Published in final edited form as:

J Public Econ. 2019 March ; 171: 18–28. doi:10.1016/j.jpubeco.2018.01.010.

Household Labor Supply and the Gains from Social Insurance*

Itzik Fadlon,

University of California, San Diego and NBER, address: Department of Economics, University of California, San Diego, 9500 Gilman Drive #0508, La Jolla, CA 92093-0508

Torben Heien Nielsen

University of Copenhagen, address: Department of Economics, University of Copenhagen, Øster Farimagsgade 5 building 26, DK-1353, Copenhagen K

Itzik Fadlon: fadlon@ucsd.edu; Torben Heien Nielsen: thn@econ.ku.dk

Abstract

The marginal gains from social income insurance programs are captured by the gap in the marginal utility of consumption across states of nature. To identify this gap in the context of the household, this paper offers a new labor-supply based approach that leverages household-level economic interactions and optimality conditions. We demonstrate that, in frameworks of efficient household allocations, spousal labor supply responses to shocks have direct implications for the gains from more generous government benefits to households. We show that this holds for both intensive and extensive margin responses under fairly general conditions. Our analysis illustrates how labor market data can be used for assessing marginal welfare gains in a general class of social insurance schemes, including the large and important programs of disability insurance and survivors benefits. Hence, household labor supply behavior and responses to shocks, which are widely studied in theoretical and empirical work, hold valuable information for the optimal design of social insurance.

*We thank David Autor, Martin Browning, Raj Chetty, Jeffrey Clemens, Julie Cullen, David Cutler, Gordon Dahl, Ariel Dora Stern, Raluca Dragusanu, Liran Einav, Torben Fischer, Eric French, Elhanan Helpman, Nathan Hendren, Simon Jäger, Xavier Jaravel, Lawrence Katz, Claus Thustrup Kreiner, David Laibson, Jessica Laird, Rafael Lalive, Camille Landais, Søren Leth-Petersen, Magne Mogstad, Joana Naritomi, Arash Nekoei, Jann Speiss, Yoram Weiss, Gal Wettstein, Danny Yagan, and Eric Zwick for helpful discussions and comments. We gratefully acknowledge financial support from the Alfred P. Sloan Foundation through the NBER Household Finance working group (under the direction of Brigitte Madrian and Stephen Zeldes) and the NBER Pre-doctoral Fellowship Program on the Economics of an Aging Workforce (under PI David Card), the Boston College Center for Retirement Research, the NBER Roybal Center for Behavior Change in Health and Savings, and EdLabs at Harvard University. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA or any agency of the Federal Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof. The research reported in this publication was also supported by the National Institute on Aging of the National Institutes of Health under Award Number P30AG034532. The content is the sole responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Correspondence to: Torben Heien Nielsen, thn@econ.ku.dk.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Keywords

Social Insurance; Evaluation of Welfare Gains; Household Labor Supply

1 Introduction

Welfare analysis and the optimal design of any social insurance program requires identifying both the benefits and the costs associated with that program. The welfare costs of government transfers are captured by the impact of households' behavioral responses to the policy on the government's budget, known as the "fiscal externality". The welfare gains from more generous social insurance are captured by the gap in the marginal utility of consumption across states of nature. This gap, which is zero in the first-best allocation in which marginal utilities are smoothed across states of nature, measures market inefficiency and quantifies the potential benefit from additional government intervention. Estimating the welfare costs is conceptually straightforward with the appropriate policy variations in the data that identify the policy-specific elasticities (Hendren 2016b). Estimating the welfare gains from social insurance is a more challenging task, and the main approach in the modern literature to tackle it is based on studying the consumption smoothing effects of government benefits (Chetty and Finkelstein 2013).

However, analyzing consumption across different states of nature involves significant challenges, particularly in the context of the household (Chetty and Finkelstein 2013; Bee et al. 2013; Pistaferri 2015). Consumption is very difficult to measure accurately due to noise and recall errors and is typically available for relatively small samples. Additionally, consumption measures are usually partial and cover a sub-set of goods, such as food expenditure. Focusing on limited aspects of expenditure can lead to misleading conclusions about actual consumption, e.g., in the presence of home production (as emphasized by Aguiar and Hurst 2005).¹ Moreover, even comprehensive and accurate data on households' overall expenditure across states of nature—which are largely uncommon—would require nontrivial assumptions or complex estimations to be translated into individuals' consumption bundles, since in the household setting individual consumption is not directly assignable and observable. Doing so must carefully take into account consumption flows of durable goods, household public goods, and, importantly, economies of scale in the household's consumption technology to account for the composition of the household and its evolution over the life-cycle and across states of nature (see, e.g., Blundell and Lewbel 1991, Browning et al. 2013, Blundell et al. 2013, and Low and Pistaferri 2015).

Recognizing these difficulties, recent influential studies have sought alternative techniques for recovering the gap in marginal utilities using information from the labor market. Specifically, Shimer and Werning (2007) show how to identify the gains from higher unemployment benefit levels using the comparative statics of reservation wages with respect

¹Recent advances in the analysis of consumption partially deal with some of these limitations, either by using survey data that cover a wider set of non-durable and services expenditure (Blundell et al. 2015; Low and Pistaferri 2015), or by creating consumption measures from income and wealth registers (Browning and Leth-Petersen 2003; Koijen et al. 2014; Kreiner et al. 2014; Kostøl and Mogstad 2015; Autor et al. 2015; Kolsrud et al. 2016; De Giorgi et al. 2016).

to benefit increases. Chetty (2008) develops a method to recover the gains from higher benefit levels using liquidity and substitution effects in the search effort of the unemployed, and Landais (2015) provides a similar technique for the gains from longer benefit durations.² These labor-supply based methods, however, have been so far limited to the context of unemployment shocks and unemployment insurance.

In this paper, we build on this recent work and offer a new labor-supply based approach for welfare analysis, in which we leverage household-level economic interactions and optimality conditions. Exploiting the interplay between the consumption allocations and labor supply decisions of household members, our approach uses only the labor supply behavior of the indirectly-affected spouse. We show that, under the assumption of household efficiency, spousal labor supply responses to shocks can be mapped to predictions about the welfare gains from providing more generous government benefits. Our method for welfare analysis is therefore general, as it merely relies on standard optimality conditions with respect to household members' choices. The simple logic that underlies our results is as follows. In each state of nature spouses work to the point where their *own* marginal utility loss from working equals *each* household member's valuation of the additional consumption from the spouses' increased earnings. Hence, the sensitivity of spousal labor supply to shocks and economic incentives can reveal how the household's valuation of additional consumption changes across different states, which captures the gains from insurance. Intuitively, the extent to which household labor supply responds to shocks as self-insurance is directly related to the degree to which households lack formal insurance and hence to the scope for welfare-improving government benefits.

Our approach to welfare analysis, which explicitly allows for economic dependencies across household members, has two main contributions to the existing literature. First, compared to current methods that use labor market data, our method is applicable to a general class of shocks and social insurance schemes in a household setting.³ Importantly, by analyzing the behavior of spouses who are indirectly affected by shocks through the sharing of household resources, we allow these shocks to have arbitrary effects on the directly-affected household members' working ability (e.g., their productivity or disutility from labor) and labor market opportunities (e.g., their employability).⁴ These include the prevalent setting of (fatal and non-fatal) severe health shocks, thereby allowing the welfare analysis of the large and important programs of disability insurance (DI) and survivors benefits using our framework. Other important applications include the assessment of the value of unemployment insurance for the long-term unemployed—whose long durations of unemployment significantly harm their employment prospects (as documented by, e.g., Kroft, Lange, and Notowidigdo 2013)—as well as any other setting in which the individuals who are directly impacted may be unresponsive to economic incentives or at a corner solution (of either working full-time or not working at all). For example, in the debate on the privatization of

²This is also in the spirit of arguments raised by Autor and Duggan (2007).

³An obvious limitation of our approach is, therefore, its inapplicability to analyzing single individuals.

⁴The labor supply behavior of directly-affected individuals in these cases may be unresponsive to shocks or economic incentives and hence cannot reveal the household's consumption preferences and need for additional income insurance. If these attenuated behavioral responses were wrongly used in welfare analysis, they would be misleadingly interpreted as low valuations of additional income insurance, while being actually driven by the directly-affected individuals' inability to adjust their labor supply accordingly.

Social Security, the value of protecting against pension-wealth losses in the 401(k) account of a working individual can be recovered by the labor supply response of his or her spouse. Similarly, spousal labor supply can be also used to evaluate the welfare losses caused by the discontinuation of employee compensations, such as health insurance, life insurance, and employer matching in retirement savings.

Second, compared to welfare analysis that relies on consumption data, our approach has several significant advantages attributable to the reliance on data from the labor market. Labor supply data are typically more precise and widely available, in particular with researchers' increasing access to register-based data on household earnings. Additionally, and of particular importance in the analysis of households, labor supply is directly assignable to individual household members, as opposed to intra-household allocations of purchases (i.e., individual-level consumption bundles) that are usually not observed (Browning et al. 2014). Therefore, unlike data on expenditure, information on labor supply does not require conversions to account for economies of scale, household composition changes (e.g., upon spousal death), or intra-household distribution of resources and bargaining power, but is rather readily usable for welfare analysis. Moreover, in applications with significant spousal labor supply responses which mitigate the consumption drop the household would otherwise experience, focusing on consumption fluctuations could lead to notable underestimations of the gains from social insurance. In these settings, it is the utility cost of reduced leisure that captures the gains from additional insurance, which is accounted for in our method that analyzes spousal labor supply.⁵ Of course, compared to the consumption approach, our approach has its own limitations. In particular, similar to other optimization-based approaches that map the marginal utility from consumption to the marginal utility from other arguments of the utility function, our approach requires that choices with respect to the analyzed argument are interior. Accordingly, our intensive-margin model requires that spouses' optimization leads them to an interior solution in hours of work, and our alternative extensive-margin model requires the presence of marginal households.

Our propositions and formulas, which rely on household optimality conditions, identify labor supply moments and preference parameters that are important for welfare analysis of social insurance. As such, they do not take a stance regarding how the interested researcher or policy maker should go about implementing them empirically. Specifically, we believe that our results can be a useful guide in both structural and reduced-form studies that aim to make quantitative statements about welfare and optimal policy. For reduced-form studies, that wish to analyze marginal changes in given policy environments, the formulas indicate which treatment effects to analyze and highlight the quasi-experimental variation needed for their identification. For structural studies, that wish to analyze fundamental policy reforms and simulate alternative policies, our formulas emphasize welfare-relevant household

⁵This is the case in Autor et al. (2015), who find that DI denials do not decrease the household income or consumption of married applicants due to significant spousal labor supply responses, and therefore conclude that their estimate for the welfare gains from DI receipt among married couples is due to increased leisure. It is also the case in Fadlon and Nielsen (2017), where we find large increases in spouses' labor supply following fatal health shocks that attenuate the consumption drop they would otherwise experience. In that application, analyzing labor supply would account for the utility cost of reduced leisure while the analysis of consumption would under-estimate the gains from additional survivors benefits.

behaviors and preference parameters that should be carefully modeled and identified. This can be useful in complex dynamic models in which full identification of all the model's primitives is particularly challenging and where transparency of sources of identification is critical. For such cases, our formulas offer a set of key moments to match for identifying the structural model and a set of critical structural parameters the chosen estimation procedure should identify.⁶

Besides the literature on social insurance, this paper relates to the large and recently revived work on spousal labor supply and its self-insurance role. This strand of the literature encompasses both structural and reduced-form studies in a variety of household shocks and social insurance settings. These include (but are not limited to) Ashenfelter (1980), Heckman and Macurdy (1980, 1982), Lundberg (1985), Maloney (1987, 1991), Spletzer (1997), Cullen and Gruber (2000), Stephens (2002), Blundell et al. (2015), Haan and Prowse (2015), and Wang (2016) in the context of wage and unemployment shocks and unemployment insurance, and Coile (2004), Meyer and Mok (2013), Autor et al. (2015), Olsson and Thoursie (2015), Dobkin et al. (2016), and Fadlon and Nielsen (2017) in the context of health and disability. Our paper emphasizes that the findings from these studies not only inform us of how households behave in response to shocks over the life-cycle, but also contain valuable normative content with direct implications for the design of social insurance in different settings. For example, in our empirical study of household responses to fatal and non-fatal severe health shocks (Fadlon and Nielsen 2017), we find increases in labor supply among survivors who experienced income losses when their spouses died, particularly when the deceased spouse had earned a significant share of the household income. Applying the formulas we develop here suggests large welfare gains from more generous survivors benefits and from conditioning benefits on the deceased spouse's work history in the Danish setting.

The remainder of the paper is organized as follows. We begin with Section 2 that sets a simple conceptual framework for the analysis of household labor supply and its responses to shocks. Then, in Section 3 we provide our core analysis and show how household labor supply can be normatively used to assess the welfare gains from social insurance. In Section 4 we discuss important extensions and generalizations to the simple model and illustrate the generality of our results, and in Section 5 we describe how to empirically implement our suggested welfare approach. Section 6 concludes.

2 A Stylized Model of Household Labor Supply

We start by analyzing a static model of household labor supply decisions. In this section, we lay out the conceptual framework and positively study the household's behavior in response to shocks. Specifically, we formalize how spousal labor supply can be used as insurance against income shocks to the household. While we consider both an intensive margin model

⁶In the context of unemployment insurance, Kolsrud et al. (2016) provide an example for using the consumption-based formulas to calibrate a structural model. Finkelstein et al. (2015), who study the value of Medicaid, offer an additional hybrid approach that is based on the analysis of marginal policy changes. They make inferences about non-marginal policy changes using additional statistical assumptions that allow them to interpolate between local estimates of the marginal impact of program generosity, which substitute for the economic assumptions in more structural estimations.

and an extensive margin model, we postpone the latter to Section 4 and lead with the intensive-margin model for illustrative purposes. The intensive-margin model allows for the most immediate comparison between the standard consumption-based representation of the welfare gains from social insurance and our proposed representation using household labor supply, and captures the intuition of our results in a transparent way.

The model that we study here—including the setup and preference specification—is the simplest possible model that demonstrates how household labor supply responses can be used to draw implications for the design of social insurance. Nonetheless, our qualitative arguments extend to much more general settings. After deriving the welfare formulas for the simple case, we discuss important generalizations to the highly-stylized static model as well as alternative assumptions about the household's preference structure (including state dependence and complementarities) and the household's behavior (including both the collective and the unitary approaches).

Setup—We study labor supply decisions of a two-person household, which consists of individuals 1 and 2. We consider a world with two states of nature: a “good” state, state g , in which member 1 works; and a “bad” state, state b , in which member 1 experiences a shock—e.g., a severe health shock—and drops out of the labor force. We employ this extreme assumption regarding member 1's labor supply for simplification, but any shock that leads to some degree of exogenous decline in this member's labor supply can be readily analyzed within the same framework. Households spend a share of μ^g of their adult life in state g and a share of μ^b in state b (with $\mu^g + \mu^b = 1$). In what follows, the subscript $i \in \{1, 2\}$ refers to the household member and the superscript $s \in \{g, b\}$ refers to the state of nature.

Household Budget Constraint—Denote by c_i^s and l_i^s the individual consumption and labor supply of member i in state s , respectively. Let A^s denote the household's state-contingent wealth and non-labor income—including transfers from any source of individually-purchased or employer-provided private insurance, transfers from relatives, and out-of-pocket expenses (such as medical bills). We denote by $\bar{z}_i^s(l_i^s)$ i 's net-of-tax labor income in state s , so that with a wage rate of w_i and a linear labor-income tax rate of τ_i^s we have $\bar{z}_i^s(l_i^s) = z_i^s(1 - \tau_i^s)$, where $z_i^s \equiv w_i l_i^s$ are gross earnings. Finally, let B^s represent benefits from the government in state s . It is possible to allow for income-testing in government benefits, but with some added analytical complication. We therefore choose to abstract from modeling this feature here, and focus on the state-contingent aspect of benefits which is at the core of our analysis of insurance against different states of nature. With this notation, the household's overall income in state s , y^s , satisfies $y^s = A^s + \bar{z}_1^s(l_1^s) + \bar{z}_2^s(l_2^s) + B^s$.

Individual Preferences—Let $U_i(c_i^s, l_i^s)$ represent i 's utility as a function of consumption, c_i^s , and labor supply, l_i^s , in state s . For simplicity, we assume for now (and relax later) that $U_i(c_i^s, l_i^s) = u_i(c_i^s) - v_i(l_i^s)$, where $u_i(c_i^s)$ is member i 's utility from consumption and $v_i(l_i^s)$ represents member i 's disutility from labor (including the utility loss from direct work costs

and the opportunity costs of lost home production). We employ the normalization $u_1(0) = v_1(0) = 0$. This lets the model incorporate the case in which the bad state is a fatal health shock (where $c_1^b = l_1^b = 0$). Additionally, we assume that the consumption utility and the labor disutility functions are well-behaved—i.e., that $u'_i(c_i^s) > 0$, $u''_i(c_i^s) < 0$, $v'_i(l_i^s) > 0$, and $v''_i(l_i^s) > 0$.

Household Preferences—We follow the collective approach to household behavior (Apps and Rees 1988; Chiappori 1988, 1992) and assume that household decisions are Pareto efficient. Therefore, denoting each member i 's Pareto weight by β_i , the household decisions can be characterized as solutions to the maximization of $\beta_1 U_1(c_1^s, l_1^s) + \beta_2 U_2(c_2^s, l_2^s)$.

7

Household Behavior—In the baseline static model where the household consumes its entire disposable income in each state of nature, there are no savings decisions involved, which we introduce in the dynamic extension to the model. Hence, in the current setting, the household's choices reduce to the labor supply and consumption allocation decisions. Formally, in each state s the household solves the problem:

$$\max U(c_1^s, l_1^s; c_2^s, l_2^s) \equiv \beta_1 U_1(c_1^s, l_1^s) + \beta_2 U_2(c_2^s, l_2^s) \text{ s.t. } c_1^s + c_2^s = y^s.$$

Optimal consumption allocation across spouses must satisfy $\beta_1 u'_1(c_1^s) = \beta_2 u'_2(c_2^s)$.

Additionally, in the intensive-margin model, the first-order condition with respect to the (interior) labor supply choice of the indirectly-affected member 2 satisfies:

$$u'_2(c_2^s) = \frac{v'_2(l_2^s)}{w_2(1 - \tau_2^s)}. \text{ Put together, the two optimality conditions imply that}$$

$$\beta_1 u'_1(c_1^s) = \beta_2 u'_2(c_2^s) = \frac{\beta_2 v'_2(l_2^s)}{w_2(1 - \tau_2^s)}.$$

This simple combination of optimality conditions with respect to consumption and labor supply choices will prove powerful for welfare analysis—it is the key source of the ability to map consumption utility to spousal labor disutility and hence of the identification of the gains from income insurance based on the indirectly-affected member's labor supply. Importantly, since we want to allow the directly-affected member 1 to be at a corner solution in state b ($l_1^b = 0$) due to, e.g., a severe disability that affects his or her ability to work, our analysis relies only on the labor supply responses of the indirectly-affected spouse.

⁷Note that with our normalization that $u_1(0) = v_1(0) = 0$, the household preferences reduce to the (weighted) utility from member 2's allocation when the bad state is member 1's death.

Spousal Labor Supply as Self-Insurance—At this point it is easy to see the self-insurance role of spousal labor supply responses to shocks. Define y_{-2}^s as the household's resources excluding those directly attributed to 2's labor supply decision—i.e., $y_{-2}^s \equiv A^s + \bar{z}_1^s(l_1^s) + B^s$ —such that the (exogenous) income loss from the shock is $L \equiv y_{-2}^g - y_{-2}^b$ (the gap in the spouse's unearned income across the two states). The household optimization conditions imply that the spouse's labor supply response to the shock $\frac{l_2^b}{l_2^g}$ is greater whenever the imposed income loss L is larger:

$$\frac{\partial \left(\frac{l_2^b}{l_2^g} \right)}{\partial L} = - \frac{\beta_i u_i''(c_i^b) \frac{\partial c_i^b}{\partial A^b}}{l_2^g \beta_2 v_2''(l_2^b) / w_2 (1 - \tau_2^b)} > 0, \quad (1)$$

when consumption in the bad state is a normal good ($\frac{\partial c_i^b}{\partial A^b} > 0$). Intuitively, when individuals experience shocks that cause them to decrease their labor supply and earn less income, their spouses can compensate for the associated income loss by increasing their own labor supply.⁸ Since the relative increase in spousal labor supply in response to shocks increases with the income loss, it can reveal the extent to which the household lacks formal insurance and needs to self-insure.⁹

3 Welfare Analysis: Implications for the Gains from Social Insurance

Using this setup, we provide in this section our main result and show how the gains from social insurance can be represented using only moments of spousal labor supply within our stylized framework. We do so in three steps. First, we derive the general formula for the gains from social income insurance—namely, the gap in marginal utilities of consumption across states of nature. Second, we briefly describe the consumption-based approach to identifying this gap. Third, we provide our alternative representation for these gains using spousal labor supply responses. To gain intuition, we derive the results within the intensive-margin model, where the comparison to the consumption-based approach is most straightforward, and we present the counterpart formula in the extensive-margin case later in the paper. Since we essentially rely only on individuals' optimality conditions in the context of the household, our approach to welfare analysis is general. We discuss the exact

⁸Of course, there are important non-financial linkages across spouses which we abstract from here since we want to highlight the financial channels that link household members. We return to mechanisms other than income loss that can drive spousal labor supply in response to shocks later in Section 4 where we discuss extensions that include state dependence and other forms of non-separabilities.

⁹Empirical work that finds evidence in support of the self-insurance role of spousal labor supply includes, among other studies, Stephens (2002) and Blundell et al. (2015) who find that wives' labor supply is an important consumption insurance device against permanent shocks to husbands' wages; Cullen and Gruber (2000) who study whether spousal labor supply is crowded out by unemployment insurance benefits and find a large crowd-out effect; Autor et al. (2015) who find similar crowd-out effects in the context of disability insurance; and Fadlon and Nielsen (2017), where we find a significant increase in survivors' labor supply following their spouse's death which is entirely driven by households that experience substantial income losses due to the shock.

identifying assumptions that underlie this approach at end of the current section. Important extensions and generalizations to the stylized model are considered in the next section.

In the design of optimal policies, the planner must weigh the gains against the associated costs. On the cost side, for example, transferring \$1 across states generates fiscal externalities that households impose on the government budget through their within-state behavioral responses to this policy change. In our case, the government's revenue could decrease since more generous social insurance will lead to decreases in spousal labor supply in the bad state. Identifying marginal costs is conceptually straightforward and much of the social insurance literature has focused on their estimation in different contexts. Therefore, we abstract from their analysis in this paper and focus only on the challenging task of identifying the gains from social insurance.

Planner's Problem—Denote the vector of tax rates on labor income by $T \equiv (\tau_1^g, \tau_2^g, \tau_1^b, \tau_2^b)$, and the vector of state-contingent benefits by $B \equiv (B^g, B^b)$. For the purpose of simplifying our formulas we analyze the case in which $\tau_2^g = \tau_2^b = \tau_2$, but the analysis readily extends to other cases. Let W^s denote the household's value function in state s such that $W^s \equiv \max U(c_1^s, l_1^s; c_2^s, l_2^s) \text{ s.t. } c_1^s + c_2^s = y^s$. Therefore, the household's expected utility is $J(B, T) \equiv \mu^g W^g + \mu^b W^b$. The social planner's objective is to choose the tax-and-benefit system that maximizes the household's expected utility subject to the requirement that expected benefits paid, $E(B^s) \equiv \mu^g B^g + \mu^b B^b$, equal expected taxes collected, $E\left(\sum_{i=1}^2 w_i \tau_i^s l_i^s\right) \equiv \mu^g (w_1 \tau_1^g l_1^g + w_2 \tau_2^g l_2^g) + \mu^b (w_1 \tau_1^b l_1^b + w_2 \tau_2^b l_2^b)$. Hence, the planner chooses the benefit levels B and taxes T that solve

$$\max_{B, T} J(B, T) \text{ s.t. } E(B^s) = E\left(\sum_{i=1}^2 w_i \tau_i^s l_i^s\right). \quad (2)$$

Welfare Gains from Social Insurance—What is the welfare gain from providing more generous benefits when the bad state occurs? To answer this question, consider transferring resources from the good state g to the bad state b through a small increase in, e.g., the tax rate τ_1^g to finance a balanced-budget increase in benefits in the bad state, B^b .

The social gain from this perturbation consists of the household's valuation of additional insurance. To construct a measure for this valuation, consider first the household's utility loss from the marginal increase in τ_1^g that finances the additional insurance. Within our

collective model, this loss is captured by $\left| \frac{\partial J(T, B)}{\partial \tau_1^g} \right| = \mu^g z_1^g \beta_i u'_i(c_i^g)$, as the household's income

in state g is reduced by z_1^g dollars, which are valued at $\beta_i u'_i(c_i^g)$ per dollar. Partially

differentiating the government's budget, this marginal increase in τ_1^g allows a balanced-

budget increase in B^b of the amount $\frac{\partial B^b}{\partial \tau_1^g} = \frac{\mu^g z_1^g}{\mu^b}$. The household's valuation per \$1 increase in B^b is given by $\frac{\partial J(T, B)}{\partial B^b} = \mu^b \beta_i u'_i(c_i^b)$, since it produces a value of $\beta_i u'_i(c_i^b)$ and is transferred to the household with probability μ^b . The utility gain from the increase in benefits when the shock occurs is, therefore, $\frac{\partial J(T, B)}{\partial B^b} \times \frac{\partial B^b}{\partial \tau_1^g} = \mu^g z_1^g \beta_i u'_i(c_i^b)$.¹⁰

Put together, the welfare benefits from a (balanced-budget) increase in B^b financed by an increase in τ_1^g are $\frac{\partial J(T, B)}{\partial B^b} \times \frac{\partial B^b}{\partial \tau_1^g} - \left| \frac{\partial J(T, B)}{\partial \tau_1^g} \right| = \mu^b z_1^g (\beta_i u'_i(c_i^b) - \beta_i u'_i(c_i^g))$. To gain cardinal interpretation for this expression, we follow the recent social insurance literature and normalize it by the welfare gain from decreasing the labor income tax rate in the good state, τ_1^g (Chetty and Finkelstein 2013). Overall, the normalized welfare benefit from our policy change is¹¹

$$MB \equiv \frac{\frac{\partial J(T, B)}{\partial B^b} \times \frac{\partial B^b}{\partial \tau_1^g} - \left| \frac{\partial J(T, B)}{\partial \tau_1^g} \right|}{\left| \frac{\partial J(T, B)}{\partial \tau_1^g} \right|} = \frac{u'_i(c_i^b) - u'_i(c_i^g)}{u'_i(c_i^g)}. \quad (3)$$

That is, the marginal welfare gain is captured by the insurance value of transferring resources from the good to the bad state, which is measured by the gap in either member's individual marginal utilities of consumption across the two states. This “rate of return” on shifting funds, which is zero in the first-best allocation in which marginal utilities are smoothed across states of nature for all household members, measures market inefficiency and quantifies the potential gain from government intervention. This expression mirrors the benefit side of Baily's (1978) and Chetty's (2006a) formula for the optimal level of social insurance applied to our household setting.

Consumption-Based Method for Identifying Welfare Gains—The main approach in the modern public finance literature for assessing the welfare gains from social insurance in (3) studies the consumption smoothing effects of government transfers. It can be transparently illustrated using the method that was developed by Baily (1978) and Chetty

¹⁰The partial differentiation of the government's budget allows us to focus on the gains from social insurance. To include the costs, we would analyze the total derivative $\frac{dB^b}{d\tau_1^g}$ which takes into account not only the required mechanical adjustments, but also the

households' behavioral responses to the policy change that have an impact on the government's budget (the “fiscal externality”). We illustrate that in our analysis of the dynamic participation model in the appendix.

¹¹Our analysis resembles the derivation of individuals' willingness to pay for additional unemployment insurance in Hendren (2016a). Deriving these welfare benefits can be also achieved by characterizing the first-order conditions of the planner's problem as in Chetty (2006a), Chetty and Finkelstein (2013), and Fadlon and Nielsen (2015).

(2006a) and was first implemented by Gruber (1997) in the context of unemployment insurance. Specifically, this method takes a quadratic approximation to the utility function to

represent the gap in marginal utilities as $MB \cong \gamma_i \times \left(\frac{c_i^g - c_i^b}{c_i^g} \right)$, where $\gamma_i = -\frac{u_i''(c_i^g)}{u_i'(c_i^g)}$ is $\hat{\gamma}$'s

coefficient of relative risk aversion evaluated at c_i^g . This formula evaluates fluctuations in the

consumption of goods across states, $\frac{c_i^g - c_i^b}{c_i^g}$ —which measure the degree of lack of

consumption smoothing—with the rate of change in the utility from marginal dollars, captured by the curvature of the utility function (γ)—which measures the utility cost of not smoothing consumption. Put differently, the benefits from insurance can be evaluated by the analysis of “quantity” fluctuations in consumption, which are then “priced” in utility terms.

Labor Supply Representation of Welfare Gains—We show next that simple, yet powerful, implications of the household's labor supply decisions allow us to rewrite the marginal benefit in (3) in terms of the indirectly-affected spouse's labor supply. By doing so, we show that the gains from additional insurance can be alternatively measured by evaluating changes in the consumption of the spouse's leisure instead of changes in the household members' consumption of goods. To build toward the proposition that summarizes this welfare result, which we formally state at the end of this subsection, we take two steps that correspond to the following two lemmas.

Lemma 1 (Spousal Labor Supply Responses to Shocks): The marginal benefit from raising B^b can be represented by

$$MB \cong \varphi \times \left(\frac{l_2^b - l_2^g}{l_2^g} \right), \quad (4)$$

where $\varphi \equiv \frac{v_2''(l_2^g)}{v_2'(l_2^g)}$.

Proof: The household's optimality conditions imply that *each* household member's (weighted) marginal utility from consumption can be mapped to the *spouse's* (weighted)

marginal disutility from labor, since $\beta_1 u_1'(c_1^s) = \beta_2 u_2'(c_2^s) = \frac{\beta_2 v_2'(l_2^s)}{w_2(1 - \tau_2^s)}$. This allows us to

represent the marginal benefit from social insurance by $MB = \frac{v_2'(l_2^b) - v_2'(l_2^g)}{v_2'(l_2^g)}$.¹² Intuitively, we

¹²Recall that we analyze the case of $\tau_2^g = \tau_2^b$ which simplifies the formula. The formula can be readily adjusted to other cases using the same equalities described in this proof.

use the household's optimality conditions to represent the degree to which households are able to smooth each member's marginal utility from consumption, $\frac{u'_i(c_i^b) - u'_i(c_i^g)}{u'_i(c_i^g)}$, using the degree to which they are able to smooth the marginal disutility from the spouse's labor, $\frac{v'_2(l_2^b) - v'_2(l_2^g)}{v'_2(l_2^g)}$. A quadratic approximation to member 2's labor disutility function around l_2^g yields the result.¹³

The intuition behind this lemma is as follows. We saw that the benefits from extra dollars of income insurance are measured by the relative utility gain from additional consumption of goods in the bad state compared to the good state. This additional insurance also decreases the need to compensate for the income loss associated with the shock through spousal labor supply as a self-insurance mechanism. In turn, it allows us to alternatively express the benefits from more generous social insurance by using the relative utility gain from additional consumption of spousal leisure. Put differently, the formula in Lemma 1 assesses the benefits of incrementally smoothing labor supply across states as a result of additional formal social insurance that reduces costly self-insurance.

Similar to the consumption representation, our approach evaluates benefits by multiplying the change in the "quantity" of spousal labor supply in response to shocks, $\frac{l_2^b - l_2^g}{l_2^g}$, by the rate of change in the spouse's disutility from additional work, ϕ , which captures the utility "price" of the labor supply quantity fluctuations across the two states. According to the formula, the welfare gains from additional benefits are higher whenever spousal responses to shocks are larger—that is, whenever the household's baseline ability to smooth the spouse's consumption of leisure across states of nature is lower. In the comparative statics of our model in equation (1) we saw that this quantity term increases with income losses and captures the self-insurance role of spousal labor supply. Therefore, intuitively, larger spousal labor supply responses—which correspond to a stronger need to self-insure—imply a greater scope for welfare-improving social insurance due to lack of adequate formal insurance. Similarly, the welfare gains from additional benefits are higher whenever ϕ is larger, as it implies that self-insurance through spousal labor supply (and the lack of smoothing spousal leisure) is more costly.

Analogous to analyzing fluctuations in the marginal utility of consumption that requires calibrating or estimating the utility parameter γ_i , analyzing fluctuations in the spouse's marginal disutility from labor requires calibrating or estimating the utility parameter ϕ . This can be done using the variety of tools that the literature has developed for identifying preference parameters (in the lab, using quasi-experiments, or estimating structural models).

¹³When the third order terms of spousal labor disutility are not small, the labor supply representation requires an additional term, analogous to the additional term that involves the coefficient of relative prudence in the consumption smoothing representation (Chetty 2006a).

Lemma 2 offers a straightforward technique for identifying φ , the curvature of the labor disutility function, using directly-estimable labor supply elasticities.

Before presenting Lemma 2, we need to introduce an additional concept from the family economics literature: the “sharing rule” interpretation of the collective model (Chiappori 1992). Under efficiency, one can think of the household’s decision within the assumptions of the stylized model as a two stage process. In the first stage, household members share non-labor income according to a given sharing rule that depends on their relative bargaining power; and in the second stage, each member optimally chooses his or her own labor supply and consumption. We denote this sharing rule in the following way: in each state s wealth and non-labor income are shared between the members such that $y_2^s \equiv \pi_2^s(w_1, w_2, A^s; B, T)$ is the amount received by member 2 and $y_1^s \equiv A^s + B^s - \pi_2^s(w_1, w_2, A^s; B, T)$ is the amount received by member 1. With these definitions, one can write 2’s program as

$$\max_{c_2^s, l_2^s} u_2(c_2^s) - v_2(l_2^s) \quad \text{s.t.} \quad c_2^s = y_2^s + \bar{w}_2 l_2^s, \quad (5)$$

where $\bar{w}_2 = w_2(1 - \tau_2)$ is the net-of-tax wage rate. We are now ready to present the second lemma:

Lemma 2 (Identification of φ): Define $\varepsilon(x, y) \equiv \frac{\partial x}{\partial y} \frac{y}{x}$. The parameter $\varphi \equiv \frac{v_2''(l_2^g)}{v_2'(l_2^g)} l_2^g$ can be mapped to spousal labor supply elasticities in the following way

$$\varphi = \frac{1 + \varepsilon(l_2^g, y_2^g) \frac{\bar{w}_2 l_2^g}{y_2^g}}{\varepsilon(l_2^g, \bar{w}_2) - \varepsilon(l_2^g, y_2^g) \frac{\bar{w}_2 l_2^g}{y_2^g}}.$$

Proof: The first-order conditions of the program (5) imply that $\bar{w}_2 u_2'(y_2^g + \bar{w}_2 l_2^g) = v_2'(l_2^g)$.

Partially differentiating the latter equation with respect to y_2^g and $\bar{w}_2$ yields

$$\frac{\partial l_2^g}{\partial y_2^g} = - \frac{\bar{w}_2 u_2''}{(\bar{w}_2)^2 u_2'' - v_2''} \quad \text{and} \quad \frac{\partial l_2^g}{\partial \bar{w}_2} = - \frac{u_2' + \bar{w}_2 l_2^g u_2''}{(\bar{w}_2)^2 u_2'' - v_2''}.^{14}$$

With some algebra, it follows that

$$\varphi \equiv \frac{v_2''(l_2^g)}{v_2'(l_2^g)} l_2^g = \frac{1 + \varepsilon(l_2^g, y_2^g) \frac{\bar{w}_2 l_2^g}{y_2^g}}{\varepsilon(l_2^g, \bar{w}_2) - \varepsilon(l_2^g, y_2^g) \frac{\bar{w}_2 l_2^g}{y_2^g}}.$$

The logic of the mapping in Lemma 2 follows a strategy similar to that of Chetty (2006b) for estimating risk aversion (i.e., the curvature of the consumption utility function). Here, we apply this logic to the curvature of the labor disutility function instead, and we extend it in a simple way to the household context using the collective approach and the sharing rule (decentralized) representation. The intuition for this result is that the extent to which a household member responds to changes in economic incentives (own wages and income) is directly linked to the rate at which preferences change (over labor or consumption).

The combination of Lemmas 1 and 2 leads to the following main proposition, which illustrates how the gains from social insurance can be assessed using solely moments of spousal labor supply within the stylized framework. We discuss its empirical implementation and required variation in Section 5.

Proposition 1 (Labor Supply Representation of Welfare Gains): The marginal benefit from raising B^b can be represented by

$$MB \cong \varphi \times \left(\frac{l_2^b - l_2^g}{l_2^g} \right), \quad (6)$$

$$\text{where } \varphi = \frac{1 + \varepsilon(l_2^g, y_2^g) \frac{\bar{w}_2 l_2^g}{y_2^g}}{\varepsilon(l_2^g, \bar{w}_2) - \varepsilon(l_2^g, y_2^g) \frac{\bar{w}_2 l_2^g}{y_2^g}}.$$

3.1 Discussion

Within-State Efficiency—The key identifying assumption that underlies our analysis is that household decisions are Pareto efficient. While individual rationality on its own implies that a household member's disutility from labor is equated to his or her marginal utility from consumption, household-level efficiency implies additionally that marginal utilities from consumption are equated across household members (with the appropriate Pareto weights). Hence, on the margin, all members of the household exhibit the same weighted returns from the consumption of additional resources, and any member not at a corner solution can reveal through labor supply responses the consumption preferences of each member of the household.

The assumption of household Pareto efficiency relies on the premise that when spouses have symmetric information about each other's preferences and consumption (because they interact on a regular basis) we would expect them to find ways to exploit any possibilities of Pareto improvements. Importantly, as emphasized by Browning, Chiappori, and Weiss

¹⁴Note the subtlety that we focus on partial derivatives of the spouse's behavior with respect to y_2^g and $\bar{w}_2$. In particular, y_2^g is held fixed when we change $\bar{w}_2$.

(2014), this does not preclude the possibility of power issues such that the allocation of resources within the household can depend on its members' respective Pareto weights. It simply assumes that no resources are left on the table.¹⁵

Since we emphasize the household as the operative economic unit, we have taken the collective approach. It has the important appealing feature of inherently accounting for a decision process that is taking place between household members, who can potentially hold different views on the decisions being made, and of allowing for intra-household distribution of powers that can matter for behavior as has been shown empirically. Moreover, as we allude to below, the collective approach addresses issues such as household members' ability to share risk among them (as in, e.g., Mazzocco 2004), which is central for our social insurance context. Note, however, that for a given household composition any approach to modeling the household's behavior and preferences which assumes efficiency (either explicitly or implicitly) would yield our results. Specifically, this includes the common unitary approach, which treats the family as one decision maker with a single utility function.

Cross-State Efficiency—Our analysis has assumed that risk is shared efficiently between household members. This implies that intra-household allocations are efficient both *ex-post* within a state of nature (as discussed above) and *ex-ante* across states of nature. In this respect, and in the context of spouses' relative bargaining power, Browning, Chiappori, and Weiss (2014) discuss the important distinction in the collective model between *ex-post* realizations and *ex-ante* distributions of income shocks. Specifically, when risk is shared efficiently, income realizations are pooled, so that one member's allocation should not suffer from his or her own bad luck. In contrast, the Pareto weights do depend on the *ex-ante* situations of spouses; so that, for example, if member 2 has larger expected income, one can expect that his or her Pareto weight will be larger, which would translate to higher levels of consumption. That is, under efficient risk sharing, these are *ex-ante* distributions, rather than *ex-post* realizations, that may affect the spouses' individual powers. Accordingly, our analysis has maintained the Pareto weights stable across states of nature and, as a result, baseline weights do not affect our welfare results for a given household composition. When the bargaining power does change across states (according to some particular process the researcher hypothesizes), or when the shock changes the household's composition (e.g., when state *b* is member 1's death), the corresponding adjustments are required to reflect the new weights on the household members' utilities.

Use of Optimality Conditions—Both the analysis of fluctuations in marginal utility of consumption (through the use of the envelope theorem) and the analysis of fluctuations in spousal marginal disutility from labor rely on the assumption that households make optimal choices. In fact, with this assumption, multiple representations for the gains from social insurance can be recovered using the marginal utilities of any single argument of the utility function since they are linked through the household's optimality conditions.¹⁶ This

¹⁵There are some cases in which the efficiency assumption fails (see discussion in Browning, Chiappori, and Weiss 2014). To model these cases, one would need to specify the underlying model of household decision making in order to identify one spouse's preferences from the other spouse's behavior.

flexibility allows researchers to use the representation most applicable given the available data and research tools (Chetty 2009; Chetty and Finkelstein 2013; Finkelstein et al. 2015). A main advantage of our proposed approach is the availability of large-scale and accurate data on labor market outcomes and the wide array of research tools the literature has developed for the analysis of household labor supply.

Compared to other optimization-based approaches, assessing the utility cost of consumption fluctuations does not preclude household members from being at a corner with respect to other choice variables. However, Proposition 1 for the intensive-margin model requires that spouses' optimization leads them to an interior solution in hours of work, so that the marginal disutility from spousal labor is linked to the marginal utility from consumption through an equality (see Finkelstein et al. 2015 for a related discussion). To allow for violations of this assumption, Proposition 2 (stated in the next section) provides the corresponding formula for labor force participation decisions. However, it is not assumption free: the requirement for identification in the stylized extensive-margin case is that the marginal entrant's value of labor disutility is interior to the support of the labor disutility distribution.¹⁷ Hence, identification would not be achieved for applications in which all spouses never work (e.g., due to significant labor market frictions) or in which all spouses work full-time prior to the shock.

4 Extensions and Generalizations

The stylized model that we analyzed is the simplest possible model that demonstrates our normative findings. However, as we mentioned earlier, the qualitative arguments that we made so far extend to much more general settings. In this section, we discuss some main variations and extensions to the simple model.

Extensive-Margin Model—In various cases, a participation model may better describe actual labor markets as it allows for common labor market frictions, such as hour requirements set by employers, which can limit employees' ability to optimize on the intensive margin.¹⁸ We describe here the setup of the extensive-margin counterpart of our simple static model and state the analogous welfare formula. Since the analysis tracks that of the intensive-margin model very closely, a complete investigation of the participation model, a proof of the corresponding formula, a discussion of its intuition, and extensions to the static case all appear in the appendix.¹⁹

¹⁶For example, in the context of health insurance, Finkelstein et al. (2015) develop a health-based approach for welfare analysis of Medicaid coverage that relies on evaluating the marginal health returns to out-of-pocket medical spending in different states of nature.

¹⁷In a search model of labor force participation (as in Appendix B.2), this is equivalent to having spouses at an interior solution for search effort.

¹⁸The choice of the appropriate model should depend on the context. In applications where there is strong evidence for such frictions, e.g., the Danish context as documented by Chetty et al. 2011, a participation model would be more appropriate. Indeed, in our own work on household responses to fatal and non-fatal severe health shocks using Danish administrative data (Fadlon and Nielsen 2017), we find that most of the estimated increases in surviving spouses' labor supply are concentrated on the extensive margin.

¹⁹The simple participation model of this section is most closely related to Kleven et al. (2009) and Immervoll et al. (2011), who study optimal taxation of couples with extensive-margin labor supply responses.

In the labor force participation version of our model $l_i^s = 1$ if i works and $l_i^s = 0$ otherwise. We adjust the household resource constraint so that $\bar{z}_i^s(l_i^s) = z_i^s \times (1 - \tau_i^s) \times l_i^s$, where z_i^s are gross earnings conditional on working, and τ_i^s are average tax rates. We also let $B^s(l_2^s)$ represent benefits from the government in state s as a function of l_2^s .²⁰ Individual preferences are adjusted such that we replace $v_i(l_i^s)$ with $v_i \times l_i^s$, where v_i represents each member i 's disutility from labor. The couple's disutilities from labor (v_1, v_2) are drawn from a continuous distribution defined over $[0, \infty) \times [0, \infty)$. We denote the marginal probability density function of v_2 by $f(v_2)$ and its cumulative distribution function by $F(v_2)$. For simplicity, we set $\beta_1 = \beta_2 = 1$, which is without loss of generality under our maintained assumption of efficient risk sharing. Together, these imply that in each state s the household maximizes $U(c_1^s, l_1^s; c_2^s, l_2^s) \equiv u_1(c_1^s) - v_1 \times l_1^s + u_2(c_2^s) - v_2 \times l_2^s$, subject to the budget constraint $c_1^s + c_2^s = y^s(l_2^s)$, where $y^s(l_2^s)$ is the household's overall income which depends on member 2's participation decision. For this reason, it is also the case that consumption bundles are a function of l_2^s , which we denote by $c_1^s(l_2^s)$ and $c_2^s(l_2^s)$ for members 1 and 2, respectively.

To characterize the welfare gain from providing more generous benefits when the bad state occurs, recall that government benefits to households, $B^s(l_2^s)$, can depend on both the state of nature and on spousal employment. Since our focus is on the value of insurance across states (as opposed to insurance within a state), we consider transferring resources from the good state g to the bad state b for a given choice of spousal employment. Specifically, we analyze a small decrease in the benefit $B^g(0)$ to finance a balanced-budget increase in $B^b(0)$.²¹ Similar to the intensive-margin case, the normalized welfare benefit from our policy change

is $MB = \frac{u'_1(c_1^b(0)) - u'_1(c_1^g(0))}{u'_1(c_1^g(0))}$. The labor supply representation of these welfare gains in the

extensive-margin model is summarized in the following proposition:

Proposition 2: Let $\bar{v}_2^s$ represent the labor disutility threshold such that the spouse works in state s if and only if $v_2 < \bar{v}_2^s$, and denote the spouse's probability of participation in state s by $e_2^s \equiv F(\bar{v}_2^s)$. The marginal benefit from raising $B^b(0)$ can be represented by

²⁰In this model we let benefits differ by 2's participation since it allows the welfare analysis to focus on the value from insurance across different states (rather than across different spousal employment statuses within states) in an analytically and conceptually simple way. Specifically, it enables us to analyze the optimality of insurance generosity across states of the world for given spousal employment.

²¹In the simple model, this perturbation concerns the distribution of benefits to low-income households across different states of nature. The analysis of other perturbations to the system will follow similar steps. We focus on this particular aspect of the policy since it captures the essence of insuring households against shocks in a mathematically simple way.

$$MB = \Phi \times \left(\frac{e_2^b}{e_2^g} \right) - 1, \quad (7)$$

where $\Phi \equiv \phi^b/\phi^g$, $\Phi \equiv \phi^b/\phi^g$, $\phi^s \equiv \frac{|\varepsilon(e_2^s, B^s(0))|}{B^s(0) \times f(v_2^s)}$, and $\varepsilon(e_2^s, B^s(0))$ is the spouse's participation elasticity with respect to the policy tool $B^s(0)$.

Proof: See Appendix A.

We discuss in the appendix how both Propositions 1 and 2 are based on capturing the value of additional social insurance using the value of additional spousal leisure. In broad terms, in Proposition 1 this value is given by the foregone welfare cost from the disutility of marginal hours of work, and in Proposition 2 it is given by the forgone welfare cost from the disutility of the marginal worker.²²

Before we proceed, it is worth highlighting an important advantage of the participation model, which relates to issues we discuss below. In particular, whenever individual preferences have a separable labor disutility component, the welfare formula remains the same. This means that individuals' consumption preferences, which we assumed to be $u_i(c_i^s)$, can take a much more general form that flexibly incorporates consumption-leisure complementarities, $u_i(c_i^s, l_i^s)$, as well as a variety of complementarities across spouses' allocations of both consumption and time, $u_i(c_1^s, l_1^s; c_2^s, l_2^s)$.

Household Resource Constraint and Consumption-Technology Economy of Scale—Our model accommodates income streams that can depend on the state of nature and on employment choices, and can therefore flexibly incorporate any state and employment contingent private or social insurance payments. In a similar fashion, it is also possible to allow for income-tested (and age-dependent) transfers and taxes and more complex taxation schemes. For example, one can account for non-linear income taxation and differential tax rules for joint filing by studying a general state-dependent function that maps labor supply into household-level net-of-tax earnings: $\bar{z}^s(l_1^s, l_2^s)$.

So far, however, we have abstracted from potentially important economies of scale in the household's consumption technology. To allow for household public goods and economies of scale in a straightforward and general way, we follow Browning et al. (2013) and introduce an arbitrary technology, G^s , that transforms income into consumption.²³ That is,

²²Note that Proposition 2 shows in one step how the marginal net benefit from social insurance can be expressed using different moments of the spouse's labor supply in a model of labor force participation. Specifically, the formula directly maps the price component Φ to labor supply elasticities. In the extensive-margin model the price term also requires calibrating the labor disutility

distribution (specifically, the ratio $\frac{f(v_2^g)}{f(v_2^b)}$, which can be done in a variety of different ways as we discuss in the appendix.

we can rewrite the household state-specific resource constraint as $c_1^s + c_2^s = G^s(y^s)$, where each member's c_i^s is measured in *unobservable* “private good equivalent” units.²⁴ Since G^s is state specific, it also accounts for potential household composition changes as a result of the shock (e.g., when state b represents a fatal health shock), as well as any post-shock changes in the household's technology (e.g., in the degree or nature of home production). With this specification, the formula for the welfare gains from a \$1 increase in insurance becomes

$$MB = \frac{G^{b'}(y^b)u'_i(c_i^b) - G^{g'}(y^g)u'_i(c_i^g)}{G^{g'}(y^g)u'_i(c_i^g)},$$

since the benefit from a \$1 transfer must be scaled by the

amount of consumption units that it produces. This highlights a key challenge in the analysis of consumption: the researcher must identify the function G^s in order to translate income or overall expenditure into individual consumption in household settings, which entails strong assumptions or complex estimations (Browning et al. 2013). The labor supply representation, however, remains the same since we can still write the marginal gains as

$$MB = \frac{v'_2(l_2^b) - v'_2(l_2^g)}{v'_2(l_2^g)},$$

which tremendously simplifies the analysis when economies of scale are

considered. This significant advantage of studying (spousal) labor supply instead of consumption arises from the fact that in household settings individual consumption is not directly assignable and observable while labor supply is.

Dynamics over the Life-Cycle and General Choice Variables—In the context of social insurance over the life-cycle, it is important to consider households' self-insurance through ex-ante mechanisms such as precautionary savings. In Appendix B, we analyze a fully-dynamic life-cycle model. This model allows for endogenous savings, as well as private and informal insurance arrangements.

Generally speaking, our formulas extend to this model with the adjustment that spousal labor supply in different states of nature are replaced with their averages taken over the periods households spend in each state (analogous to the dynamic consumption formula in Chetty 2006a and Chetty and Finkelstein 2013).²⁵ Note that even in the presence of ex-ante responses in expectation of shocks, these are still the ex-post responses to shocks that assess the gains from social insurance policies that condition benefits on shock realizations. The intuition behind this result is as follows: when forward-looking households make adjustments in anticipation of shocks (according to their expectations), their responses after the shock is experienced recover its residual uninsured risk. This leftover risk assesses the

²³This is also similar to Becker's (1965) model of household production and more recent applications and extensions to his framework, such as Kleven (2004) and Kleven and Kreiner (2007).

²⁴One important and widely used element involved in converting expenditure data into private consumption is equivalence scales. However, despite their practical importance and the sensitivity of welfare analysis to them, the main equivalence scale estimates are ad-hoc and not theoretically based (such as the modified OECD equivalence scale of 0.67 and the square-root scale of 0.71). A recent example for model-based estimates for adult equivalence scales, is Browning, Chiappori, and Lewbel (2013), who find non-negligible differences across genders. Under the assumption of equal sharing of income among the two spouses, their scale estimates are 0.80 for males and 0.72 for females.

²⁵The dynamics of the life-cycle analysis likewise enter the marginal costs of social insurance. A household in state g not only decreases its labor supply due to higher taxes in the present, but also in response to increased benefits in the hitherto unencountered state b . The prospect of higher benefits in the case that the household experiences a shock lowers its need to save for that scenario, which translates into a decrease in labor supply in state g .

insurance gap that the government can potentially fill. Put differently, the causal ex post impact of shocks (on either consumption or spousal labor supply) are the moments required for identifying households' willingness to pay for benefits (see discussions on these issues in Chetty 2008 and Hendren 2016a). We return to this point in Section 5 where we discuss the empirical implementation of our approach.

Both the static and the dynamic models can additionally incorporate a general class of arbitrary choice variables—such as time investment in home production (similar to Chetty 2006a, Chetty and Finkelstein 2013, and Finkelstein et al. 2015). The generality of our analysis to inclusion of additional choice variables and response margins stems from the fact that our results are derived using optimality conditions, which map each member's consumption utility to spousal labor disutility, and using the envelope theorem. Since these conditions hold in more complex models that maintain the efficiency assumption, the economic forces that underlie the assessment of the gains from social insurance using spousal labor supply remain similar in more general settings.²⁶

State-Dependence and other forms of Non-Separabilities—Besides income losses, there are other important ways in which households can be directly affected by shocks. Specifically, in some applications, it is important to allow household members' preferences to change across states of nature. For example, when the bad state is member 1's disability, the household may value less activities such as traveling (a consumption utility state dependence), and member 2's utility loss from time spent away from home may increase because he or she would like to take care of their sick spouse (a labor disutility state dependence). Similarly, when the bad state is member 1's death, working may become more costly if the surviving spouse experiences depression and has difficulties working; and, conversely, working may become more desirable if the surviving spouse feels lonely and wishes to seek social integration or if he or she no longer has to care for an ill spouse. Another example is when the bad state is member 1's unemployment that is accompanied by consumption drops, where state dependence would be present if there are consumption-leisure complementarities.

When such dependencies or complementarities are present, the simple formulas no longer hold. We therefore show next how to extend our approach to allow preferences to be state dependent in different dimensions. First, note that consumption utility state dependence would not alter our formulas since we mapped the identification of welfare gains from consumption to labor supply. Similarly, state dependence in the directly-affected member's labor disutility—e.g., due to a severe health shock—would not affect the results which rely on spousal labor supply. This is, indeed, one of the main motives for focusing on the indirectly-affected spouses. The type of state dependence that would affect the formulas is with respect to the indirectly-affected spouse's disutility from labor.

²⁶More precisely, when the additional arguments enter the household's utility in a way that is separable from labor supply the same formulas go through. When there are dependencies in the utility function across these arguments and labor supply, spousal marginal labor disutility still directly maps into marginal utilities from consumption, but the mapping to labor supply behavior requires adjustments. These adjustments are identical to those described below when utility is state dependent. Related analyses and alternative assumptions in the inclusion of general choice variables can be found in Chetty (2006a) in the context of welfare evaluations of social income insurance using consumption and in Finkelstein et al. (2015) in the context of welfare evaluations of health insurance using various optimization-based methods.

To illustrate how this sort of state dependence enters our formulas, let us analyze the intensive-margin model with state-dependent preferences of the form

$U^S(c_1^S, l_1^S; c_2^S, l_2^S) = u_1^S(c_1^S) - v_1^S(l_1^S) + u_2^S(c_2^S) - v_2^S(l_2^S)$. We provide an illustration for the extensive-margin model in Appendix C. With these preferences, the formula in Proposition 1 becomes:

$$MB \cong \left[\theta_l \times \varphi \times \left(\frac{l_2^b - l_2^g}{l_2^g} \right) \right] + [\theta_l - 1], \text{ where } \varphi = \frac{v_2^{g''}(l_2^g)}{v_2^{g'}(l_2^g)} l_2^g \text{ similar to before, and } \theta_l = \frac{v_2^{b'}(l_2^g)}{v_2^{g'}(l_2^g)}$$

measures the degree of state dependence by evaluating the extent to which the marginal cost of spousal labor supply (or the marginal value of spousal leisure) varies across states of nature starting from equal levels of labor supply.²⁷

Compared to Proposition 1, the first bracketed term adjusts the “price” component of self-insurance through spousal labor supply from φ to $\theta_l \times \varphi$. To gain intuition, consider the case in which $\theta_l > 1$ so that spousal leisure is more valuable in the bad state. Since the formula assesses benefits from social insurance by evaluating the gains from the consumption of leisure, this would make the transfer of resources from state g to state b more socially desirable. The reason is that it would allow for more leisure in the bad state, in which it is valued more highly, by decreasing the need to compensate for the associated income loss through spousal labor supply. The additional component in the second bracketed term captures welfare considerations that are beyond income losses, and may be present even

when households are well-insured and there is no self-insurance ($\frac{l_2^b - l_2^g}{l_2^g} = 0$). For example,

when $\theta_l > 1$ the planner would want to allow for more leisure in the bad state simply because it is valued more highly on the margin in that state (and the opposite when $\theta_l < 1$). That is, this term captures the value of insurance against “utility shocks” rather than income shocks that are associated with state transitions. Note that θ_l is a composite parameter that includes direct changes in spouses’ marginal disutility from labor as a result of a shock but also indirect changes. For example, in the presence of consumption-leisure utility complementarities, the marginal disutility from labor can vary across states of nature due to changes in consumption and the utility dependence between consumption and leisure.

Overall, state dependence affects any analysis that aims to make quantitative welfare statements, whether it relies on analyzing consumption or analyzing labor supply and whether it uses reduced-form or structural estimation techniques. There is no consensus in the literature on the magnitude (or even sign) of utility state dependence in the context of the

²⁷Evidently, similar adjustments to the welfare formulas are required in the consumption-smoothing approach (see Chetty and Finkelstein 2013). For completeness, we report the adjusted consumption-based formula in the context of our model. In the intensive-

margin model the marginal gains from social insurance become $MB \cong \left[\theta_{c_i} \times \gamma_i \times \left(\frac{c_i^g - c_i^b}{c_i^g} \right) \right] + [\theta_{c_i} - 1]$, where

$\gamma_i = -\frac{u_i^{g''}(c_i^g)}{u_i^{g'}(c_i^g)} c_i^g$ and $\theta_{c_i} = \frac{u_i^{b'}(c_i^b)}{u_i^{g'}(c_i^g)}$. Chetty and Finkelstein (2013) provide an equivalent version of this formula when state

dependence is defined differently as $\frac{u_i^{b'}(c_i^b) - u_i^{g'}(c_i^b)}{u_i^{g'}(c_i^g)}$.

significant shocks that social insurance programs aim to protect against (such as disability and unemployment). In practice, many empirical normative studies ignore this important aspect since the identification of state dependence has proven very challenging. Some papers address this challenge by offering general identification strategies, estimating the degree of state dependence in different applications, or developing tests for the presence of such dependencies in their particular empirical contexts in order to provide more grounded welfare statements.²⁸ An example for the latter can be found in Fadlon and Nielsen (2017), where we provide an intuitive test for the presence of a certain leading hypothesized type of state dependence in the context of spousal death. In that paper we are specifically interested in testing the hypothesis that the increase in survivors' labor supply can be attributable to lower costs of supplying labor following the death of a spouse, due to loneliness and the desirability of social integration or because the survivor no longer has to care for an ill spouse. In Appendix D we additionally illustrate with a simple example, using our theoretical framework, how one can reach conclusions about the relative role of different state dependence dimensions based on estimations of responses to shocks.²⁹

Heterogeneity—Extending the underlying logic of the dynamic generalization, it is possible to also allow for heterogeneity across households in elements of the resource constraint—for example, in the income loss that they experience or their degree of insurance. In this case, both our intensive and the extensive margin formulas for changes in the generosity of universal benefits would adjust to include the population average responses. See similar arguments and analyses within the consumption-based approach in, e.g., Chetty (2006a) for unemployment benefit levels and Kolsrud et al. (2016) for unemployment benefit timing.

However, when there is also heterogeneity in preferences across households, aggregation of utilities is much more complicated and requires additional assumptions or adjustments to the welfare formulas (Chetty 2006a). Related analysis within the consumption approach can be found in Andrews and Miller (2013) for the optimal levels of unemployment benefits and in Kolsrud et al. (2016) for the optimal timing. Andrews and Miller's (2013) analysis of optimal levels (that identifies the covariance between risk aversion and the consumption drop as an additional component of the welfare formula) can be applied to our intensive-margin model if arguments regarding consumption and consumption utility are modified to arguments regarding spousal labor supply and labor disutility. For the participation model, one possible approach is to apply to our dynamic search model for household extensive-margin labor supply assumptions regarding heterogeneity similar to those in Chetty (2008) and Landaís (2015). In their analysis of dynamic search models in the context of unemployment insurance, they discuss different independence assumptions across preference or behavior moments. In Appendix E we provide an example for heterogeneity in preferences (specifically, in state dependence) and derive the resulting formula under such an independence assumption.

²⁸See, for example, Viscusi and Evans (1990), Evans and Viscusi (1991), Lillard and Weiss (1997), Rust and Phelan (1997), Sloan et al. (1998), De Nardi et al. (2006), Edwards (2008), Finkelstein et al. (2013), Ball and Low (2014), and Low and Pistaferri (2015).

²⁹This example is based on the estimated household responses to non-fatal health shocks in Fadlon and Nielsen (2017), where we find small declines in spousal earnings even though households are not perfectly insured.

5 Implementation

Before we conclude, we offer in this section a brief guide for how to empirically implement our method. The formula in Proposition 1 shows that a full identification of the gains from social insurance using spousal labor supply requires two sets of responses: across states of nature and within states of nature. The two lemmas conveniently break down the formula into these two sets of household labor supply moments and point to the different sources of variation and estimation techniques required for their identification. We discuss the two sets of moments successively. While our formulas do not take a stance regarding how researchers should estimate these moments (e.g., using structural or reduced-form approaches), our following discussion on implementation focuses on reduced-form methods that rely on quasi-experimental variation.

Lemma 1, which constitutes the primary portion of our welfare result, highlights the important role of spousal labor supply responses to shock *realizations*, $\frac{l_2^b - l_2^g}{l_2^g}$, or the “shock

elasticity”. As we emphasized above, even when households can make behavioral adjustments in anticipation of shocks, the moments that enter the formula are the *ex-post* responses to actual realizations. This means that, empirically, when ex-ante responses are possible, one has to carefully choose research designs that cleanly recover the causal ex-post impact of shocks. In the presence of complex dynamics, this could be a challenging task. In such cases, identification requires constructing counterfactuals that account for life-cycle and time patterns in family labor supply which, among many other factors, are likely to depend on ex-ante expectations. Some papers have successfully done so in various contexts by using matched control groups from the pool of untreated units based on observables.

However, there are important classes of applications relevant for our analysis, where strategies that rely on unaffected households as controls may be inadequate. In particular, in the context of household labor supply responses to fatal and severe non-fatal health shocks, we show in Fadlon and Nielsen (2017) that affected and observably-similar unaffected households exhibit substantially different behavioral patterns over time, in violation of the requirement of parallel pre-trends across the two groups. We therefore take a different approach in that paper which relies on the common notion that the timing of shocks within a short period of time may be as good as random. Specifically, we employ a quasi-experimental research design that constructs counterfactuals to affected households by using households that experience the same shock but a few years in the future. The identifying assumption is that, absent the realization of the shock, the outcomes of the treatment and control groups would run parallel (which can be evaluated based on pre-trends). Among many other dimensions, this ensures the similarity of the groups in terms of their expectations. Importantly, this empirical strategy allows for behavioral adjustments by treated households in expectation of a shock, but estimates ex-post responses to the realization of the shock as necessary for the welfare formula.³⁰

³⁰Note that our analysis, which explicitly accounts for strategic intra-household interactions using the collective approach, has focused on shocks whose particular timing is likely unpredictable and whose financial consequences are potentially severe. We have done so since these are the shocks for which social insurance is mostly justified, as households’ ability to efficiently self-insure against them is particularly limited. In a concurrent paper, Hendren (2016a) studies within the unitary framework how spousal labor supply

Lemma 2 completes the characterization of the required moments. It lists the within-state labor supply responses that the researcher needs to estimate, which are the traditional elasticities with respect to (own) income and net-of-tax wages. The reduced-form literature has used a variety of sources of variation to identify these elasticities. For example, to estimate labor supply responses to liquidity from income or wealth, the literature has exploited lottery winnings, variation in severance pay policies and schedules, the timing of EITC refunds, and kinks in the schedule of unemployment insurance benefits (Imbens et al. 2001; Card et al. 2007; Chetty 2008; LaLumia 2013; Landais 2015; Cesarini et al. 2017); and, to estimate labor supply responses to net-of-tax wages, exploiting variation from tax reforms or discontinuous incentives in the income tax schedule has become a common practice (see, e.g., Chetty et al. 2013 for a meta analysis of extensive margin responses). Note that, in our context, the researcher should consider the relevant population for which estimates are required. For example, if one is interested in analyzing the gains from insurance against health shocks, and if husbands typically experience these shocks prior to their wives, then the mean labor supply elasticities to be used in the formula should reflect the gender composition of the population of indirectly-affected spouses.

6 Conclusion

This paper develops a new approach to welfare analysis of social insurance in general household settings using information from the labor market. We have shown how, when households make optimal choices, household labor supply behavior can be used to draw implications for the gains from social insurance. Intuitively, our analysis has illustrated that the degree to which households self-insure through spousal labor supply in response to shocks and differentially respond to economic incentives across states of nature reveal their lack of formal insurance and, hence, the scope for more generous government benefits in bad states of the world. Doing so, we have also highlighted the important welfare relevance of the prevalent empirical work on household labor supply and its self-insurance role in different settings.

Using labor supply rather than consumption data for welfare evaluations involves significant advantages due to the wide availability of large-scale and precise data on household labor market outcomes and due to the ability to directly assign labor supply behavior to individual members of a household. Our approach offers a way to exploit these advantages, using either reduced-form or structural estimation techniques, in a general class of shocks and social insurance schemes. These include important applications such as fatal and non-fatal severe health shocks and long run unemployment. They also include less traditional but increasingly important applications, such as assessing the value of protecting against pension-wealth losses in private savings accounts (relevant for the debate on the privatization of Social Security) and evaluating the welfare losses caused by discontinuation

responses to variation in anticipation to unemployment shocks map to households' willingness to pay for unemployment insurance. Unlike the class of unpredictable shocks that we focus on, the applicability of his approach relies on households' anticipation to shocks, on revelation of new information pertaining to the likelihood of experiencing the shock, and on the researcher's ability to empirically observe or impute such variations. Nevertheless, its reliance on ex-ante responses inherently incorporates state dependence.

of employee compensations, such as health insurance, life insurance, and employer matching in retirement savings.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Aguiar M, Hurst E. 2005; Consumption versus expenditure. *Journal of Political Economy*. 113(5):919–948.
- Andrews I, Miller C. 2013 Optimal social insurance with heterogeneity.
- Apps PF, Rees R. 1988; Taxation and the household. *Journal of Public Economics*. 35(3):355–369.
- Ashenfelter O. 1980; Unemployment as disequilibrium in a model of aggregate labor supply. *Econometrica*. 48(3):547–564.
- Autor D, Kostøl A, Mogstad M. 2015 Disability benefits, consumption insurance, and household labor supply.
- Autor DH, Duggan MG. 2007; Distinguishing income from substitution effects in disability insurance. *American Economic Review*. 97(2):119–124.
- Baily MN. 1978; Some aspects of optimal unemployment insurance. *Journal of Public Economics*. 10(3):379–402.
- Ball S, Low H. 2014; Do self-insurance and disability insurance prevent consumption loss on disability? *Economica*. 81(323):468–490.
- Becker GS. 1965 A theory of the allocation of time. *Economic Journal*. :493–517.
- Bee, A, Meyer, BD, Sullivan, JX. Improving the Measurement of Consumer Expenditures. University of Chicago Press; 2013. The validity of consumption data: Are the consumer expenditure interview and diary surveys informative?.
- Blundell, R, Dias, MC, Meghir, C, Shaw, JM. Working Paper 19007. National Bureau of Economic Research; 2013. Female labour supply, human capital and welfare reform.
- Blundell R, Lewbel A. 1991; The information content of equivalence scales. *Journal of Econometrics*. 50:49–68.
- Blundell R, Pistaferri L, Saporta-Eksten I. 2015 Consumption inequality and family labor supply. *American Economic Review*.
- Browning M, Chiappori PA, Lewbel A. 2013; Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *Review of Economic Studies*. 80(4):1267–1303.
- Browning, M, Chiappori, P-A, Weiss, Y. Economics of the Family. Cambridge University Press; 2014.
- Browning M, Leth-Petersen S. 2003; Imputing consumption from income and wealth information. *Economic Journal*. 113(488):F282–F301.
- Card D, Chetty R, Weber A. 2007; Cash-on-hand and competing models of intertemporal behavior: New evidence from the labor market. *Quarterly Journal of Economics*. 122(4):1511–1560.
- Cesarini D, Lindqvist E, Notowidigdo MJ, Östling R. 2017 The effect of wealth on individual and household labor supply: evidence from swedish lotteries. *American Economic Review*.
- Chetty R. 2006a; A general formula for the optimal level of social insurance. *Journal of Public Economics*. 90(10–11):1879–1901.
- Chetty R. 2006b; A new method of estimating risk aversion. *American Economic Review*. 96(5):1821–1834.
- Chetty R. 2008; Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy*. 116(2):173–234.
- Chetty R. 2009; Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*. 1:451–488.

- Chetty, R, Finkelstein, A. Volume 5 of *Handbook of Public Economics*. Elsevier; 2013. Chapter 3 - social insurance: Connecting theory to data; 111–193.
- Chetty R, Friedman JN, Olsen T, Pistaferri L. 2011; Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records. *Quarterly Journal of Economics*. 126(2):749. [PubMed: 21836746]
- Chetty R, Guren A, Manoli D, Weber A. 2013; Does indivisible labor explain the difference between micro and macro elasticities? a meta-analysis of extensive margin elasticities. *NBER Macroeconomics Annual*. 27(1):1–56.
- Chiappori P-A. 1988; Rational household labor supply. *Econometrica*. 56(1):63–90.
- Chiappori PA. 1992; Collective labor supply and welfare. *Journal of Political Economy*. 100(3):437–467.
- Coile, CC. Working Paper 10810. National Bureau of Economic Research; 2004. Health shocks and couples' labor supply decisions.
- Cullen JB, Gruber J. 2000; Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Economics*. 18(3):546–572.
- De Giorgi, G, Frederiksen, A, Pistaferri, L. Working Paper 22357. National Bureau of Economic Research; 2016. Consumption network effects.
- De Nardi, M, French, E, Jones, JB. Working Paper 12554. National Bureau of Economic Research; 2006. Differential mortality, uncertain medical expenses, and the saving of elderly singles.
- Dobkin, C, Finkelstein, A, Kluender, R, Notowidigdo, MJ. Working Paper 22288. National Bureau of Economic Research; 2016. The economic consequences of hospital admissions.
- Edwards RD. 2008; Health risk and portfolio choice. *Journal of Business & Economic Statistics*. 26(4):472–485.
- Evans WN, Viscusi WK. 1991 Estimation of state-dependent utility functions using survey data. *Review of Economics and Statistics*. :94–104.
- Fadlon, I, Nielsen, TH. Working Paper 21352. National Bureau of Economic Research; 2015. Household responses to severe health shocks and the design of social insurance.
- Fadlon I, Nielsen TH. 2017 Family labor supply responses to severe health shocks.
- Finkelstein, A, Hendren, N, Luttmer, EF. Working Paper 21308. National Bureau of Economic Research; 2015. The value of medicaid: Interpreting results from the oregon health insurance experiment.
- Finkelstein A, Luttmer EF, Notowidigdo MJ. 2013; What good is wealth without health? the effect of health on the marginal utility of consumption. *Journal of the European Economic Association*. 11:221–258.
- Gruber J. 1997; The consumption smoothing benefits of unemployment insurance. *American Economic Review*. 87(1):192–205.
- Haan, P, Prowse, V. Discussion Paper 8980. IZA Institute for the Study of Labor; 2015. Optimal social assistance and unemployment insurance in a life-cycle model of family labor supply and savings.
- Heckman JJ, Macurdy TE. 1980; A life cycle model of female labour supply. *Review of Economic Studies*. 47(1):47–74.
- Heckman JJ, Macurdy TE. 1982; Corrigendum on a life cycle model of female labour supply. *Review of Economic Studies*. 49(4):659–660.
- Hendren N. 2016a Knowledge of future job loss and implications for unemployment insurance.
- Hendren N. 2016b; The policy elasticity. *Tax Policy and the Economy*. 30(1):51–89.
- Imbens GW, Rubin DB, Sacerdote BI. 2001 Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players. *American Economic Review*. :778–794.
- Immervoll H, Kleven HJ, Kreiner CT, Verdelin N. 2011; Optimal tax and transfer programs for couples with extensive labor supply responses. *Journal of Public Economics*. 95(11):1485–1500.
- Kleven HJ. 2004; Optimum taxation and the allocation of time. *Journal of Public Economics*. 88(3): 545–557.
- Kleven HJ, Kreiner CT. 2007; Optimal taxation of married couples with household production. *FinanzArchiv: Public Finance Analysis*. 63(4):498–518.

- Kleven HJ, Kreiner CT, Saez E. 2009; The optimal income taxation of couples. *Econometrica*. 77(2): 537–560.
- Koijen, R, Van Nieuwerburgh, S, Vestman, R. Improving the Measurement of Consumer Expenditures. University of Chicago Press; 2014. Judging the quality of survey data by comparison with “truth” as measured by administrative records: Evidence from sweden.
- Kolsrud J, Landais C, Nilsson P, Spinnewijn J. 2016 The optimal timing of unemployment benefits: Theory and evidence from sweden.
- Kostøl A, Mogstad M. 2015; Earnings, disposable income, and consumption of allowed and rejected disability insurance applicants. *American Economic Review*. 105(5):137–41. [PubMed: 29543409]
- Kreiner, CT, Lassen, DD, Leth-Petersen, S. Working paper. University of Copenhagen; 2014. Liquidity constraint tightness and consumer responses to fiscal stimulus policy.
- Kroft K, Lange F, Notowidigdo MJ. 2013; Duration dependence and labor market conditions: Evidence from a field experiment. *Quarterly Journal of Economics*. 128(3):1123–1167.
- LaLumia S. 2013; The eite, tax refunds, and unemployment spells. *American Economic Journal: Economic Policy*. 5(2):188–221.
- Landais C. 2015; Assessing the welfare effects of unemployment benefits using the regression kink design. *American Economic Journal: Economic Policy*. 7(4):243–278.
- Lillard LA, Weiss Y. 1997; Uncertain health and survival: Effects on end-of-life consumption. *Journal of Business & Economic Statistics*. 15(2):254–268.
- Low H, Pistaferri L. 2015; Disability insurance and the dynamics of the incentive insurance trade-off. *American Economic Review*. 105(10):2986–3029.
- Lundberg S. 1985; The added worker effect. *Journal of Labor Economics*. 3(1):11–37.
- Maloney T. 1987; Employment constraints and the labor supply of married women: A reexamination of the added worker effect. *Journal of Human Resources*. 22(1):51–61.
- Maloney T. 1991; Unobserved variables and the elusive added worker effect. *Economica*. 58(230): 173–187.
- Mazzocco M. 2004; Saving, risk sharing, and preferences for risk. *The American Economic Review*. 94(4):1169–1182.
- Meyer, BD, Mok, WK. Working Paper 18869. National Bureau of Economic Research; 2013. Disability, earnings, income and consumption.
- Olsson M, Thoursie PS. 2015; Sickness insurance and spousal labour supply. *Labour Economics*. 33:41–54.
- Pistaferri L. 2015; Household consumption: Research questions, measurement issues, and data collection strategies. *Journal of Economic and Social Measurement*. 40(1–4):123–149.
- Rust J, Phelan C. 1997 How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica*. :781–831.
- Shimer R, Werning I. 2007; Reservation wages and unemployment insurance. *Quarterly Journal of Economics*. 122(3):1145–1185.
- Sloan FA, Viscusi WK, Chesson HW, Conover CJ, Whetten-Goldstein K. 1998; Alternative approaches to valuing intangible health losses: the evidence for multiple sclerosis. *Journal of Health Economics*. 17(4):475–497. [PubMed: 10180927]
- Spletzer JR. 1997; Reexamining the added worker effect. *Economic Inquiry*. 35(2):417–427.
- Stephens MJ. 2002; Worker displacement and the added worker effect. *Journal of Labor Economics*. 20(3):504–537.
- Viscusi WK, Evans WN. 1990 Utility functions that depend on health status: estimates and economic implications. *American Economic Review*. :353–374.
- Wang H. 2016 Intra-household risk sharing and job search over the business cycle.