market activity affects health. Besides the direct effect of working, for example, in an adverse environment, institutional factors also play a role. In particular, private health insurance in the United States is tied to employment. Therefore, individuals have an incentive to work for an employer that provides them with health benefits, although they may be a better match with another employer that does not provide health insurance.

The relationship between labor supply and public health insurance is less obvious but nevertheless important. Since Medicaid eligibility depends on income, individuals have an incentive to decrease their labor supply in order to qualify for this public health insurance benefit. This effect matters especially for individuals who do not receive private health insurance through their employers. Since jobs with higher wages are more likely to provide health benefits, there is a range in the middle of the earnings distribution where individuals earn too much to be eligible for Medicaid, but do not receive employer-sponsored health insurance (ESHI). In this setting, the recent U.S. health care reform, the Patient Protection and Affordable Care Act (PPACA), has significant potential effects on labor supply. Both the expansion of Medicaid eligibility thresholds and the introduction of health insurance subsidies serve to fill the gap that currently exists between Medicaid and ESHI. The changes enacted by this reform imply that individuals can choose their labor supply without necessarily considering if they can obtain public or private health insurance. This can potentially lead to higher earnings and increased productivity. If individuals do not have to work full time in order to obtain ESHI, however, this reform may also decrease labor supply. In the first two chapters of my dissertation, I analyze the effect of the PPACA on the labor supply of a group that can particularly benefit from this reform—single mothers.

Health plays an important role not only in the labor market but also in the marriage market. As in the labor market, the relationship between health and individual behavior is complex. For example, being married may improve one's health because of financial and emotional support. On the other hand, being healthy may improve one's prospects on the marriage market. In the third chapter of my dissertation, I focus on the latter effect and study whether individuals invest in their health by improving their health-related behavior

(such as exercising or quitting smoking) in order to attract a spouse.

Chapter 1

Medicaid and the Labor Supply of Single Mothers: Implications for Health Care Reform

The PPACA has the potential to significantly affect employment among the low-wage population. On the one hand, increased Medicaid eligibility and newly established health insurance subsidies will reduce the incentive to seek health insurance through employment. On the other hand, expanding the income cutoffs for eligibility will remove work disincentives for individuals who currently refrain from working because they are afraid to lose Medicaid eligibility. The effects of this law on work incentives for those marginally attached to the workforce are therefore ambiguous. At the same time, this population is important to policymakers because of its vulnerability and relative deprivation.

In this chapter, I aim to determine the work incentive effects of the PPACA provisions that apply specifically to single mothers, a group that is characterized by low attachment to the labor force and that will be particularly affected by these reforms. Single mothers and their children are also the main beneficiaries of Medicaid. In contrast to married women, they do not have the option to obtain health insurance coverage through the employers of their spouses. Moreover, they often lack the qualifications necessary to find a job with ESHI.

I estimate a structural model of labor supply that incorporates Medicaid and ESHI. In estimating the model, I rely on exogenous variation from recent expansions in Medicaid eligibility. After obtaining estimates of the preference parameters governing employment choice, I use them to simulate single mothers' labor supply and take-up of ESHI under health care reform. The structural approach allows me to analyze the effects of health insurance subsidies, a policy that has not yet been introduced at the national level.

My contributions in this chapter are fourfold. First, I model and estimate the joint impact of Medicaid and ESHI on single mothers' labor supply. Potential ESHI coverage can affect the direction of the effect of Medicaid on employment. On the one hand, many low-income individuals are not qualified enough for jobs that provide health benefits. Moreover, they are only eligible for Medicaid if their incomes fall below the applicable threshold. This induces work disincentives. Expanding Medicaid eligibility or introducing health insurance subsidies would therefore relax this constraint and potentially increase labor supply. On the other hand, if the income threshold increases sufficiently, full-time workers with ESHI coverage may become eligible for Medicaid or subsidies. Since Medicaid and subsidized health insurance are cheaper and may be more generous than ESHI, these

workers may drop ESHI coverage (crowding-out of ESHI) and potentially reduce their labor supply.

Second, I distinguish between full-time and part-time employment. The literature typically considers only the participation decision. As argued above, workers with low initial labor supply might increase their labor supply when the Medicaid eligibility threshold increases. Others might work full-time prior to health care reform in order to qualify for ESHI coverage. Introducing health insurance subsidies allows these individuals to reduce their labor supply and drop ESHI coverage while obtaining subsidized health insurance. Therefore, not allowing for an intensive margin would mask these changes in labor supply.

Third, I allow individual differences in the valuation of health insurance to affect the labor supply effects of Medicaid expansions. When analyzing the effects of Medicaid expansions on labor supply, it is important to account for individuals' valuation of health insurance, which in turn depends on their health status. For example, a healthy person might change her behavior less in response to Medicaid expansions than someone with chronic medical conditions that require expensive health care. Specifically, I allow for health-related heterogeneity in medical expenditure and preferences over health insurance coverage.

Fourth, the exogenous variation used to identify the impact of Medicaid expansions on labor supply comes from changes in Medicaid policies after the 1996 welfare reforms. These reforms affected mostly parents instead of children who benefited from earlier expansions. Since the recent health care reform extends Medicaid eligibility to even broader groups, the analysis of the more current Medicaid expansions is of particular policy interest. Moreover, some states implemented Medicaid thresholds that are close to the one specified by PPACA, which allows for more realistic policy simulations.

I specify a discrete static labor supply model. Single mothers maximize their utility, which is derived from consumption, leisure, and own and children's health insurance coverage subject to a budget constraint that incorporates out-of-pocket medical expenditure and government transfers. The possible choices are nonwork, part-time without ESHI, full-time work without ESHI, and full-time work with ESHI. Depending on earnings and state- and year-specific Medicaid thresholds, single mothers and their children are eligible for Medicaid.

To estimate the choice model, I develop a two-stage estimator. In the first stage, I estimate a wage equation and two medical expenditure equations for mothers and children. The goal is to obtain conditional distributions of the utility function arguments. The wage regression accounts for selection into the labor force and into jobs with or without ESHI using a bivariate selection correction mechanism. Both the wage and medical expenditure regressions are estimated allowing for random individual effects, thus taking advantage of the

panel structure in my data. In the second stage, I estimate the parameters of the utility function using draws from the conditional distributions of earnings and medical expenditures as regressors. I allow both observed and unobserved heterogeneity in the utility parameters. The unobserved heterogeneity is modeled as a multivariate normal distribution. Hence, I use a multinomial logit model with choice probabilities that are integrated over the normal distribution. To estimate the parameters of this distribution and the fixed preference parameters, I use maximum simulated likelihood. Standard errors are block-bootstrapped over the two stages of the estimator.

Individual data come from the Medical Expenditure Panel Survey (MEPS) from 1996 to 2008. The MEPS is the ideal data set for the question at hand since it combines detailed information in health insurance coverage and medical expenditures with labor market variables. The medical expenditure data are very reliable since individuals' responses have been cross-checked with medical providers. Moreover, the MEPS provides a panel, albeit a short one, so I can account for individual effects in my two-stage estimator. I use a restricted version of the MEPS, which allows me to match individual data with state- and year-specific policy variables such as Medicaid thresholds and welfare rules.

The estimated preference parameters indicate that single mothers with medical conditions are significantly more likely not to work in order to be eligible for Medicaid or to work full-time with ESHI. This result implies that women who need health insurance most avoid the gap between Medicaid and ESHI where they would be uninsured. It is likely that these women would otherwise choose a different location on the hours distribution to increase their productivity or have more time for leisure or child care. Hence, these women benefit particularly from PPACA since it allows them to enter the labor force and to work in jobs without ESHI coverage.

Based on these estimates, I simulate single mothers' labor supply choices. First, I check the results by simulating labor supply under current Medicaid rules. The simulated choice probabilities match the observed ones closely, thereby providing evidence that the model is valid. Then I simulate labor supply under two policy counterfactuals: expansion of Medicaid only and both expansion of Medicaid and introduction of health insurance subsidies. While the PPACA corresponds to the latter simulation, it is important to analyze how the components of this reform affect labor supply decisions independently. The simulation results show that health care reform increases labor force participation among single mothers by about 6 percent. Moreover, labor supply at the intensive margin grows by about 5 percent. Finally, health care reform leads to crowding-out of ESHI of about 40 percent in this population. These results are heterogeneous across subgroups, however, with single mothers with medical conditions reacting most strongly to the reform, as expected based on the preference parameter estimates. I also

use these simulations to estimate the welfare implications of this reform. In particular, I calculate the compensating variation of this reform, that is, the dollar amount a single mother would be willing to forgo in order to obtain the reform. On average, this amount is about three times larger than the per-family increase in public expenditures, where the increase comes mostly from the new health insurance subsidies. Hence, I conclude that this reform is welfare improving.

Chapter 2

The Effect of Recent Medicaid Expansions on the Labor Supply of Single Mothers

One drawback of the structural approach in the first chapter is the functional and distributional assumptions necessary to estimate policy-invariant preference parameters. They are needed to simulate labor supply under a reform that is not yet observed, but it is difficult to test directly if these assumptions are valid. To check the robustness of my simulation results in Chapter 1, I perform reduced-form regressions of single mothers' labor supply on Medicaid eligibility. Exogenous variation comes again from recent expansions in Medicaid thresholds for parents. Instead of estimating the effect of the PPACA, which would require more structure, I simply estimate the effect of existing Medicaid reforms. I can then compare these estimation results to the policy counterfactual I perform in Chapter 1 that only includes Medicaid expansions but no health insurance subsidies. Besides providing a robustness check for the results in Chapter 1, this chapter is also of broader interest in its own right. No study exists to date that analyzes the overall effect of recent Medicaid expansions on labor supply, and this chapter fills this gap in the literature.

Since Medicaid eligibility depends on earnings, which in turn are determined by labor supply, simply regressing labor supply on Medicaid coverage or eligibility yields inconsistent estimates due to endogeneity of the Medicaid eligibility variable. In order to avoid this problem, I use two different identification strategies. First, I use Medicaid and Children's Health Insurance Program (CHIP) income eligibility thresholds as instruments for Medicaid and CHIP coverage. To the extent that these thresholds are exogenous, they are valid instruments. As with all public policies, there is a concern that policymakers choose Medicaid thresholds in response to individual behavior, rendering these thresholds endogenous. In this case, however, it is plausible that policymakers do not react to single mothers' labor supply when choosing Medicaid thresholds. Medicaid is a program intended to provide health insurance coverage to low-income families, but not as a labor market program.

Second, I use a triple-differences strategy by comparing labor market outcomes across states, over time, and between single mothers and childless single women. This

identification strategy does not require exogeneity of Medicaid thresholds with respect to single women's labor supply. Hence, it is valid even if the exogeneity assumption described above fails. Instead, it relies on the assumption that variation in labor supply across single mothers and childless single women is independent of Medicaid thresholds when controlling for state and year fixed effects. Childless single women constitute an ideal control group since they face a similar labor market environment as single mothers, but they are never eligible for Medicaid. Under current policies, only parents of Medicaid-eligible children qualify for this benefit. Since Medicaid coverage is identically equal to zero for childless women, I cannot use it as the explanatory variable. Instead, I simulate Medicaid eligibility by state, year, and demographic group, and use this simulated measure interacted with an indicator for the presence of children as the regressor of interest.

A second contribution of this chapter is the analysis of migration patterns between states in response to changes in Medicaid thresholds. An important concern when using cross-state variation in policies is potential migration in response to these policies. If single mothers select their location based on Medicaid generosity, for example, the identifying assumption that Medicaid eligibility thresholds are exogenous would be invalid. To test if migration because of differences in Medicaid generosity is a problem in this context, I compare year-to-year changes in Medicaid thresholds of single mothers and childless single women who did or did not migrate between states. Again, childless women are an ideal comparison group since they are never eligible for Medicaid and therefore have no incentive to migrate because of differences between these policies. Otherwise, the incentives for single women with and without children to move between states should be similar because of cross-state differences in labor markets. Therefore, we expect to see no systematic difference in migration patterns between these two groups in response to Medicaid threshold changes.

Using a large sample of single women with and without children from the March CPS from 1996 to 2007, I first provide extensive descriptive evidence on the effect of Medicaid expansions on single mothers' labor supply. This mostly graphical suggests that single mothers adjust their labor supply in response to the incentives that are implied by Medicaid income tests. Overall, expanded Medicaid thresholds have a negative effect on labor supply. This adjustment mostly occurs at the extensive margin, but to some extent single mothers also reduce hours worked in order to qualify for Medicaid.

In the main estimation results both identification strategies yield the same result that expansions in Medicaid eligibility do not decrease labor supply. I use various measures of labor supply for both the extensive and intensive margins, including labor force participation, full-time versus part-time work, weeks worked per year, and hours worked per week.

The results for all outcomes are consistent in showing that an increase in Medicaid eligibility thresholds does not have a negative effect on labor supply.

To test if single mothers migrate between states in order to benefit from more generous Medicaid policies, I use CPS data on the state of residence in the previous year. This allows me to compare migration of single women with and without children in a difference-in-differences framework when the destination state has a different Medicaid threshold from the initial state. There is no systematic difference between these groups when comparing the unconditional distributions of differences in Medicaid thresholds pre- and postmigration. Hence, I conclude that single mothers do not choose their locations based on Medicaid thresholds.

In summary, this chapter provides evidence that the structural assumptions in Chapter 1 are valid since I obtain qualitatively similar results. Medicaid expansions alone have moderately negative effects on single mothers' labor supply. This result is of interest to policymakers since it shows that the health insurance subsidies in PPACA are necessary to avoid decreases in labor supply in this population. Moreover, I show that migration in response to Medicaid thresholds does not bias the estimation results. This conclusion is important since it backs an important assumption when using policy variation across states. It also shows that Medicaid policies do not impact individual behavior enough to cause location changes.

Chapter 3

The Marriage Market Ratio and Investment in Health

Health also plays an important role in the marriage market: healthy individuals may be more likely to get married or find a more desirable spouse. Therefore, individuals can affect their marriage market outcomes by improving their health. This is the topic of the third chapter: I test whether individuals invest more in their health when it is harder to find a spouse. Sociologists and researchers in public health have long recognized the relationship between marriage and health. The correlation between marital status and health is often given a causal interpretation termed *marital protection*. However, most existing studies merely uncover a correlation between marital status and good health or lower mortality. Causal effects are difficult to determine because individuals select into marriage. The effect of marriage on health decreases when marital selection is accounted for, but the positive correlation persists. Therefore, the consensus is that both selection and a causal protective effect of marriage on health matter.

If being married improves one's health and people sort into marriage based on their premarital health status, an interesting question arises: To what extent do people decide

to improve their health in order to attract a potential spouse and to increase the likelihood of marriage? The third chapter attempts to answer this question. Depending on the result, this has important implications for the findings of the existing literature claiming the presence of a causal protective effect of marriage on health. If individuals improve their health in order to get married, any positive correlation between marital status and health is confounded by preexisting health status. This effect is not necessarily the same as marital selection because the decision to be healthy is made consciously to attract a potential spouse.

To formalize the relationship between marriage market conditions and health investment, I develop a simple model of premarital health behaviors. Intuitively, the incentive to invest in one's health in order to attract a potential spouse is highest when it is not certain to find a spouse, but the probability is sufficiently high. Once it becomes very unlikely to find a spouse, the cost of investing in one's health outweighs the (uncertain) benefit of marriage. On the other hand, when marriage is certain, the incentive for health investments is smaller because it is more likely to find a spouse in any case. The likelihood to become married, in turn, depends on the relative numbers of men and women. Hence, the model predicts that the relationship between the relative numbers of individuals of own and opposite sexes and healthy behavior has an inverted U-shape.

A simple way to capture relative numbers of men and women on the marriage market is to use sex ratios, which are defined by the ratio of men to women at marriageable ages. Instead of simply using the ratio of cohort sizes, I develop a marriage market ratio that explicitly accounts for the empirical distribution of age differences between married men and women. A novel contribution is to distinguish between different racial compositions of couples. For example, the distribution of age differences among mixed-race couples exhibits a larger variance than that of same-race couples. My marriage market ratio accounts for these differences and therefore correctly reflects the marriage market conditions of black women, for example, who have the option to marry a black or a white man of different ages.

Using data from the National Health Interview Survey and the Behavioral Risk Factor Surveillance System, I test the hypothesis that the marriage market ratio affects the health-related behaviors of smoking, drinking, and exercising. The results indicate that men and women of all races are less likely to smoke and black men and women are less likely to drink when facing unfavorable marriage market conditions. The results for physical activity, however, are mixed.