Unobserved Inputs in Household Production

Sergey Mityakov, Thomas A. Mroz *

Abstract

With few exceptions, empirical household production studies ignore unobserved inputs. We demonstrate that without additional assumptions, the estimable impacts of the observed inputs cannot provide informative estimates of their marginal products due to contaminating variation in unobserved inputs; not even the sign of marginal impacts can be ascertained. Instrumental variables cannot solve this problem since every candidate for an instrument affecting an observed input, including experimental assignments, would also affect unobserved inputs choices through the budget constraint, invalidating this variable as an instrument. We show that under certain additional assumptions an appropriately specified empirical model can provide bounds for true marginal products. Our main point is that unless one is willing to make assumptions of this nature, estimated effects would have no useful interpretation. Almost all existing empirical studies of health, child development, and job-training programs fail to account for this issue, rendering their conclusions incomplete and possibly misleading.

Keywords: Household Production; Model Construction and Evaluation

*Mityakov: Department of Economics, Clemson University, 228 Sirrine Hall, Clemson, SC 29634. Email: smityak@clemson.edu. Mroz: Department of Economics, Andrew Young School of Policy Studies, Georgia State University, Atlanta, GA 30303 tmroz@gsu.edu

1 Introduction.

To make informed policy recommendations, economists need to understand how inputs to household production functions affect measurable outcomes like health and children's test scores. Two key issues make this a difficult task. First, a household's choices of inputs likely depend on unobserved to the researcher baseline characteristics and households' unobserved abilities to make use of the inputs. As a consequence, households' choices of the levels of the inputs are likely to be statistically endogenous determinants for the estimation of household production functions. Researchers have used a variety of approaches to address this issue, such as better measures of productivity, experimentally assigned inputs, instrumental variables, and natural experiments.

The second issue, which is the focus of this paper, arises because one almost never can observe all of the inputs to production chosen by households. In general the estimable impact of an observed input on the health outcome would confound the true marginal effect of that observed input with the marginal effects of unobserved inputs. Instrumental variables, experiments, or other approaches used to solve the omitted variables problem in the case of unobserved fixed household characteristics mentioned above will not solve the omitted variable problem in this case. This happens because the unobserved inputs are optimally chosen by households. Any "exogenous" variation in observed inputs would be associated with changes in unobserved inputs, since the unobserved inputs are chosen subject to the same budget constraint as the observed inputs. Any exogenous or endogenous variation in any observed input would be correlated with changes in unobserved inputs even when observed and unobserved inputs are separable in the production function.

We use a model of utility maximization subject to a budget constraint in conjunction with a household production function to derive precise interpretations of estimated effects of observable inputs on household's outcomes. The economic model provides considerable guidance for researchers about the types of variables one needs to include in a "hybrid" household production function in order to justify these interpretations. In general, these estimated effects do not correspond exactly to standard *ceteris paribus* marginal effects of the observed inputs in the home production. Often, however, the estimated effects will provide a bound on the magnitude of the true marginal effect. We also show how one can improve on those bounds by controlling for the readily available but often ignored information on pure consumption goods when estimating the hybrid production function. These bounds arise solely from a theoretical model describing the behavior of an optimizing economic agent. We demonstrate the potential importance of these bounds through a simple simulation exercise.

We show that the least informative bound for the marginal product is closely related to the "policy effect" one would estimate in an experimental setting. In particular, if one were to randomly assign the production input under consideration while also controlling for incomes and prices as we describe below, then the estimated "policy effect" would be identical to the bound we present. Averaging these conditional effects over the sample distribution of prices and incomes would yield the simple policy effect that only examines mean differences between the treated and untreated in the experimental study. Given this close relationship, our bound can be used to calculate the "policy effect" for different target populations without the need to conduct numerous experimental studies.

While we frame our discussion in terms of health production to make the analysis more specific, the estimation and interpretation issues we derive apply to any household maximization problem with production. This includes, for example, estimates of the impacts of school inputs and parents' behaviors on children's developmental outcomes. Given that one almost never observes all of the relevant inputs to the household production function, our findings suggest that most of the estimates in the literature on household production, not just those in the health economics literature, need to be reinterpreted. Unless researchers bring external information to bear in their analyses of household production functions, they will be unable to validate most interpretations of their estimates of the impacts of observed inputs as interesting technological relationships.

2 Background

Early work on the estimation of production functions with missing inputs mostly focused on the case where there was a fixed unobserved input that was not varied as part of the optimization process. The motivation for these types of formulations came from an assumption that there could be unobserved, firm specific managerial factors affecting input choices and output levels (Hoch, 1955; Mundlak, 1961). In general, longitudinal data with firm specific fixed effects could be used to obtain consistent estimates of the marginal impacts of the observed inputs to the production process. More recently the industrial organization literature has explored structural methods to control for time-varying unobserved productivity shocks that could affect a firm's input choices.[1] Such approaches, however, typically would not work in the case when the missing input itself is a choice variable, which is the focus of this paper.

Rosenzweig and Schultz (1983) took the analysis of production functions

[1] Olley and Pakes(1996), Levinsohn and Petrin(2003), Ackerberg, Caves and Frazer(2006).

with missing inputs to a more fundamental level. All inputs are chosen optimally as a part of a household utility maximization process, but the researcher does not observe the chosen levels for a subset of the inputs. They discuss a commonly used approach, the "hybrid production function," to deal with the unmeasured inputs. In that approach, the researcher estimates a relationship where output is a function of the observed inputs, the prices of the unobserved inputs, and the household's level of exogenous income. They demonstrate that the estimated impact of an observed input on health outcomes in this hybrid specification does not measure the true marginal impact of the observed input holding constant the levels of the other observed inputs and the levels of the unobserved inputs. Unobserved inputs that are chosen as part of the household's utility maximization, subject to a budget constraint, result in consequences well beyond those addressed in the early literature that only had fixed, unobserved inputs affecting the choices of the variable inputs and output levels.

Todd and Wolpin (2003) discuss production functions for cognitive achievement and point out that the inclusion of proxy variables like income and prices for unobserved inputs could lead to more biased measures of the impacts of the observable inputs than an empirical approach that excludes these proxy variables (see, also, Wolpin, 1997). They discuss various approaches one might use when not all of the relevant inputs can be observed and assumptions needed for these approaches to obtain asymptotically unbiased estimates of the marginal effects of the observed inputs. A major conclusion of their study is that instrumental variables approaches will be unlikely to resolve problems arising from omitted inputs in the production function. This happens because the omitted inputs are chosen by the families and so would typically be correlated with the observed inputs. In this situation, any instrument that has power to predict the observed input should also predict the unobserved inputs. It could not be a valid instrument.

3 Preliminary Modeling Issues

A common shortcoming of the studies discussed above is their failure to provide an exact link between the theoretical model and the specification of the empirical model. In this section we fill in that gap. In the subsequent section we use the results from this analysis to specify and interpret feasible empirical specifications of health production functions that are consistent with a theoretical model of household utility optimization. Throughout most of the analysis in this and the subsequent section, we assume that there are only two purchased inputs used in the health production function, X and Z, and that utility only depends on the amount of health produced by the household, H, and the consumption of a composite commodity C. We extend the analysis to the general case where the are multiple observed and unobserved inputs in section 4.4.

Let the function $H = F(X, Z)$ be a household's health production function. The standard demand functions for the two health inputs are given by $X(p_X, p_Z, p_C, I)$ and $Z(p_X, p_Z, p_C, I)$ where the p's are the prices of the three purchased goods and I is exogenously determined income. We assume that one could estimate nonparametrically the two demand functions and the health production function $F(X, Z)$ if the health outcome H, the two inputs X and Z, the prices of the three goods, and exogenous income I were observed by the researcher. Since prices and incomes do not enter the production function directly, they are potential candidates to use as instrumental variables to control for the possible endogeneity of X and Z. The problem we want to address is what one might be able to learn about the effect of X on H when there is only information on H, the prices, income, and the quantity of the input X. That is, the levels of the input Z and the consumption goods C are not observed.

A seemingly obvious approach would be to substitute the demand function for Z into the production function and then estimate this form of the "hybrid" production function. This demand function, by definition, will depend on the household's preferences and the form of the health production function. This approach, however, will in general result in an unidentified model[2]. To see this, substitute the demand function for the unobserved input into the production function. This yields $H = F(X, Z(p_X, p_Z, p_C, I))$. When the form of the demand function is unknown, this becomes some general function of observed inputs, prices, and income: $H = G(X, p_X, p_Z, p_C, I)$.

Since X depends on exactly the same set of variables determining Z, (p_X, p_Z, p_C, I), there is an exact functional relationship among the five arguments in the function $G(\cdot)$. A nonparametric model for estimating the function G could admit almost any estimate of the effect of X on H through the function G by offsetting changes in the impacts of p_X, p_Z, p_C, and I on H. This nonparametric expression of an identification problem is similar to perfect multicollinearity in a linear regression model[3]. Like in the linear regression

[2]This might not be an issue if one can impose the exact functional form of the health production function $F(X, Z)$ and has precise information about the functional forms for the demand function $Z(p_X, p_Z, p_C, I)$. Identification here would come from functional form assumptions.

[3]The function $F(\cdot)$ does contain some separability restrictions that are not imposed on a general function like $G(\cdot)$. However, given the nonidentification result discussed above, it will be impossible to exploit these separability restrictions to uncover the marginal effect of X. For example, instead of the function $H = F(X, Z(p, I))$ one can always substitute an observationally equivalent function $H = F(X, (Z(p, I))) + \phi(X) - \phi(X(p, I))$ for any function $\phi(\cdot)$, where $X(p, I)$ is the true demand function for X. Since $\phi(\cdot)$ is arbitrary one can estimate any effect of X on H while satisfying the separability restrictions.

model, this identification problem can only be overcome by the imposition of some, hopefully valid, set of constraints. Economic theory, however, provides little guidance for the types of constraints one might impose in order to obtain the true marginal impact of the input X on the health outcome.

This non-identification problem is distinct from the endogeneity of inputs issue arising from unobservable productivity in the industrial organization literature on estimating production functions. That literature explores structural approaches to control for time varying productivity differentials that are not due to variations in optimally chosen unobserved inputs[4]. Here, all inputs, both observed and unobserved, are choice variables in the individual optimization problem. The key issue here arises because variations in the observed input X, even under random experimental assignment, typically will be associated with variations in the choice of the unobserved input Z.

Rosenzweig and Schultz's (1983) analysis of the hybrid production function differs from the one presented here by its exclusion of the price of the observed input (p_X) as a determinant of the health outcome. In general this would be valid only when the unconditional demand for Z does not depend on p_X. Variations in the observed input X would then arise from variations in p_X, which would not be perfectly determined by variations in p_Z, p_C, and I. The Rosenzweig and Schultz formulation for the hybrid production function could more generally be derived when all households face the same price p_X. But in this case, there would be no variation in the input X that did not arise from variations in p_Z, p_C, and I, resulting again in a non-identified specification. Without strong and mostly ad hoc assumptions, the form of the hybrid

[4]The control-function approaches suggested in Olley and Pakes(1996), Levinsohn and Petrin(2003), and Ackerberg, Caves and Frazer(2006), for example, would not be feasible when the "productivity shock" is itself a chosen input.

production function discussed by Rosenzweig and Schultz cannot be derived from a standard model of utility maximization or used to uncover empirically the impacts of observed health inputs.

The conditional demand function approach discussed in Liu et al (2009) can overcome the basic identification issue inherent in the unrestricted form of the hybrid production function G. In particular, consider the demand function for the unobserved input Z conditional on the optimally chosen level of the observed input X. Using standard rationed demand analysis, this conditional function can be written as $Z = q_z(p_C, p_Z, I^*, X)$, where $I^* = I - p_X X$ is the amount of income the household has left to allocate between the consumption good C and the unobserved input Z. In general, the conditional demand for Z will depend on the amount of X chosen by the household even holding the level of I^* fixed. Substituting this constrained demand for Z into the true production function yields $H = F(X, q_z(p_C, p_Z, I^*, X))$. Without assumptions on the form of the function $q_z(\cdot)$, the estimable conditional hybrid production function becomes $H = G_C(X, p_C, p_Z, I^*)$. In this situation, the effect of X on H, through the function G_C and conditional on p_C, p_Z, and I^*, should be nonparametrically identified.

It is crucial that one conditions on the value of I^* instead of its components in order for this particular effect of X to be identified. The estimate of the partial effect of X on H obtained through the conditional hybrid production function G_C, however, does not have a simple and straightforward interpretation. In the next section we derive interpretations of this type of effect.

4 Basic Model

We begin this section with a detailed analysis of the case where there is only one observed input and one unobserved input to the household production function. The setup, intuition and analytic derivations for this simple case carry over to the multidimensional case. In the last subsection we extend the analysis to situations with multiple inputs, and the appendix contains complete derivations when there are multidimensional unobserved and observed inputs.

4.1 Preferences and Technology

Assume consumers derive utility U from health H and some other consumption goods C. For simplicity, H and C are assumed to be one-dimensional. Health is produced with several inputs. We denote as X inputs which are observed and as Z the unobserved inputs. Assume preferences are given by a general utility function

$$U = U(C, H; \tau), \tag{1}$$

where τ is an arbitrary vector of household-specific taste parameters. The household health production is given by a function F with standard properties

$$H = F(X, Z; \rho), \tag{2}$$

where ρ represents productivity parameters that could vary from household to household[5]. The household budget constraint is:

$$p_X X + p_C C + p_Z Z = I. \tag{3}$$

[5] The taste (τ) and productivity (ρ) parameters do not affect the comparative static analysis presented below, so we often drop them in the derivations to save notation. They are, however, crucial determinants of the household's optimal choices. In empirical analyses the presence of these unobserved preference and productivity parameters means that all observed household inputs must be treated as endogenous variables.

Throughout we consider an interior solution and assume that the corresponding second order conditions are satisfied.

As a straightforward extension, one can interpret this static optimization formulation as part of a dynamic optimization model with two stage budgeting. To see this explicitly, one can write the utility function as

$$U = U(C, H; \tau) = W(C_t, H_{t-1}, H_t, \tau_{1_t}) + \frac{1}{1+\beta} V(H_t, \tau_{2_t}, \Omega_t) \tag{4}$$

where H_{t-1} is the individual's health when entering the current period. It is not a contemporaneous choice variable. $W(.)$ represents the current period utility function that depends on current consumption, the health stock inherited from the previous period, the amount of health stock at the end of the period and possibly a subset of the original taste parameters τ. $V(.)$ is the expected maximal value of future utility that the household will receive from the next period onwards, discounted at rate β. It includes the information set available to the household at time period t, Ω_t. Income, I, in this instance should be interpreted as the household's optimal expenditure level in time period t. The household's technology parameters in the production function, ρ, can also depend upon the the level of the inherited capital stock.

4.2 Interpreting Estimated Effects of Observed Inputs

Consider the following econometric problem. We would like to estimate the marginal product of input X on health production: $\frac{\partial F}{\partial X}$. The information available is structured in the following way. The levels of H and X are observed; prices p_X, p_C, and p_Z are observed. Income I is observed. The levels of other goods C and the health input Z are not observed. Our research goal is to understand which effects we are able to estimate and whether we can use these to place informative bounds on the marginal effects of the observed health inputs.

The estimated effect of the observed input X on a health measure H when conditioning on an arbitrary set of controls Y would measure:

$$\left.\frac{dH}{dX}\right|_Y = \frac{\partial F}{\partial X} + \frac{\partial F}{\partial Z}\left.\frac{dZ}{dX}\right|_Y \tag{5}$$

Here $\left.\frac{dZ}{dX}\right|_Y$ is the derivative which indicates the change in the unobserved inputs Z when X changes by dX given the set of control variables Y.

SERGEY EDIT 2016 We assume that the data are rich enough so that individual level heterogeneity parameters that are fixed for a given individual (denoted by ρ, τ in the previous section) can be perfectly controlled for by an appropriate (potentially non-parametric) estimation technique. For example, this would be the case when one has the data on same individual facing different prices p_X and observed consuming X and $X + dX$. Thus, all results, we derive in our paper, apply at each point of individual level heterogeneity parameters $(\rho_0,\ \tau_0)$: e.g. in formula (5) above all partial derivatives are evaluated at the point (X_0, Z_0) which is the optimal chose for the individual with individual heterogeneity parameters $(\rho_0,\ \tau_0)$ at particular levels of prices p_X, p_Z, p_C and income (current period optimal expenditures) I. With this in mind, we omit $(\rho_0,\ \tau_0)$ from all formulas to simplify notation.

The bias term $\frac{\partial F}{\partial Z}\left.\frac{dZ}{dX}\right|_Y$ in (5) arises from individual optimal choice of unobserved production input Z in response to the changes in prices and income that bring about the change in the optimal choice of observed production input X. Instrumental variables approaches, which are useful to deal with unobserved heterogeneity due to omitted variables *fixed* at the individual level, are likely to be invalid to deal with this bias due to unobservability of the *optimally chosen* variables Z. As we show below any factor that changes the optimal choice of X is likely to involve a concurrent change in unobserved optimally chosen unobserved variables Z rendering all potential candidates for instruments invalid. END SERGEY EDIT 2016

In this regard a major issue for an empirical analysis of the effect of X on H is the choice of an appropriate set of controls Y to minimize the "bias term" in the above expression, $\frac{\partial F}{\partial Z}\frac{dZ}{dX}\Big|_{Y=const}$.

As we argued above, using all available information $Y = (p_C, p_Z, p_X, I)$ results in an unidentified model, as X itself is fully explained by those same variables. One needs to put restrictions on the set of conditioning variables to obtain an identified econometric model. Once effects are identified, one can provide an economic interpretation of the estimable effect of X on H.

Using the conditional demand function for the unobserved input discussed above, consider the following optimization problem conditional on the level of observed input X:

$$\begin{aligned}&\max_{C,Z} U(C, F(X,Z))\\&\text{s.t. } p_C C + p_Z Z = I^* \equiv I - p_X X\end{aligned} \tag{6}$$

The conditional demand function for unobserved health input Z associated with this problem is:

$$Z = q_Z(p_C, p_Z, I - p_X X, X) \tag{7}$$

We assume that the data are rich enough so that we observe relevant variations in X while holding total expenditure on other goods C and unobserved input Z, $I^* = I - p_X X$, constant. Then, if we regress the observed health level H on the observed level of health input X (which does not enter the utility function directly) and the total expenditures on all goods other than X, I^*, (controlling for prices p_Z, p_C) we would estimate the following effect[6]:

[6]This effect is identified by variations in X induced by changes in p_x and I that leave I^* constant.

$$\left.\frac{dH}{dX}\right|_{I^*=I-p_X X=const} = \frac{\partial F}{\partial X} + \frac{\partial F}{\partial Z} \left.\frac{dZ}{dX}\right|_{I^*=I-p_X X=const} \tag{8}$$

The estimated effect is the sum of the effect of interest, the marginal product of input X in health production $\frac{\partial F}{\partial X}$, as well as a bias term related to the fact that as we change the level of input X the individual might change the level of unobserved health input Z, even when prices p_Z and p_C and total expenditures on C and Z stay constant, i.e. $\frac{\partial F}{\partial Z} \left.\frac{dZ}{dX}\right|_{I^*=I-p_X X}$[7].

The key question we ask is what is the direction and size of the bias. Assuming that both the observed and unobserved inputs have positive marginal products, the estimated effect will be biased in the direction towards zero (negatively biased) whenever the derivative of the conditional demand for Z with respect to the observed input X is negative. To examine whether this would be the case, we need to compute how the unobservable input Z changes when we change the observed input X holding the combined expenditure on Z and C fixed, $\left.\frac{dZ}{dX}\right|_{I^*=I-p_X X=const}$. That is, we need to understand the derivative of the conditional demand function $Z = q_Z(p_C, p_Z, I^*, X)$ with respect to the observed input X holding I^* fixed. Theorem 1 provides one possible answer to this question

Theorem 1 *Suppose health H and other goods C are normal goods and the degree of complementarity in health production between the beneficial observed input X and the beneficial unobserved input Z is sufficiently small (the cross derivative F_{ZX} is small if positive or negative: i.e. the increase in one of the inputs lowers the marginal effect of the other or barely increases it).*[8] *Then the*

[7] The set of controls Y is this case is: (p_Z, p_C, I^*)

[8] Theorem 1 considers one possible set of conditions under which we can derive an informative bound for the marginal product of an observed input X. In Theorem 2 we derive necessary and sufficient conditions that do not require knowledge about substitutability between X and Z but instead rely upon some information about the behavior of consumption

regression of observed health H on the observed health input X holding prices p_C and p_Z and total expenditure on C and Z ($I^* = I - p_X X$) constant, would underestimate the true value of the marginal product of X in health production. The estimable effect of the productive input might even be negative.

The Appendix section 9.1 contains a complete derivation of results for a wide set of cases and the interpretations for other effects estimated by a hybrid production model. Here we outline the main result for Theorem 1. The key equation describing the change in the demand for the unobserved input due to a change in the observed input holding I^* constant is:

$$Bias = Bias_1 = F_Z \frac{dZ}{dX} = \frac{U_H F_Z^2 F_X}{\Delta} \left[\frac{\partial}{\partial H} \left(\log \frac{U_C}{U_H} \right) - \frac{F_{ZX}}{F_Z F_X} \right] \tag{9}$$

The partial derivative with respect to H of the term in parentheses will be positive whenever C is a normal good[9], and Δ is negative by the second order conditions. One's ability to unambiguously sign the overall bias therefore depends on the substitutability of the two inputs in producing H. If both X and Z are beneficial inputs and the two inputs are substitutes or only weak complements, then the conditional demand for Z will fall with an increase in X. From (8), this implies that the identified effect of X on Z will underestimate the true marginal impact of a beneficial input X.

Intuitively, an increase in the observed beneficial health input X increases the level of health H if Z is kept constant. Such an increase in H can be thought of as an increase in health endowment in the conditional (on

---

choice C in response to a change in X.

[9]In particular, $\frac{\partial}{\partial H}\left(\log \frac{U_C}{U_H}\right) = \frac{\eta_C}{-\epsilon^*_{CC}} \frac{1-s_C}{H}$, where ϵ^*_{CC} is the compensated own price elasticity, η_C is the income elasticity of demand for good C and s_C is the share of income spent on C. This term will be large when the income elasticity for the consumption good is high and also when the compensated own price elasticity for the C is small. See Lemma 1 in Appendix 10.2.

X) demand system[10]. Since consumption is a normal good, this endowment/"income" increase would tend to be partially reallocated to an increase in consumption C. This effect is represented by the first term inside the brackets in equation (9). If this effect were to dominate the second term inside the brackets, the consumer would unambiguously choose a higher C and hence a lower level of the unobserved input Z because the income I^* allocated between C and Z is held constant.

If the change in X makes Z less productive ($F_{ZX} < 0$), then this reallocation towards consumption would be reinforced because the shadow price of health, $\frac{p_Z}{F_Z(X,Z)}$, would rise with the increase in X. The income and substitution effects operate in the same direction. This substitution effect is represented by the second term inside the brackets in equation (9). In fact, if this second term is strong enough, the beneficial input X might even appear to be harmful in estimation[11]. When the increase in X increases the marginal product of Z ($F_{ZX} > 0$), the income effect (due to higher endowment of health) and substitution effect (due to fall in the shadow price of health) work in opposite directions and the sign of the bias cannot be determined.

One can also sign this bias term in the case when X and/or Z are harmful, but have no direct impact on utility (e.g. X may be dangerous working conditions for which the person is compensated, so that $p_X < 0$)[12]. The sign

[10]Due to an increase in X by dX health endowment would rise to $H_1 = F(X + dX, Z)$.

[11]This finding parallels the famous Peltzman (1975) argument about the impact of mandatory seatbelt laws on automobile accidents. In this argument, mandatory seatbelt use (an exogenous increase in the beneficial input X) could result in a perceived, dramatic decline in the marginal product of safe driving and lead to a decline in safe driving (unobserved Z). That decline could be so large that the total incidence of accidents could increase after the introduction of the law.

[12]However, in the case of harmful health inputs, a more relevant assumption would be that those inputs could also affect utility directly, e.g. an individual consumes alcohol because he receives utility from it despite the fact that it is bad for her health. We discuss this extension in the appendix.

of this term does not depend on the sign of F_Z. When the observed input X adversely affects health then the bias will be positive provided the term $\frac{F_{ZX}}{F_X F_Z}$ is negative or small if positive. Thus, we establish the following:

Corollary 1 *Suppose health H and other goods C are normal goods. Assume the observed health input X and the unobserved health input Z have no direct effect on utility. Suppose that $\frac{F_{ZX}}{F_Z F_X} < 0$ or small if positive. Then the regression of observed health H on observed health input X holding prices p_C, p_Z and total expenditure on C and Z ($I^* = I - p_X X$) constant, would underestimate the true value of the marginal product of a beneficial health input X and overestimate the marginal product of a harmful health input X (underestimate the adverse impact). The bias may be large enough so that the estimated effect would be opposite in sign to the true marginal effect of X.*

The interpretation of the condition $\frac{F_{ZX}}{F_Z F_X} < 0$ is quite straightforward. In the case when both X and Z are beneficial inputs it means that $F_{ZX} < 0$, i.e. an increase in one of the inputs *decreases* the marginal effect of the other input. The same condition holds in the case when both X and Z are harmful. The intuitive interpretation will be different though. Suppose that X is smoking and Z is illegal drug use, then $F_{ZX} < 0$ would mean that the increase in smoking *raises* the marginal damage from illegal drug use. When one of the two inputs is beneficial and the other harmful, then the relevant condition for the bound to hold is $F_{ZX} > 0$, i.e., the increase in the amount of beneficial input (e.g. jogging) *decreases* the marginal damage of the harmful input (e.g. smoking).

These substitutability conditions illustrate a one possible set of sufficient conditions needed to characterize the relation between the actual marginal product and the estimable impact of observed health input. One can instead use the budget constraint in consumer problem (6) to provide a necessary and sufficient condition for the estimated effect to provide an informative bound for the true marginal product of observed input X. This involves understanding how consumption C changes in response to changes in X. Theorem 2 summarizes our findings:

Theorem 2 *Consider the regression of observed health H on the observed health input X holding prices p_C and p_Z and total expenditure on C and Z ($I^* = I - p_X X$) constant.*

i. Suppose in response to the increase in the observed input X the choice of consumption good C increases: $dC > 0$. This is a necessary and sufficient condition for the estimated effect of observed input X to be lower than its true marginal product.

ii. $dC < 0$ is a necessary and sufficient condition for the estimated effect of observed input X to be an upper bound for its true marginal product.

Proof: This is a special case of Theorem 5 (proved below), for the single-dimensional X and Z case. ■

This result does not depend upon substitutability patterns between observed (X) and unobserved (Z) inputs, nor does it require any assumptions about the sign of the marginal products of X or Z. Further this result does not depend upon consumption C or health H being a normal good. This flexibility comes at a cost. In Theorem 2 one imposes a restriction not on the model fundamentals but rather on the optimal choice of an endogenous variable (consumption good). From a practical perspective, however, researchers might find it easier to verify whether this assumption holds. Strictly speaking, the budget constraint logic behind Theorem 2 does not require Z or X to be single-dimensional. In section 5, where we analyze the multidimensional inputs case, we discuss a multidimensional extension of Theorem 2.

To summarize, to obtain an identified model in the case when some inputs in the health production function are unobserved one can estimate the following regression model:

$$\begin{aligned} H &= F(X, Z(p_C, p_Z, I - p_X X, X, \rho, \tau), \rho) \\ &\equiv h(X, p_C, p_Z, I - p_X X, \rho, \tau). \end{aligned} \tag{10}$$

As implied by economic theory, the regression function $h(\cdot)$ should contain all of the observed health inputs, the prices of all the unobserved health inputs and pure consumption goods, and the income the household has to allocate after it purchases the observed health inputs[13]. Theorem 1 and Corollary 1 describe conditions when the estimated effect $\frac{\partial h}{\partial X}$ is likely to be a lower bound for the true marginal effect of the observed input X.

Unlike the attenuation bias one finds for measurement error problems in empirical models, the attenuation bias we derive here follows solely from economic theory. The bias arises from a researcher's uncertainty about the actual amount of the unobserved input Z used by the household. This theoretical result provides a "bound" when interpreting a correctly specified hybrid production function when one does not "include" a relevant health input but does account for all other relevant factors, including taste and productivity shifters. It provides the theoretical underpinnings for the specification and interpretation of the empirical hybrid health production function.

Note that our empirical specification differs from the ones suggested in the literature. Todd and Wolpin (2003) argued that including income as a proxy for omitted inputs is likely to confound the estimates of the effects of observed

[13] Since the theoretical hybrid production function depends upon unobserved tastes and productivities (τ and ρ), in the empirical analysis all observed household inputs should be treated as endogenous.

inputs. We argue, however, that a properly adjusted income measure should always be included in the regression for the estimated effects to have meaningful economic interpretations. Rosenzweig and Schultz (1983) suggest dropping the prices of included inputs p_X, but this also results in a specification incompatible with economic theory unless one is willing to believe quite restrictive forms for the demand function for input Z.

4.3 Conditioning on Consumption of Non-health Inputs.

A natural question in light of the downward bias result discussed above is whether one should control for the observable part of the consumption vector C, which does not affect health production function per se. In the Appendix section 10.5 we investigate this issue. Though the general direction is ambiguous, one might be able to reduce the bias by controlling for such inputs in some plausible cases. This would provide a more informative bound for the true marginal effect. Theorem 3 summarizes this discussion:

Theorem 3 *Assume that the observed health input X has no direct effect on utility. Further assume that health H does not affect the marginal rate of substitution between two pure consumption goods C and W: $\frac{\partial}{\partial H}\left(\log \frac{U_W}{U_C}\right) = 0$. (This is true when consumption goods C and W are weakly separable from health in the utility function). Controlling for W in the estimation of the hybrid health production function would result in a smaller bias (in absolute value) for the estimated marginal product of observed health input X.*

Theorem 3 can provide information about the true direction of the marginal effect of X and the type of bound one estimates. If the estimated effect is positive and larger than the estimate without controlling for W, then the true impact of X is positive and one has uncovered a (more informative) lower bound for its magnitude. Similarly, if the estimated effect is negative and lower than the estimate without controlling for W, then the true impact is negative and it provides a lower bound on $X's$ detrimental impact. If one has a priori information on the direction of the true marginal effect of X, the bias reduction described in Theorem 3 can allow one to determine whether the estimated effect is an upper or lower bound. Unlike Theorem 1, this result does not require any knowledge about the substitution patterns in the production function.

To see more intuitively the rationale behind this result, suppose we could observe and control for all of the household's pure consumption goods. Then, all remaining income in I^* would be spent on the unobserved input. A non-parametric specification of the regression model would control exactly for $Z = I^*/p_Z$, and there would be no bias from the "omitted input" in the estimation of the effect of the observed input X on health[14]. When we can control for only a subset of the pure consumption goods, we are able to restrict somewhat the possible levels for the expenditure on the unobserved input. And when the marginal rate of substitution between the two pure consumption goods is unrelated to the level of the health output (and consequently to the level of the observed health input), the remaining budget set shrinks without inducing a relative shift between the two consumption goods that is related directly to health. This allows one to obtain a tighter bound without changing the direction of the bias.

To summarize, one cannot estimate the true marginal product of an ob-

[14] This is similar to an approach used in the industrial organization literature, e.g., Olley and Pakes(1996), where one conditions on investment demand to hold constant the unobserved firm fixed effect when estimating production function parameters.

served health input X when some essential health inputs are unobserved. However, one may be able to bring external information to bear and uncover informative bounds using the approaches we outlined above as described in Theorems 1, 2, and 3. If anything, it is crucial to include the prices of omitted inputs and consumption goods in the regression model to obtain an econometric specification consistent with economic theory. The failure to adjust income properly and include it as a regressor in the hybrid production function makes it nearly impossible to interpret estimated effects and to assess how they might differ from marginal effects on health production. Controlling for the consumption of other goods (not necessarily health inputs) can reduce the size of the bias and make the bound more informative.

4.4 Multidimensional Inputs

The extension to the case of multiple observed inputs $X = (x_1, ..., x_N)'$ with a single unobserved input $Z = (z_1)$ is nearly identical to the analytic results derived above. When there are multiple unobserved inputs $Z = (z_1, ..., z_K)'$, however, an increase in an observed input x_i could impact the marginal rates of technical substitution among the unobserved inputs Z. This could lead to increases in the demand for some of the unobserved inputs and decreases in the demands for others. Consequently, even if all observed inputs are pairwise substitutes with each of the unobserved inputs, there could be an increase in the demand for some unobserved input due to an increase in an observed input. Substitutability of the inputs in this situation would not be sufficient for the estimable effect of the observed input to provide an informative bound. However, if there is in addition weak separability of the unobserved inputs in the production process, then the estimable effects do provide informative bounds. We summarize our findings in the following theorem.

Theorem 4 *Suppose health H and other goods C are normal goods. Assume that unobserved health inputs $Z = (z_1, ..., z_K)'$ are weakly separable in production from observed health inputs $X = (x_1, ..., x_N)'$: $H = F(X, Z) \equiv \Phi(X, g(Z))$. Assume that health inputs X and Z have no direct effect on utility. Consider running a regression of observed health H on the observed health input x_i holding constant other observed health inputs x_{-i}, prices p_C and p_Z, and total expenditure on C and Z ($I^* = I - p_X' X$).*

- *If the health input x_i is beneficial and the degree of complementarity between x_i and Z is sufficiently small (the cross derivative Φ_{gx_i} is small if positive or negative) then the beneficial effect of x_i estimated in the regression above would be lower than the true value of the marginal product of x_i in health production. The estimable effect of the productive input might even be negative.*

- *If the health input x_i is harmful and the degree of substitutability between x_i and Z is sufficiently small (the cross derivative Φ_{gx_i} is positive or small in absolute value if negative) then the effect of x_i estimated in the regression above would be higher than the true value of the marginal product of x_i in health production (i.e., the adverse effect of x_i would be underestimated). The estimable effect of the harmful input might even be positive.*

Proof: In the Appendix 9.1 we show that in the case when health inputs X and Z have no direct effect on utility the total bias is given by:

$$Bias = U_H \frac{\partial}{\partial H} \left[\log \left(\frac{U_C}{U_H} \right) \right] F_{Z'} \Delta^{-1} F_Z F_{x_i} - U_H F_{Z'} \Delta^{-1} F_{Z x_i} \qquad (11)$$

The first term in this bias is always nonpositive since $F_{Z'} \Delta^{-1} F_Z \leq 0$ due to the second order conditions, and $\frac{\partial}{\partial H}\left[\log\left(\frac{U_C}{U_H}\right)\right] \geq 0$ whenever consumption and health are normal goods (see Technical Lemma).

Under the separability assumption we can write the second term as:

$$U_H F_{Z'} \Delta^{-1} F_{Zx_i} = U_H \Phi_g \Phi_{gx_i} G_{Z'} \Delta^{-1} G_Z. \quad (12)$$

As before $G_{Z'} \Delta^{-1} G_Z \leq 0$ since Δ is a negative semidefinite matrix. Thus, the contribution of this term to the bias is opposite in sign to Φ_{gx_i}. If x_i is a beneficial health input, $F_{x_i} > 0$, then the $Bias$ will be negative provided Φ_{gx_i} is negative or positive but sufficiently small: i.e. x_i and the unobserved health inputs (in the aggregate) are substitutes or the degree of complementarity is not strong. If x_i is a detrimental health input, $F_{x_i} < 0$, then the $Bias$ will be positive provided Φ_{gx_i} is positive or negative but not large in absolute value. ■

The intuition behind these results is nearly identical to that outlined in the one-dimensional case in our discussions of Theorem 1 and Corollary 1. An increase in health H due to an increase in some beneficial health input x_i can be thought of as an increase in health endowment in the conditional on X demand framework. Facing such an increase in real income, an individual would rationally spend some of that higher health endowment on the (normal) consumption good C. Thus, the observed increase in health H will be lower than suggested by the pure increase in x_i, provided the increase in x_i does not reduce the shadow price of health by a large amount..

5 Consumption behavior and sign of the bias in the multidimensional case.

As in the single unobserved input case discussed in Theorem 2, one can use knowledge of how consumption C would change in the conditional on X demand function, in response to an increase in an observed input by the amount dx_i to establish whether or not the estimable effect of an observed input provides a directional bound on its true marginal product. The only condition we require is that unobserved inputs $Z = (z_1, ..., z_K)$ do not have direct effects on utility (i.e., x_i can affect utility directly but the z_k's cannot).

In this case consumer optimization problem (6) admits the following equivalent restatement with the help of a cost function:

$$\begin{aligned} &\max_{C,Z} U(C, X, H) \\ &\text{s.t. } p_C C + K(X, H) = I^* \equiv I - p'_X X \end{aligned} \tag{13}$$

where $K(X, H)$ is a cost function for household production holding the observed inputs X fixed. It is a solution the following problem:

$$\begin{aligned} &\min_{Z} p'_Z Z \\ &\text{s.t. } F(X, Z) = H \end{aligned} \tag{14}$$

This approach allows to effectively reduce the conditional on X multidimensional problem in C and Z to a two-dimensional choice of C and H (conditional on X).

Now consider changing one of the observed health inputs x_i by dx_i while holding other observed inputs x_{-i}, prices of consumption and unobserved inputs (p_C, p_Z) and income in conditional demand problem I^* constant. Totally differentiating the budget constraint we get:

$$p_C dC + K_{x_i} dx_i + K_H dH = 0 \tag{15}$$

From the envelope theorem we have that: $K_H = \lambda_H$ and $K_{x_i} = -\lambda_H F_{x_i}$ where $\lambda_H > 0$ is the Lagrange multiplier from the cost minimization problem. Thus, we obtain:

$$F_{x_i} - \frac{dH}{dx_i} = \frac{p_C}{\lambda_H} \frac{dC}{dx_i} \tag{16}$$

This equation allows us to establish the necessary and sufficient condition on the estimated effect of health $\frac{dH}{dx_i}$ to be lower or upper bound for the true marginal product of x_i: F_{x_i}. This establishes the following result which is analogous to Theorem 2 discussed in Section 3 above.

Theorem 5 *Assume that observed and unobserved health inputs are multidimensional. Further assume that unobserved health inputs Z have no direct utility effects. Consider the regression of observed health H on the observed health inputs $X = (x_1, ..., x_N)$ holding prices p_C and p_Z and total expenditure on C and Z ($I^* = I - p'_X X$) constant.*

i. Suppose in response to the increase in one observed input x_i, while keeping other observed inputs x_{-i} constant, the choice of consumption good C increases: $dC > 0$. This is a necessary and sufficient condition for the estimated effect of observed input x_i to be lower than its true marginal product F_{x_i}.

ii. Similarly $dC < 0$ is a necessary and sufficient condition for the estimated effect of observed input x_i to be an upper bound for its true marginal product.

The intuition behind this result is straightforward. If consumption were to increase following the change dx_i, then aggregate spending on the unobserved inputs Z would decline (since I^* is held constant): $\sum_{k=1}^{K} p_{z_k} dz_k < 0$. The total effect on health H from an increase in x_i would be the sum of main effect of x_i, $F_{x_i} dx_i$, and the combined responses due to the change in unobserved inputs Z: $dH = F_{x_i} dx_i + \sum_{k=1}^{K} F_{z_k} dz_k$. From the cost minimization problem we know that at the optimal choice of unobserved health inputs Z their marginal product are proportional to their prices: $p_{z_k} = \lambda_H F_{z_k}, \forall k = 1, ..., K$ (with the Lagrange multiplier λ_H being the proportionality factor). Hence lower overall spending on Z would imply that the contribution of changes in Z to the total effect on health H would be negative: $\sum_{k=1}^{K} F_{z_k} dz_k < 0$.

Note that this derivation does not require any additional assumptions on the substitutability between X and Z, nor does it require consumption C or health H to be normal goods. Health inputs X and Z could be either beneficial or harmful. We do require that unobserved health inputs Z do not have direct effects on utility, because this condition is needed to reduce a multidimensional problem in C and Z to a choice of C and H via the cost function $K(X, H)$.

The derivation above implicitly assumes that consumption good C is one-dimensional. If there were multiple consumption goods $C = (c_1, ..., c_M)$, one could investigate empirically the direction of the bias under the following additional assumption. Suppose consumption goods C are separable from the inputs X in the utility function: $U(Q(c_1...c_M), X, H)$. If the conditional demand for any consumption good c_m increased(decreased) in response to an increase dx_i, then the estimable effect of X_i in the hybrid health production function would be a lower (upper) bound for its true marginal effect.

To sum up, to interpret the estimates of the hybrid production function as relevant economic quantities a researcher must bring additional information about the underlying economic problem. Theorem 4 in the previous section presents a set of conditions when one understands well the technological relations (in particular substitution patterns) between observed and unobserved health inputs. This information, for example, might come from medical literature. In other situations the researcher might be able to justify a separability assumption like the one just described. In that case one could rely on Theorem 5 and empirically determine whether she/he has estimated a upper or lower bound on the true marginal effect of the observed health input by looking at observed changes in endogenous variables (e.g. some consumption goods). Our main point is more general than the specific assumptions we mention in this and previous sections. A failure to impose and test these conditions, or to establish alternative identifying assumptions of a similar nature, would yield estimates which cannot be interpreted as the *ceteris paribus* effects of household production inputs.

6 Instruments and Experimental Effects.

Suppose, for example, that one specifies the empirical health production function as only a function of a single input X. This misspecified model incorrectly excludes all terms involving the unobserved inputs Z; these are subsumed in the "error term." Since the observed input X is chosen jointly with the unobserved inputs, the error term in this specification usually will be correlated with the observed input[15]. To control for endogeneity bias in the estimation of the misspecified model, one might consider using variations in variables like the price of the observed input or an experimental assignment as an instrumental variable. However, as noted by Todd and Wolpin (2003), any determinant of the observed input typically will also influence the demand for the unob-

[15]Throughout this discussion, we assume that the unobserved taste and productivity parameters, τ and ρ, are independent of all prices and incomes.

served input[16], violating one of the requirements for the variable to be a valid instrument[17]. The marginal effect of X will not be identified.

Adding variables like household income and prices of the unobserved inputs typically will also fail to yield interpretable estimates of the effects of the observed input on health. This happens for two reasons. First, the conditional demand function for the unobserved input as described in equation (10) depends on income after removing expenditures on the observed input, not total household income. This specification issue, however, may be not too severe provided expenditures on the observed input are small.

Second, the correct specification of the hybrid production function also includes the prices of all consumption goods. Unless consumption good prices are independent of the health input prices and incomes, or perfectly explained by them, the empirical model will be misspecified. If one uses an instrumental variables approach to control for the endogeneity of the observed inputs for the estimation, with consumption good and input prices included only as instruments, then a rejection of the overidentification restrictions would be indicative of an incorrectly specified empirical model in this instance.

An experimental assignment of the observed input to a random sample of households also will not provide an asymptotically unbiased estimator of the marginal effect of X, but it can provide an estimator of the bound of the

[16] In some special cases the demand for Z might not depend on the price of the observed input. Provided all the determinants of the unobserved input, including its price and household income are independent of the price of X, estimation in this instance will be quite similar to the case of experimental assignment of the input X discussed below. To simplify this discussion, we also assume the conditions for Theorem 1 to hold are satisfied.

[17] In a similar way, government policies and regulations which prescribe a certain level of consumption of observed input X, e.g. a ban on smoking, would not be valid instruments since a change in the amount of observed input X consumed (a reduction to zero in the case of a ban) would, in general, also result in changes in the choices of the unobserved inputs Z. For example, in the case of smoking ban, people might switch to other stimulants which substitute for smoking (e.g. chewing tobacco or illegal substances); this would also contaminate the estimation of the actual marginal effect of smoking on health.

impact of X described in equation (8). Suppose households perfectly comply with their experimental assignment of X, where X is provided at zero cost and the X assignments are made independent of all other factors affecting household decisions and outcomes (in particular p_z, p_c, I, and ρ and τ). Given X, households will optimally choose the unobserved inputs Z and consumption goods C to maximize utility, and equation (10) evaluated at $p_x = 0$ will be the correctly specified hybrid production function.

Assume first that there is no individual level heterogeneity in preferences and health production function (ρ and τ are the same for everybody). In this case nonparametric regression of H on X controlling for observables (namely p_z, p_c, I[18]) would be able to identify exactly the same bound as presented in equation (8). In this simplest case the choice of the input Z is exactly determined by prices p_z, p_c, income I, and level of observed health input X. But since different (assigned) values of X imply different values of Z being chosen even for the same prices and income, the estimated effect of X on H would be contaminated by associated changes in Z. This yields the same bias described in Theorem 1 above[19].

In the more realistic case when there is individual level heterogeneity (ρ, τ), such a nonparametric regression would identify only some "average" of the bounds described in (8), where the average is taken with respect to the joint distribution of ρ and τ. In the case when estimated nonparametric regression does not condition on prices and income, the average will be taken over the whole joint distribution of p_z, p_c, I, and ρ and τ. The bound defined in Theorems 1 and 2 is identical to the "policy effect" in Heckman(1992) and

[18]In this case $I^* = I$ since $p_X = 0$.

[19]This just highlights the fact that the choice of Z depends on the experimentally assigned value of X, $Z = Z(X, I, p_z, p_c)$.

discussed by Todd and Wolpin (2003).

In the absence of a random experimental assignment for X, the regression model needs to include the prices of all consumption goods (not just omitted health inputs) and the adjusted income, I^*. As a result, the estimated effect of X will be a function of these prices and I^*. As discussed above it will be identical to the estimate from the experimental assignment which conditions on these variables. Experimental assignment studies usually do not condition on these prices and incomes and thus are likely to provide estimates of limited use outside of the population studied under the experiment. Our findings suggest that experimental assignment estimates could be more widely applicable if one were to condition on more prices and (adjusted) incomes.

In our simulation exercises below, we present average effects for X at each value of the observed input X, where we average with respect to the distribution of all prices and incomes that could have given rise to that particular observed value of X. In general, these conditional distributions will vary with the level of X, unless the input X is assigned so that it is independent of all prices and incomes and all unobserved tastes and productivities.

7 Simulations Illustrating the Bounds

Though no finite set of simulations can provide complete information about the consequences on the estimation of marginal products when one does not observe all of the relevant inputs to the production function, these simulations do demonstrate three key implications of the lack of complete information. First, we show that the bias described in Theorem 1 could be considerable. Second, we illustrate that including non-health related information into the estimation of the hybrid production function has the potential to provide a more informative bound (Theorem 2). Third, we demonstrate that estimation approaches that either ignore the missing inputs or proxy for them using household income have the potential to provide severely biased estimates of the impacts of the observed inputs.

We present three simulations to illustrate how the theoretical bounds on the marginal products might work in practice. The first simulation allows for observable heterogeneity across households arising from variations in observable prices and incomes to affect input demands as well as unobserved to the researcher heterogeneity affecting the health production function. The second simulation removes all unobserved heterogeneity (unobserved by the researcher); we use it to illustrate that the biases arise even in an idealized world. The third simulation mimics a random experimental assignment of the observed input, where the household can alter its choice of the unobserved input and consumption goods in response to the experimental assignment. In each simulation we generate one million observations on prices, incomes and heterogeneity and solve for each household's optimal commodity and health input demands and the resulting health outcomes. We use various approaches for estimating the "marginal effect" of an observed health input[20].

Suppose the utility function depends on two consumption goods, c_1 and c_2, and health H. It has the following Stone-Geary form:

$$U_i = (c_{1i} - 3)^{.1}(c_{2i} - 4)^{.5}(H_i - 5)^{.4} \tag{17}$$

Health is produced by the CES production function

$$H_i = (\rho_i x_i^{0.75} + z_i^{0.75})^{0.8} \tag{18}$$

[20]These simulations cannot be representative, as alternative choices for the data generating processes and parameters described below will yield different "true" impacts and "estimable" impacts. They only demonstrate the potential issues and solutions we discuss.

where x and z are the two health inputs and ρ_i is the heterogeneity known to the household when it makes its input choices but unmeasured by the researcher. The budget constraint is given by

$$p_{xi}x_i + p_{c1i}c_{1i} + p_{c2i}c_{2i} + p_{zi}z_i = I_i. \tag{19}$$

All prices and incomes follow uniform distributions generated from a five-variate normal copula[21] and are independent of the heterogeneity in the production function, ρ. These functional form assumptions for the utility and production functions correspond to those specified in Theorems 1 and 2.

When estimating the production functions and hybrid functions, we do not impose the known functional forms imposed by the above specifications of the utility and production functions. In all of the "estimations" with these simulated data, we use a third degree fully-interacted polynomial in the log-explanatory variables as the approximate functional form in the (hybrid) production function. Any interaction term containing a choice variable we treat as endogenous, using a fully-interacted fourth degree polynomial in the logarithms of the exogenously determined prices and incomes as instruments[22].

The left hand panel in Figure 1 presents the calculated partial derivative of health with respect to a unit change in the observed health input for the

[21] The normal components are correlated at 0.4, except for the component used to generate p_z which has a -0.4 correlation with each of the other four components. Each of the four prices follow a U(1,2) while incomes follow a U(50,150). The uniform random variables have correlations of approximately +/-0.38. Unobserved heterogeneity follows a independent uniform distribution on (1,2) when it is present; otherwise it is fixed at 1.5.

[22] In the simulations with no unobserved heterogeneity, for the true production function and the correctly specified hybrid models, the R^2 values are always 0.9967 or larger, revealing that that the third order polynomials approximate well the true functions. For the incorrectly specified functions, such as the regression of the observed health outcome only on a polynomial in the observed input or only on a polynomial in the observed input and income, the R^2 values can be as low as 0.7669. The first stage regressions almost always have reported R^2 values of 1.0000 when there is no unobserved heterogeneity, indicating that the fourth degree polynomials approximate well the true functional forms.

estimated true production function and three of the correctly specified hybrid production functions derived above. There is unobserved heterogeneity in the production function; it impacts all effects presented in Figure 1[23]. We construct the average derivative about each value of X using the distribution of the exogenous prices and incomes around each point[24].

The two overlapping lines at the top of this panel measure the "true" average analytic derivative. One comes from evaluating the analytic formula and the other from estimating the third order approximation when both inputs are observed. They do not coincide exceptionally well at the lowest levels of the observed input due to combined effect of the estimation by instrumental variables and the presence of the unobserved heterogeneity in the health production function that households take into account when they make their choices. These lines represent the true effects we would like to measure in an ideal world. The lowest line in the left panel measures the average effect of X from the correctly specified hybrid function from equation (10), and it measures the bound on the true marginal effect described in (8). It differs from the true derivative by the term described in (9). Even though the demand for the observed input is three times larger than the demand for the unobserved input (average budget share of 30% versus 10%), the bias is considerable over most of the range of the observed input[25]. We label this the minimum information bound.

The line just above the lower bound reveals how the bound tightens when

[23]The mean effect of the observed input X on health is 0.242(sd: 0.047); the mean effect of the unobserved input Z is 0.242 (sd:0.073); and the cross derivative is -0.0015(sd: 0.0006).

[24]We do this empirically; for each value of X we construct the average of the derivatives in the range(X-0.5, X+0.5) in the one million observation data set. Since the joint distributions of these characteristics vary with X, the slopes of the lines in the figures do not measure how the marginal product of X varies as X increases.

[25]The small numbers on the horizontal axis indicate the 5^{th}, 10^{th}, 25^{th}, 50^{th}, 75^{th}, 90^{th}, and 95^{th} percentiles of the observed input.

one controls for the demand for a consumption good as discussed in Theorem 2. It provides a tighter bound. Since its average budget share is low (only 13%), controlling for it provides little new information. The line closer to the true average marginal effect line demonstrates the implication of instead controlling for a good comprising a larger average budget share (47%). Controlling for the good with a large budget share substantially limits the size of the remaining budget allocated between the unobserved input and the other consumption good. In this example it provides a substantively more precise bound.

The right panel of Figure 1 explores less theoretically motivated approaches for estimating the impact of the observed input on health. The top two lines in this figure repeat the true average marginal effect line and the least informative bound line from the left panel. The long dashed line measures the average effect when the empirical production function only depends on the observed input. Over most of the range it lies further from the true values than the least informative bound.[26] The line with short dashes provides the average effect when one also includes the log income in the production function specification. This specification supports Todd and Wolpin's (2003) contention that including income to capture the missing inputs could lead to more severe biases.

The dotted line adds the prices of the two consumption goods and the missing input to the specification. It corresponds to Rosenzweig and Schultz's(1983) hybrid production function. For this simulation it does yield improvements in the estimation of the effects of the observed input over the other two misspecified models, but it almost never provides a better estimated average effect than the least informative bound that we derive. Figure A1 in the appendix

[26] While it consistently underestimates the true effect here, for other specifications of the distributions of the prices and inputs it can instead overestimate the true effect.

repeats theses graphs for the case when there is no unobserved heterogeneity in the health production function, and nearly all of the above discussion carries over to this more idealized case. The only significant exception is that the estimate of the true marginal effect is nearly indistinguishable from the true calculated effect. This happens because there is no unobserved heterogeneity so all R^2's are nearly identical to 1.00. In summary, these figures reveal that the bias due to not observing an input can be large and that the bound can be tightened considerably by incorporating information about non-health related expenditures. None of the less theoretically motivated estimators provide better estimates of "effects" than the loosest, minimum information bound model over most of the range of the observed input.

In Figure 2 we display estimates of the estimated average impacts of the observed input when the observed health input is experimentally assigned[27]. The left hand panel displays the bounds using correctly specified hybrid models with and without controls for pure consumption goods. The results are nearly identical to those presented in Figure 1. The amount of bias for the estimated marginal effect can be fairly substantial and the bias can be reduced by incorporating non-health related information. The right hand panel displays estimates using econometric models that are standard in the evaluation of experimental effects. These are exactly the same empirical models examined in the right hand panel of Figure 1. Because the observed input is assigned independently of all prices, incomes, and heterogeneities, the joint distribution of these other health determinants are independent of the assigned input level, and all of the average effects from using different control variables are identical at each particular value of X. As discussed in the previous section, the average

[27]We randomly reallocate the observed demands for the input X from the model with heterogeneity used to produce Figure 1 to observations to make these assignments.

effect for each of these estimators of the experimental effect are equivalent to the minimum information bound described in equation (8).

It is crucial to recognize that the marginal effects, bounds, and experimental effects presented in Figures 1 and 2 are averages over the joint distribution of prices and incomes about each point on the X axis. In Figure 3 we demonstrate how these effects change when one uses the experimentally assigned input X but only considers those individuals with incomes in the lowest quartile and who face prices of the unobserved input in the highest price quartile[28]. This subset of observations contains those who are least likely to choose large values for the unobserved input. The two solid lines at the top of Figure 3 are true marginal effects calculated by using information on both the observed and unobserved inputs. The thin solid line repeats the true average marginal effects from Figure 2, while the thick solid line is the true average marginal effect for those with low income and high price for the unobserved input. The true marginal effect for the restricted sample exceeds that for the entire sample due to the substitutability of inputs imposed by the form of the production function and the lower choices of the unobserved input for this group. The lowest line repeats the average minimum information lower bound from the correctly specified hybrid production model[29], which as we saw in Figure 2 equals the estimated effect obtained when one uses any of the experimental model estimators and averages over the entire distribution of prices and income.

The two dashed lines (almost completely overlapping) represent the average minimum information bound for those with low income and high prices for the input Z and the average of the experimental effect for low income and

[28]The model is estimated using all observations, exactly as was done for Figure 2.

[29]Averaged over the entire distribution of prices and incomes at each point X.

high price for Z[30]. The last two lines (dashes with dots) are the estimated bounds from the correctly specified hybrid model with controls for the choice of the second pure consumption good(c_2), with the thicker line being the average bound for individuals with low income and high prices and the thinner line averaging over the entire distribution of incomes and prices as in Figure 2. Even when focusing on this subset of low-income/high p_z individuals, conditioning on the "irrelevant" input c_2 provides a tighter bound for the true marginal effect than the bound obtained when one does not include c_2 in the hybrid production function.

8 Summary

This paper demonstrates how one can use economic theory to specify empirical models of household production functions and provide possibly interesting and useful economic interpretations for effects of inputs estimated using these correctly specified hybrid production functions. Provided observed and unobserved inputs are not strongly complementary, theoretical analysis reveals that the estimated effect of an observed productive input would typically yield a lower bound on the marginal product of the observed input. For a "bad" observed input (e.g., smoking), the estimated impact would provide a lower bound on its true, marginal detrimental effect. We also discuss bounds on the true *ceteris paribus* marginal effects that use only information on how consumption would change with an exogenous increase in an observed input.

These bounds follow from theoretical *ceteris paribus* derivations; they do not depend upon any assumptions about endogeneity of inputs or the form of

[30]For these two experimental effect to differ from the lowest bound estimate, it is necessary for the estimation with the experimental data to include income and the price of Z as regressors in the empirical model.

a statistical model. We further show that one can improve on those bounds by including seemingly irrelevant information about pure consumption goods, provided particular separability assumptions for the utility function are reasonable. This also allows one to establish the sign of the bias and determine whether the estimated effect represents an upper or lower bound on the true marginal product of the observed input. We demonstrate the potential empirical relevance of these bounds and improvements on them using data generated by a simple simulation exercise.

The least informative bound that we identify for the marginal effect is identical to the "policy effect" one would estimate in an experimental setting holding prices and incomes constant. If the policy effect is the only effect of interest, then our analysis describes how one can estimate this effect without needing to rely upon an experimental study. Given the estimated relationship, it would be straightforward to solve for a "policy effect" in many alternative environments by integrating over different joint distributions of prices and incomes. Experimental studies that ignore such characteristics in the estimation would be less informative, unless other target populations face a similar distribution of prices and incomes to that for the experimental subjects.

The least informative bound, since it is the "policy effect," would provide sufficient information in many situations. But a more complete understanding of how households produce health, education, and other outputs could help researchers uncover and develop more effective policy tools. Our derivation of the bounds on the marginal products and how one can improve those bounds is an essential step in this direction.

We take no stand on whether or not the specific conditions we derive are likely to hold in practice. Our main point is that unless one is able to come up with and justify assumptions of this or a similar nature, one cannot claim that the estimated effects of observed household production inputs provide any resemblance to their marginal effects. This problem pervades all of the existing studies of household production including health, child development, human capital, and educational investments.

To reiterate, our main criticism of the household production literature concerns the interpretation of existing estimates, not necessarily their point values. Under some conditions[31], the point estimates from correctly specified hybrid production functions may not differ appreciably from those found in the literature. However, for these estimates to have meaningful economic content, such as the bounds interpretation we derive, one needs to invoke and justify additional assumptions about the substitution patterns among observed and unobserved inputs.

Acknowledgements

For valuable suggestions and advice we thank Michael Grossman, Pierre-Andre Chiappori, Bill Dougan, Kevin Tsui, seminar participants at Instituto Technologico Autonomo de Mexico's Seminario Aleatorio, UNC-Greensboro, the Andrew Young School of Policy Studies, UNC-Chapel Hill, the Federal Reserve Bank of Atlanta and the University of Virginia, and conference participants at 2011 North American Meeting of the Econometric Society for helpful comments. NIH Grant R01-HD47213 provided support for this paper.

[31]For example, this could be the case when income effects are small, expenditure on observed inputs is low, and there is little variability in consumption goods' prices and prices of omitted inputs.

9 References

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer, 2006, "Structural Identification of Production Functions," working paper.

Heckman, James J., 1992 "Randomization and Social Policy Evaluation", in C. Manski and I. Garfinkel, eds.), Evaluating Welfare and Training Programs, Cambridge, MA: Harvard University Press, pp. 201-230.

Heckman, James J., 2000, "Causal Parameters and Policy Analysis in Economics: A Twentieth Century Retrospective," Quarterly Journal of Economics, vol. 115, pp. 45-97.

Heckman James J. and Edward Vytlacil, 2005, "Structural Equations, Treatment Effects, and Econometric Policy Evaluation" Econometrica, vol.73(3), pp.669-738.

Hoch, Irving, 1955, "Estimation of production function parameters and testing for efficiency," (abstract), Econometrica Vol. 23, No.2, pp.325-6.

Lancaster, K., 1966, "A new approach to consumer theory," Journal of Political Economy Vol. 74, No. 1, pp. 132-157.

Levinsohn, James and Amil Petrin, 2003, "Estimating Production Functions Using Inputs to Control for Unobservables," Review of Economic Studies Vol. 70, No. 2, pp. 317-342.

Liu, Haiyong, Thomas Mroz, and Linda Adair, 2009, "Parental Compensatory Behaviors and Early Child Health Outcomes in Cebu, Philippines," Journal of Development Economics, Vol. 90, No. 2, November, pp. 209-230.

Michael, Robert T. and Gary S. Becker, 1973, "On the New Theory of Consumer Behavior," The Swedish Journal of Economics, Vol.75(4): 378-396.

Mundlak, Yair, 1961, "Empirical Production Function Free of Management Bias," Journal of Farm Economics, Vol. 43, No. 1 (February), pp. 44-56.

Olley, Steven and Ariel Pakes, 1996, "The Dynamics of Productivity in the Telecommunications Equipment Industry," Econometrica vol. 64(6): 1263-98.

Peltzman, Sam, 1975. "The Effects of Automobile Safety Regulation," Journal of Political Economy, vol. 83, No. 4 (August): pp. 677-726.

Rosenzweig, Mark R. and Schultz, T. Paul, 1983. "Estimating a Household Production Function: Heterogeneity, the Demand for Health Inputs, and Their Effects on Birth Weight," Journal of Political Economy, vol. 91(5): 723-46.

Todd, Petra E. and Kenneth I. Wolpin, 2003, "On the specification and estimation of the production function for cognitive achievement," The Economic Journal, Vol. 113, February, pp. F3-F33.

Wolpin, Kenneth I, 1997, "Determinants and Consequences of the mortality and health of infants and children," in Handbook of Population and Family Economics, M. Rosenzweig and O. Start, eds., Chapter 10, pp. 483-557.

Figure 1

Average Effects* of Observed Input with Heterogeneity

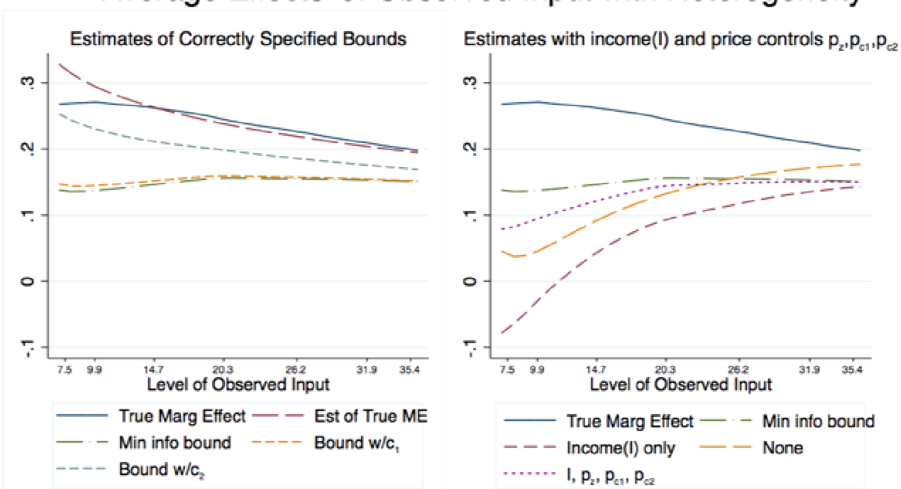

X-axis ticks at 5th, 10th, 25th, 50th, 75th, 90th, and 95th percentiles of observed input
* averaged wrt distribution of prices, incomes, and unobservables about chosen input level

Figure 2

Average Effects* of Experimentally Assigned Input with Heterogeneity

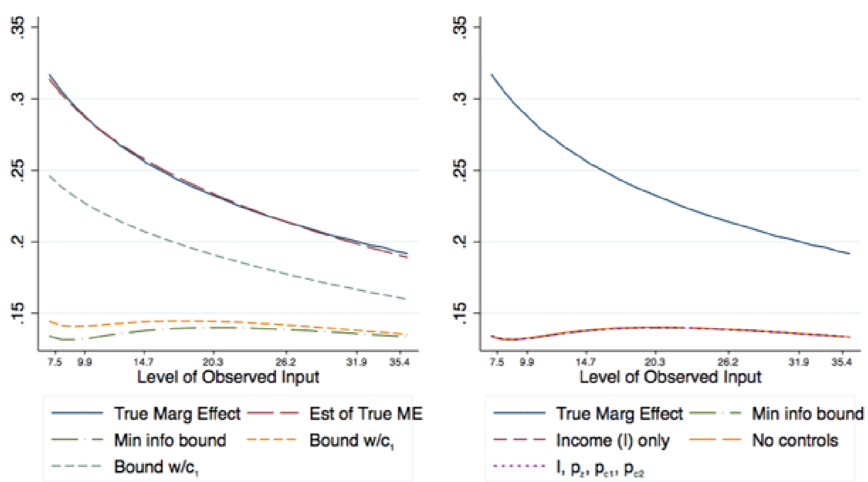

X-axis ticks at 5th, 10th, 25th, 50th, 75th, 90th, and 95th percentiles of observed input
* averaged wrt distribution of prices, incomes, and heterogeneity about assigned input

Figure 3

Average Effects* of Experimental Input; Low Income, High p_z

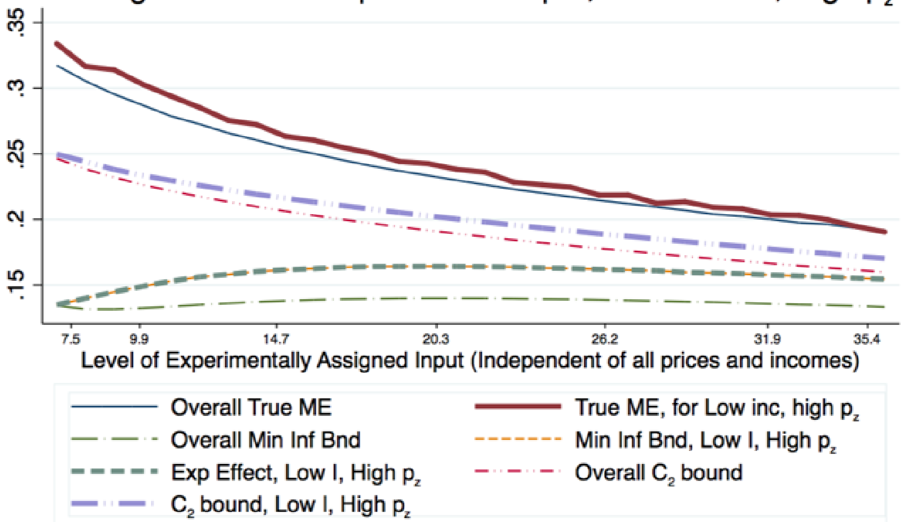

X-axis ticks at 5th, 10th, 25th, 50th, 75th, 90th, and 95th percenitiles of the experimentally assigned input. Heterogeneity in health production function.
* A Subset of the experimental data in Figure 2; averages taken wrt distribution of prices and incomes about each X-point as in Figures 1 and 2.

Figure A1

Average Effects* of Observed Input Without Heterogeneity

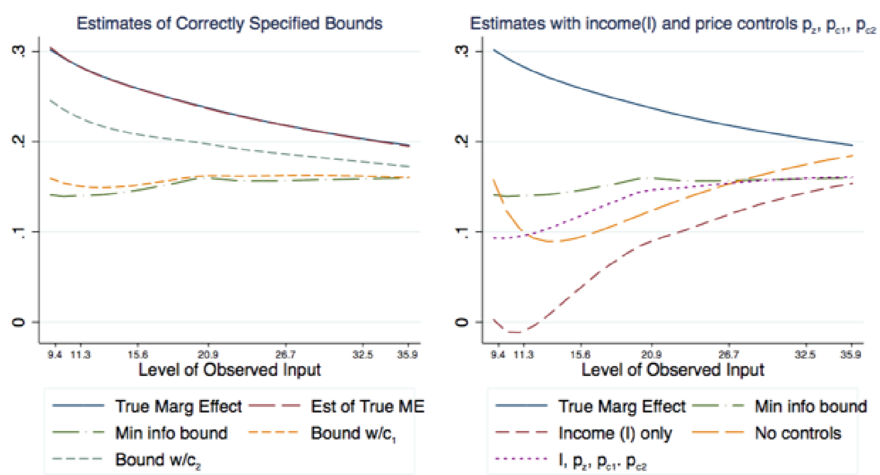

X-axis ticks at 5th, 10th, 25th, 50th, 75th, 90th, and 95th percentiles of observed input
* averaged wrt distribution of prices and incomes about chosen input level

10 For Online Appendix

10.1 Derivation of the bias

In this section of the appendix we derive an expression for the bias of the estimated marginal effect of observed inputs X in the most general case when both X and unobserved input Z have direct effects on utility. One can derive biases in the case when X and/or Z are affecting only health as special cases of this problem. We consider the general case when there are multiple observed and unobserved health inputs: i.e. $X = (x_1, ..., x_N)'$ and $Z = (z_1, ..., z_K)'$ are vectors. However, for simplicity we assume that consumption is represented by scalar aggregated consumption good C. The consumer's problem (conditional on X) in this case can be written as:

$$\begin{aligned} &\max_{C,Z} U(C, X, Z, F(X, Z)) \\ &\text{s.t. } p_C C + p_Z' Z = I^* = I - p_X' X \end{aligned} \tag{20}$$

where p_X and p_Z are vectors of prices conforming to X and Z respectively.

Expressing C from the budget constraint: $C = I^R - t_X' X - t_Z' Z$, where $t_X = \frac{p_X}{p_C}$, $t_Z = \frac{P_Z}{P_C}$, and $I^R = \frac{I}{p_C}$, we can express constrained optimization problem (20) with respect to C and Z as the following unconstrained optimization with respect to Z only:

$$\max_{Z} U(I^R - t_X' X - t_Z' Z, X, Z, F(X, Z)) \tag{21}$$

Optimality conditions for this problem are:

$$-t_Z U_C + (U_H F_Z + U_Z) = 0^{32} \tag{22}$$

[32]Here and in what follows to save space we denote column vector $\frac{\partial F}{\partial Z}$ as F_Z and row

and the following matrix is negative semidefinite:

$$\Delta \equiv \frac{\partial}{\partial Z'}\left(-t_Z U_C + (U_H F_Z + U_Z)\right) \leq 0. \tag{23}$$

Consider varying the observed input X by vector dX of infinitesimal changes: $dX = (dx_1, dx_2, ..., dx_N)'$. To assess the bias, we need to determine the sign of the change in the unobserved level of input Z, dZ. Totally differentiating first order condition (22) we obtain:

$$\frac{\partial}{\partial Z'}\left(-t_Z U_C + (F_Z U_H + U_Z)\right) dZ + \frac{\partial}{\partial X'}\left(-t_Z U_C + (F_Z U_H + U_Z)\right) dX = 0 \tag{24}$$

The first term in this sum is ΔdZ where Δ is negative semidefinite matrix from (23). The second term can be derived from utility function in (21). Thus we get:

$$\begin{aligned}\Delta dZ = [&t_Z U_{CX'} + t_Z F_{X'} U_{CH} - U_{ZX'} - U_{ZH} F_{X'} - \\ &- F_Z U_{HX'} - F_Z F_{X'} U_{HH} - U_H F_{ZX'}]\end{aligned} \tag{25}$$

Expressing price ratios vector t_Z from the first order condition (22): $t_Z = \frac{U_H}{U_C} F_Z + \frac{U_Z}{U_C}$ and substituting into (25) we get:

$$\begin{aligned}\Delta dZ = \Big[&\tfrac{1}{U_C} U_Z U_{CX'} + \tfrac{U_H}{U_C} F_Z U_{CX'} + \tfrac{U_{CH}}{U_C} U_Z F_{X'} + \tfrac{U_{CH} U_H}{U_C} F_Z F_{X'} - \\ &- U_{ZX'} - U_{ZH} F_{X'} - F_Z U_{HX'} - F_Z F_{X'} U_{HH} - U_H F_{ZX'}\Big] dX\end{aligned} \tag{26}$$

For the later analysis it is instructive to combine the terms in (26) as:

vector $\left(\frac{\partial F}{\partial Z}\right)'$ as $F_{Z'}$, matrix $\frac{\partial^2 F}{\partial Z \partial X'}$ as $F_{ZX'}$, etc

$$\begin{aligned} B_1 \equiv \left(\tfrac{U_{CH} U_H}{U_C} - U_{HH} \right) F_Z F_{X'} - U_H F_{ZX'} = \\ = U_H \tfrac{\partial}{\partial H} \left[\log \left(\tfrac{U_C}{U_H} \right) \right] F_Z F_{X'} - U_H F_{ZX'} \end{aligned} \tag{27}$$

$$B_2 \equiv \tfrac{U_H}{U_C} F_Z U_{CX'} - F_Z U_{HX'} = U_H F_Z \tfrac{\partial}{\partial X'} \left[\log \left(\tfrac{U_C}{U_H} \right) \right] \tag{28}$$

$$B_3 \equiv \tfrac{1}{U_C} U_Z U_{CX'} - U_{ZX'} = -U_C \tfrac{\partial}{\partial X'} \left(\tfrac{U_Z}{U_C} \right) \tag{29}$$

$$B_4 \equiv \tfrac{U_{CH}}{U_C} U_Z F_{X'} - U_{ZH} F_{X'} = -U_C \tfrac{\partial}{\partial H} \left(\tfrac{U_Z}{U_C} \right) F_{X'} \tag{30}$$

When one runs a regression of health outcome H on observed health inputs X, the point estimate for a particular health input X_i would measure: $\frac{\partial F}{\partial x_i} + F_{Z'} \frac{dZ}{dx_i}\Big|_{x_{-i}, I - p'_X X}$, where $\frac{dZ}{dx_i}\Big|_{x_{-i}, I - p'_X X}$ represents a vector of changes in unobserved health inputs Z associated with a change in x_i while keeping other observed health inputs x_{-i} fixed and total income $(I - p'_X X)$ spent on consumption C and unobserved health inputs Z constant.

Using our derivation above, the change dZ can be calculated from equations (26)-(30) while setting the i-th element of dX to dx_i and the remaining elements to zero. Denote such a vector of infinitesimal changes as dX_i. With this in mind, the total bias term in this case can be written as:

$$\begin{aligned} BIAS = F_{Z'} \tfrac{dZ}{dx_i}\Big|_{x_{-i}, I - p'_X X} = F_{Z'} \Delta^{-1} \left(B_1 + B_2 + B_3 + B_4 \right) \tfrac{dX_i}{dx_i} \equiv \\ \equiv BIAS_1 + BIAS_2 + BIAS_3 + BIAS_4 \end{aligned} \tag{31}$$

where[33]

$$\begin{aligned} BIAS_1 \equiv F_{Z'}\Delta^{-1}B_1\tfrac{dX_i}{dx_i} = \\ = U_H\tfrac{\partial}{\partial H}\left[\log\left(\tfrac{U_C}{U_H}\right)\right]F_{Z'}\Delta^{-1}F_ZF_{x_i} - U_HF_{Z'}\Delta^{-1}F_{Zx_i} \end{aligned} \tag{32}$$

$$BIAS_2 \equiv F_{Z'}\Delta^{-1}B_2\tfrac{dX_i}{dx_i} = U_HF_{Z'}\Delta^{-1}F_Z\tfrac{\partial}{\partial x_i}\left[\log\left(\tfrac{U_C}{U_H}\right)\right] \tag{33}$$

$$BIAS_3 \equiv F_{Z'}\Delta^{-1}B_3\tfrac{dX_i}{dx_i} = -U_CF_{Z'}\Delta^{-1}\tfrac{\partial}{\partial x_i}\left(\tfrac{U_Z}{U_C}\right) \tag{34}$$

$$BIAS_4 \equiv F_{Z'}\Delta^{-1}B_4\tfrac{dX_i}{dx_i} = -U_CF_{Z'}\Delta^{-1}\tfrac{\partial}{\partial H}\left(\tfrac{U_Z}{U_C}\right)F_{x_i} \tag{35}$$

The first bias term, $Bias_1$, results from the presence of X and Z in the production function. The second term $Bias_2$ is present when X also affects utility function directly but Z is affecting only health production. The third term $Bias_3$ appears when Z has a direct impact on utility. And the fourth term is present when both X and Z have direct impacts on utility.

In the special case when X and Z are one dimensional total bias can be written as:

$$\begin{aligned} Bias = Bias_1 + Bias_2 + Bias_3 + Bias_4 = \\ = \tfrac{U_HF_Z^2F_X}{\Delta}\left(\tfrac{\partial}{\partial H}\left(\log\tfrac{U_C}{U_H}\right) - \tfrac{F_{ZX}}{F_ZF_X}\right) + \tfrac{U_HF_Z^2}{\Delta}\tfrac{\partial}{\partial X}\left(\log\tfrac{U_C}{U_H}\right) + \\ + \tfrac{U_ZF_XF_Z}{\Delta}\tfrac{\partial}{\partial H}\left(\log\tfrac{U_C}{U_Z}\right) + \tfrac{U_ZF_Z}{\Delta}\tfrac{\partial}{\partial X}\left(\log\tfrac{U_C}{U_Z}\right) \end{aligned} \tag{36}$$

10.2 Technical Lemma

Lemma 1 *If C is a normal good then* $\frac{\partial}{\partial H}\left(\log\frac{U_C}{U_H}\right) \geq 0.$

Proof: Consider an arbitrary point (C, H). Set the ratio of prices $\frac{p_C}{p_H}$ equal to the ratio of marginal utilities $\frac{U_C}{U_H}$ at this point. Then this point will

[33]Here we use the fact that $F_{X'}dX_i = \frac{\partial F}{\partial x_i}dx_i \equiv F_{x_i}dx_i$

be a solution to the individual utility maximization problem for income level $I = p_C C + p_H H$ at these prices.

Consider the following thought experiment: increase income I by some dI and change the price of C by some dp_C in such a way that the individual's choice of C does not change but the chosen level of H changes. Taking the first differential of the demand functions for C and H yields:

$$0 = dC = \frac{\partial C}{\partial p_C} dp_C + \frac{\partial C}{\partial I} dI \tag{37}$$

$$dH = \frac{\partial H}{\partial p_C} dp_C + \frac{\partial H}{\partial I} dI \tag{38}$$

Solve for dI from equation (37) and substitute this into (38)

$$\frac{dH}{dp_C} = \frac{\partial H}{\partial p_C} - \frac{\partial H}{\partial I} \frac{\frac{\partial C}{\partial p_C}}{\frac{\partial C}{\partial I}} \tag{39}$$

or, equivalently

$$\frac{dH}{dp_C} = \frac{H}{p_C} \left[\epsilon_{HC} - \eta_H \frac{\epsilon_{CC}}{\eta_C} \right] \tag{40}$$

where the ϵ's are the (uncompensated) price elasticities of demand and η are income elasticities.

From the Cournot aggregation condition, $s_H \epsilon_{HC} + s_C \epsilon_{CC} + s_C = 0$, where s_C and s_H are the budget shares of C and H, $\epsilon_{HC} = -\frac{s_C}{s_H}(\epsilon_{CC} + 1)$. Substituting this relation into equation (40) yields:

$$\frac{dH}{dp_C} = \frac{H}{p_C} \left[-\frac{s_C}{s_H} - \epsilon_{CC} \left(\frac{\eta_H}{\eta_C} + \frac{s_C}{s_H} \right) \right] \tag{41}$$

Using the Engel aggregation condition, $s_H \eta_H + s_C \eta_C = 1$ yields:

$$\frac{dH}{dp_C} = -\frac{H}{p_C s_H}\left[s_C + \frac{\epsilon_{CC}}{\eta_C}\right] = -\frac{H\epsilon^*_{CC}}{p_C s_H \eta_C} \tag{42}$$

where $\epsilon^*_{CC} = \epsilon_{CC} + s_C\eta_C$ is compensated own price elasticity.

In the above derivation we kept C constant allowing H to vary, hence $\frac{\partial}{\partial H}\left(\log\frac{U_C}{U_H}\right) = \frac{d}{dH}\left(\log\frac{U_C}{U_H}\right)$. Since the ratio of marginal utilities equals the price ratio at the optimal choice point we obtain:

$$\frac{\partial}{\partial H}\left(\log\frac{U_C}{U_H}\right) = \frac{d(p_C/p_H)}{dH}\frac{p_H}{p_C} = \frac{dp_C}{dHp_C} = \frac{\eta_C}{-\epsilon^*_{CC}}\frac{s_H}{H}. \tag{43}$$

Since the own price compensated elasticity, ϵ^*_{CC}, is negative, the sign of the expression above is the same as sign of η_C. When C is a normal good, this term is always positive.

■

10.3 Health Inputs with Direct Utility Effects

In the main text we considered the case when the health inputs under consideration $X = (x_1, ..., x_N)'$ had no direct impact on utility. In the empirical analysis often one is concerned with the case when some of the health inputs are detrimental to health, $F_{x_i} < 0$ for some i, but individuals still consume them since they derive utility from them. In this section we allow for the observed input to have a positive direct effect on utility while having a (potentially negative) effect on health. For example, X could be smoking or binge drinking.[34]

[34] In our derivations above we also assumed that unobserved health inputs $Z = (z_1, ..., z_K)'$ have no direct utility effects either. In this section we try to relax this assumption as well.

An individual in this case would maximize the following utility function

$$U(C, X, F(X, Z; \rho); \tau) \tag{44}$$

subject to the same budget constrain as above. As before we are interested in assessing the size of the bias in the estimation of the marginal effects of the observed inputs $X = (x_1, ..., x_N)$: $\frac{\partial F}{\partial Z} \frac{dZ}{dX_i}\Big|_{I^*=I-p_X X=const, x_{-i}=const}$.

As we show in Appendix 10.1 in this case the bias for the estimated effect of a particular observed health input x_i could be written as:

$$\begin{gathered} Bias = Bias_1 + Bias_2 = \\ = U_H \frac{\partial}{\partial H}\left[\log\left(\frac{U_C}{U_H}\right)\right] F_{Z'} \Delta^{-1} F_Z F_{x_i} - U_H F_{Z'} \Delta^{-1} F_{Z x_i} + \\ + U_H F_{Z'} \Delta^{-1} F_Z \frac{\partial}{\partial x_i}\left[\log\left(\frac{U_C}{U_H}\right)\right] \end{gathered} \tag{45}$$

As discussed above, the first term $Bias_1$ typically[35] has the opposite sign to $\frac{\partial F}{\partial x_i}$. In the case when x_i is a "bad" input ($\frac{\partial F}{\partial x_i} < 0$) this term will be positive, and the estimated marginal effect would underestimate the true detrimental impact of x_i or even cause it to appear to be a "good" input.

However, compared to the baseline case we have an additional term in the total bias that relates to the relative substitutability of x_i with pure consumption goods C and health H in the utility function: $\frac{\partial}{\partial x_i}\left(\log \frac{U_C}{U_H}\right)$.[36] In general, the sign of this term has to be assessed by the researcher on a case by case basis.

In the case of x_i being smoking, for example, the term $\frac{\partial}{\partial x_i}\left(\log \frac{U_C}{U_H}\right)$ would

[35] In the case when Z is one dimensional it suffices to put some restrictions on the degree of complementarity/substitutability between x_i and Z. In the multidimensional case additional separability conditions have to be imposed to establish unambiguously the direction of the bias. See Theorems 1 and 3 for more details.

[36] As before $F_{Z'} \Delta^{-1} F_Z \leq 0$ due to second order conditions.

be negative when, as people smoke more, they value health H more (at the margin) than other consumption goods C, keeping the levels of health and those consumption goods constant[37]. In this case the last term in the bias will also be positive. The total bias will be positive for a bad input such as smoking, and estimable effects would still provide a bound on the true marginal effect. For the case of good input which has a direct effect on utility, the reverse condition would be needed for the estimable effect to be a bound.

When this assumption is violated then the total bias might still be opposite in sign to F_{x_i} if the contribution from this term does not dominate $Bias_1$. In this case the estimated upper (lower) bound for $\frac{\partial F}{\partial x_i} < 0$ ($\frac{\partial F}{\partial x_i} > 0$) would be more precise. However, in general the sign of the total bias for the estimable effect cannot be interpreted as a bound without incorporating additional information.

In the case when both the observed and unobserved inputs have dual impacts, there are two additional terms in the bias (see equations (29) and (30) in Appendix 10.1), which relate to changes in substitutability between pure consumption goods C and the unobserved inputs Z as consumption goods (i.e. ignoring Z's impact on utility through health production) when the levels of health H and observed inputs X change. In general, the signs of those terms have to be assessed on a case by case basis. However, if the utility function is separable in (X, H) and (C, Z), these additional bias terms would equal zero and the above interpretations would hold.

[37]I.e., when x_i increases we do not take into account its impact on H through the production function.

10.4 Interpretation of Other Estimated Effects

In the main text we interpreted the effects of observed health input in the correctly specified hybrid production function in equation (10). In this section we derive interpretation for estimated "effects" of other explanatory variables, which are included in regression (10).

The estimable effect of I^* measures the derivative of observed health H with respect to $I - p'_X X$,

$$\frac{\partial h}{\partial I^*} = \frac{\partial F}{\partial Z}\frac{\partial Z}{\partial I^*} \tag{46}$$

The effect of adjusted income I^* is the combination of the marginal product of missing input(s) and their (conditional on X) income effects. This effect is not guaranteed to be positive. In fact, if missing inputs negatively affect health (e.g. smoking) and are normal goods (conditionally on X) then the estimated effect of I^* might be negative.

Interpretations of the impacts of the prices of other missing inputs can be simply derived. For example, the effect of the price p_Z of unobserved health inputs Z would measure:

$$\frac{\partial h}{\partial p_Z} = \frac{\partial F}{\partial Z}\frac{\partial Z}{\partial p_Z} \tag{47}$$

The effect of the price of unobserved non-health input C would measure

$$\frac{\partial h}{\partial p_C} = \frac{\partial F}{\partial Z}\frac{\partial Z}{\partial p_C} \tag{48}$$

When there is more than one missing health input then those estimable effects would measure the sum of marginal products of all the unobserved health inputs each weighted by the price derivative of the conditional demand for it with respect to the corresponding price.

10.5 Proof of Theorem 2.

In this Appendix section we investigate whether it is better to control for the observable consumption of non-health production goods in order to minimize the bias of the estimated marginal product of observed health inputs. For simplicity consider first the case when the unobserved input Z has no direct impact on utility.

In particular, we now assume that part of consumption C is observable. Slightly abusing the notation let C be the consumption input which is not observable and W be the consumption input which is observable. X is the observable health input, and Z is the unobservable health input. Consider estimating the marginal impact of the observable health input X on H controlling for the value of the observable non-health input W. The bias, as before, can be inferred from:

$$\left.\frac{dH}{dX}\right|_{I^*=const,W=W^*} = \frac{\partial F}{\partial X} + \left.\frac{\partial F}{\partial Z}\frac{dZ}{dX}\right|_{I^*=const,W=W^*} \tag{49}$$

We would like to analyze how the term $\left.\frac{\partial F}{\partial Z}\frac{dZ}{dX}\right|_{I^*=const}$ changes depending on whether or not one controls for W (with I^* being different in those two cases).

If one does not control for W, then as X changes both dZ and dW are potentially non-zero. When we control for W then $dW = 0$. Without loss of generality we consider the bias for an arbitrary dW and zero it out as needed.

Since neither X or Z affect utility directly, the individual's problem in this case can be written as:

$$\begin{aligned} &\max_{C,Z,W} U(C, W, F(X, Z)) \\ &\text{s.t. } p_C C + p_W W + p_Z Z = I^* \equiv I - p_X X \end{aligned} \tag{50}$$

To simplify the derivation normalize $p_C = 1$. Then expressing C from the budget constraint and substituting it into the objective function we can equivalently rewrite the consumer's optimization problem as:

$$\max_{Z,W} V(Z, W, F(X, Z); I^*) \equiv \max_{Z,W} U(I^* - p_W W - p_Z Z, W, F(X, Z)) \tag{51}$$

The first order conditions can be written as usual:

$$\begin{aligned} V_W &= 0 \\ V_Z + V_H F_Z &= 0 \end{aligned} \tag{52}$$

The second order condition in this case requires that the following matrix of second derivatives is negative semidefinite:

$$\begin{pmatrix} V_{WW} & V_{WZ} + V_{WH} F_Z \\ (V_{WZ} + V_{WH} F_Z)' & V_{ZZ} + 2V_{HZ} F_Z + V_{HH} F_Z^2 + V_H F_{ZZ} \end{pmatrix} \leq 0 \tag{53}$$

Consider changing the observable health input X by some amount dX while keeping I^*, expenditure on other goods, constant. Totally differentiating first order conditions yields:

$$V_{WW} dW + (V_{WZ} + V_{WH} F_Z) dZ = -V_{WH} F_X dX \tag{54}$$

$$\begin{aligned} &(V_{ZZ} + 2V_{ZH} F_Z + V_{HH} F_Z^2 + V_H F_{ZZ}) dZ + (V_{ZW} + V_{HW} F_Z) dW = \\ &= -(V_{ZH} F_X + V_{HH} F_X F_Z + V_H F_{ZX}) dX \end{aligned} \tag{55}$$

The term in front of dZ in the previous equation is $\Delta_{22} \leq 0$ (i.e., this is the $(2, 2)$ element of the negative semidefinite matrix in (53)).

From (54) we can express dW as:

$$dW = -\frac{V_{WH}F_X dX + (V_{WZ} + V_{WH}F_Z)dZ}{V_{WW}}, \tag{56}$$

and substituting this into (55) yields:

$$F_Z \frac{dZ}{dX} = F_Z \frac{\frac{(V_{ZW}+V_{HW}F_Z)}{V_{WW}} V_{WH}F_X - (V_{ZH}F_X + V_{HH}F_X F_Z + V_H F_{ZX})}{\Delta_{22} - \frac{(V_{ZW}+V_{HW}F_Z)^2}{V_{WW}}}. \tag{57}$$

When we do control for W then we have an expression for the bias which is similar to what we had before (modulus our new notation). In this case the regression of health on observable health input X would estimate:

$$\left.\frac{dH}{dX}\right|_{I^*=const,W=const} = \frac{\partial F}{\partial X} + \frac{\partial F}{\partial Z}\left.\frac{dZ}{dX}\right|_{I^*=const,W=const}, \tag{58}$$

where $\left.\frac{dZ}{dX}\right|_{I^*=const,W=const}$ could be derived using the same equations as above with $dW = 0$. In this case, we have

$$F_Z \frac{dZ}{dX} = F_Z \frac{-(V_{ZH}F_X + V_{HH}F_X F_Z + V_H F_{ZX})}{\Delta_{22}} \tag{59}$$

In order to estimate relative magnitudes of the bias one would need to compare expressions (57) and (59).

In the case when we do not control for W the denominator in (57) is smaller in absolute value than the denominator in (59):

$$\Delta_{22} < \left(\Delta_{22} - \frac{(V_{ZW} + V_{HW}F_Z)^2}{V_{WW}}\right) < 0^{38} \tag{60}$$

This effect, as it works through the denominator, tends to amplify the bias in

[38] Note that second order conditions imply that: $V_{WW}\Delta_{22} - (V_{ZW} + V_{HW}F_Z)^2 > 0$

the case when we do not control for W.

In order to the understand total effect on the bias term we need to compare the numerators as well. The term $B_1 \equiv -(V_{ZH}F_X + V_{HH}F_XF_Z + V_HF_{ZX})F_Z$ is contained in both expressions. Earlier we established that this term is likely to be opposite in sign to F_X (see Theorem 1 and Corollary 1).

In the case when we do not control for W we also have an additional term in the bias:

$$B_2 \equiv \frac{(V_{ZW} + V_{HW}F_Z)}{V_{WW}} V_{WH}F_XF_Z \tag{61}$$

This term, however, has an indeterminate sign. To further analyze this term, it is useful to return to the original function U. Using definition (51) we find:

$$\begin{aligned} V_Z &= -p_Z U_C \\ V_H &= U_H \\ V_W &= -p_W U_C + U_W \end{aligned} \tag{62}$$

The first order conditions above can then be written then as:

$$-p_W U_C + U_W = 0, -p_Z U_C + U_H F_Z = 0 \tag{63}$$

or

$$p_W = \frac{U_W}{U_C}, p_Z = \frac{U_H F_Z}{U_C} \tag{64}$$

Thus, we obtain:

$$\begin{aligned} V_{ZW} = p_Z p_W U_{CC} - p_Z U_{CW} = \tfrac{U_H U_W F_Z}{U_C} \left(\tfrac{U_{CC}}{U_C} - \tfrac{U_{CW}}{U_W} \right) = \\ = \tfrac{U_H U_W F_Z}{U_C} \tfrac{\partial}{\partial C} \left(\log \tfrac{U_C}{U_W} \right) \end{aligned} \tag{65}$$

$$V_{HW} = -p_W U_{CH} + U_{WH} = -U_W \left(\frac{U_{CH}}{U_C} - \frac{U_{WH}}{U_W} \right) = \\ = -U_W \frac{\partial}{\partial H} \left(\log \frac{U_C}{U_W} \right) \tag{66}$$

Using these one can write the bias term B_2 in equation (61) as:

$$B_2 = \frac{U_W^2 F_X F_Z^2}{V_{WW}} \frac{\partial}{\partial H} \left(\log \frac{U_C}{U_W} \right) \left[\frac{\partial}{\partial H} \left(\log \frac{U_C}{U_W} \right) - \frac{U_H}{U_C} \frac{\partial}{\partial C} \left(\log \frac{U_C}{U_W} \right) \right] \tag{67}$$

As before it is possible to show that $\frac{\partial}{\partial C} \left(\log \frac{U_W}{U_C} \right) > 0$ for a normal good W, but the sign of the other term is indeterminate as well as the sign of the whole term B_2.

However, we can determine a sign in the following special case. Assume that health does not affect the marginal rate of substitution between consumption goods W and C: $\frac{\partial}{\partial H} \left(\log \frac{U_W}{U_C} \right) = 0$ (e.g. preferences are weakly separable in health and non-health goods). Then B_2 would vanish and the total bias will be determined only by the common $B_1 > 0$ term and the denominators in (60). In this situation the bias will be larger (and the bound less precise) when one does not control for the observed part of consumption W. The result is likely to hold also when $\frac{\partial}{\partial H} \left(\log \frac{U_W}{U_C} \right)$ is sufficiently close to zero.

The case when Z also has a direct impact on utility can be analyzed similarly. One should define $V(Z, W, H; I^*) \equiv U(I^* - p_Z Z - P_W W, Z, W, H)$ and the derivation would go unchanged until the equation for the bias term B_2 in equation (61). The exact analogue of condition (67) is more involved since, when Z has a direct impact on utility, as V_{ZW} has more complex form. However, the formula for the derivative V_{HW} will be unchanged (in terms of partial derivatives of U), as will the first order condition with respect to W. Hence, V_{HW} (and hence B_2) would vanish under the same condition as before, namely, $\frac{\partial}{\partial H} \left(\log \frac{U_W}{U_C} \right) = 0$ and the total bias will again be larger in absolute value in the case when one does not control for W. Consequently, when one estimates a home production function, controlling for the chosen amount of a pure consumption good can result in a smaller bias and a tighter bound for the estimated marginal product of an observed input to the production function when not all of the chosen inputs can be observed.