Women's empowerment and family health: Estimating LATE with mismeasured treatment

Authors: Rossella Calvi, Arthur Lewbel, Denni Tommasi

Boston College Working Papers in Economics, 2018

Originally posted on: http://ideas.repec.org/p/boc/bocoec/959.html

Women's Empowerment and Family Health: Estimating LATE with Mismeasured Treatment

Rossella Calvi*[1], Arthur Lewbel†[2], and Denni Tommasi‡[3]

[1]Department of Economics, Rice University

[2]Department of Economics, Boston College

[3]ECARES, Université Libre de Bruxelles

May 2018

Abstract

We study the causal effect of women's empowerment on family health in India. We define treatment as a woman having primary control over household resources and use changes in inheritance laws as an instrument. Due to measurement difficulties and sharing of goods, treatment cannot be directly observed and must be estimated using a structural model. Treatment mismeasurement may therefore arise from model misspecification and estimation errors. We provide a new estimation method, MR-LATE, that can consistently estimate local average treatment effects when treatment is mismeasured. We find that women's control of substantial household resources improves their and their children's health.

JEL Codes: D13, D11, D12, C31, I32.

Keywords: causality, LATE, structural model, collective model, resource shares, bargaining power, health.

*Rice University, Backer Hall 274, 6100 Main Street, Houston, TX 77005, USA (e-mail: rossella.calvi@rice.edu).

†Boston College, Maloney Hall 315, 140 Commonwealth Avenue, Chestnut Hill, MA 02467, USA (e-mail: arthur.lewbel@bc.edu).

‡ECARES, Avenue F. D. Roosevelt 50, CP 114, B-1050 Brussels, Belgium (e-mail: dtommasi@ulb.ac.be).

We thank Samson Alva, Erich Battistin, Jonathan de Quidt, Bram De Rock, Laurens Cherchye, Flávio Cunha, Frank DiTraglia, Yingying Dong, Jeremy Fox, Mette Gørtz, Rachel Heath, Nikolaj Harmon, Søren Leth-Petersen, Benjamin Solow, Takuya Ura, Kenneth Wolpin, Yi Zhang, seminar participants at University of Copenhagen, ECARES, Rice University, Tilburg University, Stockholm University, Penn State, UT Austin, and participants at the Winter Meeting of the Econometric Society, NEUDC, SEA Conference, the Annual Conference on Economic Growth and Development at ISI-Delhi, the Texas Econometrics Camp, and the Economic Demography Workshop for their suggestions. All errors are our own.

1 Introduction

There is a large literature stressing the importance of female intra-household decision making power in developing countries (see, e.g., seminal work by Thomas (1990; 1994; 1997), based on related issues raised by Becker (1965; 1974; 1981) and by Sen (1983; 1988; 1989)). For example, numerous studies document that income or assets accruing to women or believed to be controlled by women are more likely than those of men to be allocated to expenditures that benefit children as well as themselves, such as food and health care (see, e.g., Haddad et al. (1997); Duflo (2003); Quisumbing and Maluccio (2003); Smith (2003); Rubalcava et al. (2009); LaFave and Thomas (2017)). Control over household resources, however, is difficult to directly observe because consumption is measured at the household level and goods can be shared. As a consequence, most of these studies focus on estimating the effect of randomized treatments or of other proxies that happen to be directly observed and are believed to empower women.[1]

The goal of this paper is to overcome these measurement issues and directly study the causal effect of women's control over resources on family health in India. We define treatment as a woman controlling substantial household resources and estimate this treatment indicator using a structural model of intra-household allocation. Our outcomes are health measures of household members and our instrument is a plausibly exogenous change in inheritance laws in India.

Note that a typical causal analysis might look at the impact of the change in law itself on health, but this would tell us nothing about how other changes in women's control over resources might impact health. In contrast, a typical structural analysis of this problem would require not only modeling the intra-household resource allocation process, but also structurally modeling how inheritance laws affect intra-household allocations and how control of household resources affects health. Many might find such models implausible. This is an example of the commonly noted trade-off that analyses based on natural experiments often answer questions of perhaps limited interest, while structural analyses depend on strong assumptions regarding the underlying true behavior. We seek to combine the best of both worlds. We know little about how female control over household resources may affect health outcomes, and so we address that question using LATE (local average treatment effect) estimation. A great deal more is known, both theoretically and empirically, about the economics of household consumption allocations. Hence, we first employ a structural modeling approach to estimate our treatment indicator. We then estimate a LATE of this treatment on a variety of health outcomes, using inheritance law changes as an instrument.

A drawback of this general procedure is that, due to estimation errors, which include possible model misspecification, our observed (estimated) treatment indicator is likely to be mismeasured for some households. More generally, even when treatment is observed rather than estimated, it may sometimes be mismeasured due, e.g., to reporting errors or to people who for whatever reason choose not to take the treatment they were assigned. To deal with these issues, we propose a new estimation method, which we call MR-LATE (for Mismeasurement Robust LATE), that can identify

[1] See section 2 for more details.

and consistently estimate LATE even when the endogenous binary treatment indicator contains measurement errors. Unlike Battistin et al. (2014) or DiTraglia and García-Jimeno (2017), MR-LATE does not require re-survey data or homogeneity of treatment effects. Unlike Ura (2016), who obtains bounds, we (under certain conditions) point identify and estimate LATE.

We apply our novel estimator to account for the inevitable estimation and specification errors that arise from using an intra-household allocation model to estimate treatment. The end result is that instead of asking what the health impact is of a particular policy intervention (e.g., the inheritance law change), we can address a more general question: What is the impact of empowering women (by giving them control of household resources) on family health? Taking the standard caveats about the external validity of LATE into account, the answer will then allow us to assess the potential impacts of a wide variety of policy interventions that affect control of household resources.

Our analysis relies on the *collective* households framework (Chiappori (1988, 1992)), which allows to structurally estimate a measure of women's decision making power from household consumption data. This measure is based on *resource shares* (i.e., the fraction of a household's resources allocated to each member), which exactly capture members' control over household resources and provide good measures of intra-household bargaining power (Browning et al. (2013)). We estimate resource shares for women and men in India using detailed expenditure data from the 2005-2006 National Sample Survey (NSS) of Consumer Expenditure and a methodology developed in Dunbar et al. (2013). We then define a household to be treated if the woman controls (roughly) more than half of household resources.[2]

Due to the lack of NSS data on health outcomes, we use the structural estimates to perform an out-of-sample prediction on the 2005-2006 Indian National Family Health Survey (NFHS-3), which includes the same socio-economic characteristics of individuals and households as the NSS dataset as well as detailed information about women's, men's and children's health status. Using our MR-LATE estimator, we then study the causal effect of living in households with relatively highly empowered women on household members' health status. We exploit changes to the Indian inheritance law and NFHS information on women's religion, year of marriage, and state of residence to construct our instrumental variable. Specifically, we focus on women's exposure to the Hindu Succession Act (HSA) amendments that equalized women's inheritance rights to men's in several Indian states between 1976 and 2005. In order to benefit from the reform, a woman needed to be Hindu, Buddhist, Jain or Sikh, and unmarried at the time of the reform in her state.

We find that accounting for specification, estimation, and/or measurement error in the estimate of treatment is empirically important, with some substantial differences between the standard 2SLS estimator for LATE (Imbens and Angrist (1994)), which cannot control for such errors, and our MR-LATE estimation method. The results of our empirical analysis indicate that women's empowerment positively affects women's and children's health outcomes, but does not affect men's health outcomes. The estimated effects are sizable. Our most conservative estimates indicate that, for compliers, the average treatment effect on women's body mass index is 7.7 and that women

[2]For simplicity, our analysis focuses on nuclear households with children.

in treated households are 72 percent less likely to be underweight and 52 percent less likely to be anemic. Also, mother's empowerment substantially reduces the occurrence of cough, fever and diarrhea in children, and mildly improves their height-for-age and weight-for-age z-scores. Our MR-LATE estimates indicate that living in households where the mother controls most of household resources reduces children's likelihood to be sick with fever by 43 to 65 percent, with cough by 66 to 89 percent, and with diarrhea by 45 to 65 percent. Our structural model of treatment indicates that women control about 60 percent of resources in households where they control a majority of resources, and about 40 percent when they do not control a majority. These large differences in resource control help to understand why our estimated treatment effects are so sizable.

Beyond our particular application, MR-LATE should be generally useful for many other instances where treatment is observed with error, due for example to reporting or recording errors. However, we want to emphasize its potential for applications like ours in which the treatment itself is estimated. There are many examples of potential treatments of clear economic significance, that are rarely analyzed causally because the treatments themselves cannot be directly observed and so must be estimated. Examples of such treatments are measures of expectations, ability, opportunity, utility, risk aversion, or welfare. One might be interested, for example, in quantifying the effect of high discount rates on individuals' likelihood of dropping out of school, or the effect of risk aversion on investment decisions, or the effect of abilities such as non-cognitive skills on future earnings. Our paper provides a practical method of doing causal analyses in many such settings.

The remainder of the paper is organized as follows. The next section is a literature review. This is followed by our identification results and the derivation of our associated MR-LATE estimation method, including a Monte-Carlo study of its small sample properties. We then implement our study of women's intra-household power on family health outcomes, reporting the results of both our structural model of treatment and its causal effects. These are followed by some concluding remarks. Proofs and additional material are in an Appendix.

2 Literature Review

Our empirical analysis consists of a structural model of treatment (women's control over household resources) and a causal LATE analysis of the effect of this treatment on a range of individual health outcomes. Our structural treatment model is based on the collective household framework pioneered by Becker (1965, 1981), Apps and Rees (1988) and Chiappori (1988, 1992). In these models, each household is characterized as a collection of individuals, each of whom has a well defined utility function, and who interact to generate Pareto efficient allocations. Recent advances in this literature permit the recovery of resource shares (or sharing rule), defined as each member's share of total household consumption (Lewbel and Pendakur (2008), Browning et al. (2013), Dunbar et al. (2013)). Our specific model is based on Dunbar et al. (2013) and Calvi (2016), who obtain resource shares from estimates of Engel curves (demand equations holding prices constant) of clothing items that are consumed exclusively by women, men or children. Partly due to data

limitations, in our application we focus on estimating resource shares for men and women only, as in Lewbel and Pendakur (2008) and Browning et al. (2013), while treating children as public goods (see, e.g., Blundell et al. (2005)).

From a policy perspective, our results contribute to the general literature on estimating the effect of women's empowerment on adults' and children's outcomes. An application with a similar motivation to ours is Maitra (2004), who uses relative educational attainment of the parents and self-reported measures of autonomy and decision power to show that a woman's control over household resources (her ability to keep money aside) has a significant effect on health care usage and child mortality in Indian households. More broadly, economic analyses of bargaining power within the household have made use of varying proxies of resource control, such as assets (e.g., in Beegle et al. (2001); Frankenberg et al. (2001); Quisumbing and Maluccio (2003); LaFave and Thomas (2017)), or unearned income (e.g., Schultz (1990)). Thomas (1990) documents that child health in Brazil tends to improve if additional non-labor income is in the hands of women rather than men. He estimates that income in the hands of a mother has, on average, twenty times the impact of the same income in the hands of a father with respect to children's survival probabilities. Duflo (2003), studying elderly benefits in South Africa, concludes that the same transfer has drastically different impacts on the health of female grandchildren depending on whether it is paid to the grandmother or to the grandfather.

Legal reforms aimed at improving women's property or inheritance rights have been also used to assess the effects of changes in bargaining power within the household. In the Indian context, Deininger et al. (2013) find evidence of an increase of women's likelihood of inheriting land following the introduction of Hindu Succession Act (HSA) amendments that equalized women's inheritance rights to men's in several Indian states between 1976 and 2005. Roy (2008) documents that women's exposure to the HSA reforms improved women's autonomy within their marital families. Deininger et al. (2013), Roy (2015) and Bose and Das (2015) find that it increased female education, while Heath and Tan (2014) claim that it increased women's labor supply, especially into high-paying jobs. Calvi (2016) shows that women's exposure to the HSA reforms led to improved health outcomes and increased access to household resources.[3] We recall that our analysis uses the HSA reforms as the instrument for LATE estimation, where an imperfect measure of treatment is obtained by estimation of a structural model of control over household resources.

Papers empirically documenting substantial measurement (misclassification) errors in observed treatments include Bollinger (1996), Angrist and Krueger (1999), Kane et al. (1999), Card (2001), Black et al. (2003), and Hernandez et al. (2007). In our application, these measurement errors in treatment come (in part) from treatment being structurally estimated, but as these references show, measurement errors are also common even in applications where treatment is directly observed.

A few previous papers have considered alternative techniques for dealing with such measure-

[3]Parallel to these studies, an extensive literature studying the effects of Conditional Cash Transfer programs in developing countries has shown that providing women with a large amount of cash in their hands can significantly increase, e.g., the budget shares of expenditures on clothing for children and lower shares of alcohol (Attanasio and Lechene, 2002), increase children's health and education, and livestock (Rubalcava et al., 2009) and improve child development (Macours et al., 2012; Tommasi, 2016).

ment errors in treatment. Homogeneous treatment effects, corresponding to estimation of constant coefficients of a mismeasured binary regressor, have been estimated using instruments by many authors, including Aigner (1973), Kane et al. (1999), Black et al. (2000) and Frazis and Loewenstein (2003). When treatment is mismeasured, point-identification and associated estimators of average treatment effects (without assuming treatment effects are homogeneous) are provided by Mahajan (2006), Lewbel (2007) and Hu (2008). These papers obtain identification exploiting both the assumption that the true treatment is exogenously determined and an assumed instrument to deal with the measurement error. Under more general conditions, bounds on average treatment effects with misclassified treatment are provided by Klepper (1988), Manski (1990), Bollinger (1996), Kreider and Pepper (2007), Molinari (2008), Imai et al. (2010), and Kreider et al. (2012).

The causal effect we focus on identifying and estimating is the local average treatment effect (LATE) of Imbens and Angrist (1994), which is applicable when the true treatment is endogenous, an exogenous binary instrument is available, and treatment effects may be heterogeneous. Identification of LATE with misclassified treatment has recently received some attention. Ura (2016), for example, considers estimation of LATE with mismeasured treatment and standard LATE instrument assumptions, but only obtains set identification bounds.[4] Our MR-LATE methodology obtains point-identification making use of two rather than a single misclassified treatment indicator. Battistin et al. (2014) also use two measures of the misclassified treatment to obtain point-identification of LATE and an associated semiparametric estimator. However, they require re-survey data (that is, multiple observations of the same individuals), which are often not available. DiTraglia and García-Jimeno (2017) and Yanagi (2017) also obtain point-identification of LATE with mismeasured treatment. The former, however, requires that treatment effects be homogeneous, while the latter requires the availability of two instrumental variables with specific properties, one for the endogenous treatment and the other to deal with the measurement error. We achieve point-identification without these requirements.

Our estimation problem has the standard LATE structure that a randomized instrument is correlated with treatment, and the true treatment affects an outcome. But in our case we must overcome the issue that the observed treatment does not equal the true treatment. A similar structure arises in models where outcomes of interest and randomized treatment are not available in the same data set. In these models, a randomized treatment (corresponding to our instrument) affects an intermediate outcome called a *statistical surrogate* (see Prentice (1989)), corresponding to our mismeasured treatment indicator. The surrogate then affects (or at least strongly correlates with) the outcome of interest. These estimators require that the surrogate satisfy a strong conditional independence assumption (see, e.g., Rosenbaum (1984); Begg and Leung (2000); Frangakis and Rubin (2002); VanderWeele (2015)). Athey et al. (2016) overcome this limitation by observing multiple surrogates, each of which may not satisfy the required conditional independence. They assume that

[4]Specifically, the instrument must satisfy the exclusion restriction and weakly increase the true treatment (Imbens and Angrist, 1994; Angrist et al., 1996). In the case of a continuous misclassified treatment, Lewbel (1998), Song et al. (2015), Hu et al. (2015) and Song (2015) use instruments and further exclusion restrictions to obtain identification and estimation of average marginal effects with classical or nonclassical measurement error in a nonparametric or semiparametric context.

there exists a single latent, unobserved surrogate that has the desired properties, and combine the observed surrogates to model the impact on the outcome of the underlying latent surrogate. In a roughly analogous way, we exploit multiple (two) mismeasures of treatment to model the impact on the outcome of an underlying latent (true) treatment. Beyond this analogy, however, the details of their estimator and underlying assumptions differ substantially from ours.

More generally, this paper contributes to the long-standing debate on the relative benefits and limitations of structural modeling vs. causal, randomization based analyses. Proponents of randomization question the validity of results obtained from complex structural modeling assumptions (Angrist and Pischke, 2010; Imbens, 2010). Advocates of structural models stress the insights that can be obtained when one allows economic theory to guide the empirical work (Wolpin, 2013). Recent contributions in the econometrics literature have started to formally unify the two camps in order to overcome these divisions (Vytlacil (2002), Heckman et al. (2006); Heckman and Vytlacil (2007); Pearl (2009)). Heckman (2010), for example, proposes to combine the best features of both the structural and the causal modeling approaches in what he calls a *third way* of policy analysis. Similarly, Lewbel (2016) argues that combining the strengths of both approaches can often be the best strategy for identification. We contribute to this literature by using structure provided by economic theory to significantly expand the set of causal questions that researchers can answer, when combined with statistical techniques (like MR-LATE) that account for the inevitable estimation and specification errors of structural models.

3 LATE With Mismeasured or Misspecified Treatment

This section is organized in three parts. First, we describe our theoretical framework and derive the new MR-LATE estimator. Second, we provide examples for identification of LATE. Finally, we implement a Monte Carlo experiment to assess the finite sample behavior of the estimator.

3.1 Set Up and Identification

We first introduce some notation. We ignore additional covariates for now, as everything immediately extends to conditioning on covariates X.

Let D be the true binary treatment variable that affects an outcome of interest. D is *not* observed and cannot be consistently estimated. Let Z be an unconfounded (e.g., randomized) binary instrument that is correlated with D and will satisfy the standard (Imbens and Angrist (1994)) assumptions of an instrument for LATE estimation. Let the random binary variables D_0 and D_1 denote the potential treatments $D_z = D(z)$ for possible realizations z of Z. By definition,

$$D = (1 - Z) D_0 + Z D_1. \tag{1}$$

Let Y be an observed outcome of interest and let random variables Y_0 and Y_1 be the potential

outcomes $Y_d = Y(d)$ for possible realizations d of D. Then,

$$Y = (1-D)Y_0 + DY_1 = Y_0 + (Y_1 - Y_0)[(1-Z)D_0 + ZD_1]. \tag{2}$$

Assumption 1. *Y and D satisfy the standard Imbens and Angrist (1994) LATE assumptions:*

i. $0 < E(D) < 1$, $0 < E(Z) < 1$ *and* $Z \perp (Y_1, Y_0, D_1, D_0)$.

ii. (Y_1, Y_0, D_1, D_0, Z) *are independent across individuals and have finite means.*

iii. There are no defiers, so $\Pr(D_0 = 1 \text{ and } D_1 = 0) = 0$.

Let C denote a complier, i.e., someone who has $D_0 = 0$ and $D_1 = 1$. If D was observed then, under the conditions listed in Assumption 1 above, the Imbens and Angrist (1994) LATE would be identified by the instrumental variable estimand:

$$\frac{cov(Y,Z)}{cov(D,Z)} = E(Y_1 - Y_0 \mid C) = LATE. \tag{3}$$

Since we do not observe D, however, we cannot implement this standard approach.

Instead of D, what we observe is a binary treatment indicator T, which could be a proxy for or an estimate of D, or could correspond to reported values of D that are mismeasured for some observations (later we will make use of two such D indicators, but just consider one for now). In our empirical application, for instance, T will be an estimate of D that is based on a structural model. Thus, in our application, T will not equal D for some individuals either because of estimation error or because the structural model may be misspecified. Our goal is to point-identify and estimate LATE, even though D is not directly observed and cannot be consistently estimated.[5] Define random variables T_0 and T_1 as potential observed treatments (or potential estimated treatments) so that $T_d = T(d)$ for possible realizations d of D. Then by definition

$$T = (1-D)T_0 + DT_1. \tag{4}$$

The variables T_0 and T_1 can be interpreted as indicators of whether treatment is correctly measured or not. In particular, if $T_0 = 0$ and $T_1 = 1$, then treatment is not mismeasured. This shows the two possible types of measurement or classification error: if $T_0 = 1$, then a true $D = 0$ is misclassified as treated, and if $T_1 = 0$, then a true $D = 1$ is misclassified as untreated.

Assumption 2. *T is such that the following conditions are satisfied:*

i. $Z \perp (Y_1, Y_0, D_1, D_0, T_1, T_0)$.

ii. $(T_1, T_0) \perp (Y_1, Y_0) \mid C$.

iii. $E(T_1 - T_0 \mid C) \neq 0$.

Assumption 2-i just combines the LATE unconfoundedness assumption that $Z \perp (Y_1, Y_0, D_1, D_0)$ with the assumption that the instrument is also independent of the potential measurement errors,

[5] Although we observe T and not D, people's behavior is still based on their actual D. This means that introducing measurement error does not change the no defiers assumption. If we had incorrectly assumed behavior was based on T, and estimated LATE using T in place of D, then what would appear to be defiers could exist. That would be just one of multiple sources of bias in LATE estimates that ignore the measurement error.

and hence of (T_1, T_0). The standard assumption that Z is randomized by experimental or quasi-experimental design is sufficient to make 2-i hold. Assumption 2-ii says that, for compliers, the potential mismeasures (T_1, T_0) are independent of the potential outcomes (Y_1, Y_0). Combined with unconfoundedness, this corresponds to the standard assumption that measurement errors are unrelated to outcomes.[6] Finally, Assumption 2-iii is a minimal relevance condition saying that, at least for compliers, T provides some information regarding D. This assumption implies that, for compliers, the correlation between D and T is nonzero.

Let $p_d = E(T_d \mid C)$. By definition, p_1 is the probability that a complier would have their treatment correctly observed if they were assigned the true treatment $D = 1$. That is, p_1 is the probability that a complier would have $T = 1$ if they were assigned $D = 1$. In contrast, p_0 is the probability that a complier would have their treatment *incorrectly* observed (meaning $T = 1$) if they were assigned the true treatment $D = 0$. Note that Assumption 2-iii ensures that $p_1 - p_0$ is nonzero.

Define q and λ as follows

$$q = \frac{p_1}{p_1 - p_0} \tag{5}$$

$$\lambda = \frac{cov(YT, Z)}{cov(T, Z)} = \frac{E(YT \mid Z = 1) - E(YT \mid Z = 0)}{E(T \mid Z = 1) - E(T \mid Z = 0)}. \tag{6}$$

If one were to ignore measurement error in T, one would estimate LATE as in Imbens and Angrist (1994) by an instrumental variables regression of Y on T using Z as the instrument, which would asymptotically equal $cov(Y, Z)/cov(T, Z)$. The constant λ, defined by equation (6), is similar to but not the same as this asymptotic value of the LATE estimator. Instead, λ equals the asymptotic value of an instrumental variables regression of YT on T using Z as the instrument. Abadie (2002) also makes use of this λ, though for a different purpose. In the following Theorem, we prove that λ is a mixture of the potential outcomes for compliers.

Theorem 1. *Let Assumptions 1 and 2 hold. Then:*

$$\lambda = E[qY_1 + (1 - q)Y_0 \mid C]. \tag{7}$$

Proof. See Appendix A.1. □

Elements of Theorem 1, and results related to Theorem 1, appear in some earlier work, including Abadie (2002), Ura (2016), and references therein. Our primary novelty is in how we make use of the relationship given by Theorem 1.

Assume we observe two different mismeasures of treatment, called T^a and T^b. These could be, for instance, two proxies or two different estimates of D. Recalling that $p_d = E(T_d \mid C)$, let $p_d^a = E\left(T_d^a \mid C\right)$ and $p_d^b = E\left(T_d^b \mid C\right)$, where T_d^a and T_d^b are the potential mismeasured treatments associated with T^a and T^b. Similarly, define q^a, q^b, λ^a, λ^b, ρ, and our MR-LATE estimator as

[6] A sufficient condition for Assumption 2-ii to hold is that $(T_1, T_0) \perp (Y_1, Y_0, D_1, D_0)$, meaning that the measurement errors are independent of the potential outcomes and potential treatments. If this stronger assumption holds then q defined below satisifes $q = E(T_0)/E(T_1 - T_0)$. We do not require this plausible but stronger condition.

follows:

$$q^a = \frac{p_1^a}{p_1^a - p_0^a}, \quad q^b = \frac{p_1^b}{p_1^b - p_0^b},$$

$$\lambda^a = \frac{cov\left(T^a Y, Z\right)}{cov\left(T^a, Z\right)}, \quad \lambda^b = \frac{cov\left(T^b Y, Z\right)}{cov\left(T^b, Z\right)}, \text{ and}$$

$$\text{MR-LATE} = \rho = \lambda^a - \lambda^b.$$

This Corollary follows immediately:

Corollary 1. *Let Assumption 1 hold, and let Assumption 2 hold with $T = T^a$ and with $T = T^b$. Then:*

$$\textit{MR-LATE} = \left(q^a - q^b\right) E\left[Y_1 - Y_0 \mid C\right] = \left(q^a - q^b\right) LATE \tag{8}$$

where $LATE = cov\left(Y, Z\right) / cov\left(D, Z\right)$.

Corollary 1 has some straightforward implications:

Corollary 2. *Let Assumption 1 hold, and let Assumption 2 hold with $T = T^a$ and with $T = T^b$. If $q^a - q^b = 1$, then MR-LATE $= LATE = cov\left(Y, Z\right) / cov\left(D, Z\right)$. A sufficient condition for MR-LATE to equal LATE is $p_0^a = p_1^b = 0$.*

The fact that MR-LATE equals the true LATE when $p_0^a = p_1^b = 0$ follows from equations (5) and (8). More generally, MR-LATE provides a good approximation to LATE when p_0^a and p_1^b are close to zero. Having $p_0^a = 0$ means that, among compliers, the probability that $T_0^a = 1$ is zero. In other words, all compliers who have $D = 0$ must also have $T^a = 0$. Therefore, p_0^a equals zero if, among compliers, the treatment measure T^a has only one kind of measurement error, never mistaking the actually untreated as treated (even if it often mistakes the actually treated as untreated). Having $p_1^b = 0$ is a little more peculiar, since it requires that all compliers who have $D = 1$ must also have $T^b = 0$. This means that, among compliers, $1 - T^b$ never mistakes the actually treated as untreated (even if it is frequently wrong about mistaking the actually untreated as treated).

The restrictions that $p_0^a = p_1^b = 0$ can also be thought of as rough analogs to the standard no defiers assumption. No defiers assumes a zero probability of certain combinations of D and Z realizations, while $p_0^a = 0$ assumes a zero probability of certain combinations of D and T^a realizations, and similarly for $p_0^b = 0$. The $p_0^a = p_1^b = 0$ restrictions are stronger than no defiers, but it is not unreasonable to assume that mismeasurement of D would have a stronger connection to the true D than a mere instrument would.

A simple summary is as follows: MR-LATE equals the true LATE if, among compliers, when D is zero then T^a is zero, and when D is one then T^b is zero. More generally, p_0^a and p_1^b will be close to zero, making MR-LATE close to the true LATE (meaning $q^a - q^b$ is close to one), if T^a is rarely one

when D is zero, and if T^b is rarely one when D is one. This is analogous to how one might interpret the usual LATE estimator as being close to but not equal to a true causal effect, if the probability of defiers in the populations is small but nonzero.

When $q^a - q^b$ is not close to or equal to one, Corollary 1 might alternatively be used for set identification.

Corollary 3. *Let Assumption 1 hold, and let Assumption 2 hold with $T = T^a$ and with $T = T^b$. If $q^a - q^b > 0$, MR-LATE signs LATE. If $q^a - q^b \geq 1$, LATE lies between 0 and MR-LATE. A sufficient condition for set identification is $p_1^a > p_0^a$ and $p_0^b > p_1^b$.*

Corollary 3 states that if, for compliers, the share of actual treated in T^a is larger than the share of misclassified actual untreated, and analogously, if the share of actual untreated in T^b is larger than the share of misclassified actual treated, then MR-LATE is informative regarding the sign and the magnitude of LATE. Later, we discuss in more detail the situations under which the (mild) conditions that make MR-LATE at least close to LATE hold, and the sufficient conditions for set identification.

3.2 The MR-LATE Estimator and Examples

The numerical computation of the MR-LATE estimator is extremely easy. Assume we have independent, identically distributed observations of the vector $\left(Y_i, Z_i, T_i^a, T_i^b\right)$ for individuals $i = 1, ..., n$ and that $E\left(Y_i^2\right)$ is finite. Consider a linear instrumental variables regression of $Y_i T_i^a$ on a constant and on T_i^a, using as instruments a constant and Z_i. Let the estimated coefficient of T_i^a in this regression be $\widehat{\lambda}^a$. Note that $\widehat{\lambda}^a$ does not equal the ordinary Imbens and Angrist (1994) LATE estimator, since $\widehat{\lambda}^a$ regresses $Y_i T_i^a$ on T_i^a instead of regressing Y_i on T_i^a (as noted earlier, Abadie (2002), Ura (2016) among others have similar constructions, though used differently).

Similarly, let $\widehat{\lambda}^b$ be the estimated coefficient of T_i^b in a linear instrumental variables regression of $Y_i T_i^b$ on a constant and on T_i^b, again using as instruments a constant and Z_i. The MR-LATE estimator is then given by $\widehat{\rho} = \widehat{\lambda}^a - \widehat{\lambda}^b$. With our identifying assumptions, consistency of this estimator follows immediately from applying the law of large numbers to each expectation in these regressions. Similarly, root-n asymptotic normality follows mechanically from the Lindeberg-Levy central limit theorem and the delta method.

One convenient way to write the estimator is to consider the following moments:

$$\begin{aligned}
E\left(Y_i T_i^b - \alpha^b - \lambda^b T_i^b\right) &= 0 \\
E\left(\left(Y_i T_i^b - \alpha^b - \lambda^b T_i^b\right) Z_i\right) &= 0 \\
E\left(Y_i T_i^a - \alpha^a - \left(\rho + \lambda^b\right) T_i^a\right) &= 0 \\
E\left(\left(Y_i T_i^a - \alpha^a - \left(\rho + \lambda^b\right) T_i^a\right) Z_i\right) &= 0
\end{aligned}$$

for some constants α^a, α^b, λ^b, and ρ. These moments correspond to the two instrumental variables regressions that comprise MR-LATE. One could therefore estimate the constants α^a, α^b, λ^b, and ρ

using GMM (which would actually just reduce to method of moments estimation), and the standard GMM asymptotic distribution formula would then deliver standard errors for $\widehat{\rho}$. Alternatively, one could just bootstrap the two instrumental variables regressions that define $\widehat{\lambda}^a$ and $\widehat{\lambda}^b$ and hence $\widehat{\rho}$.

Before proceeding further, we discuss two examples (related to our later empirical application) that illustrate the use of the MR-LATE estimator.

Example 1: Incompletely Measured Treatment. Due to growing attention regarding the status of women in developing countries, in household surveys of such countries a common type of question to ask is, "Who usually makes decisions about [X] in your household?" For example, the National Family Health Survey used in this paper asks this, with [X] referring to decisions regarding own health care, household purchases, and visits to family or relatives. Possible answers to this type of survey question are "the wife" or "the husband," but other ambiguous responses are also common. These other answers can include no response, or "someone else," or "both." Let $P = 1$ if the answer is "the wife," let $P = -1$ if the answer is "the husband," and let $P = 0$ denote any other answer. Assume we also have some outcome Y and some randomized binary instrument Z that is positively correlated with P.

Suppose we define treatment to be $D = 1$ if the wife makes most of the decisions regarding [X], otherwise $D = 0$ if the husband makes most of these decisions. This means we observe D for households that have $P = 1$ or $P = -1$, but we do not observe who is ultimately deciding [X] when $P = 0$. Thus, for those households with $P = 0$, whatever value we assigned to D would be subject to measurement error. One common procedure with this type of data is to construct a treatment proxy T where $T = 1$ if $P = 1$, otherwise $T = 0$, and apply the usual LATE estimator using T in place of D. Alternatively, one might discard all the observations that have $P = 0$, and apply the standard LATE estimator to the remaining observations. However, either of these estimators is likely biased in unknown ways, due to misclassification errors or to correlated selection when $P = 0$. Instead of these typical estimators, suppose we let $T^a = 1$ if $P = 1$ ($T^a = 0$ otherwise), and let $T^b = 1$ if $P = -1$ ($T^b = 0$ otherwise). With these definitions of T^a and T^b, MR-LATE can correctly identify and consistently estimate the LATE that corresponds to observing the true D for all households without error. This follows directly from Theorem 1 and Corollary 1 discussed above.

A drawback of using individual self-reported measures of power, as in this example, is that each response concerns a relatively specific type of decision (e.g., own health care, household purchases, and visits to family or relatives) and is likely not representative of women's power more generally. Thus, in this application the LATE we obtain might not be particularly informative regarding empowerment. It is also possible that some of the $P = 1$ and $P = -1$ responses are mismeasured, meaning that a wife or husband may incorrectly claim that they make most of the decisions about [X]. In this case p_0^a and p_1^b will not equal zero as is required for MR-LATE to point-identify the true LATE, though if such misreports are rare, MR-LATE will still be close to the true LATE.

Example 2: Threshold Crossing Model. Partly to deal with the problems listed in the previous paragraph, it is common to combine responses from multiple empowerment related questions

into a single index of the wife's relative decision making power within a household (see, e.g., Smith (2003), Roy (2008), Upadhyay et al. (2014) and references therein). Alternatively, as we do in our empirical application, one might estimate a power measure using a structural model of household behavior. One may then use MR-LATE to deal with potential measurement or estimation errors in these power indices or measures, as follows.

Define the indicator function $\mathbb{I}(\cdot)$ to equal one if its argument is true, and zero otherwise. Let R^* be the true measure of a woman's relative decision making power, or control over resources. Suppose that the true treatment D is determined by a typical threshold crossing model, so $D = \mathbb{I}(R^* \geq e)$ for some unobserved random threshold e. This means that the wife has primary decision making power, or controls a majority of household resources, corresponding to $D = 1$, if R^* is sufficiently large.

Assume we cannot observe R^* perfectly, either because of measurement or response errors, or because of specification or estimation errors. Instead, what we observe, construct, or estimate, is the variable R, which is related to the true R^* by $R = R^* + \varepsilon$, where ε is an unknown error due to mismeasurement, misspecification, or estimation error. For simplicity in exposition, assume $\varepsilon + e$ is independent of Z and R^*.[7]

Let κ^a and κ^b be two constants chosen by the researcher, with $\kappa^a > \kappa^b$, and define treatment measures T^a and T^b as follows:

$$T^a = \mathbb{I}(R \geq \kappa^a) \quad \text{and} \quad T^b = \mathbb{I}(R < \kappa^b). \tag{9}$$

Note that $R \geq \kappa^a$ implies $R^* \geq \kappa^a - \varepsilon$. Thus, by construction,

$$p_0^a = \Pr(R^* \geq \kappa^a - \varepsilon \mid R^* < e).$$

Therefore, if κ^a is larger than the maximum value that $\varepsilon + e$ can take on, then $p_0^a = 0$. More generally, p_0^a is near zero if the chance that $\varepsilon + e$ is greater than κ^a is small. The intuition is straightforward: an individual is untreated (having $D = 0$) when their true R^* is sufficiently small. So, if we define T^a to equal one only when the observed or estimated R is very large, then the probability of having $T^a = 1$ when $D = 0$ is very small, meaning that p_0^a is near or equal to zero. In section A.2 in the Appendix, we provide a graphical illustration of this construction.

We could guarantee that p_0^a is zero, as desired, by taking κ^a to be infinite (or, if R is bounded, taking κ^a to be greater than the largest value that R can take on). Then, however, T^a would equal zero for every observation, and so would be useless as a measure of treatment. Using the notation of Theorem 1, if κ^a is infinite, then $p_0^a = p_1^a = \Pr(R^* \geq \kappa^a - \varepsilon \mid R^* \geq e) = 0$, which would lead to a violation of Assumption 2-iii.

In practice, we have a trade-off in the selection of κ^a. Recall $q^a = p_1^a / \left(p_1^a - p_0^a\right)$, and we want q^a to be as close as possible to one. This means choosing κ^a to make the probability p_0^a as small as

[7] The assumption that e and ε are independent of Z and R^* can be relaxed. We only assume this here to simplify the exposition. Specifically, if e or ε correlates with Z or R^*, then the correct expressions for p_0^a and p_1^b will need to condition on compliers, and will therefore be more complicated than the expressions we derive below, though the corresponding intuition regarding identification will remain unchanged.

possible relative to the probability p_1^a. The larger κ^a is, the closer p_0^a is to zero as desired. However, when κ^a is very large, p_1^a becomes small too. This is because the larger κ^a is, the less informative T^a becomes as a measure of treatment (e.g., the lower is the correlation between T^a and the true D). Empirically, to get q^a to be as close to one as possible, we will want to choose a moderate value of κ^a, which makes p_0^a close to zero and p_1^a relatively large (see the Monte Carlo analysis for further analysis of this point).

A comparable construction applies to T^b, where

$$p_1^b = \Pr(R^* < \kappa^b - \varepsilon \mid R^* \geq e),$$

so $p_1^b = 0$ if κ^b is less than the minimum value that $\varepsilon + e$ can take on, and p_1^b is near zero if the chance of $\varepsilon + e$ being less than κ^b is small.

Once κ^a and κ^b are chosen, the estimation of MR-LATE is as described earlier, with $T^a = \mathbb{I}(R \geq \kappa^a)$ and $T^b = \mathbb{I}\left(R < \kappa^b\right)$. It is interesting to contrast this with ordinary LATE estimation. In this context, one would typically construct $T = \mathbb{I}(R \geq c)$, where c is one's best guess of the mid-point of $\varepsilon + e$, thereby constructing T to be as close as possible to the true unknown D. However, estimation of LATE using this T in place of the true unknown D will generally be substantially biased. The problem with replacing the unknown D with the known T in an ordinary LATE estimation is that compliers who have R close to c are precisely those who are most likely to be misclassified.

Point-identification, where MR-LATE equals the true LATE in this example, requires that $\varepsilon + e$ be bounded from both above and below, and that κ^b and κ^a are chosen to lie outside these bounds (but still within the range of R^* and R). In some contexts, we may have sufficient information to know these bounds: for instance, the threshold e might be an observable policy variable, and the measurement error ε might be rounding errors of known maximum possible magnitude. However, even when such information is not available, we can still obtain bounds on the true LATE by selecting κ^a such that $p_1^a > p_0^a$ and κ^b such that $p_0^b > p_1^b$ as in Corollary 3. The selection of κ^b and κ^a as described in the previous paragraph will generally satisfy these inequalities, since they make p_1^a large relative to p_0^a and p_0^b large relative to p_1^b.

To illustrate this point, suppose we observed the answer to m survey questions of the type discussed in Example 1 above. Give each answer a value of one if the response is that the wife controls that decision, a value of minus one if the response is that the husband controls that decision, and a value of zero for any other response. Now define our estimated index R to equal the sum of these responses across all of the questions. Then, if we took $\kappa^a = k$ for some integer k, with $k < m$, we would be assuming that it is extremely unlikely (p_0^a is very close to zero) that the husband has most of the power if the household reports that the wife makes $m - k$ or more of the m decisions. Similarly, taking $\kappa^b = -k$ means assuming it is extremely unlikely (p_1^b is very close to zero) that the wife really has most of the power if they report that the husband makes $m - k$ or more of the decisions. We might not be certain that these choices of κ^b and κ^a will point-identify LATE, but we can choose large enough values of k to be confident that MR-LATE is at least close to LATE. Moreover, we can

have a very high degree of confidence that at least the conditions for set identification, $p_1^a > p_0^a$ and $p_0^b > p_1^b$, are satisfied.

Likewise, in our later empirical application, where R^* is the share of household resources controlled by a woman, we cannot be certain that our choices of κ^b and κ^a will point identify LATE, but we will choose large enough values to be confident that p_0^a and p_1^b are at least small relative to p_1^a and p_0^b, giving set identification and making MR-LATE at close to LATE. More generally, we will want to choose κ^a and κ^b to make use of whatever economic, distributional, or other information is available. We provide some examples later.

3.3 Monte Carlo Simulations

In this section, we implement some Monte Carlo experiments to illustrate when and how a practitioner could apply MR-LATE, and to check the finite sample properties of our new estimator. The experiments correspond to Examples 1 and 2 discussed above.

In our data generating process (hereafter DGP), we construct unobserved potential outcomes Y_0 and Y_1 and the corresponding observed outcome Y as follows:

$$Y_0 = 0.5 + X + S + V_0$$
$$Y_1 = 1.5 + X + S + V_1$$
$$Y = (1-D)\,Y_0 + DY_1$$

where X is an observed covariate, while S, V_0 and V_1 are random unobserved errors. The unobserved true treatment indicator D is given by $D = \mathbb{I}(R^* \geq e)$, where the unobserved index R^* is

$$R^* = 0.1X + 0.1Z + 0.1S + U \tag{10}$$

with Z being the observed binary instrument and U is an additional error. The exogenously determined variables, errors, and parameter values in our simulations are all set to values that resemble our empirical application. In particular, we let $X \sim N(0,1)$, $S \sim N(0, 0.1)$, $Z = \mathbb{I}(\sim U(0,1) > 0.9)$, $V_0 \sim N(0,1)$, $V_1 \sim N(0,1)$, and $U \sim N(0, 0.04)$. To ease comparison, we here rescale R^* such that the relevant threshold for treatment is equal to zero, that is $e = 0$.[8] Unless noted otherwise, each sample consists of 10,000 observations, and the DGP is simulated 1,000 times.

In our first experiment (which corresponds to Example 1 above) we assume the econometrician just observes Y, X, Z, and P, where P takes on values of -1, 0, and 1 each with probability $1/3$, based on low, medium and high values of R^*. So, $P = D = 1$ in one third of the sample, $D = 0$ and $P = -1$ in another third of the sample, and $P = 0$ in the remaining sample, regardless of the value of D.

We compare the performances of five different estimators. The first two estimators are infeasible since they assume D is observed without error. First, we estimate an ordinary least squares (OLS)

[8] In our empirical application, the relevant threshold is 50 percent.

Table 1: Simulations: Example 1

	D is known		D is unknown		
	OLS	2SLS	OLS	2SLS	MR-LATE
D	1.015	1.004	0.712	0.903	1.021
sd	0.014	0.101	0.020	0.096	0.084
Bias	0.015	0.004	-0.288	-0.097	0.021
MSE	0.000	0.010	0.083	0.019	0.007

Notes: In each simulation, the true value is set to 1.000. Results are based on 1,000 simulations for 10,000 observations each.

regression model of Y on a constant, D, and X. Due to the correlation between treatment and potential outcomes, this OLS is inconsistent. In particular, the variable S causes a violation of the unconfoundedness assumption. Second, we estimate a two stage least squares (2SLS) regression model of Y on a constant, D, and X, using Z as an instrument for D. The coefficient of D in this 2SLS regression is the standard LATE estimator, which is consistent but infeasible because it uses D observed without error.

The remaining estimators we consider are feasible. First, we let $T = \mathbb{I}(P = 1)$ and estimate a linear regression of Y on a constant, T, and X, using OLS. Next, we estimate the same linear model using 2SLS, taking Z to be an instrument for T. The coefficient of T in this 2SLS regression is the standard LATE estimator, using T in place of the unobserved true treatment indicator D. Third, we construct $T^a = \mathbb{I}(P = 1)$ and $T^b = \mathbb{I}(P = -1)$, and apply the MR-LATE estimator. MR-LATE is given by $\widehat{\lambda}^a - \widehat{\lambda}^b$, where, for $j = a, b$, $\widehat{\lambda}^j$ is the 2SLS coefficient of T^j, obtained by regressing YT^j on a constant, T^j, and X, using Z as an instrument for T^j.

Table 1 shows our first simulation results. As expected, both of the OLS estimators are biased due to the correlation between treatment and potential outcomes, with the second OLS estimator behaving particularly poorly (bias of 29 percent), since it is also biased due to the measurement error in T relative to D. The standard 2SLS LATE estimator, which is infeasible because it uses the true unobserved D, has bias near zero. This is our benchmark for comparison of the feasible estimators. The standard 2SLS LATE estimator that uses T in place of D, which is feasible, has about 10 percent bias and an MSE almost double that of the infeasible LATE. This estimator is inconsistent due to the measurement error in T. Finally, our MR-LATE estimator, which is both feasible and consistent, has a relatively small bias (2 percent). Surprisingly, in this problem MR-LATE performs so well that it has a smaller mean squared error than even the infeasible 2SLS LATE that uses the true D.

In our next experiment, we assume that the practitioner observes multiple measures $P_1, \ldots, P_m$ relating to treatment, each constructed analogously to the variable P defined above. Following Example 2 of section 3.2, let $R = \sum_{j=1}^{m} P_j$ be an index intended to approximate the true unobserved R^*. We now only consider feasible estimators. We first let $T = \mathbb{I}(R \geq 0)$ and estimate a linear regression of Y on a constant, T, and X, using OLS. Next, we estimate the same linear model using 2SLS, taking Z to be an instrument for T. The coefficient of T in this 2SLS regression is again the standard LATE estimator that uses T in place of the unobserved true treatment indicator D. Third, we let $T^a = \mathbb{I}(R \geq k)$ and $T^b = \mathbb{I}(R < -k)$, for a few different choices of the integer k, and apply the

Table 2: Simulations: Example 2

	Panel A: $R = \sum_{j=1}^{6} P_j$					
			MR-LATE			
	OLS	2SLS	$k=0$	$k=1$	$k=2$	$k=3$
D	0.758	0.899	0.899	1.012	1.013	1.015
sd	0.019	0.092	0.092	0.083	0.081	0.081
Bias	-0.242	-0.101	-0.101	0.012	0.013	0.015
MSE	0.059	0.019	0.019	0.007	0.007	0.007
	Panel B: $R = R^* + \varepsilon$ (known bounds)					
	$\kappa = 0$			$\kappa = 0.05$		
	OLS	2SLS	MR-LATE	OLS	2SLS	MR-LATE
D	1.015	0.999	0.999	0.827	1.032	1.016
sd	0.014	0.097	0.097	0.018	0.105	0.085
Bias	0.015	-0.001	-0.001	-0.173	0.032	0.016
MSE	0.000	0.009	0.009	0.030	0.012	0.007
	Panel C: $R = R^* + \varepsilon$ (unknown bounds)					
			MR-LATE			
	OLS	2SLS	$\kappa = 0$	$\kappa = 0.01$	$\kappa = 0.05$	$\kappa = 0.1$
D	0.780	1.063	1.063	1.042	1.021	1.031
sd	0.020	0.109	0.109	0.105	0.091	0.088
Bias	-0.220	0.063	0.063	0.042	0.021	0.031
MSE	0.049	0.016	0.016	0.013	0.009	0.009

Notes: Results are based on 1,000 simulations of 10,000 observations each.

MR-LATE estimator. Since R can only take on integer values from $-m$ to m, k must be an integer less than m. Specifically, we set $m = 6$ and consider $k = 0, 1, 2$, and 3.[9] Results of this experiment are reported in Panel A of Table 2. Note that for $k = 0$, MR-LATE is numerically identical to the 2SLS estimator. As the table shows, we find that MR-LATE is not sensitive to the choice of k and performs better than the other estimators for all $k > 0$.

For our last set of experiments, we again assume the econometrician can observe Y, X, Z, and R, but now R is given by $R = R^* + \varepsilon$, where ε is either normal with variance 0.04 (the same variance as U), or ε is distributed as a truncated normal ranging between $-\kappa$ and κ. We may interpret ε either as measurement error in R^*, or as the specification and estimation error in a structural model estimate of R^*.[10] Once again, feasible estimators include ignoring the measurement error in R and applying OLS or 2SLS using $T = \mathbb{I}(R \geq 0)$ in place of D. We compare these to MR-LATE, using $T^a = \mathbb{I}(R \geq \kappa)$ and $T^b = \mathbb{I}(R < -\kappa)$ for a few different values of κ. In Panel B of Table 2, we report results where ε is bounded, with bounds κ and $-\kappa$ known to the econometrician. In these cases, MR-LATE point-identifies LATE. In Panel C of Table 2, ε is normal so bounds are nonexistent. In

[9] For this particular DGP, R turns out to be highly correlated with R^*, resulting in the proxy T having about a 0.90 correlation with D. So this example represents a case of relatively little measurement error. Our empirical application likely has more measurement error than this.

[10] It is worth pointing out that we set the variance of the measurement error such that it is sufficiently large to generate misclassification, but sufficiently low to make T^a and T^b informative about D. This is analogous to a common real world situation where misclassification of the treatment indicator matters but does not lead to a sample where the share of actual treated in T^a is smaller than the share of misclassified actual untreated, and analogously, where the share of actual untreated in T^b is smaller than the share of misclassified actual treated. Hence, in our simulations, the sufficient condition of Corollary 3 is satisfied and MR-LATE can be used (at least) to set identify LATE.

this case MR-LATE does not point-identify LATE, but can be arbitrarily close asymptotically for large values of κ. Both with known bounds and nonexistent bounds, we find that MR-LATE performs considerably better than OLS and 2SLS.

As discussed in section 3.1, in the case of unknown or nonexistent bounds, the practitioner faces a trade-off when choosing κ. The larger the chosen value of κ, the closer p_0^a and p_1^b are to zero, but the farther $\widehat{\lambda}^a$ and $\widehat{\lambda}^b$ are to their limiting values at any given sample size. The numerical results presented here are reassuring in showing that as long as κ is not chosen to be either very small or very large, one obtains a good approximation of the true LATE (much better than the standard 2SLS estimator) over a wide range of possible choices of κ. When bounds don't exist, MR-LATE is biased, though the bias can be reduced by increasing κ as the sample size grows. We examine the performance of MR-LATE as the sample size increases (results are in Table A1 in the Appendix). We find that while the bias in MR-LATE does indeed shrink as the sample size increases, this happens quite slowly. The MR-LATE bias, however, remains very small compared to the bias in the standard 2SLS LATE estimator using T in place of D.

4 Women's Intra-household Power and Family Health

We apply the MR-LATE estimator to study the impact of women's intra-household empowerment on family members' health outcomes in India. As primary caregivers, the greater is a woman's decision power and control over resources, the more effective her care for herself and her children can be (Smith (2003)). We therefore define our unobserved true treatment indicator D to equal one if the wife has primary control of resource allocation decisions in the household, and zero otherwise. Formally, we define $D = \mathbb{I}(R^* \geq e)$, where R^* is a woman's unobserved decision making power (or the share of household resources that she controls) and e is a threshold that might vary across households for unobserved reasons.

We examine the impact of treatment D on a variety of health related outcomes Y. For adults, we consider body mass index (hereafter BMI), an indicator for being underweight, and an indicator for being anemic. For children, we consider height-for-age and weight-for-age z-scores, recent occurrences of diarrhea, fever and cough, and an indicator for whether a child has been vaccinated against one or more diseases. Our instrument Z is based on inheritance law reforms that equalized women's inheritance rights to men's in several Indian states between 1976 and 2005.

As discussed in section 3.2, there exist a variety of indicators of women's status and control over resources that might be used to measure treatment, including self-reports of decision making power. However, these measures are quite crude and imprecise, usually focused on very few specific decisions. These self-reports are examples of what we called P_j in the previous section. A typical LATE study might use these survey responses (or an index of them) as direct measures of treatment, despite the fact that they are at best crude proxies for women's overall decision making power. In fact, in our dataset we find a quite low correlation between these self-reports and the instrument

Z.[11] Thus, we find that the use of these self-reports leads to weak and uninformative estimates of treatment effects, regardless of whether they are taken to equal treatment for ordinary LATE estimation, or if they are used for MR-LATE estimation. These empirical results are reported in Appendix A.3.

For our primary analysis, we instead employ a structural model that makes use of both economic theory (of collective household decision making), and detailed household expenditure data, to estimate the percentage of resources in the household that a woman controls (that is, her resource share). MR-LATE estimation then allows us to consistently estimate a LATE associated with treatment defined in terms of resource shares, even though this estimated treatment may suffer to some extent from both model misspecification error and estimation error.

In contrast to the structural model used to estimate resource control, which is grounded in economic theory, we do not model health production as a function of control over resources. One could imagine many possible mechanisms linking the unobserved resource control treatment indicator D to health outcomes Y. For example, mothers may have different priorities than fathers regarding expenditures on health related goods, or circumstances that contribute to women having greater power might also affect the health care of family members. Previous research (discussed in section 2) provides indirect evidence of linkages between D and Y, by establishing, e.g., how variables that affect women's power correlate with health outcomes of wives and children. Our goal is therefore to directly quantify a treatment effect of D on Y, using a causal rather than structural model, exploiting plausibly exogenous variation in an instrument Z that correlates with D.

Note that this analysis assumes that D defined this way is a relevant measure of treatment for health outcomes. To the extent that health outcomes are the results of many health decisions, it is reasonable to assume that the family member with the most power will generally determine most of these decisions. This is roughly analogous to voter models, where outcomes are primarily determined by the party with the most power. For example, a woman who has primary control over the household's resources may be able to make timely decisions to treat herself or a sick child after discovering an illness, or to more easily make use of health services and follow through with treatment recommendations.[12]

An alternative possibility is that the magnitude of R^*, not just D, is more relevant for determining Y. To test this, and for comparison with our main results, we estimate a model that linearly regresses Y on R (our estimate of R^*) and other covariates. The results are reported in Appendix A.3. We find there is little to no evidence of a linear relationship between women's resource shares and health outcomes, suggesting that what matters is indeed who has control over a substantial fraction of household resources, i.e., D, and not the resource shares themselves. One could also estimate a more flexible, non-linear model, where R^* is replaced with R and then instrumented using Z. Note, however, that the estimation of such models might still lead to biased estimates. In fact, we have

[11] This is in line with results in Heath and Tan (2014), who document no significant link between self-reported participation in household decisions and these inheritance reforms.

[12] Moreover, changes in behavior may occur when women control a majority of household resources. Bertrand et al. (2015), for instance, study the consequences of relative income within households in the United States and document substantial changes in social and economic outcomes following the violation of the prescription that "a man should earn more than his wife."

no reasons to believe that errors in R satisfy the classical measurement error assumptions. One important reason for discretizing the treatment the way we do is to successfully tackle estimation issues related to non-classical measurement error.

Note that by definition, LATE averages over all random variables that affect outcomes, conditional on treatment. In our model where the treatment D is defined as $\mathbb{I}(R^* \geq e)$, this means that the average treated outcome includes averaging over R^* for all compliers who have $R^* \geq e$, while the average control outcome includes averaging over R^* for all compliers who have $R^* < e$. As reported below, with e centered at 50 percent, we estimate (based on R, since R^* is unobserved) that the treated group has an average R^* of around 60 percent, while the control has group has an average R^* of approximately 40 percent. So the average woman in the treatment group has much more power than the average woman in the control group, which is important to recognize for interpreting the results.

4.1 Modeling Resource Shares

A key difficulty in observing or calculating a woman's resource share R^* is that most goods in a household can be shared or consumed jointly to some extent by household members. For example, home heating is almost completely shared, while cooking fuel is jointly consumed just among household members who are eating together. Other goods, like food, are consumed individually, but it is difficult to track exactly who eats what within the household. Define a good to be *private* if it is not shared or consumed jointly. Define a good to be *assignable* if it appears in just one (known) household member's utility function, and so is only consumed by that household member. If all goods were private and assignable, then we could potentially directly observe R^*. The difficulty is that most goods are not assignable or private. In addition, expenditure surveys generally collect consumption data at the level of households, not of individuals.

We apply the structural methodology developed by Dunbar et al. (2013) (hereafter DLP) to obtain R, an estimate of R^*, which we then use to construct the treatment indicators T^a and T^b required for MR-LATE estimation. Assume we observe the household's total expenditure on all goods, M, and we observe the household's expenditures on (at least) one private, assignable good for each decision maker in the household. Let scalars w_w and w_m denote the household's budget shares (fraction of total expenditures M) spent on the observed private, assignable goods, which in our data are women's clothes and men's clothes, respectively.

We cannot just use w_w and w_m as measures of R^* and $1-R^*$, because men and women may have very different tastes for clothing. For example, a wife might consume fewer household resources than her husband, but still consume more clothes than him, because she derives more utility from clothing consumption than her husband does. Following DLP, we instead identify and estimate a separate clothing Engel curve for each decision maker. The Engel curve for person t is defined to be that person's demand for clothing, expressed as a function of the total resources person t controls: R^*M for women and $(1-R^*)M$ for men. Then, given w_w, w_m and M, we implicitly invert these

Engel curves to solve for R^*.

Let $X = (X_1, ..., X_K)$ denote a vector of observable attributes of households and their members. Household attributes X may affect the preferences of each household member and may also affect the household's bargaining process or social welfare function, and as a result may directly affect resource shares.[13] We employ the commonly used piglog (price independent generalized logarithmic) functional form for these Engel curves, which is

$$\begin{cases} w_w &= R^*\delta_w + R^*\beta \ln(R^*M) \\ w_m &= (1-R^*)\delta_m + (1-R^*)\beta \ln((1-R^*)M) \end{cases} \tag{11}$$

where $\beta = \beta(X)$, $\delta_w = \delta_w(X)$, $\delta_m = \delta_m(X)$, and $R^* = R^*(X)$ (see Appendix A.4 for details).[14] Note that the demand functions for other goods (those that are not private and assignable) are more complicated, but are not required to estimate the resource shares. DLP prove that the functions $\beta(X)$, $\delta_w(X)$, $\delta_m(X)$, and $R^*(X)$ are identified in this model.[15] For our empirical application, we assume the functions $\beta(X)$, $\delta_w(X)$, $\delta_m(X)$, and $R^*(X)$, are all linear in their arguments.[16] In particular, we specify

$$R^*(X) = \theta_0 + \theta_1 X_1 + ... + \theta_K X_K. \tag{12}$$

We append an error term to the equations in system (11), yielding a two equation system that we estimate using non-linear Seemingly Unrelated Regressions (SUR).[17]

Recognizing that the resulting estimates may suffer from both misspecification and estimation errors, we let R_i denote the resulting estimate of $R^*(X_i)$ for each woman i, and apply MR-LATE estimation using these estimates R_i.

4.2 Estimating the Unobserved Treatment and MR-LATE

For our empirical analysis, we employ two different datasets from India. One, the 62^{nd} round of the NSS Consumer Expenditure Survey (hereafter NSS), contains detailed consumption data that we use for estimating the above model of resource shares. The other, the third round of the National

[13] In the collective household model literature, covariates that only affect the household's bargaining process but not the tastes of the household members are known as *distribution factors*. A feature of the DLP approach is that it does not require observation of distribution factors. However, if any of our covariates are distribution factors, then they would affect R^* but not the other parameters.

[14] Jorgenson et al. (1982) Translog demand system and the Deaton and Muellbauer (1980) Almost Ideal Demand System have Engel curves of the piglog form, and piglog Engel curves were also used in empirical collective household models estimates by DLP. DLP estimate resource shares for children as well as the adults in the household. Due to data availability (see section 4.3), we can only estimate resource shares of the mother and father. Our framework still allows for caring preferences and for the possibility that mothers and fathers may value differently the well-being of their children.

[15] The identification depends partly on the assumption that $\beta(X)$ is the same for men and women. DLP call this the SAP (similar across people) assumption, and provide empirical evidence supporting this restriction.

[16] We do not include Z as an element of X (particularly in the equation for R^*) for two reasons. First, doing so could induce spurious correlation between the estimated treatment indicator and the instrument. Second, the NSS expenditure data does not include information on women's year of marriage, which is required to construct an exact measure of exposure to the inheritance law reforms and hence Z in the NSS dataset. However, we acknowledge that this might lead to a violation of Assumption 2. We therefore repeat our analysis by including a measure of women's *eligibility* to the amendments, defined as the interaction between an indicator variable for being Hindu, Buddhist, Sikh or Jain, and an indicator variable equal to one if a woman was 14 or younger at the time of the amendment in her state and to zero if she was 23 or older (see Heath and Tan (2014) and Calvi (2016)). Results are confirmed and available upon request.

[17] The non-linear SUR is iterated until the estimated parameters and the covariance matrix converge. The result is asymptotically equivalent to maximum likelihood with multivariate normal errors.

Family Health Survey (hereafter NFHS), reports the health outcomes we use for our causal treatment effects estimation. Both surveys were conducted between 2005 and 2006, and the covariate vector X (of attributes of households and their members) is observed in both datasets.

Let $\widehat{\theta}$ denote the estimate of θ in equation (12), which we obtain as described in the previous subsection, using the NSS data. Then, for each individual i drawn from the NFHS data, we take the estimates $\widehat{\theta}$ obtained from the NSS data, and use them to construct an estimate of the share of resources controlled by the woman in individual i's household as

$$R_i = \widehat{\theta}_0 + \widehat{\theta}_1 X_{1i} + ... + \widehat{\theta}_K X_{Ki}.$$

Our goal is to estimate a LATE of $D = \mathbb{I}(R^* \geq e)$ for a range of health outcomes Y. We separately consider health outcomes for mothers, fathers, and children. So, e.g., when i is a child and Y_i is an indicator of whether the child has been vaccinated, the treatment effect we wish to estimate is the change in i's probability of being vaccinated if he/she is exposed to highly empowered mothers, corresponding to $D = 1$. We wish to estimate this treatment effect, even though the mother's true resource share R_i^* is unobserved.

As discussed in Example 2 in section 3.2, we apply our MR-LATE estimator by constructing two mismeasures of treatment, i.e., $T_i^a = \mathbb{I}(R_i \geq \kappa^a)$ and $T_i^b = \mathbb{I}\left(R_i < \kappa^b\right)$, where κ_a and κ_b are chosen bounds with $\kappa_a > \kappa_b$. For $j = a, b$, the estimation procedure consists of regressing $Y_i T_i^j$ on a constant, T_i^j, and X_i using 2SLS (with Z_i being the excluded instrument).[18] The MR-LATE parameter is then obtained as the difference between the estimated coefficients of treatment in these two 2SLS regressions, that is $\widehat{\rho} = \widehat{\lambda}^a - \widehat{\lambda}^b$.

The way we choose bounds κ_a and κ_b is as follows. We choose a percentage $\mathscr{K}$, and let κ_a be the value such that $\mathscr{K}/2$ percent of the sample has R in the interval $[50, \kappa_a]$ and κ_b is the value such that $\mathscr{K}/2$ percent of the sample has R in the interval $[\kappa_b, 50]$. This is consistent with (but does not require) having e being centered around 50 percent, implying that households with $D_i = 1$ are usually ones in which the mother has control over a majority of household resources.[19] Essentially, $\mathscr{K}$ corresponds to how much misclassification error T^a and T^b would contain if R was exactly equal to R^*. Recall that we are unable to directly observe or estimate κ_a and κ_b. While we cannot be sure that the sufficient condition for point identification holds exactly, it is quite likely that the sufficient conditions for set identification are indeed satisfied. Therefore, following Corollary 3, we can interpret the MR-LATE estimates as providing upper bounds for the true LATE.

We apply this procedure separately for a few different health outcomes for men, women, and children, using a few different values of $\mathscr{K}$. In the special case of $\mathscr{K} = 0$, MR-LATE becomes numerically identical to the standard Imbens and Angrist (1994) 2SLS LATE estimator, using the mismeasured $T_i = \mathbb{I}(R_i \geq 50)$ in place of the unobserved true $D_i = \mathbb{I}\left(R_i^* \geq e_i\right)$. Generally, this $\mathscr{K} = 0$ estimator will only be consistent if there is no measurement or estimation error in R_i and if e_i exactly

[18] Results are robust to including R in the linear model for Y alongside with the mismeasures of treatment and are available upon request.

[19] We recall from section 3.2 that the relevant threshold e_i can vary across households for unobserved reasons. In effect, random variation in e is observationally equivalent to measurement error and therefore addressed by our estimation method.

equals 50 percent for all households.

What we mainly require for consistency of MR-LATE is that $\mathscr{K}$, and hence κ_a and κ_b, be large enough so that any complier who has $D = 0$ will also have $T^a = 0$, and any complier who has $D = 1$ will also have $T^b = 0$. However, as discussed in section 3.3, the larger the values of κ_a and κ_b, the less informative T_a and T_b are of the true unobserved treatment D. As a result, an excessively small value of $\mathscr{K}$ will lead to bias relative to the true LATE, while an excessively large value of $\mathscr{K}$ will lead to imprecision (high variance) of the estimates.

4.3 Data

We implement our empirical analysis using the 2005-2006 India National Sample Survey (NSS) of Consumer Expenditure (62^{nd} round) and the 2005-2006 India National Family Health Survey (NFHS-3). Table A2 and A3 in the Appendix present some descriptive statistics.[20]

NSS data. The 2005-2006 NSS Consumer Expenditure Survey contains detailed data on household expenditures, socio-economic characteristics, and other particulars of household members. We select households consisting of a mother, a father, and one to four children.[21] Among other items, households are asked to report how much they spent on clothing and footwear. Given the detailed breakdown of clothing expenditure, it is possible to identify the expenditures on some items of clothing that can be specifically assigned to women and to men, thereby allowing us to construct expenditures on private assignable clothing for each decision maker.[22] Table A2 in the Appendix contains some descriptive statistics. For clothing items the NSS reports expenditures that occurred in the past 365 days. For simplicity and consistency with other data, we convert these annual expenditures into monthly figures. We consider covariates that characterize each individual, the household, and the environment of the household. Specifically, our covariates include the gender composition of children, wife's age, the age gap between spouses, the average age of children, and dummies for the number of children, geographic region, religion (Hindu, Buddhist, Sikh or Jain), for living in rural areas, for female and male higher education, and for belonging to a Scheduled Caste, Scheduled Tribe, or other backward classes.

NFHS data. The 2005-2006 National Family Health Survey provides a range of health indicators for women aged 15 to 49, for men aged 15 to 54, and for children born in the 5 years

[20] As shown in table A2, with the exception of a few variables, the household socio-economic characteristics are on average quite similar in the two samples. The main differences are related to the definitions of completed schooling and land ownership in the two surveys. Moreover, the NFHS covers the 29 states in India, while the NSS includes both the 29 states and the 7 union territories of India. Any errors introduced by the use of two different samples will take the form of estimation error in R, and so should be accounted for by the MR-LATE estimator.

[21] More precisely, we select households with one woman and one man above age 15 (with one of these designated as the head of household), and from 1 to 4 children under 15. We exclude households in the top or bottom 1 percent of expenditure, and we exclude households that report having performed any ceremony during the month prior to the survey, as unusual purchases of clothing items and non-standard expenditure patterns may occur for festivities and ceremonies. The final estimation sample contains 7,480 households.

[22] We define expenditure on women's assignable clothing as the sum of expenditures on saree, chaddar, dupatta, and shawl. For men's assignable clothing, we combine expenditure on dhoti, lungi, salwar, pajamas, and shirts. Notice that Tommasi and Wolf (2018) shows that if the data exhibit relatively flat Engel curves in the consumption of the private assignable goods, then the DLP model can be weakly identified. However, households in our dataset display a large variation in the consumption of private assignable goods (see Figure A1 in the Appendix). Hence, we do not appear to have a weak identification problem with our data.

prior to the date of interview. The survey also contains many covariates, comparable to those we observe in the NSS data. As above, we select households consisting of a mother, a father, and one to four children. We consider women, men and children datasets separately, observing a few different health measures for each individual. The health measures for adults include anthropometrics like BMI (weight in kilograms divided by height in meters squared) and measures of anemia. A cut-off BMI point of 18.5 is used to define thinness or evidence of undernutrition. Anemia is a condition in which the number of red blood cells, or their oxygen-carrying capacity, is insufficient. Although its primary cause is iron deficiency, it often coexists with, and hence serves as an indicator of, a number of other health issues such as malaria, parasitic infection, and nutritional deficiencies. For children the health related measures we observe include weight-for-age and height-for-age z-scores (standard deviations from the reference median based on the 2006 WHO Child Growth Standards). A z-score greater than 2 indicates over-nourishment with respect to the corresponding anthropometric measurements. Deficits on these indicators (measured by their values less than -2 standard deviations below the median) are known as stunting and underweight respectively. Another child health measure we observe is mothers' reports of whether a child was sick with fever, cough or diarrhea in the past two weeks. Finally, we observe child vaccination records, which we use to construct an additional indicator variable equal to one if a child has ever received any vaccine to prevent diseases.[23]

The Hindu Succession Act and its Amendments. We exploit changes in the Indian inheritance law to construct a plausibly unconfounded instrumental variable Z. A woman's right to inherit land and other property is often claimed to play a significant role in determining women's position within the household (World Bank, 2014). Inheritance rights in India differ by religion and, for most of the population, are governed by the Hindu Succession Act (HSA). The HSA was first introduced in 1956 and only applied to Hindus, Buddhists, Sikhs, and Jains, in all states other than Jammu and Kashmir.[24] Before then, the traditional systems (Mitakhshara and Dayabhaga) were strongly biased in favor of sons (Agarwal, 1995). Gender inequalities, however, remained even after the introduction of the HSA. On one hand, in the case of a Hindu male dying intestate (without leaving a will) all his separate or self-acquired property devolved equally upon sons, daughters, widow, and mother. On the other hand, the deceased's daughters had no direct inheritance rights to joint family property, whereas sons were given direct right by birth to belong to the coparcenary. In the decades following the introduction of the HSA, state governments passed amendments that equalized inheritance rights for daughters and sons (Kerala in 1976, Andhra Pradesh in 1986, Tamil Nadu in 1989, and Maharashtra and Karnataka in 1994). A national-level ratification of the amendments occurred in 2005. However, these amendments only applied to Hindu, Buddhist, Sikh or Jain women who

[23] We observe whether a child received a BCG vaccine (against tuberculosis), one to three DPT vaccines (against diphtheria, pertussis, and tetanus), and one to four polio vaccines (at birth and one to three years after).

[24] The HSA did not apply to individuals of other religions, such as Muslims, Christians, Parsis, Jews, and other minority communities. Most laws for Christians formally granted them equal rights as of 1986. However, gender equality for Christian women was and is not the practice, as the Synod of Christian Churches has being arranging legal counsel to help draft wills to disinherit female heirs. The inheritance rights of Muslim women in India are governed by the Muslim Personal Law (Shariat) Application Act of 1937, under which daughters inherit only a fraction of what sons inherit (Agarwal, 1995).

were not yet married at the time of the amendment.

For each individual in our NFHS sample, we construct our instrumental variable Z as the indicator of whether the inheritance law reform applied to the woman in that individual's household. Whether Z equals one or not depends on the woman's religion, state of residence, and year of marriage, since exposure to the reform varies by these characteristics. We find that $Z = 1$ for 18 percent of women in the sample. Due to the gender age gap at marriage (on average 5 years), the percentage of men married to HSA exposed women is larger (28 percent). All specifications presented in our analysis include woman's cohort, religion, state, cohort-religion and state-religion fixed effects, together with state specific time trends up to degree four. The exclusion restriction needed for identification is that, once these fixed effects and time trends are included, being Hindu, Buddhist, Jain or Sikh and unmarried at the time of implementation has an effect on health outcomes only through women's higher control of household resources.[25]

Previous works have evaluated the HSA amendments using difference-in-difference methods (see, e.g., Roy (2008, 2015), Deininger et al. (2013), Heath and Tan (2014), Calvi (2016)). That type of analysis considers exposure to the HSA amendments as treatment. Our goal is not to estimate the treatment effect of this particular policy, but more broadly to estimate the health effect of living in a household where a woman controls a substantial fraction of resources, using exposure to these inheritance rights reforms as an instrument.

4.4 Empirical Results

In this section, we summarize our estimates of the resource share R, associated treatment measures, and the results of our causal analysis of the effect of women's empowerment D on health outcomes Y. Estimates of the Engel curves of women's and men's private assignable clothing, used to construct R, are reported in Table A11 in the Appendix.

Table 3 contains descriptive statistics of the predicted resource shares R in the NSS and NFHS samples. Our tables present R in percentage form, so e.g., the reported mean of R of 48 means the wife is estimated to control 48 percent of the household's resources. The reported summary statistics vary somewhat across the two samples, because they entail averages over the empirical distributions of the covariates $(X_1, ..., X_K)$ in each sample. It is therefore reassuring that the estimated means and standard deviations of R in the two samples are very similar, indicating that the samples are highly comparable. It is also reassuring that the minima and maxima of the estimated resource shares do not fall outside the zero to 100 percent range for all households, despite them being modeled as linear (and hence not bounded) functions of household characteristics X. In both samples, the share of resources controlled by the mother is slightly lower than that for fathers, with the mean of R (the estimated percentage of household resources controlled by the mother) equaling 48 percent

[25] This assumption is plausible. Despite other factors and policies may have differentially affected young Indian women, we do not expect these to vary by religion. Moreover, Calvi (2016) demonstrates that the potential endogeneity of women's time of marriage is not a concern. She also shows that the HSA amendments changed neither sorting in the marriage market nor total household expenditures. Finally, the existence of other changes *resulting from* women's increased power within their marital families due to their HSA exposure (e.g., higher labor force participation; see Heath and Tan (2014)) does not determine a violation of the exclusion restriction.

Table 3: Estimated Resource Shares and Woman's Power

	Obs.	Mean	St. Dev.	Min.	Max.
NSS Sample :					
Woman's Resource Share (R)	7,440	48.27	11.64	13.54	86.92
$T = \mathbb{I}(R \geq 50)$	7,440	0.41	0.49	0.00	1.00
NFHS Sample:					
Woman's Resource Share (R)	22,767	48.13	11.81	6.50	86.97
$T = \mathbb{I}(R \geq 50)$	22,767	0.40	0.49	0.00	1.00

Note: Household level data. R is in percentage form.

in both data sets. Interestingly, the estimates accord with our ex ante expectations. For instance, the average R is particularly high in the North-East states (61 percent), which is consistent with the presence of a number of matrilineal societies and cultures in these regions (Khasi and Garo societies, for example). In contrast, North Indian women seem to have a much lower control over resources (40 percent). Finally, highly educated women (who have completed high school) are found to have a substantially higher command over resources (55 percent) relative to low educated women (47 percent).

Our best estimate of D is $T = \mathbb{I}(R \geq 50)$. In the NFHS sample, women who have $T = 1$ have an average R of 60, while those having $T = 0$ have an average R of 40. So, while we cannot know the average fraction of resources controlled by the truly treated and untreated, i.e., $E(R^* \mid D)$, our estimates of $E(R \mid T)$ indicate that the treated group controls a considerably larger fraction of household resources than the control group. It is important to note that our estimate R of the true R^* refers to resources controlled by the woman, not necessarily those consumed by the woman. For example, mothers and fathers may value differently the well-being of their children, and so they might allocate different fraction of the resources they control to children.[26]

Figure A2 in the Appendix shows the empirical distribution of R in the NFHS sample (the distribution in the NSS is very similar). In section A.3 of the Appendix, we document a clear positive relationship between our estimated R and some self-reported measures of decision making power that are available in the NFHS dataset. Moreover, we show that women are more likely to self-report participating in household decisions in households that have $T = 1$ vs. $T = 0$, even after conditioning on individual and household level controls, fixed effects, and state time trends. These results verify that our expenditure based estimates of control of resources, R and T, do indeed correspond to measures of decision making power in the household.

As discussed in section 3.2, if we ignored specification and estimation error in R, we would apply the usual Imbens and Angrist (1994) LATE estimator (corresponding to $cov(Y,Z)/cov(D,Z)$), by replacing the the true unknown D with our best guess $T = \mathbb{I}(R \geq 50)$. However, this will likely lead to biased estimates due to measurement, estimation, and specification errors in R. We therefore apply our MR-LATE estimator to account for these errors.

[26] Using a more recent round of the NSS of Consumer Expenditure that includes a richer set of assignable goods, Calvi (2016) estimates separate resource shares for men, women and children. For households with children Calvi estimates that women' resource shares are on average only 67 percent of men's, which suggests that, relative to fathers, a higher fraction of resources controlled by mothers is diverted to children.

Table 4: Adult's Health

	Women			Men		
	BMI	Pr(BMI≤18.5)	Pr(Anemic)	BMI	Pr(BMI≤18.5)	Pr(Anemic)
MR-LATE ($\mathscr{K}=0$)	9.7989	-0.9175	-0.5572	2.4074	-0.2778	0.0860
	(1.9869)	(0.2288)	(0.2531)	(2.4452)	(0.2971)	(0.2318)
MR-LATE ($\mathscr{K}=1$)	9.2903	-0.8836	-0.5194	2.0658	-0.2345	0.0556
	(2.1346)	(0.2200)	(0.2420)	(2.7973)	(0.2785)	(0.2181)
MR-LATE ($\mathscr{K}=5$)	9.1945	-0.8482	-0.5239	3.0378	-0.2629	0.0925
	(4.1294)	(0.2225)	(0.2487)	(4.3572)	(0.2926)	(0.2195)
MR-LATE ($\mathscr{K}=10$)	12.4103	-1.1476	-0.6254	3.7843	-0.3852	0.0558
	(6.2875)	(0.2715)	(0.3124)	(7.4517)	(0.3633)	(0.2630)
MR-LATE ($\mathscr{K}=20$)	7.7153	-0.7232	-0.5151	-1.0725	-0.5217	-0.1373
	(9.1580)	(0.2915)	(0.3730)	(15.7390)	(0.4454)	(0.3185)

Notes: Estimates are obtained using the NFHS-3 data. The women sample includes married women of age 15 to 49 in nuclear households with up to 4 children. The men sample includes married men of age 15 to 54 in nuclear households with up to 4 children. All specifications include an indicator variables for being Hindu, Buddhist, Sikh or Jain, for region of residency, for number of children, rural areas, for being part of Scheduled Castes, Scheduled Tribes or Other Backward Classes, land ownership, woman's and man's high school completion, the fraction of female children, woman's and man's ages and average age of children 0-14. All specifications include state-religion and cohort-religion fixed effects, and state specific time trends (up to degree four). Anemia includes severe and moderate anemia. Bootstrap standard errors in parentheses.

Table 5: Children's Health

	Weight-for-age (z-score)	Height-for-age (z-score)	Pr(Cough)	Pr(Fever)	Pr(Diarrhea)	Pr(Any Vaccination)
MR-LATE ($\mathscr{K}=0$)	2.0547	2.8140	-0.6649	-0.6114	-0.4549	-0.1296
	(1.4246)	(1.8384)	(0.2973)	(0.4423)	(0.2185)	(0.2797)
MR-LATE ($\mathscr{K}=1$)	2.0093	2.7298	-0.7264	-0.6565	-0.5153	-0.2065
	(1.3814)	(1.7865)	(0.3176)	(0.4551)	(0.2255)	(0.2883)
MR-LATE ($\mathscr{K}=5$)	1.7328	3.0385	-0.7306	-0.6273	-0.4923	-0.2159
	(1.3947)	(1.6276)	(0.2868)	(0.4259)	(0.1981)	(0.3159)
MR-LATE ($\mathscr{K}=10$)	2.4247	3.0848	-0.8890	-0.6361	-0.5121	-0.1029
	(1.7819)	(1.8369)	(0.3115)	(0.4158)	(0.2023)	(0.4178)
MR-LATE ($\mathscr{K}=20$)	2.2458	2.8007	-0.6878	-0.4312	-0.6457	0.0141
	(1.6342)	(1.8806)	(0.3053)	(0.3345)	(0.1789)	(0.5104)

Notes: Estimates are obtained using the NFHS-3 data. The sample includes children 0 to 5 in nuclear households with up to 4 children. All specifications include an indicator variables for being Hindu, Buddhist, Sikh or Jain, for region of residency, for number of children, rural areas, for being part of Scheduled Castes, Scheduled Tribes or Other Backward Classes, land ownership, parents' high school completion, the fraction of female children, parents' ages, the child's age and gender. All specifications include state-religion and cohort-religion fixed effects for the mother, and state specific time trends (up to degree four). Bootstrap standard errors in parentheses.

To apply MR-LATE, we construct bounds κ_a and κ_b based on choosing a misclassification percentage $\mathcal{K}$ as described in section 4.2. The percentages $\mathcal{K}$ we consider are $\mathcal{K}$ equal to 0, 1, 5, 10, and 20. Table A4 in the Appendix reports the bounds κ^a and κ^b that correspond to each of these values of $\mathcal{K}$. The values of κ^a and κ^b vary across the subsamples of women, men and children in the NFHS due to variation in the distribution of covariates.

Table 4 reports the resulting MR-LATE estimates for adult health outcomes, while estimates for children's health outcomes are reported in Table 5. Bootstrapped standard errors are reported in parentheses. The MR-LATE estimates for $\mathcal{K} = 0$ are numerically identical to the standard LATE estimator that ignores errors in R. Even quite small deviations of $\mathcal{K}$ from zero substantially change the MR-LATE estimates for some outcomes, showing that accounting for errors in R appears to be empirically important.

Overall, our MR-LATE estimates indicate that a woman's control of household resources exerts a positive and significant effect on her own health. Women with high control over household resources have a much higher BMI and face a lower likelihood to be underweight or anemic. The estimated effects are sizable: our most conservative estimates indicate that the average treatment effect on women's body mass index is 7.7 and that women in treated households are 72 percent percent less likely to be underweight and 52 percent less likely to be anemic. This is not surprising. As the sufficient conditions for set identification are likely to be satisfied, the MR-LATE estimates provide upper bounds for the true LATE. Moreover, even when point-identification is achieved and MR-LATE equals LATE, the large magnitudes are consonant with us comparing women who on average have control of about 60 percent of household resources with women who on average control about 40 percent of household resources.

We find mother's control over household resources also positively affects children's health. A mother's high control over resources decreases her children's likelihood of being sick with cough, fever or diarrhea in the two weeks prior to the survey by 66, 43 and 45 percent, respectively. It also leads to higher children's height-for-age and weight-for-age, though these effects are not significantly different from zero. By contrast, we do not find any positive (or negative) effect of a wife's control of resources on her husband's health.

Table A5 in the Appendix shows results of the first stage of the MR-LATE estimates for the different values of κ^a and κ^b considered above, together with the corresponding F-statistics. We include household level and individual level characteristics, fixed effects and state specific time trends in all specifications. Even conditioning on several sources of unobserved heterogeneity, the instrument Z is positively and significantly correlated with T^a and T^b. The first stage F-test statistics are largely above 10 for $\mathcal{K} = 0$, 1, 5, 10, which is consistent with not having a weak instrument problem. The F-test statistics, however, do fall below 10 for our largest bounds, corresponding to $\mathcal{K} = 20$. This is consistent with our predictions, since T^a and T^b become less correlated with the true D and $1-D$ once the bounds κ^b to κ^a become overly large.

Overall, our empirical results indicate that policies aimed at empowering women within households, such as strengthening their inheritance and property rights, tend to increase their control

over household resources. This increased control over resources leads to improvements in women's and children's overall health, while having little effect on men's health.

4.5 External Validity

Instead of going to the trouble of estimating resource shares, we could have simply calculated the ATE of the change in inheritance laws Z on the outcomes Y. However, our interest is not in these particular inheritance policies. Rather, we wish to learn about the likely impact of any policy that changes women's power within the household, as measured by control over resources. This is value of defining treatment D the way we do.

Given our assumptions, MR-LATE consistently estimates the average treatment effect for compliers (those for whom treatment D and the inheritance policy Z are the same binary random variable). As with ordinary LATE estimation, the question then remains: how representative are compliers of the general population and, hence, how close are our estimates to the population ATE of empowering women (by giving them control of household resources) on family health?

In Appendix A.5, we discuss the conditions under which our compliers are close to representative of the general population. Exploiting the fact that our treatment is defined according to a threshold crossing model, we can cast our problem within the marginal treatment effect (MTE) framework of Heckman and Vytlacil (1999), where the relevant threshold e_i is a source of unobserved heterogeneity across households. We then make use of some recent results in Kowalski (2016) (and references therein) to implement a test of external validity of our LATE estimates. While this analysis abstracts from the issue of measurement or estimation error in R, and therefore needs to be interpreted with caution, the results in Appendix A.5 suggest that our estimated treatment effects are unlikely to vary much with the choice of instrument, lending some empirical support for the external validity of our estimates.

5 Conclusion

We propose a novel approach to study the effects of intra-household women's empowerment on the health status of family members in India, using a change in inheritance laws as an instrument. Our model looks at the effect on health outcomes Y of a treatment D, defined as a woman having relatively high control over household resources. The treatment is based on an unobservable continuous variable R^*. We rely on a structural model of household consumption to obtain R, an estimate of R^*, and use R to construct estimated treatment indicators. Due to measurement, estimation, or specification errors in the structural model for R, estimated treatment indicators may not equal or consistently estimate the true treatment indicator D. To account for these several possible sources of error, we propose a new mismeasurement-robust LATE estimator, called MR-LATE, which uses two estimated treatment indicators, T^a and T^b, along with an outcome Y and an instrument Z, to obtain consistent estimates of the same LATE that would be obtained if we could observe the true

treatment indicator D.

Our empirical results suggest that policies aimed at increasing women's bargaining power within households should lead to improvements in overall family health. In particular, exogenously increasing a woman's control of resources within a household substantially improves both her own health and her children's, without damaging her husband's health.

Our application emphasizes the potential use of our new estimator in situations where treatment is not observed and must be estimated. However, we wish to stress that MR-LATE can be useful in other applications where a binary treatment indicator is simply observed with error, due for instance to mis-reporting.

More broadly, our analysis highlights potential advantages of combining both structural and causal, reduced form methodologies in conducting empirical analysis. The MR-LATE estimator specifically accounts for the fact that structural estimation generally suffers from multiple errors, including specification errors. But by exploiting structure, we can estimate causal effects of substantial economic interest and relevance. This may be particularly useful for constructing causal tests and benchmarks of economic models, since the researcher can directly focus on treatments that are motivated by theory (in our example, women's control of household resources), instead of only calculating the treatment effects of less relevant proxies that happen to be directly observed.

References

ABADIE, A. (2002): "Bootstrap Tests for Distributional Treatment Effects in Instrumental Variable Models," *Journal of the American Statistical Association*, 97, 284–292. [9], [11]

AGARWAL, B. (1995): "A Field of One's Own," *Cambridge Books*. [24]

AIGNER, D. J. (1973): "Regression with a binary independent variable subject to errors of observation," *Journal of Econometrics*, 1, 49 – 59. [6]

ANGRIST, J. D. (2004): "Treatment Effect Heterogeneity in Theory and Practice," *The Economic Journal*, 114, C52–C83. [43]

ANGRIST, J. D., G. W. IMBENS, AND D. B. RUBIN (1996): "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, 91, 444–455. [6]

ANGRIST, J. D. AND A. B. KRUEGER (1999): "Chapter 23 - Empirical Strategies in Labor Economics," Elsevier, vol. 3, Part A of *Handbook of Labor Economics*, 1277 – 1366. [5]

ANGRIST, J. D. AND J.-S. PISCHKE (2010): "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics," *Journal of Economic Perspectives*, 24, 3–30. [7]

APPS, P. F. AND R. REES (1988): "Taxation and the Household," *Journal of Public Economics*, 35, 355 – 369. [4]

ATHEY, S., R. CHETTY, G. IMBENS, AND H. KANG (2016): "Estimating Treatment Effects using Multiple Surrogates: The Role of the Surrogate Score and the Surrogate Index," *arXiv preprint arXiv:1603.09326*. [6]

ATTANASIO, O. AND V. LECHENE (2002): "Tests of Income Pooling in Household Decisions," *Review of Economic Dynamics*, 5, 720–748. [5]

BARGAIN, O. AND O. DONNI (2012): "Expenditure on Children: A Rothbarth-type Method Consistent with Scale Economies and Parents' Bargaining," *European Economic Review*, 56, 792–813. [42]

BARGAIN, O., O. DONNI, AND P. KWENDA (2014): "Intrahousehold distribution and poverty: Evidence from Cote d'Ivoire," *Journal of Development Economics*, 107, 262–276. [42]

BATTISTIN, E., M. D. NADAI, AND B. SIANESI (2014): "Misreported schooling, multiple measures and returns to educational qualifications," *Journal of Econometrics*, 181, 136 – 150. [3], [6]

BECKER, G. S. (1965): "A Theory of the Allocation of Time," *The Economic Journal*, 493–517. [2], [4]

——— (1974): "A Theory of Social Interactions," *Journal of Political Economy*, 82, 1063–1093. [2]

——— (1981): *A Treatise on the Family*, Harvard University Press. [2], [4]

BEEGLE, K., E. FRANKENBERG, AND D. THOMAS (2001): "Bargaining Power Within Couples and Use of Prenatal and Delivery care in Indonesia," *Studies in family planning*, 32, 130–146. [5]

BEGG, C. B. AND D. H. LEUNG (2000): "On the use of surrogate end points in randomized trials," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 163, 15–28. [6]

BERTANHA, M. AND G. W. IMBENS (2014): "External Validity in Fuzzy Regression Discontinuity Designs," Tech. rep., National Bureau of Economic Research. [43]

BERTRAND, M., E. KAMENICA, AND J. PAN (2015): "Gender Identity and Relative Income Within Households," *The Quarterly Journal of Economics*, 130, 571–614. [19]

BJÖRKLUND, A. AND R. MOFFITT (1987): "The Estimation of Wage Gains and Welfare Gains in Self-selection Models," *The Review of Economics and Statistics*, 42–49. [42]

BLACK, D., S. SANDERS, AND L. TAYLOR (2003): "Measurement of Higher Education in the Census and Current Population Survey," *Journal of the American Statistical Association*, 98, 545–554. [5]

BLACK, D. A., M. C. BERGER, AND F. A. SCOTT (2000): "Bounding Parameter Estimates with Non-classical Measurement Error," *Journal of the American Statistical Association*, 95, 739–748. [6]

BLUNDELL, R., P.-A. CHIAPPORI, AND C. MEGHIR (2005): "Collective Labor Supply with Children," *Journal of Political Economy*, 113, 1277–1306. [5]

BOLLINGER, C. R. (1996): "Bounding mean regressions when a binary regressor is mismeasured," *Journal of Econometrics*, 73, 387 – 399. [5], [6]

BOSE, N. AND S. DAS (2015): "Women's Inheritance Rights, Household Allocation and Gender Bias," *Unpublished*. [5]

BRINCH, C., M. MOGSTAD, AND M. WISWALL (forthcoming): "Beyond LATE With a Discrete Instrument," *Journal of Political Economy*. [42], [43]

BROWNING, M., P.-A. CHIAPPORI, AND A. LEWBEL (2013): "Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power," *Review of Economic Studies*, 80, 1267–1303. [3], [4], [5], [40]

CALVI, R. (2016): "Why Are Older Women Missing in India? The Age Profile of Bargaining Power and Poverty," *Mimeo*. [4], [5], [21], [25], [26]

CARD, D. (2001): "Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems," *Econometrica*, 69, 1127–1160. [5]

CARNEIRO, P., J. J. HECKMAN, AND E. J. VYTLACIL (2011): "Estimating Marginal Returns to Education," *The American economic review*, 101, 2754–2781. [42]

CHERCHYE, L., B. DE ROCK, A. LEWBEL, AND F. VERMEULEN (2015): "Sharing Rule Identification for General Collective Consumption Models," *Econometrica*, 83, 2001–2041. [42]

CHIAPPORI, P.-A. (1988): "Rational household labor supply," *Econometrica: Journal of the Econometric Society*, 63–90. [3], [4]

——— (1992): "Collective Labor Supply and Welfare," *Journal of Political Economy*, 100, pp. 437–467. [3], [4]

DEATON, A. AND J. MUELLBAUER (1980): "An Almost Ideal Demand System," *American Economic Review*, 70, 312–26. [21]

DEININGER, K., A. GOYAL, AND H. NAGARAJAN (2013): "Women's Inheritance Rights and Intergenerational Transmission of Resources in India," *Journal of Human Resources*, 48, 114–141. [5], [25]

DITRAGLIA, F. J. AND C. GARCÍA-JIMENO (2017): "Mis-classified, Binary, Endogenous Regressors: Identification and Inference," Tech. rep., National Bureau of Economic Research. [3], [6]

DUFLO, E. (2003): "Grandmothers and Granddaughters: Old-age Pensions and Intrahousehold Allocation in South Africa," *The World Bank Economic Review*, 17, 1–25. [2], [5]

DUNBAR, G. R., A. LEWBEL, AND K. PENDAKUR (2013): "Children's Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi," *American Economic Review*, 103, 438–71. [3], [4], [20], [40]

FRANGAKIS, C. E. AND D. B. RUBIN (2002): "Principal stratification in causal inference," *Biometrics*, 58, 21–29. [6]

FRANKENBERG, E., D. THOMAS, ET AL. (2001): *Measuring Power*, Citeseer. [5]

FRAZIS, H. AND M. A. LOEWENSTEIN (2003): "Estimating linear regressions with mismeasured, possibly endogenous, binary explanatory variables," *Journal of Econometrics*, 117, 151 – 178. [6]

HADDAD, L. J., J. HODDINOTT, H. ALDERMAN, ET AL. (1997): *Intrahousehold Resource Allocation in Developing countries*, Johns Hopkins University Press. [2]

HEATH, R. AND X. TAN (2014): "Intrahousehold Bargaining, Female Autonomy, and Labor Supply: Theory and Evidence from India," Mimeo, Department of Economics, University of Washington. [5], [19], [21], [25]

HECKMAN, J., S. URZUA, AND E. VYTLACIL (2006): "Understanding Instrumental Variables in Models with Essential Heterogeneity," *Review of Economics and Statistics*, 88, 389–432. [7]

HECKMAN, J. J. (2010): "Building Bridges between Structural and Program Evaluation Approaches to Evaluating Policy," *Journal of Economic Literature*, 48, 356–98. [7]

HECKMAN, J. J. AND E. VYTLACIL (2005): "Structural Equations, Treatment Effects, and Econometric Policy Evaluation," *Econometrica*, 73, 669–738. [42]

HECKMAN, J. J. AND E. J. VYTLACIL (1999): "Local Instrumental Variables and Latent Variable Models for Identifying and Bounding Treatment Effects," *Proceedings of the national Academy of Sciences*, 96, 4730–4734. [29], [42], [43]

——— (2007): "Econometric Evaluation of Social Programs, part I: Causal Models, Structural Models and Econometric Policy Evaluation," *Handbook of econometrics*, 6, 4779–4874. [7], [42]

HERNANDEZ, M., S. PUDNEY, AND R. HANCOCK (2007): "The welfare cost of means-testing: pensioner participation in income support," *Journal of Applied Econometrics*, 22, 581–598. [5]

HU, Y. (2008): "Identification and estimation of nonlinear models with misclassification error using instrumental variables: A general solution," *Journal of Econometrics*, 144, 27 – 61. [6]

HU, Y., J.-L. SHIU, AND T. WOUTERSEN (2015): "Identification and estimation of single-index models with measurement error and endogeneity," *The Econometrics Journal*, 18, 347–362. [6]

IMAI, K., L. KEELE, AND T. YAMAMOTO (2010): "Identification, Inference and Sensitivity Analysis for Causal Mediation Effects," *Statist. Sci.*, 25, 51–71. [6]

IMBENS, G. W. (2010): "Better LATE Than Nothing: Some Comments on Deaton (2009) and Heckman and Urzua (2009)," *Journal of Economic Literature*, 48, 399–423. [7]

IMBENS, G. W. AND J. D. ANGRIST (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62, 467–475. [3], [6], [7], [8], [9], [11], [22], [26]

JORGENSON, D., L. LAU, AND T. M. STOKER (1982): *The Transcendental Logarithmic Model of Aggregate Consumer Behavior*, Greenwich: JAI Press, 97–238, welfare 1, ch. 8, pp. 203-356. [21]

KANE, T. J., C. E. ROUSE, AND D. STAIGER (1999): "Estimating Returns to Schooling When Schooling is Misreported," NBER working paper 7235. [5], [6]

KLEPPER, S. (1988): "Bounding the effects of measurement error in regressions involving dichotomous variables," *Journal of Econometrics*, 37, 343 – 359. [6]

KOWALSKI, A. (2016): "How to Examine External Validity Within an Experiment," *Mimeo*. [29], [42], [43], [56]

KOWALSKI, A., Y. TRAN, AND L. RISTOVSKA (2016): "MTEBINARY: Stata module to compute Marginal Treatment Effects (MTE) With a Binary Instrument," Statistical Software Components, Boston College Department of Economics. [43], [56]

KREIDER, B., C. GUNDERSEN, AND D. JOLLIFFE (2012): "Identifying the effects of food stamps on child health outcomes when participation is endogenous and misreported," . [6]

KREIDER, B. AND J. V. PEPPER (2007): "Disability and Employment: Reevaluating the Evidence in Light of Reporting Errors," *Journal of the American Statistical Association*, 102, 432–441. [6]

LAFAVE, D. AND D. THOMAS (2017): "Extended Families and Child Well-being," *Journal of Development Economics*, 126, 52–65. [2], [5]

LEWBEL, A. (1998): "Semiparametric Latent Variable Model Estimation with Endogenous or Mismeasured Regressors," *Econometrica*, 66, 105–121. [6]

——— (2007): "Estimation of Average Treatment Effects with Misclassification," *Econometrica*, 75, 537–551. [6]

——— (2016): "The Identification Zoo - Meanings of Identification in Econometrics," *Forthcoming on Journal of Economic Literature*. [7]

LEWBEL, A. AND K. PENDAKUR (2008): "Estimation of Collective Household Models with Engel Curves," *Journal of Econometrics*, 147, 350–358. [4], [5], [42]

LISE, J. AND S. SEITZ (2011): "Consumption Inequality and Intra-household Allocations," *Review of Economic Studies*, 78, 328–355. [42]

MACOURS, K., N. SCHADY, AND R. VAKIS (2012): "Cash Transfers, Behavioral Changes, and Cognitive Development in Early Childhood: Evidence from a Randomized Experiment," *American Economic Journal: Applied Economics*, 4, 247–73. [5]

MAHAJAN, A. (2006): "Identification and Estimation of Regression Models with Misclassification," *Econometrica*, 74, 631–665. [6]

MAITRA, P. (2004): "Parental Bargaining, Health Inputs and Child Mortality in India," *Journal of Health Economics*, 23, 259 – 291. [5]

MANSKI, C. F. (1990): "Nonparametric Bounds on Treatment Effects," *The American Economic Review*, 80, 319–323. [6]

MENON, M., K. PENDAKUR, AND F. PERALI (2012): "On the Expenditure-dependence of Children's Resource Shares," *Economics Letters*, 117, 739–742. [42]

MOLINARI, F. (2008): "Partial identification of probability distributions with misclassified data," *Journal of Econometrics*, 144, 81 – 117. [6]

PEARL, J. (2009): "Causal inference in statistics: An Overview," *Statistics Surveys*, 3, 96–146. [7]

PRENTICE, R. L. (1989): "Surrogate endpoints in clinical trials: definition and operational criteria," *Statistics in medicine*, 8, 431–440. [6]

QUISUMBING, A. R. AND J. A. MALUCCIO (2003): "Resources at Marriage and Intrahousehold Allocation: Evidence from Bangladesh, Ethiopia, Indonesia, and South Africa," *Oxford Bulletin of Economics and Statistics*, 65, 283–327. [2], [5]

ROSENBAUM, P. R. (1984): "The consquences of adjustment for a concomitant variable that has been affected by the treatment," *Journal of the Royal Statistical Society. Series A (General)*, 656–666. [6]

ROY, S. (2008): "Female Empowerment Through Inheritance Rights: Evidence from India," Mimeo, London School of Economics, London. [5], [13], [25]

——— (2015): "Empowering women? Inheritance rights, female education and dowry payments in India," *Journal of Development Economics*, 114, 233–251. [5], [25]

RUBALCAVA, L., G. TERUEL, AND D. THOMAS (2009): "Investments, Time Preferences, and Public Transfers Paid to Women," *Economic Development and cultural change*, 57, 507–538. [2], [5]

SCHULTZ, T. P. (1990): "Testing the Neoclassical Model of Family Labor Supply and Fertility," *Journal of Human resources*, 599–634. [5]

SEN, A. (1983): "Economics and the Family," *Asian Development Review*, 1. [2]

——— (1988): *Family and Food: Sex-Bias in Poverty*, New York: Columbia University Press. [2]

——— (1989): "Cooperation, Inequality, and the Family," *Population and Development Review*, 15, 61–76. [2]

SMITH, L. C. (2003): *The Importance of Women's Status for Child Nutrition in Developing Countries*, vol. 131, Intl Food Policy Res Inst. [2], [13], [18]

SONG, S. (2015): "Semiparametric estimation of models with conditional moment restrictions in the presence of nonclassical measurement errors," *Journal of Econometrics*, 185, 95 – 109. [6]

SONG, S., S. M. SCHENNACH, AND H. WHITE (2015): "Estimating nonseparable models with mismeasured endogenous variables," *Quantitative Economics*, 6, 749–794. [6]

THOMAS, D. (1990): "Intra-household Resource Allocation: An Inferential Approach," *Journal of human resources*, 635–664. [2], [5]

——— (1994): "Like Father, like Son; Like Mother, like Daughter: Parental Resources and Child Height," *The Journal of Human Resources*, 29, 950–988. [2]

——— (1997): "Incomes, Expenditures, and Health outcomes: Evidence on Intrahousehold Resource Allocation," *Intrahousehold Resource Allocation in Developing Countries*, 142–64. [2]

TOMMASI, D. (2016): "How Cash Transfers Improve Child Development," Ssrn working paper. [5]

TOMMASI, D. AND A. WOLF (2018): "Estimating household resource shares: A shrinkage approach," *Economics Letters*, 163, 75 – 78. [23]

UPADHYAY, U. D., J. D. GIPSON, M. WITHERS, S. LEWIS, E. J. CIARALDI, A. FRASER, M. J. HUCHKO, AND N. PRATA (2014): "Women's Empowerment and Fertility: A Review of the Literature," *Social Science & Medicine*, 115, 111–120. [13]

URA, T. (2016): "Heterogeneous Treatment Effects with Mismeasured Endogenous Treatment," Mimeo, Duke University. [3], [6], [9], [11]

VANDERWEELE, T. (2015): *Explanation in causal inference: methods for mediation and interaction*, Oxford University Press. [6]

VYTLACIL, E. (2002): "Independence, Monotonicity, and Latent Index Models: An Equivalence Result," *Econometrica*, 70, 331–341. [7]

WOLPIN, K. I. (2013): *The limits of inference without theory*, MIT Press. [7]

WORLD BANK (2014): "Women, Business and the Law 2014. Removing Restrictions to Enhance Gender Equality," *World Bank Report*. [24]

YANAGI, T. (2017): "Inference on Local Average Treatment Effects for Misclassified Treatment," mimeo. [6]

A Appendix

This Appendix contains five main sections. The proof of Theorem 1 is in Appendix A.1. Appendix A.2 provides a graphical illustration of the MR-LATE estimator. Appendix A.3 presents some additional results and robustness checks. Appendix A.4 summarizes the derivation of our model presented in section 4.1 for estimating resource shares from demand equations for private assignable goods. In Appendix A.5, we present a brief discussion of the external validity of our results. Additional figures and tables are in Appendix A.6.

A.1 Proof of Theorem 1

Substituting equation (1) into equation (4) gives

$$T = T_0 + (T_1 - T_0)\,D = T_0 + (T_1 - T_0)\,[(1-Z)\,D_0 + Z D_1] \tag{A1}$$

Multiplying equation (2) by equation (A1) gives

$$YT = [Y_0 + (Y_1 - Y_0)\,[(1-Z)\,D_0 + Z D_1]]\,[T_0 + (T_1 - T_0)\,[(1-Z)\,D_0 + Z D_1]]$$

Using assumption 2-i, this makes

$$\begin{aligned} E(YT \mid Z=1) &= E\,[[Y_0 + (Y_1 - Y_0)\,D_1]\,[T_0 + (T_1 - T_0)\,D_1]] \\ &= E\,[T_0 Y_0 + (Y_1 T_1 - Y_0 T_0)\,D_1] \end{aligned}$$

where the last equality uses $D_1^2 = D_1$. Similarly,

$$E(YT \mid Z=0) = E\,[T_0 Y_0 + (Y_1 T_1 - Y_0 T_0)\,D_0]$$

So,

$$E(YT \mid Z=1) - E(YT \mid Z=0) = E\,[(Y_1 T_1 - Y_0 T_0)\,(D_1 - D_0)]$$

Given the no defiers assumption, either $D_1 - D_0 = 0$ or $D_1 - D_0 = 1$, and someone is a complier if and only if they have $D_1 - D_0 = 1$. The probability of being a complier is $\Pr(D_1 - D_0 = 1) = E\,[D_1 - D_0]$. We therefore apply the standard LATE logic:

$$\begin{aligned} E(YT \mid Z=1) - E(YT \mid Z=0) &= E\,[(Y_1 T_1 - Y_0 T_0)\,(D_1 - D_0)] \\ &= E\,[Y_1 T_1 - Y_0 T_0 \mid D_1 - D_0 = 1]\Pr(D_1 - D_0 = 1) \\ &= E\,(Y_1 T_1 - Y_0 T_0 \mid C)\,E\,(D_1 - D_0)\,. \end{aligned}$$

Let $p_d = E\,(T_d \mid C)$. Then, using Assumption 2-ii, the above further simplifies to

$$E(YT \mid Z=1) - E(YT \mid Z=0) = E\,(p_1 Y_1 - p_0 Y_0 \mid C)\,E\,(D_1 - D_0)\,.$$

Replacing Y with one gives

$$\begin{aligned} E(T \mid Z=1) - E(T \mid Z=0) &= E(p_1 - p_0 \mid C)\, E(D_1 - D_0) \\ &= (p_1 - p_0)\, E(D_1 - D_0). \end{aligned}$$

And therefore

$$\frac{E(YT \mid Z=1) - E(YT \mid Z=0)}{E(T \mid Z=1) - E(T \mid Z=0)} = \frac{E(Y_1 p_1 - Y_0 p_0 \mid C)}{p_1 - p_0}$$

which equals equation (7), thereby proving the Theorem.

A.2 A Graphical Illustration

Assume that $supp(\varepsilon) \subset \left[\kappa^b - e, \kappa^a - e\right)$. Then it follows that for $T = T^a$ we have $p_1^a = 1$ with $p_0^a = 0$, and for $T = T^b$ we have $p_1^b = 0$ and $p_0^b = 1$, and so $\lambda^a - \lambda^b = E[Y_1 - Y_0 \mid C]$. Given corollary 2, LATE can be point identified. Figure A3 provides a graphical representation of this. If there was no measurement error, the true treatment and control groups would coincide with the respective observed groups. All individuals on the black line on the right hand side of e, would have a R^* larger than the threshold value; otherwise, they would be on the black line on the left hand side of e. One could construct a treatment proxy $T = \mathbb{I}(R \geq c)$, where R is an estimate of R^* and c is one's best guess of the midpoint between $\varepsilon + e$. This approach, however, will not identify the treatment effect of interest. To achieve point identification of LATE in presence of measurement error or misclassification error, we need to have two treatment indicators, T^a and T^b, such that $q^a = p_1^a / \left(p_1^a - p_0^a\right) = 1$ and $q^b = p_1^b / \left(p_1^b - p_0^b\right) = 0$. By knowing the bounds κ^a and κ^b, we are able to define a T^a such that for all individuals on the red line on the left hand side of κ^a, $p_0^a = 0$. That is, with probability 0, these individuals, who are observed in the control group, belong to the true treatment group. Analogously, we are able to define also a T^b such that for all individuals on the blue line on the right hand side of κ^b, $p_1^b = 0$. That is, with probability 0, these individuals, who are observed in the treatment group, belong to the true control group.

A.3 Additional Results: Alternative Measures of Power

In this section, we first empirically compare our structurally-motivated measure of bargaining power R, estimated from household expenditure data, with some more typical proxies of power, namely, women's self-reports of control over various household decisions and mobility. We then present MR-LATE estimates that are instead based on an index of power constructed from these self-reported proxies, which are available in the NFHS dataset. Finally, we explore the possibility of a continuous relationship between R and the health outcomes of family members

Comparing Self-reported and Structural Measures of Power. The NFHS data contains questions of the form, "Who usually makes decisions about [X] in your household?". Specifically, women are asked to report who has the final say over their own health care, household purchases,

and visits to family or relatives. We construct indicator variables equal to 1 if the answer to these questions is “respondent alone" or “respondent and husband/partner jointly" and 0 if the answer is “husband/partner".[27] Several women in our sample report having no say in household decisions: 29 percent of women say they do not participate in decisions over their own health, 25 percent report having no say in determining visits to family and friends, and 33 percent claim to have no say in large household purchases. In addition, women are asked whether they are allowed to go alone to places outside the village, to the health facility or to the market. Many women report an inability to go places alone, especially to places outside the village or community (51 percent). One out of three women report not being allowed to go to the market or to a health facility alone. We combine the above information (three questions on women’s mobility) with the responses to the three questions on women’s participation in household decisions to construct an index of women’s autonomy. Specifically, we give each answer a value of one if the response is that wife controls that decision (or if she can go alone to different places), a value of minus one if the response is that the husband controls that decision (or if she cannot go alone), and a value of zero for any other response. We then define our estimated index to equal the sum of the responses across all of the questions. Figure A5 in the Appendix shows the distribution of this index.

Panels (a) to (c) of Figure A4 display the results of non-parametric regressions of women’s reported participation in household decisions on our estimated resource share R. Panel (d) shows the non-parametric relationship between our index of women’s autonomy and R. In all cases, the presence of positive relationships emerges clearly. We also examine the link between the self-reported decision making and our binary structural treatment variable $T = \mathbb{I}(R \geq 50)$ conditional on individual and household level controls, fixed effects, and state time trends. The estimation results are in Table A6 (in the Appendix). Overall, women are significantly more likely to report participating in decisions in treated households, i.e., in households where we estimate, based on expenditures, that they have substantial control over resources.[28] Thus, these results corroborate the theory underlying our structural model of treatment: the larger is R (meaning the higher is the likelihood that a woman has control over household resources), the higher are her self reported decision making and bargaining powers within the household.

MR-LATE Estimates Based on Self-reported Measures of Power. Here we consider estimation of treatment effects where the treatments are women’s self-reported measures of power. As discussed in Example 1 in section 3.2, we could discard all the observations where responses to power questions are ambiguous, and apply standard LATE estimation to the remaining observations. But if giving other answers is correlated with the instrument or with potential outcomes, the resulting LATE estimates could then be biased in unknown ways. And even if a resulting LATE estimate is not biased, discarding observations would entail a loss of efficiency. So instead of discarding any observations, we use MR-LATE estimation to make use of these incompletely observed treatments.

[27] We exclude women who answer “other/someone else" (less than 1 percent in any question).

[28] We also repeated this exercise using principal components of the self reported responses. The findings do not change. Results are available upon request.

This is done by setting $T^a = 1$ if the answer is "the wife," otherwise $T^a = 0$, and similarly letting $T^b = 1$ if the answer is "the husband," otherwise $T^b = 0$. Thus the difference between T^a and $1-T^b$ is in how the ambiguous responses are handled.

Following Example 2 of section 3.2, we can also construct MR-LATE estimates based on an index constructed by summing the self-reported power responses. This index is based on the answers to three questions about women's participation in household decision and three questions on women's mobility, for a total of $m = 6$ questions. This index therefore takes on integer values in the range -6 to 6.

Using the notation of Example 2, if we took $\kappa^a = k$ for some integer k, with $k < 6$, we would be assuming that it is extremely unlikely that the husband really has most of the power if the household reports that the wife makes $6-k$ or more of the decisions (meaning p_0^a is very close to zero). Similarly, taking $\kappa^b = -k$ for some small integer k means assuming it is extremely unlikely that the wife has most of the power if they report that the husband makes $6-k$ or more of the decisions (meaning p_1^b is very close to zero).

Table A7 and A8 report the MR-LATE estimates for models where T^a and T^b are constructed as described in the previous paragraphs, following Examples 1 and 2 of section 3.2, respectively.[29] The empirical results of these alternative estimates appear unreliable, with peculiar magnitudes and some large standard errors. The problem with these estimates are that self reports of power appear to be rather noisy indicators of true power, and as a result, the treatment indicators T, T^a, and T^b based on these self reports are not significantly correlated with the instrument Z. This can be seen in the first-stage MR-LATE F-statistics, which are almost all below 10. So, while the results here illustrate alternative ways that MR-LATE estimates can be constructed, we find that more stable and reliable treatment effect estimates are obtained when treatment is defined using our structural model estimates of women's power.

Linear Model. In section 4 we discussed why we expect Y to depend on D. Here we consider the alternative possibility of Y depending continuously on R^*. We do this by linearly regressing each health status measure Y_i on a constant, R_i, and X_i, using 2SLS where Z_i is the excluded instrument for R_i. These estimates should be interpreted cautiously, since we have no reason to expect that the true relationship of Y to R^* and X is linear, or that estimation errors in R relative to R^* satisfy the classical measurement error assumptions needed for validity of linear 2SLS estimation. But if these assumptions do hold, then the estimated coefficient of R in this regression will be a consistent estimate of the average marginal effect of R^* on Y.

Tables A9 and A10 in the Appendix contain the estimation results of these linear regressions for women's and children's health outcomes.[30] For women, we restrict the sample to those living in households where the woman has control of 40 to 60 percent of the resources, which is where we see a positive and significant correlation between R and Z (see columns 1 and 2). In general, while

[29] For save space, we report results for the NFHS women's sample only. Results for men's and children's samples are available upon request.

[30] We do not report results for the men's sample, because (possibly due to the smaller sample size), no significant first stage estimates could be obtained using the men's sample.

the estimated coefficients of these regressions have the expected signs, we do not find marginal effects that are statistically different from zero. While the correlation between R and Z is statistically significant, based on low first-stage F-statistics we cannot rule out the possibility of a weak instrument problem. The lower significance of these estimates relative to our main MR-LATE results also suggests that the linear model may be misspecified.

A.4 Derivation of Household Demand Equations of Private Assignable Goods

Here we summarize the derivation of our structural model, based on Browning et al. (2013) (BCL) and Dunbar et al. (2013) (DLP), for estimating resource shares from the demand equations of private assignable goods. Consider a household comprised of T types of individuals indexed $t = 1, ..., T$. Recall M is the total expenditures of the household, i.e., the household's total budget, X denotes a vector of observable attributes of households and their members, $\widetilde{Z}$ denotes a vector of possible distribution factors (if any), and Q_1,...,Q_T are quantities of each private assignable good consumed by household member t. Let S be a vector of quantities of all other goods the household consumes. Unlike Q_1,...,Q_T, the goods S may be shared and hence jointly consumed to some extent. In particular, $S = \sum_{t=1}^{T} S_t$ where S_t is the vector of quantities of these goods consumed by member t. The purchased quantities of these goods are given by $A(X)S$, where the matrix $A(X)$ summarizes the extent to which these goods are shared.

Let P_1,...,P_T be the market prices of the private assignable goods, let P_S be the vector of market prices of goods S, and let P denote the vector of all of these prices.

The household chooses what to consume using the program

$$\max_{Q_1,...,Q_T,S_1,...S_T} \widetilde{V}\left[V_1(Q_1,S_1,X),...,V_T(Q_T,S_T,X) \mid \widetilde{Z},X,P/M\right] \tag{A2}$$

$$\text{such that} \quad S = \sum_{t=1}^{T} S_t \quad \text{and} \quad M = P_S' A(X) S + \sum_{t=1}^{T} P_t Q_t$$

where $V_t(Q_t,S_t,X)$ for $t = 1, ..., T$ is the utility function of household member t, and the function $\widetilde{V}$ describes the social welfare function or bargaining process of the household. A function $\widetilde{V}$ exists because the household is Pareto efficient.

What makes Q_1,...,Q_T be private is that they are not shared. What makes them assignable is that the econometrician can observe who consumes each. In particular, each member t has quantity Q_t in his or her utility function, and does not have Q_ℓ for all $\ell \neq t$ in his or her utility function. The square matrix $A(X)$ is what is called by BCL a linear consumption technology function over goods. Having $A(X)$ differ from the identity matrix is what allows goods in S to be partly shared and/or consumed jointly. In particular, $A(X)S$ equals the quantity vector of these goods that the household actually purchases, while $S = \sum_{t=1}^{T} S_t$ is total quantity vector of these goods that the household consumes. These quantities are not the same due to sharing and joint consumption. The smaller an element of $A(X)S$ is relative to the corresponding element of S, the more that good is shared or

jointly consumed. See BCL for details.

Household attributes X may affect preferences, and so appear inside the utility functions V_t. These X variables can also affect the extent to which goods are shared through $A(X)$, and they can directly affect the bargaining process or social welfare function given by $\widetilde{V}$ (by, e.g., affecting the relative bargaining power of members). As a result, resource shares may also depend on X. The difference between X and distribution factors $\widetilde{Z}$ is that the vector $\widetilde{Z}$ appears in the model only as arguments of $\widetilde{V}$, and so only directly affects the allocation of resources within the household, but not the tastes of the individual household members or the jointness of consumption.

Applying duality theory and decentralization welfare theorems, it follows from BCL that the household's program above is equivalent to a program where each household member t chooses what to consume using the program

$$\max_{Q_t,S_t} V_t(Q_t,S_t,X) \quad \text{such that} \quad \eta_t(P,M,X,\widetilde{Z})M = P_S'A(X)S_t + P_tQ_t \tag{A3}$$

where $\eta_t = \eta_t(P,M,X,\widetilde{Z})$ is the resource share of member t, that is, η_t is the fraction of total household resources M that are allocated to member t. This member then chooses quantities Q_t and the vector S_t subject to a linear budget constraint. The vector $P_SA(X)$ equals the vector of shadow prices of goods S. These shadow prices for the household may be lower than market prices, due to sharing. Being private and assignable, the shadow price of each Q_t equals its market price P_t. Let $\widetilde{M}_t = \eta_t M$ denote the shadow budget for member t. As shown in BCL, the resource share functions $\eta_t(P,M,X,\widetilde{Z})$ for each member t in general depend on the function $\widetilde{V}$ and on the utility functions V_1...,V_T.

BCL show that the more bargaining power a household member has (i.e., the greater is the weight of his or her utility function in $\widetilde{V}$), the larger is their resource share η_t. Resource shares η_t all lie between zero and one, and resource shares sum to one, that is, $\sum_{t=1}^{T}\eta_t = 1$.

As in DLP, we will not work with the household demand functions of all goods (which, as shown in BCL, are rather complicated). Instead, we only make use of the demand functions of the private assignable goods Q_t, which are simpler. Since equation (A3) is an ordinary utility function maximized under a linear budget constraint (linear in shadow prices and a shadow budget), the solution to equation (A3) is a set of Marshallian demand equations for Q_t and S_t.

Let $h_t\left(\widetilde{M}_t,P,X\right)$ be the Marshallian demand function of person t' for their private assignable good, that is, $h_t\left(\widetilde{M}_t,P,X\right)$ is the quantity person t in a household with member attributes X would demand of their assignable good if they had a budget equal to their shadow budget $\widetilde{M}_t$ and faced the within-household shadow price vector that corresponds to the market price vector P. Since each Q_t is private and assignable, the quantity Q_t that member t chooses to consume equals the quantity of this good that the household buys. It therefore follows from the above that the household's quantity demand of each private assignable good Q_t is given by

$$Q_t = h_t\left(\eta_t(P,M,X,\widetilde{Z})M,P,X\right) \quad \text{for} \quad t = 1,...,T. \tag{A4}$$

The interpretation of this equation is that the total resources allocated to member t are $\eta_t M$ (the share η_t of total household budget M) and the function h_t is that member's Marshallian demand function for this good. Since the good is private and assignable, the household's demand for the good just equals that member's own demand for the good. It is important to note that only private assignable goods have the simple form given by equation (A4). The demand functions for other goods are much more complicated, as in BCL.

Let $\widetilde{h}_t\left(\widetilde{M}_t, P, X\right) = P_t h_t\left(\widetilde{M}_t, P, X\right)/\widetilde{M}_t$ denote the Marshallian demand function written in budget share form. That is, $\widetilde{h}_t\left(\widetilde{M}_t, P, X\right)$ is the fraction of the total budget $\widetilde{M}_t$ that is spent on the good t. DLP assume data are drawn from single price regime (that is, Engel curve data), so P is a fixed constant that can be dropped from the model. They provide empirical and theoretical evidence that η_t does not depend on M.[31] This allows them to rewrite equation (A4) as $w^t = \eta_t(X, \widetilde{Z})\widetilde{h}_t\left(\eta_t(X, \widetilde{Z})M, X\right)$ for $t = 1, ..., T$, where $w^t = P_t Q_t/M$ is the household's budget share of good t, that is, the fraction of the household's total budget M that is spend on buying Q_t. DLP provide a class of functional forms for the utility functions $\widetilde{V}$ that make $\widetilde{h}_t$ linear in the log of its first argument, so $w^t = \eta_t(X, \widetilde{Z})[\delta^t(X) + (lnM + ln\eta_t(X, \widetilde{Z}))\beta(X)]$ for some functions $\delta^t(X)$ and $\beta(X)$. The assumption that $\beta(X)$ does not depend on t is what DLP call the SAP (similar across people) assumption.

A.5 External Validity of LATE

Here we use a latent variable model and the marginal treatment effects framework developed by Björklund and Moffitt (1987), Heckman and Vytlacil (1999, 2005, 2007), Carneiro et al. (2011), Brinch et al. (forthcoming) and Kowalski (2016) to shed light on the external validity of our results.[32] Specifically, we wish to clarify the relationship between the local average treatment effect (LATE) and the average treatment effect (ATE) in our empirical application.

In our application, we are interested in estimating the treatment effect of D on individual health outcomes Y, with covariates (observable individual and household characteristics) X. Given potential outcomes Y_1 and Y_0, define the functions h_1 and h_0, and corresponding errors U_0 and U_1, by $h_0(X) = E(Y_0 \mid X)$, $h_1(X) = E(Y_1 \mid X)$, $U_0 = Y_0 - h_0(X)$, and $U_1 = Y_1 - h_1(X)$.

Under the standard monotonicity assumption for LATE estimation, the determination of treatment D can be represented by standard threshold crossing model

$$D = \mathbb{I}(R^* - e \geq 0)$$

where R^* is an underlying latent variable and e is an unobserved threshold that can vary across households. It is assumed that $(U_0, U_1, e) \perp Z \mid X$, which implies validity of the instrument Z. As-

[31] Lise and Seitz (2011), Lewbel and Pendakur (2008), Bargain and Donni (2012), Bargain et al. (2014) and DLP all use this restriction in their identification results, and supply some theoretical arguments for it. Cherchye et al. (2015) and Menon et al. (2012) provide empirical support for this restriction.

[32] Brinch et al. (forthcoming) in particular show how a discrete (binary) instrument can be used to identify the MTE under functional structure that allows for treatment heterogeneity among individuals with the same observed characteristics and self-selection based on the unobserved gain from treatment.

sume also that e and R^* are continuously distributed. Instrument relevance requires that $E(R^* \mid X, Z)$ varies with Z. In our application, R^* is a continuous measure of women's control of household resources and Z is a binary variable capturing women's exposure to inheritance law reforms that improved their ability to inherit property.

The threshold crossing model for D means that households are treated if their unobserved threshold e is less than or equal to R^*. Variation in e can be interpreted as meaning that different households have different levels of R^* that are needed for the wife to have substantial control, making $D = 1$. Under monotonicity, in households where e is low enough ($e \leq R^*_L$), wives control substantial resources even if they are not exposed to the plausibly exogenous changes in women's inheritance rights (that is, they are *always-takers*, with $D = 1$ and $Z = 0$). In households where e is high enough ($e > R^*_H$), husbands control substantial resources even if their wives are exposed to the reforms (that is, they are *never-takers*, with $D = 0$ and $Z = 1$). Under the standard no-defiers assumption, the remaining households correspond to the group of *compliers* ($R^*_L < e \leq R^*_H$), whose treatment status is determined by women's exposure to the inheritance law amendments ($D = Z$).

As in Heckman and Vytlacil (1999), define the marginal treatment effect as the effect on the marginal individual entering treatment. That is, MTE is the average impact for the marginal individual receiving treatment among those with $e = R^*$, so

$$MTE(R^*) = E(Y_1 - Y_0 \mid e = R^*) \tag{A5}$$

We may without loss of generality normalize e to be uniformly distributed, by taking a suitable monotonic transform of R^* and e. Under the normalization $e \sim U(0,1)$, Heckman and Vytacil show that LATE equals the weighted average of MTE over the interval $(R^*_L, R^*_H]$, with weights equal to $\frac{1}{R^*_H - R^*_L}$. A sufficient condition for LATE to be externally globally valid (so LATE equals ATE) is if $MTE(R^*)$ is constant for all R^*.

Following Kowalski (2016) (analogous tests are proposed in Angrist (2004), Bertanha and Imbens (2014), and Brinch et al. (forthcoming)), we implement a test of global external validity using a difference-in-difference regression.[33] Since we cannot directly observe R^* and D, we perform the test using R and $T = \mathbb{I}(R \geq 50)$ instead. This test must therefore be interpreted with caution, since it does not account for the measurement errors that our MR-LATE estimator is designed for. To implement the test, we regress individuals' health outcomes Y on the covariates X in the sample of households where women are not exposed to the inheritance law reforms (i.e., $Z = 0$). Then, using the estimated coefficients, we predict outcomes $\widehat{Y}$ for all individuals with $Z = 1$ and $Z = 0$, and estimate the following model:

$$\widehat{Y} = \lambda_{TZ} TZ + \lambda_T T + \lambda_Z Z + \lambda \tag{A6}$$

Ignoring the measurement error from using R and T in place of R^* and D, if $MTE(R^*)$ is constant for all R^* (implying external validity) then $\lambda_{TZ} = 0$. To implement the test we assume $MTE(R^*)$

[33] We implement the test using the `mtebinary` Stata routine (Kowalski et al. (2016)).

is linear in R^* for simplicity, and we compute standard errors based on 200 bootstrap replications. Table A12 reports the estimates for λ_{TZ} and the corresponding bootstrap standard errors. For all health outcomes Y, in the three NFHS samples (women, men, and children), we cannot reject $\lambda_{ZT} = 0$. So under the caveat that the test is based on T, we do not reject the hypothesis that our LATE is externally valid, meaning that the treatment effects we identify in our empirical application equal ATE, and do not depend on our specific choice of instrument.

A.6 Additional Figures and Tables

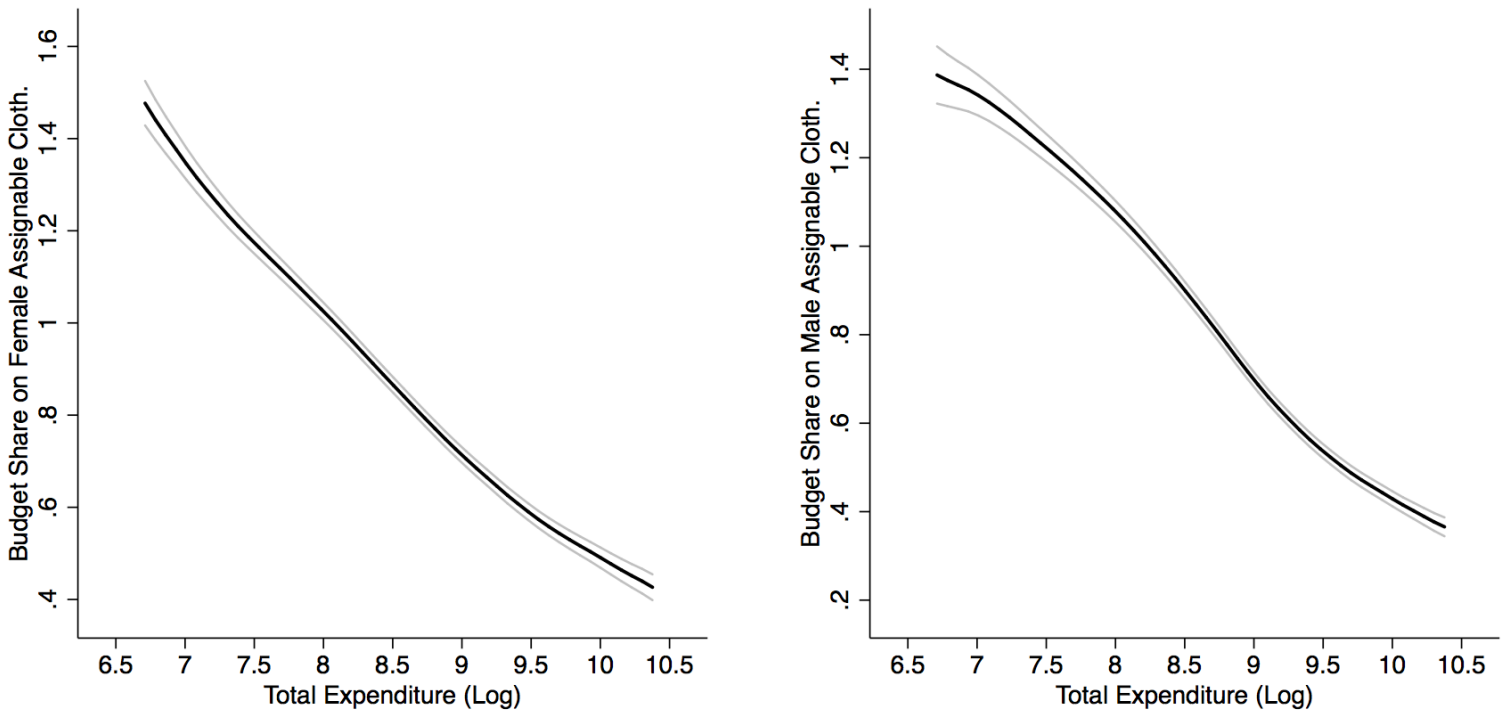

(A) Woman's Assignable Clothing (B) Man's Assignable Clothing

Figure A1: Non-parametric Engel curves

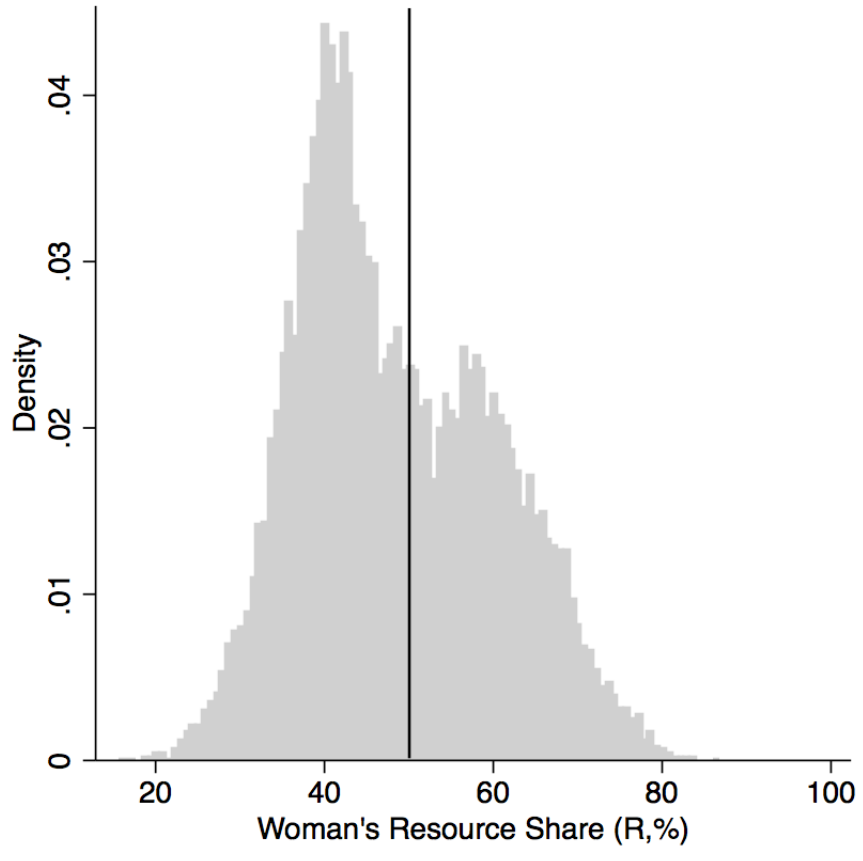

Figure A2: Estimated of Women's Resource Shares (NFHS Sample)

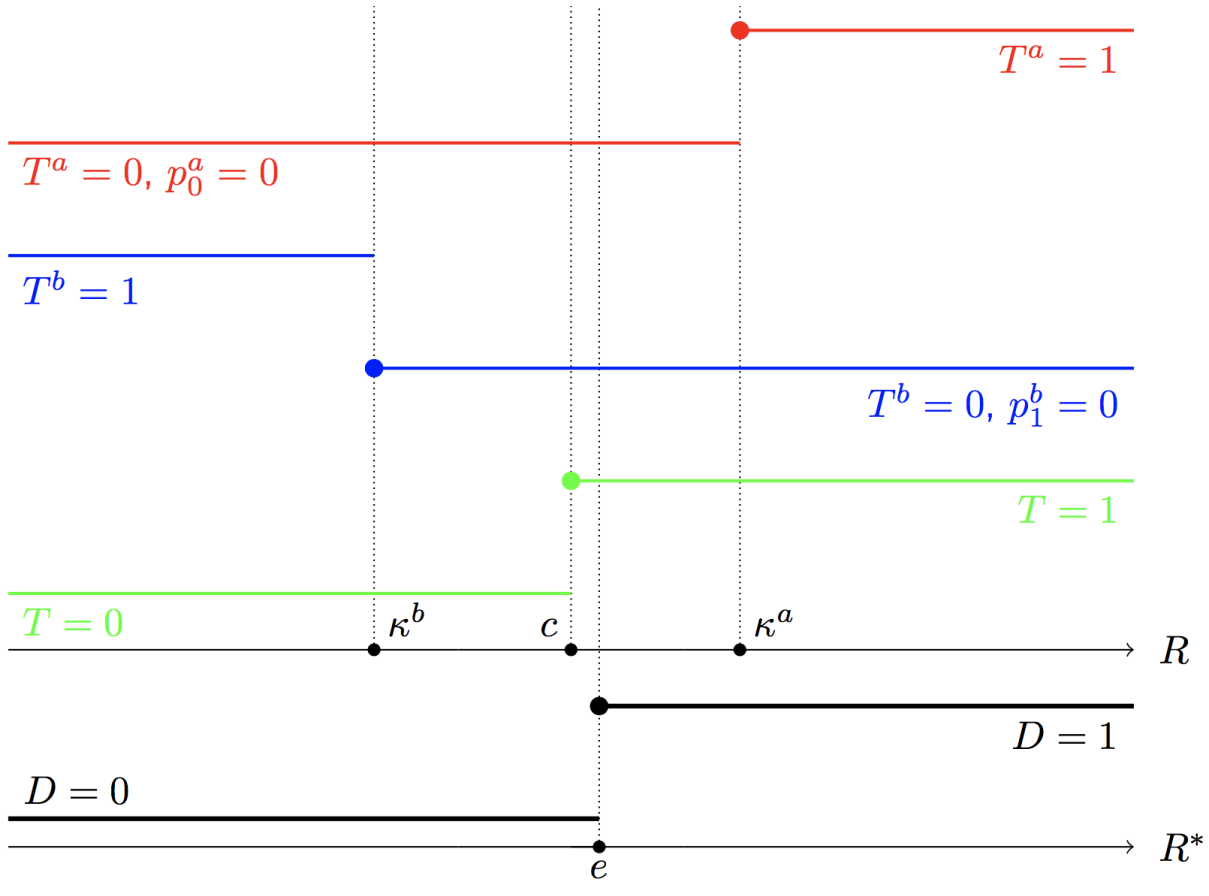

Figure A3: Illustrative Example

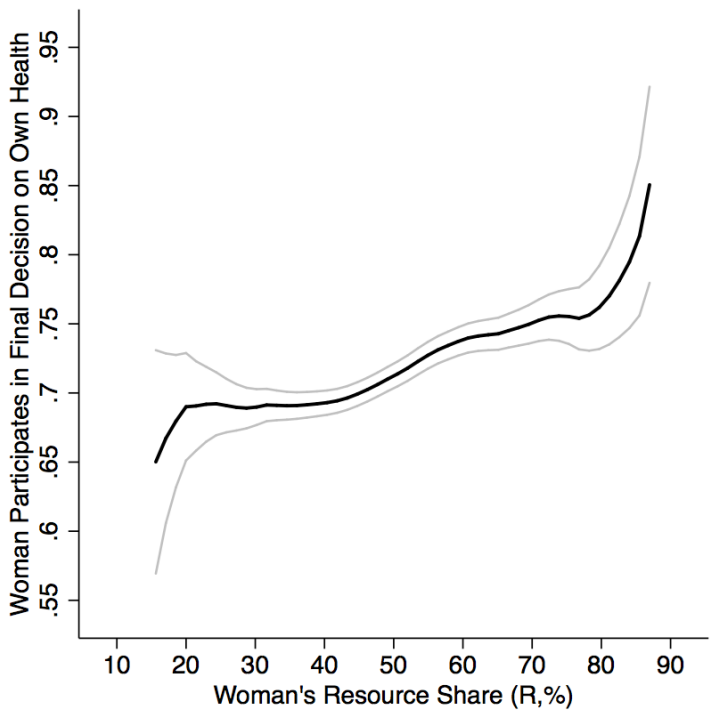

(A) Woman's Health

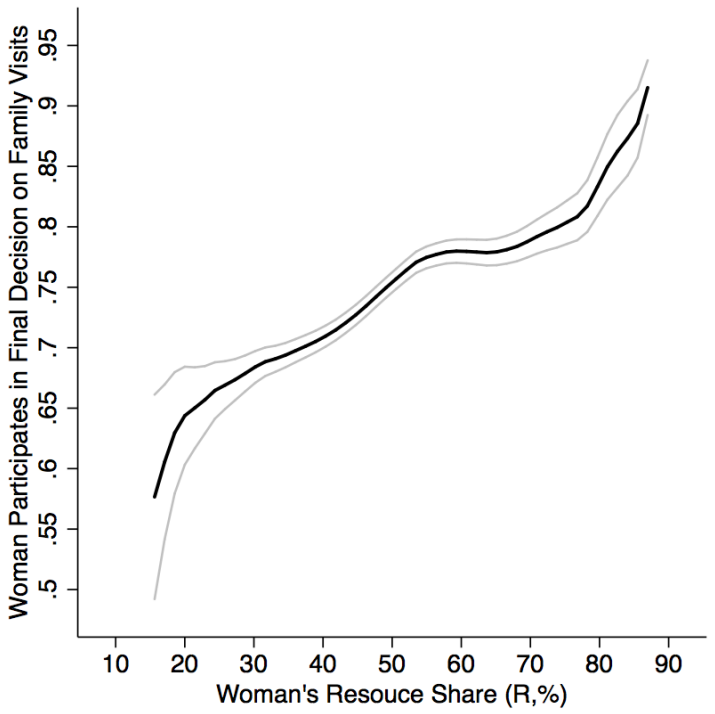

(B) Visits to Family and Relatives

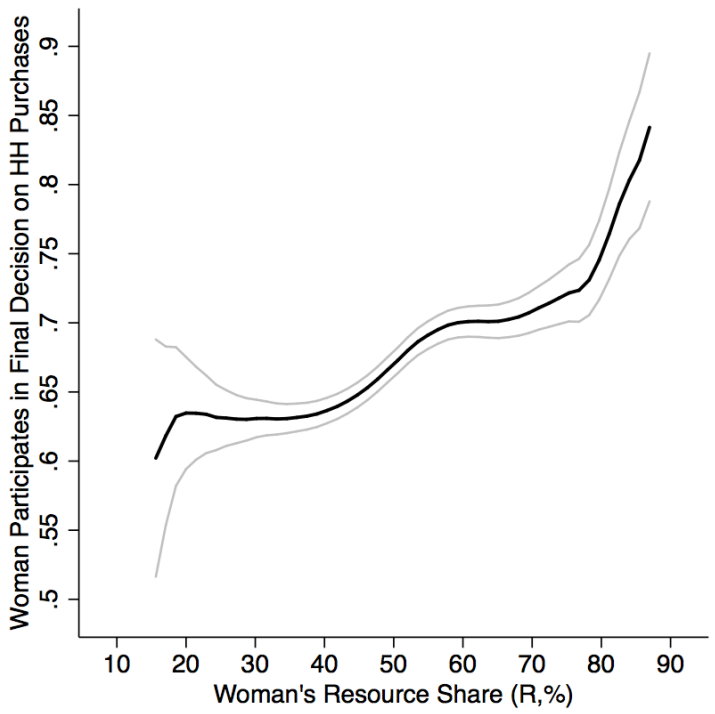

(C) Household Purchases

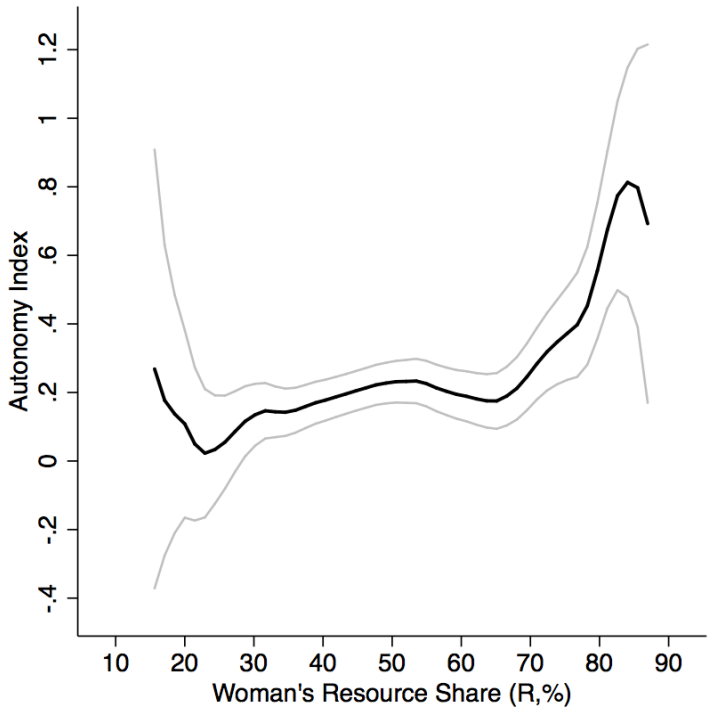

(D) Autonomy Index

Figure A4: Structurally Recovered Bargaining Power and Household Decision Making

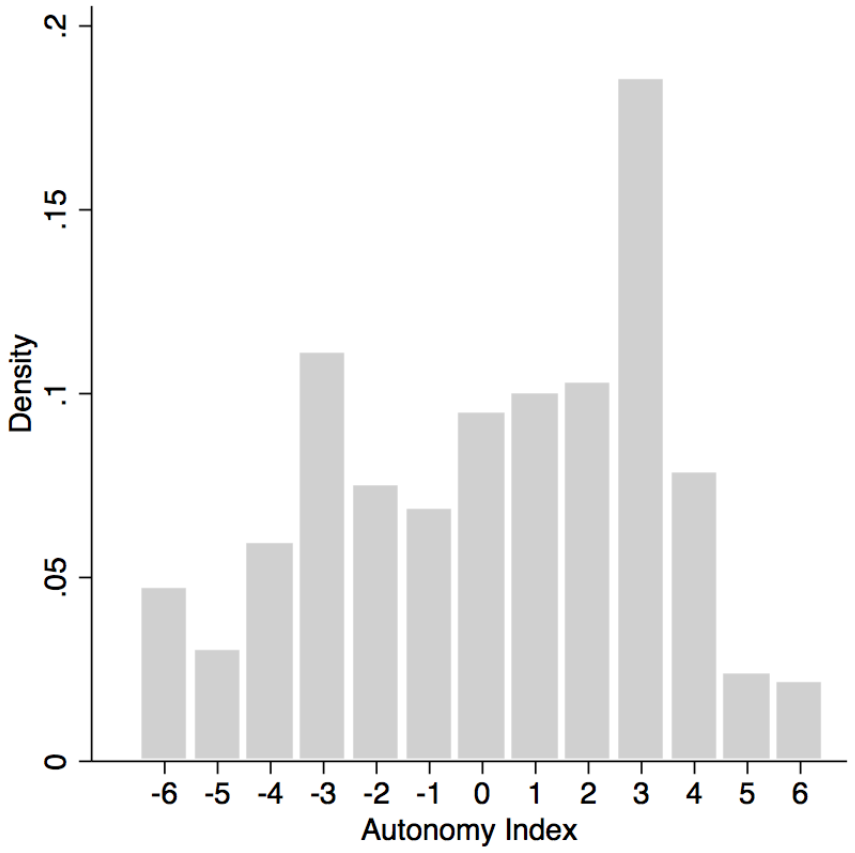

Figure A5: Autonomy Index (NFHS Sample)

Note: This index is based on women's response to 3 questions about decision making and 3 questions about mobility. We give each answer a value of one if the response is that wife controls that decision (or if she can go alone to places), a value of minus one if the response is that the husband controls that decision (or if she cannot go alone to places), and a value of zero for any other response. We then define our estimated index R to equal the sum of the responses across all of the questions.

Table A1: OLS vs. 2SLS vs. MR-LATE when ϵ is Bounded at Known κ

	Panel A: N = 1,000								
	$\kappa = 0$			$\kappa = 5$			$\kappa = 10$		
	OLS	2SLS	MR-LATE	OLS	2SLS	MR-LATE	OLS	2SLS	MR-LATE
T	1.016	0.992	0.992	0.827	1.030	1.015	0.656	1.132	1.030
sd	0.048	0.320	0.320	0.057	0.342	0.286	0.065	0.393	0.309
Bias	0.016	-0.008	-0.008	-0.173	0.030	0.015	-0.344	0.132	0.030
MSE	0.003	0.103	0.103	0.033	0.118	0.082	0.123	0.172	0.096
	Panel B: N = 10,000								
	$\kappa = 0$			$\kappa = 5$			$\kappa = 10$		
	OLS	2SLS	MR-LATE	OLS	2SLS	MR-LATE	OLS	2SLS	MR-LATE
T	1.015	0.999	0.999	0.827	1.032	1.016	0.658	1.122	1.028
sd	0.014	0.097	0.097	0.018	0.105	0.085	0.020	0.119	0.094
Bias	0.015	-0.001	-0.001	-0.173	0.032	0.016	-0.342	0.122	0.028
MSE	0.000	0.009	0.009	0.030	0.012	0.007	0.118	0.029	0.010
	Panel C: N = 500,000								
	$\kappa = 0$			$\kappa = 5$			$\kappa = 10$		
	OLS	2SLS	MR-LATE	OLS	2SLS	MR-LATE	OLS	2SLS	MR-LATE
T	1.015	0.999	0.999	0.826	1.028	1.016	0.657	1.117	1.026
sd	0.002	0.014	0.014	0.002	0.015	0.012	0.003	0.017	0.013
Bias	0.015	-0.001	-0.001	-0.174	0.028	0.016	-0.343	0.117	0.026
MSE	0.000	0.000	0.000	0.030	0.001	0.000	0.118	0.014	0.001

Notes: In each Panel, the true value is set to 1.000. Results are based on 1,000 simulations for varying number of observations N. We simulate various measurement errors and the availability of a strong and exogenous instrument. When $\kappa = 0$ means that there is no measurement error. Whereas, measurement error with $\kappa = 0.05$ (0.10) means that we estimate R that are on average +/- 5 (10) percent of the true value. This means that roughly 10 (20) percent of sample is wrongly observed to either treatment or control.

Table A2: NSS Consumer Expenditure Data and NFHS Household Data

	2005-2006 NSS Sample				2005-2006 NFHS Sample			
	Obs.	Mean	Median	St. Dev.	Obs.	Mean	Median	St. Dev.
Woman's Assignable Clothing Budget Share	7,480	0.91	0.77	0.74				
Man's Assignable Clothing Budget Share	7,480	0.92	0.70	0.87				
Total Expenditure (Rupees)	7,480	5968.30	4179.26	5064.75				
Number of Children	7,480	2.16	2.00	0.93	23,699	2.20	2.00	0.95
$\mathbb{I}$(1 child)	7,480	0.27	0.00	0.44	23,699	0.26	0.00	0.44
$\mathbb{I}$(2 children)	7,480	0.41	0.00	0.49	23,699	0.39	0.00	0.49
$\mathbb{I}$(3 children)	7,480	0.23	0.00	0.42	23,699	0.24	0.00	0.42
$\mathbb{I}$(4 children)	7,480	0.10	0.00	0.30	23,699	0.11	0.00	0.31
Fraction of Female Children	7,480	0.45	0.50	0.37	23,699	0.47	0.50	0.37
Woman's Age	7,477	30.96	30.00	6.46	23,697	30.22	30.00	7.32
Gender Age Gap (Man - Woman)	7,473	5.20	5.00	4.27	23,662	5.80	5.00	5.01
Children's Avg. Age	7,480	6.87	7.00	3.55	23,699	6.46	6.33	3.58
$\mathbb{I}$(Hindu, Buddhist, Jain, Sikh)	7,480	0.80	1.00	0.40	23,695	0.77	1.00	0.42
$\mathbb{I}$(Sch. Caste, Sch. Tribe, Oth. Back. Caste)	7,480	0.64	1.00	0.48	22,808	0.68	1.00	0.47
$\mathbb{I}$(Own Land)	7,447	0.69	1.00	0.46	23,694	0.35	0.00	0.48
$\mathbb{I}$(Woman Completed High School)	7,480	0.14	0.00	0.35	23,699	0.09	0.00	0.28
$\mathbb{I}$(Man Completed High School)	7,480	0.22	0.00	0.42	23,699	0.13	0.00	0.34
$\mathbb{I}$(Rural)	7,480	0.48	0.00	0.50	23,699	0.54	1.00	0.50
$\mathbb{I}$(North)	7,480	0.30	0.00	0.46	23,699	0.30	0.00	0.46
$\mathbb{I}$(East)	7,480	0.21	0.00	0.40	23,699	0.16	0.00	0.37
$\mathbb{I}$(North-East)	7,480	0.14	0.00	0.35	23,699	0.20	0.00	0.40
$\mathbb{I}$(South)	7,480	0.22	0.00	0.41	23,699	0.21	0.00	0.41
$\mathbb{I}$(West)	7,480	0.12	0.00	0.33	23,699	0.13	0.00	0.34

Notes: Budget shares are multiplied by 100. Woman's assignable clothing includes expenditures on saree, shawls, chaddar, and dupatta; man's assignable clothing includes expenditures on dhoti, lungi, pajamas, salwar, and shirts. North India includes Jammu & Kashmir, Himachal Pradesh, Punjab, Uttaranchal, Haryana, Delhi, Rajasthan, Uttar Pradesh, and Madhya Pradesh. East India includes West Bengal, Bihar, Jharkhand, Orissa, A & N Islands, and Chattisgarh. North-East India includes Sikkim, Arunachal Pradesh, Assam, Manipur, Meghalaya, Mizoram, Nagaland, and Tripura. South India includes Karnataka, Tamil Nadu, Andhra Pradesh, Kerala, Lakshadweep, and Pondicherry. West India includes Gujarat, Goa, Maharashtra, Daman & Diu, and D & N Haveli.

Table A3: 2005-2006 NFHS Individual Data

	Women ($N = 19,738$)			Men ($N = 10,785$)			Children ($N = 15,038$)		
	Mean	Median	St. Dev.	Mean	Median	St. Dev.	Mean	Median	St. Dev.
Body Mass Index (BMI)	21.00	20.24	3.97	21.27	20.73	3.54			
$\mathbb{I}$(BMI$\leq$18.5)	0.29	0.00	0.46	0.24	0.00	0.42			
$\mathbb{I}$(Anemic)	0.15	0.00	0.36	0.09	0.00	0.28			
Weight for Age (z-score)							-1.75	-1.83	1.19
Height for Age (z-score)							-1.60	-1.62	1.55
$\mathbb{I}$(Cough in last 2 weeks)							0.18	0.00	0.38
$\mathbb{I}$(Fever in last 2 weeks)							0.15	0.00	0.36
$\mathbb{I}$(Diarrhea in last 2 weeks)							0.09	0.00	0.28
$\mathbb{I}$(Any Vaccination)							0.91	1.00	0.28
$\mathbb{I}$(HSA Exposed)	0.18	0.00	0.39	0.28	0.00	0.45	0.18	0.00	0.38
$\mathbb{I}$(1 child)	0.24	0.00	0.43	0.24	0.00	0.42	0.20	0.00	0.40
$\mathbb{I}$(2 children)	0.40	0.00	0.49	0.41	0.00	0.49	0.39	0.00	0.49
$\mathbb{I}$(3 children)	0.24	0.00	0.43	0.24	0.00	0.43	0.27	0.00	0.45
$\mathbb{I}$(4 children)	0.11	0.00	0.32	0.11	0.00	0.31	0.14	0.00	0.34
Fraction of Female Children	0.47	0.50	0.36	0.47	0.50	0.36	0.50	0.50	0.36
Women's Age	29.58	29.00	5.86	29.44	29.00	5.71	26.31	26.00	4.45
Gender Age Gap (Men - Women)	5.84	5.00	4.28	5.65	5.00	3.68	5.58	5.00	4.22
Children's Avg. Age	6.39	6.25	3.53	6.37	6.00	3.58	3.57	3.33	2.05
$\mathbb{I}$(Hindu, Buddhist, Jain, Sikh)	0.79	1.00	0.41	0.78	1.00	0.41	0.74	1.00	0.44
$\mathbb{I}$(Sch. Caste, Sch. Tribe, Oth. Back. Caste)	0.68	1.00	0.47	0.72	1.00	0.45	0.71	1.00	0.45
$\mathbb{I}$(Own Land)	0.36	0.00	0.48	0.34	0.00	0.47	0.34	0.00	0.47
$\mathbb{I}$(Woman Completed High School)	0.08	0.00	0.27	0.09	0.00	0.28	0.07	0.00	0.25
$\mathbb{I}$(Man Completed High School)	0.13	0.00	0.34	0.13	0.00	0.34	0.11	0.00	0.31
$\mathbb{I}$(Rural)	0.54	1.00	0.50	0.52	1.00	0.50	0.56	1.00	0.50
$\mathbb{I}$(North)	0.30	0.00	0.46	0.25	0.00	0.43	0.33	0.00	0.47
$\mathbb{I}$(East)	0.16	0.00	0.37	0.10	0.00	0.30	0.16	0.00	0.37
$\mathbb{I}$(North-East)	0.20	0.00	0.40	0.19	0.00	0.39	0.23	0.00	0.42
$\mathbb{I}$(South)	0.21	0.00	0.41	0.32	0.00	0.47	0.17	0.00	0.38
$\mathbb{I}$(West)	0.13	0.00	0.33	0.14	0.00	0.34	0.11	0.00	0.32
$\mathbb{I}$(Child is Female)							0.48	0.00	0.50
Child's Age							2.18	2.00	1.39

Notes: $\mathbb{I}$(Anemic) includes moderate anemia (7.0-9.9 g/dl for women and 9.0-11.9 g/dl for men) or severe anemia (less than 7.0 g/dl for women and less than 9.0 g/dl for men). $\mathbb{I}$(Any Vaccination) includes vaccinations against polio, measles, DPT or BCG. Women of age 15 to 49, men of age 15 to 54 and children of age 0 to 5.

Table A4: Bounds κ^a and κ^b

	Women		Men		Children	
	κ^a	κ^b	κ^a	κ^b	κ^a	κ^b
$\mathscr{K}=0$	50.00	50.00	50.00	50.00	50.00	50.00
$\mathscr{K}=1$	50.20	49.80	50.19	49.83	50.18	49.82
$\mathscr{K}=5$	51.04	48.95	51.05	49.04	50.92	49.09
$\mathscr{K}=10$	52.22	47.99	52.30	48.12	51.98	48.16
$\mathscr{K}=20$	54.73	45.94	54.99	46.10	54.26	46.18

Note: NFHS data.

Table A5: First Stage Estimates (MR-LATE)

	Women		Men		Children	
	T^a	T^b	T^a	T^b	T^a	T^b
	Panel A: $\mathscr{K}=0$					
$\mathbb{I}$(HSA)	0.0781	-0.0781	0.0706	-0.0706	0.0778	-0.0778
	(0.0121)	(0.0164)	(0.0137)	(0.0161)	(0.0085)	(0.0258)
First Stage F-stat.	58.9640	58.9640	37.2133	37.2133	13.8997	13.8997
	Panel B: $\mathscr{K}=1$					
$\mathbb{I}$(HSA)	0.0790	-0.0828	0.0692	-0.0763	0.0695	-0.0816
	(0.0117	(0.0122)	(0.0113)	(0.0136)	(0.0292)	(0.0217)
First Stage F-stat.	61.4723	65.7576	36.4079	42.6378	10.4811	15.3693
	Panel C: $\mathscr{K}=5$					
$\mathbb{I}$(HSA)	0.0719	-0.0806	0.0692	-0.0732	0.0705	-0.0892
	(0.0096)	(0.0117)	(0.0117)	(0.0138)	(0.0187)	(0.0193)
First Stage F-stat.	52.4354	59.4068	38.6722	37.4487	14.3894	15.2712
	Panel D: $\mathscr{K}=10$					
$\mathbb{I}$(HSA)	0.0573	-0.0633	0.0496	-0.0598	0.0655	-0.0743
	(0.0109)	(0.0134)	(0.0119)	(0.0137)	(0.0108)	(0.0196)
First Stage F-stat.	34.7818	35.3245	21.1301	23.6691	13.5925	12.3206
	Panel E: $\mathscr{K}=20$					
$\mathbb{I}$(HSA)	0.0259	-0.0658	0.0156	-0.0602	0.0606	-0.0809
	(0.0072)	(0.0128)	(0.0075)	(0.0186)	(0.0193)	(0.0222)
First Stage F-stat.	7.0816	34.3959	2.1093	21.7688	15.7231	15.9062

Note: NFHS data. Bootstrap standard errors in parentheses. All specifications include individuals and household controls, state-religion fixed effects, mother's cohort-religion fixed effects and state specific time trends (up to degree four).

Table A6: Self-reported Decision Making and Woman's Control of Resource

	𝕀(Woman Participates in Final Decisions on)			Autonomy
	Household Purchases	Visits to Family and Relatives	Own Health	Index
$T = \mathbb{I}(R \geq 50)$	0.0245	0.0303	0.0195	0.228
	(0.0147)	(0.0125)	(0.0160)	(0.0994)
Mean Dependent Variable	0.6642	0.7130	0.7400	65.8703

Note: NFHS data. The sample includes married women of age 15 to 49 in nuclear households with up to 4 children. Bootstrap standard errors in parentheses. All specifications include individuals and household controls, state-religion fixed effects, mother's cohort-religion fixed effects and state specific time trends (up to degree four).

Table A7: Self-reported Measures of Power and MR-LATE: Example 1

	𝕀(Woman Participates in Final Decisions on)								
	Own Health			Household Purchases			Visits to Family and Relatives		
	BMI	Pr(BMI≤18.5)	Pr(Anemic)	BMI	Pr(BMI≤18.5)	Pr(Anemic)	BMI	Pr(BMI≤18.5)	Pr(Anemic)
MR-LATE	154.4886	-14.4061	-7.3242	77.0713	-10.5959	-8.0653	1.6255	-0.3637	0.6134
	(181.9458)	(6.4603)	(4.5583)	(161.4285)	(4.8256)	(3.7369)	(28.1084)	(0.6721)	(0.4610)
First stage F-statistics:									
T^a	0.1570	0.1570	0.1570	4.0934	4.0934	4.0934	1.1688	1.1688	1.1688
T^b	0.0125	0.0125	0.0125	0.0218	0.0218	0.0218	4.8848	4.8848	4.8848

Note: NFHS data. The sample includes married women of age 15 to 49 in nuclear households with up to 4 children. Bootstrap standard errors in parentheses. All specifications include individuals and household controls, state-religion fixed effects, mother's cohort-religion fixed effects and state specific time trends (up to degree four). $T^a = 1$ if the answer is "the wife," otherwise $T^a = 0$. Similarly, one can let $T^b = 1$ if the answer is "the husband," otherwise $T^b = 0$.

Table A8: Self-reported Measures of Power and MR-LATE: Example 2

	BMI	Pr(BMI≤18.5)	Pr(Anemic)
MR-LATE ($k=0$)	18.1815	-1.7050	-1.0356
	(6.7474)	(0.5801)	(0.3344)
First stage F-statistics:			
T^a	3.9385	3.9385	3.9385
T^b	3.9385	3.9385	3.9385
MR-LATE ($k=1$)	20.6218	-2.6167	-1.0582
	(20.4949)	(1.4365)	(0.4026)
First stage F-statistics:			
T^a	3.6756	3.6756	3.6756
T^b	0.5468	0.5468	0.5468
MR-LATE ($k=2$)	20.3989	-2.1133	-0.8088
	(19.1140)	(0.9096)	(0.4782)
First stage F-statistics:			
T^a	1.6341	1.6341	1.6341
T^b	1.3387	1.3387	1.3387
MR-LATE ($k=3$)	41.1914	-3.3276	-2.5234
	(86.8981)	(1.0959)	(1.0992)
First stage F-statistics:			
T^a	0.1286	0.1286	0.1286
T^b	2.1703	2.1703	2.1703

Note: NFHS data. The sample includes married women of age 15 to 49 in nuclear households with up to 4 children. Bootstrap standard errors in parentheses. All specifications include individuals and household controls, state-religion fixed effects, mother's cohort-religion fixed effects and state specific time trends (up to degree four). T^a and T^b are based on the answers to six questions about women's participation in household decision and women's mobility. $T^a = \mathbb{I}(R \geq \kappa^a)$, with $\kappa^a = k$, and $T^b = \mathbb{I}\left(R < \kappa^b\right)$, with $\kappa^b = -k$.

Table A9: Woman's Health: Linear Model

	First Stage		Second Stage		
	Woman's Resource Share (R)		BMI	Pr(BMI≤18.5)	Pr(Anemic)
$\mathbb{I}$(HSA)	-0.0337	0.0702			
	(0.0443)	(0.0278)			
Woman's Resource Share (R, %)			2.962	-0.954	-0.608
			(3.048)	(0.408)	(0.318)
Observations	19,738	10,765	10,765	10,765	10,765

Note: NFHS data. Bootstrap standard errors in parentheses. All specifications include individuals and household controls, state-religion fixed effects, mother's cohort-religion fixed effects and state specific time trends (up to degree four). Column 1: the sample includes married women of age 15 to 49 in nuclear households with up to 4 children. Column 2 to 6: the sample includes married women of age 15 to 49 in nuclear households with up to 4 children with $R \in (40, 60)$.

Table A10: Children's Health: Linear Model

	First Stage	Second Stage					
	Woman's Res. Share (R)	W-for-age (z-score)	H-for-age (z-score)	Pr (Cough)	Pr (Fever)	Pr (Diarrhea)	Pr (Vacc.)
$\mathbb{I}$(HSA)	0.107						
	(0.0501)						
Woman's Resource Share (R,%)		1.481	2.037	-0.483	-0.445	-0.330	-0.0938
		(0.909)	(1.154)	(0.297)	(0.274)	(0.201)	(0.219)

Note: NFHS data. Bootstrap standard errors in parentheses. All specifications include individuals and household controls, state-religion fixed effects, mother's cohort-religion fixed effects and state specific time trends (up to degree four). The sample includes children 0 to 5 in nuclear households with up to 4 children.

Table A11: Engel Curves Estimation Results (NSS Sample)

	$R(X)$	$\widehat{\delta}_w(X)$	$\widehat{\delta}_m(X)$	$\widehat{\beta}(X)$
$\mathbb{1}$(2 children)	-0.0554	-0.2660	-0.2260	0.0260
	(0.0166)	(0.3370)	(0.3350)	(0.0409)
$\mathbb{1}$(3 children)	-0.0395	-0.5610	-0.3470	0.0455
	(0.0211)	(0.3880)	(0.3860)	(0.0471)
$\mathbb{1}$(4 children)	-0.0824	0.0203	0.1510	-0.0180
	(0.0270)	(0.5900)	(0.5920)	(0.0721)
Fraction of Female Children	-0.0146	-0.4530	-0.5150	0.0544
	(0.0178)	(0.3600)	(0.3620)	(0.0437)
Gender Age Gap (Man - Woman)	0.0618	1.6710	1.3630	-0.1110
	(0.1320)	(2.6590)	(2.7870)	(0.3390)
Woman's Age	-0.571	2.3910	1.6600	-0.1510
	(0.1200)	(3.1000)	(3.1200)	(0.3810)
Children's Avg. Age	-0.1570	-3.4250	-2.6360	0.2270
	(0.2410)	(5.3780)	(5.3550)	(0.6520)
$\mathbb{1}$(Hindu, Buddhist, Jain, Sikh)	0.0951	1.520	1.038	-0.141
	(0.0214)	(0.3430)	(0.3480)	(0.0418)
$\mathbb{1}$(Sch. Caste, Sch. Tribe, Oth. Back. Caste)	-0.0313	-0.0360	-0.0002	-0.0180
	(0.0158)	(0.3090)	(0.3130)	(0.0380)
$\mathbb{1}$(Own Land)	0.0060	-0.1300	-0.0490	0.0299
	(0.0167)	(0.3270)	(0.3250)	(0.0393)
$\mathbb{1}$(Woman Completed High School)	0.0610	-0.2490	-0.2380	0.0368
	(0.0259)	(0.4840)	(0.4860)	(0.0563)
$\mathbb{1}$(Man Completed High School)	0.0254	0.0243	-0.2200	0.0272
	(0.0207)	(0.4020)	(0.4060)	(0.0477)
$\mathbb{1}$(Rural)	-0.0115	1.545	1.740	-0.194
	(0.0155)	(0.3200)	(0.3200)	(0.0390)
$\mathbb{1}$(North)	-0.0588	0.1130	0.987	-0.0423
	(0.0270)	(0.5250)	(0.5230)	(0.0630)
$\mathbb{1}$(East)	0.128	0.2750	0.0353	-0.0094
	(0.0276)	(0.5440)	(0.5320)	(0.0656)
$\mathbb{1}$(North-East)	0.19	-1.505	-2.055	0.150
	(0.0321)	(0.5970)	(0.5810)	(0.0713)
$\mathbb{1}$(South)	-0.0631	1.521	1.216	-0.139
	(0.0263)	(0.5590)	(0.564)	(0.0681)
Constant	0.623	6.743	6.997	-0.718
	(0.0476)	(1.0060)	(1.0070)	(0.1240)

Note: NSS data. Robust standard errors in parenthesis. Age variables are are divided by 100 to ease computation.

Table A12: Difference-in-Difference Test for External Validity

	$\widehat{\lambda}_{TZ}$	Bootstrap St. Error
Women:		
Body Mass Index (BMI)	0.6920	1.0184
Pr(BMI≤18.5)	0.0070	0.0347
Pr(Anemic)	0.0190	0.0270
Men:		
Body Mass Index (BMI)	1.0120	1.3337
Pr(BMI≤18.5)	-0.0190	0.2031
Pr(Anemic)	-0.0790	0.0755
Children:		
Height-for-age (z-score)	1.4550	8.4582
Weight-for-age (z-score)	-0.0270	9.1306
Pr(Cough)	-0.5960	0.6306
Pr(Fever)	-0.1570	0.4526
Pr(Diarrhea)	-0.1310	0.2607
Pr(Any Vaccination)	-0.2210	0.3046

Note: NFHS data. Test based on Kowalski (2016) and Kowalski et al. (2016) and implemented in Stata with `mtebinary`. $\lambda_{TZ} = 0$ implies no treatment effect heterogeneity and global external validity (under the assumption that T is the correct treatment).