Economic Theory as a Guide for the Specification and Interpretation of Empirical Health Production Functions

Sergey Mityakov, Thomas Mroz*,†

Abstract

We employ a static model of utility maximization with health production to derive precise interpretations of estimated effects of observable inputs on health outcomes when some other inputs are not

*We would like to thank Michael Grossman for valuable suggestions and advice about data. We also would like to thank Pierre-Andre Chiappori, Bill Dougan and Kevin Tsui, seminar participants at Instituto Technologico Autonomo de Mexico's Seminario Aleatorio, UNC-Greensboro, and UNC-Chapel Hill, and conference participants at 2011 North American Meeting of the Econometric Society for helpful comments. NIH Grant R01-HD47213 provided support for this paper. We thank the "Russia Longitudinal Monitoring survey, RLMS-HSE", conducted by HSE and ZAO Demoscope together with Carolina Population Center, University of North Carolina at Chapel Hill and the Institute of Sociology RAS (RLMS-HSE sites http://www.cpc.unc.edu/projects/rlms-hse, http://www.hse.ru/org/hse/rlms) for making these data available. All remaining errors are our own.

†Contact: John E. Walker Department of Economics, Clemson University, 222 Sirrine Hall, Clemson, SC 29634. Email: smityak@clemson.edu, tmroz@clemson.edu

observed. We show that if omitted or unobserved inputs are not properly accounted for, then estimated marginal products of health inputs cannot be easily interpreted. Using a general theoretical model, we propose empirical specifications to control for the omitted inputs. The resulting "effects" one can estimate using such specifications do not correspond exactly to the marginal products of the observed inputs on health. One can, however, establish some likely bounds on the "true" marginal products of the observed inputs when one uses empirical specifications compatible with economic theory. In particular, when some key health inputs are omitted from the regression, the estimated effect of an observed, health improving input will often be negatively biased. We also discuss approaches for obtaining more informative bounds if one believes that particular separability assumptions in the utility function are reasonable. We present some empirical evidence to demonstrate our methodology using Russian Longitudinal Monitoring System of HSE (RLMS-HSE) data and data from the 2008 Behavioral Risk Factor Surveillance System.

1 Introduction.

To make informed recommendations, health policy analysts need to understand how inputs to health production functions affect measurable health outcomes. Two key issues make this a difficult task. First, individuals' choices of health inputs likely depend on unobserved to the researcher baseline health characteristics and individuals' unobserved abilities to make use of the inputs. As a consequence, consumers' choices of the levels of health inputs are likely to be statistically endogenous determinants for the estimation of health production functions. Researchers have long recognized how the failure to control for the endogeneity of the health inputs can lead to biased and inconsistent estimates of the marginal effects of health inputs on health outcomes. They have used a variety of approaches to address these issues, such as better measures of health and individual productivity, experimentally assigned health inputs, and instrumental variables, natural experiments, and regression discontinuity models.

The second issue arises because one almost never can observe all of the inputs chosen by individuals to affect their health outcomes. Suppose, for example, that the health production function depends on two inputs, but the researcher can only observe one of those inputs. Only if the two inputs are neither complements nor substitutes would it be possible to estimate the marginal effect of the observed input on some transformation of the health outcome without knowledge of the level of the unobserved input or the process used to determine the choice of the second input. In general the estimable impact an observed input on the health outcome would confound the ceteris paribus effect of that observed input with the effects of unobserved inputs. In this paper we explore the consequences of not observing all of the relevant inputs to a health production function. The estimation and interpretation issues we derive apply to any maximization problem where purchased goods might provide utility indirectly through the production of household commodities as in Lancaster (1966) or Michael and Becker (1975), including studies of the impacts of schools and parents behaviors on children's developmental outcomes.

We use a static model of utility maximization subject to a budget constraint in conjunction with a health production function to derive precise interpretations of estimated effects of observable inputs on health outcomes. The economic model provides considerable guidance for researchers about the types of variables one needs to include in a "hybrid" health 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 health production function. Often, the estimated effects will provide a bound on the magnitude of the true ceteris paribus effect. These bounds arise solely from a theoretical model describing the behavior of an optimizing economic agent. Economic theory, rather than an assumption about the form of measurement error in a statistical model, provides the basis for the theoretical bounds that we derive. We close the paper with two empirical examples of how estimates of the effects of health inputs depend crucially on the inclusion or exclusion of the required control variables suggested by the theoretical model.

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 (See, Olley and Pakes(1996), Levinsohn and Petrin (2003), and Ackerberg, Caves and Frazer (2006)). 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 with missing inputs to a more fundamental level. In their analysis, 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 to deal with the unmeasured inputs and label this the hybrid production function. 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. The estimated impact depends on all of the marginal impacts of the unmeasured inputs as well as the parameters of the household's utility function. 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), in a discussion of production functions for cognitive achievement, 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 ignores these variable that proxy for the unobserved inputs (see, also, Wolpin, 1997). They present a detailed classification of the types of approaches one might use when not all of the relevant inputs can be observed and discuss the assumptions needed with these approaches to obtain asymptotically unbiased estimates of the marginal effects of the observed inputs. They also outline several specification tests that researchers could apply to help them uncover which sets of assumptions might not be rejected by the data. A major conclusion of their study is that instrumental variables approaches will be unlikely to help 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. They conclude with the somewhat pessimistic advice, "It is therefore important to have data that contains a large set of inputs spanning both family and school domains."

Liu, Mroz, and Adair (2009) use a more formal derivation of Rosenzweig and Schultz's (1983) hybrid production function to explore possible biases in the estimation of marginal effects due to there being optimally chosen unobserved inputs. Their analysis, like the one presented in this paper, assumes that all relevant prices and incomes are observed, and they demonstrate how one can substitute conditional or rationed demand functions into the structural production function to control for the levels of the unobserved inputs. By differentiating the resulting hybrid production function with respect to the observed input, they provide an exact expression for the functional effect of observed inputs on the health production function.

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 preliminary 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.

Let the function $H = F(X, Z)$ be the health production function. The standard demand functions for the two health inputs are given by $X = X(p_X, p_Z, p_C, I)$ and $Z = Z(p_X, p_Z, p_C, I)$ where the $p.$'s are the prices of the three purchased goods and I is exogenously determined income. Throughout this discussion we assume that one could estimate nonparametrically the two demand functions and the health production function $F(X, Z)$ if 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 estimate if 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 over C and H and the form of the health production function. This approach, however, will in general result in an unidentified model. This might not be an issue if one actually imposes the exact functional form of the health production function $F(X, Z)$ and has precise information about the functional forms for the demand function $Z(p_X, p_Z, p_C, I)$, but in general all estimated effects will be arbitrarily determined.

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 functional form of the demand function is unknown, this becomes $H =$

$G(X, p_X, p_Z, p_C, I)$ where G is the hybrid production function derived using standard economic concepts.

Since the input X depends on exactly the same set of variables determining Z (i.e., each input demand is a function of p_X, p_Z, p_C, and I), there is an exact functional relationship among the five arguments in the function $G(\cdot)$. This implies that 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 through the function G. This is a nonparametric expression of the identification problem, and it is similar to perfect multicollinearity in a linear regression model[1]. Like in the linear regression 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

[1]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.

one might impose in order to obtain the true impact of the input X on the health outcome.

This non-identification problem is different than 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[2]. Here, all inputs, both observed and unobserved, are choice variables in the individual optimization problem.

Rosenzweig and Schultz's (1983) presentation 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

[2]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.

could more generally be derived when all households face the same price for the input 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 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 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 condition on the value of I^* instead of its components in order for this particular effect of X to be identified. Liu et al's (2009) failure to do that in their empirical model likely limits one's ability to interpret their hybrid production function estimates, though many of their other estimated effects do retain a straightforward interpretation. 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 using standard price theory tools.

4 Basic Model

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 individual-specific taste parameters.

The household health production is given by function F with standard properties

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

where ρ represents productivity parameters that could vary from individual to individual[3]. The household budget constraint is:

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

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

[3]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 the empirical analysis the presence of these unobserved preference and productivity parameters means that all observed household inputs must be treated as endogenous variables.

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 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 a set of controls Y would measure:

$$\left.\frac{dH}{dX}\right|_Y = \frac{\partial F}{\partial X} + \frac{\partial F}{\partial Z}\left.\frac{dZ}{dX}\right|_Y \tag{4}$$

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

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

$$\begin{aligned} &\max_{C,Z} U(C, F(X, Z)) \\ &\text{s.t. } \; p_C C + p_Z Z = I^* \equiv I - p_X X \end{aligned} \tag{5}$$

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

$$Z = q_Z(p_C, p_Z, I - p_X X, X) \tag{6}$$

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 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:

$$\left.\frac{dH}{dX}\right|_{I^*=I-p_X X=const} = \frac{\partial F}{\partial X} + \frac{\partial F}{\partial Z}\left.\frac{dZ}{dX}\right|_{I^*=I-p_X X=const} \qquad (7)$$

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. $\frac{\partial F}{\partial Z}\left.\frac{dZ}{dX}\right|_{I^*=I-p_X X}$[4].

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

[4]Thus the set of controls Y is this case is: (p_Z, p_C, I^*)

changes when we change the observed input X holding the combined expenditure on Z and C fixed, $\frac{dZ}{dX}\Big|_{I^*=I-p_X X=const}$. 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. The following Theorem provides an answer:

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). Then the regression of observed health H on the 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 X in health production. The estimable effect of the productive input might even be negative.*

The appendix contains a complete derivation of results for a wide set of cases, and 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:

$$Bias = Bias_1 = F_Z \frac{dZ}{dX} = \frac{U_H F_Z^2 F_X}{\Delta} \left[\frac{\partial}{\partial H} \left(\log \frac{U_C}{U_H} \right) - \frac{F_{ZX}}{F_Z F_X} \right] \tag{8}$$

The partial derivative with respect to H of the term in parentheses will be positive whenever C is a normal good[5], 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 (7), this implies that the identified effect of X on Z will underestimate the true marginal impact of a beneficial input X.

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 maybe dangerous working

[5] 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 9.2.

conditions for which person is compensated, so that $p_X < 0$)[6]. The sign 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 (underestimate the adverse impact) the marginal product of a harmful health input X. 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

[6]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 next section.

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 would suffice 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 is $F_{ZX} > 0$, i.e., the increase in the amount of beneficial input (e.g. jogging) would *decrease* the marginal damage of the harmful input (e.g. smoking).

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

$$\begin{aligned} H &= F(X, Z(p_C, p_Z, I - p_X X, X, \rho, \tau), \rho) \\ &\equiv h(X, p_C, p_Z, I - p_X X, \rho, \tau) \end{aligned} \tag{9}$$

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[7]. Theorem 1 and Corollary 1 describe conditions when 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. Standard controls for endogenous explanatory variables, like instrumental variables estimation, cannot eliminate theoretical biases of this type.

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

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

observed inputs. We argue, however, that a properly adjusted income measure should always be included in the regression for the estimated effects to have a meaningful economic interpretation. 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 Health Inputs with Direct Utility Effects

So far we assumed that the observed input under consideration X has no direct impact on utility. In the empirical analysis often one is concerned with the inputs which are detrimental to health, $F_X < 0$, 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.

An individual in this case would maximize the following utility function

$$U(C, X, F(X, Z; \rho); \tau) \tag{10}$$

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 effect of the observed input X: $\frac{\partial F}{\partial Z}\frac{dZ}{dX}\Big|_{I^*=I-p_X X=const}$.

As we show in Appendix 9.1 in this case the bias could be written as:

$$\begin{aligned} Bias = F_Z \tfrac{dZ}{dX} = Bias_1 + Bias_2 = \\ = \tfrac{U_H F_Z^2 F_X}{\Delta}\left[\tfrac{\partial}{\partial H}\left(\log \tfrac{U_C}{U_H}\right) - \tfrac{F_{ZX}}{F_Z F_X}\right] + \tfrac{U_H F_Z^2}{\Delta}\tfrac{\partial}{\partial X}\left(\log \tfrac{U_C}{U_H}\right) \end{aligned} \quad (11)$$

As discussed above, the first term $Bias_1$ typically has the opposite sign to F_X. In the case when X is a "bad" input ($F_X < 0$) this term will be positive and, the estimated marginal effect would underestimate the true detrimental impact of X 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 with pure consumption goods C and health H in the utility function: $\frac{\partial}{\partial X}\left(\log \frac{U_C}{U_H}\right)$. In general, the sign of this term has to be assessed by the researcher on a case by case basis.

In the case of smoking, for example, the term $\frac{\partial}{\partial X}\left(\log \frac{U_C}{U_H}\right)$ would 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[8]. 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. For example, as people exercise more they would need to value health more relative to other consumption goods.

When this assumption is violated then the total bias might still be opposite in sign to F_X if the contribution from this term does not dominate $Bias_1$. In this case the estimated upper (lower) bound for $F_X < 0$ ($F_X > 0$) 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 Appendix 9.1). These relate to changes in substitutability between pure consumption goods C and

[8]i.e. when X increases we do not take into account its impact on H through production function.

the unobserved input Z as a consumption good (i.e. ignoring its impact on utility through health production function) when the levels of health H and observed input 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.

4.4 Interpretation of Other Estimated Effects

Other effects measured by the regression (9) also have non-standard interpretations. The estimable effect of I^* measures the derivative of observed health H with respect to $I - p_X X$,

$$\frac{\partial h}{\partial I^*} = \frac{\partial F}{\partial Z}\frac{dZ}{dI^*} = \frac{\partial F}{\partial Z}\frac{\partial Z}{\partial I^*} \tag{12}$$

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:

$$\frac{\partial h}{\partial p_Z} = \frac{\partial F}{\partial Z}\frac{\partial Z}{\partial p_Z} \tag{13}$$

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

$$\frac{\partial h}{\partial p_C} = \frac{\partial F}{\partial Z}\frac{\partial Z}{\partial p_C} \tag{14}$$

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.

To summarize, one cannot estimate the true marginal product of an observed health input X when some essential health inputs are unobserved. However, one can often impose some bounds using the approach we outlined above as described in Theorem 1. 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.

As an empirical matter, failing to adjust income properly, as in Rosenzweig and Schultz's (1983) specification of the hybrid production function, may be of less importance when spending on X is small. In this case I and I^* are likely to be nearly identical. In the empirical analysis we present below we do find that adjusting income for expenditures of observed health inputs has a smaller effect than including all of the other relevant prices. In other situations, however, this may not be the case. Both the theoretical analysis and the empirical analysis indicate that one should control for a more complete set of prices. These should include prices of pure consumption goods as well as the prices of the unobserved inputs.

5 Extensions

5.1 Consumption of Similar Goods in the Household

Consider the problem when a particular health input is consumed by several members of the household (such as caloric intake), and we want to analyze the marginal product of such an input on a health outcome. For example, we might want to examine the impact of calorie consumption on a child's health. Assume that we observe the caloric intake for the child (X_K) and for the other household members (X_A). We assume that the price of this input will be the same for all household members, and let $U(C, X_A, F(X_K, Z))$ be utility function of the household.

The household's optimization problem is

$$\begin{aligned} &\max U(C, X_A, F(X_K, Z)) \\ &\text{s.t. } p_X(X_K + X_A) + p_C C + p_Z Z = I \end{aligned} \tag{15}$$

and we are interested in the marginal product of X_K on child's health: $\frac{\partial F}{\partial X_K}$. A crucial feature of this formulation is the fact that the price of X does not vary across the two components of the household's total consumption of $X = X_K + X_A$. Following the analysis above, the conditional (on X_K and on X_A) demand function for the unobserved input Z is $Z = Z(p_Z, p_C, I - p_X(X_K + X_A), X_K, X_A)$. To apply the previous logic it is necessary to control not only for the consumption of the child but also for the consumption of other household members, and one would need to control for the potential endogeneity of both X_K and X_A. In this instance p_X will not be a sufficient instrument as this price also determines both the X_A and X_K consumptions. Additional instrumental variables that affect only X_A or only X_K in the hybrid production function would be required. This might be nearly impossible in most situations, as a valid instrument cannot be a determinant for any of the other conditional demands.

Given such stringent data requirements, a seemingly natural approach would be to treat X_A as a missing input. In this instance the conditional demand system would not condition on X_A. The levels of C, Z, and X_A are jointly determined given adjusted income $I - p_X X_K$ available to spend on them at prices p_C, p_Z, and p_X. The conditional demand function for the unobserved health input Z will be:

$$Z = \tilde{Z}(p_Z, p_C, p_X, I - p_X X_K, X_K) \tag{16}$$

Note that this conditional demand depends directly on p_X as this price determines the others' consumption of X (X_A). This implies the following hybrid production function:

$$F(X_K, \tilde{Z}(p_Z, p_C, p_X, I - p_X X_K, X_K) = h(p_Z, p_C, p_X, I - p_X X_K, X_K) \tag{17}$$

This equation, however, is non-identified as X_K is a function of the other variables included in the hybrid production function. As above, one can estimate any effect for the X_K by offsetting it through changes in the impacts of the three prices and the adjusted income. When one has access only to the data on the X consumption of the person whose health is affected by this input and not the aggregate numbers for the household, then one cannot say anything meaningful about the marginal product of the jointly consumed good X.

There may, however, be some realistic and justifiable additional restrictions one can impose on the demands for X_K and X_A. Continuing the above example, suppose we want to measure the impact of X on a child's health when X is the caloric intake of the child. If, as seems likely, children and adults differ in the types of food products they consume[9], then the price of the caloric consumption will differ between children and adults within the same household. In this case we can subdivide X into two related but distinct goods X'_K and X'_A representing consumptions of kids' and adults' food

[9]It is likely that children do not receive calories from consumption of alcohol and only children consume baby food or milk.

baskets respectively with differing prices p_K and p_A. In this case we can obtain an identified model through the following hybrid production function:

$$F(X'_K, \tilde{Z}(p_Z, p_C, p_A, I - p_K X'_K, X'_K) = h(p_Z, p_C, p_A, I - p_K X'_K, X'_K) \quad (18)$$

Since the price of X differs between children and adults, the corresponding hybrid function will not be fundamentally under-identified, and one can use the analysis from above to interpret the effects of X'_K on health.

5.2 Conditioning on Consumption of Non-health Inputs.

A natural question in light of the discussion above is whether it is better to control for the consumption of goods by other household members, X_A, or to ignore them in the estimation of the hybrid production function. This is essentially the same question as 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 9.3 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 2 summarizes this discussion:

Theorem 2. *Assume that the observed health input X has no direct effect on utility. Further assume that health H does not affect 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). Then controlling for W in the estimation of the hybrid health production function would result in a smaller downward bias for the estimated marginal product of observed health input X.*

To see more intuitively the rationale behind this result, suppose we could observe and control for all of the household's consumptions of pure consumption goods. Then, all of the remaining income in I^* would be used for expenditure on the unobserved input. A nonparametric 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[10]. When we can control for only a subset of the pure consumption goods, we are able to restrict somewhat the possible

[10]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.

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 a 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.

Theorem 2 also provides a possible solution to the problem caused by the attenuation type of bias discussed above. In particular, our analysis revealed that the bias could potentially be strong enough so that estimated effect might be opposite in sign to the true marginal product of observed input X. Thus, in general one cannot rely on the signs of estimated effects to infer whether a particular input is beneficial or adverse. However, one can infer this direction of an input's effect under the conditions of Theorem 2. Namely, if there is a consumption good W which is observed and together with C is weakly separable from health H, then one could estimate the effect of observed input X with and without controls for W. Since the bias in the latter case is larger in absolute value (and opposite in sign to F_X) one can infer the sign of F_X by simply comparing the estimated effects of X from the two models. In particular, if the estimated effect of X increases when W is added in the regression, this means that bias is negative and hence X is a beneficial input.

6 Empirical Examples

We use two different data sets to demonstrate the importance of the primary specification issues derived from the theoretical model. In particular, we show that appropriately defined explanatory variables can lead to substantively important changes in the estimated "effects" of the observable health inputs. The first data set is cross sectional, and we use it to explore the impacts of recent health inputs on a health stock measure. The second data set allows us to examine changes in health status over time as a function of recent health inputs. Both cross sectional and longitudinal analyses are common in the health economics literature. In each data set we observe that the theoretically indicated alterations of the empirical specifications substantially change estimated effects. We discuss all estimated effects as bounds on marginal effects, unlike nearly all of the existing literature in health economics.

6.1 BRFSS: Data Description

We use data on men aged 25 to 55 from Behavioral Risk Factor Surveillance System (BRFSS) survey conducted by Center of Disease Control (CDC) for the year 2008. This is a comprehensive dataset on health outcomes (physical health, mental well being, bmi, disability, and the incidence of several diseases) as well as possible health inputs such as visits to physicians, dentists, eye exams, smoking, and alcohol consumption. We focus on prime aged married white males who do not appear to suffer from debilitating illnesses that might affect their ability to work. Our primary outcome variable is a self-reported health measure. This dataset codes health status on a 1-5 scale with one being excellent and 5 being poor. To define our health outcome variable, we invert this scale so that higher values correspond to better health.

We consider three health inputs: dental services, alcohol consumption and tobacco smoking. The dental services variable measures how recently a person had his last dental cleaning (in years), with higher values coded to correspond to more recent dental cleaning. We use the number of times a person binge drinks per month (defined as the consumption of more than

5 drinks on one occasion) as the measure of alcohol consumption. Tobacco smoking is a categorical variable indicating whether a person smokes often, occasionally or not at all. We assume a person who smokes occasionally consumes 7.5 packs per month, whereas person who smokes often consumes 30 packs per month.

One limitation of this dataset is its information on income. The survey only records each respondent's income category (below $10,000; $10,000 to $15,000; $15,000 to $20,000; $20,000 to $25,000; $25,000 to $35,000; $35,000 to $50,000; $50,000 to $75,000; and $75,000 or more). To simplify the analysis, we use this categorical variable to construct our income measure by imputing for each individual the midpoint of his income category. This is a major shortcoming of this dataset.

A second deficiency of this dataset is that it contains a stock health measure, while it only provides information on relatively recent health inputs. Theoretically, the health stock should depend on all previous health inputs, but these are not available. The second dataset we examine allows us to look at the more theoretically appropriate value added to the health stock. We include this analysis of BRFSS data primarily because such types of analyses do appear in the literature.

We merge the BRFSS dataset with the region level data on prices of dentists, beer, wine, and cigarette taxes. We also include prices of other goods, such as apartment rents, and price indices for health care, groceries, housing, utilities and energy, as indicated by the theoretical model. Table 1 contains summary statistics for all variables used in our analysis of the BFRSS data.

6.2 BRFSS: Estimation Results

Following the theoretical results of the previous section, we consider a regression model containing all three observed inputs in health production as well as the total expenditure available for all other health inputs and goods and the prices of omitted health inputs and other consumption goods.

$$H_i = \alpha + \beta_1 Dental_i + \beta_2 Drinking_i + \beta_3 Smoking_i + \gamma Income_i^* + \mu p_{Z_i} \quad (19)$$

Here $Income_i^* = Income_i - Drinking_i P_{Drink_i} - Dental_i P_{Dent_i} - Smoking_i P_{Cigs_i}$ measures the income spent on all other goods except three health goods under consideration. p_{Z_i} are the location specific prices of omitted health inputs Z as well as consumption goods C. We include among those price indices for health care, groceries, housing, utilities, energy, apartment rent, and an overall price index.[11] We assume $Income_i$ and all prices are exogenous.

Since health inputs as well as residual income ($Income^*$) might be endogenous, in this analysis we estimate this regression by two-step feasible GMM with Dentist Services, Alcohol, Smoking and $Income^*$ being instrumented by local beer, wine, and dentist prices, total (categorical) income, and state level cigarette taxes. Note that the economic model implies that the prices of the three included health inputs should be valid instruments for these variables, while total income would be an instrument for $Income^*$. The prices of commodities not appearing in the hybrid production function are included as explanatory variables as required by economic theory[12]. We

[11] One cannot include all the relevant prices for lack of information. Non-varying prices of omitted inputs do not present any problems. A non-varying price of an included input, on the other hand, can present a severe problem. The constant price cannot be used as an instrument to control for endogeneity of this input, and if it is unobserved one cannot construct the appropriate $Income^*$ value leading to more severe interpretation issues.

[12] Because the economic model implies that one should subtract expenditures on the observed health inputs, $p_{X_i}X_i$, from total income when defining residual income, we also include all second order interactions of all prices, exogenous variables, and total income in the set of instrumental variables to help capture the level of the observed health input expenditures in the $Income^*$ term.

also control for an individual's age, age squared and dummies for four levels of educational attainment (less than high school, high school, some college, and college degree). Table 1 contains summary statistics for this data set.

The estimates of various forms for the hybrid health production function for the BRFSS data are in Table 2. The estimates in the first column do not include the income measure or the prices for all of the other goods. They do instrument for the possible endogeneity of the observed inputs as discussed above. This "naive" baseline would suggest a considerable positive effect of dental cleaning on subjective health evaluation, while alcohol and smoking have negative effects. In the second column we modify this baseline by including total household income in the regression. The impact of dental cleaning drops by more than half, but is still positive and somewhat statistically significant, while the effect of alcohol becomes quite negative and significant. The "effect" of smoking becomes positive but insignificant. In the third column we adjust income by expenditures on observed health inputs, and the results barely change. This is likely to happen because of the crude income measure in the BRFSS dataset. Finally, in column 4 of Table 2 we include prices of other goods in the regression model. This is the only specification in Table 2 compatible with the economic model. The estimated effect of dentist visits becomes indistinguishable from zero and is quite small and negative. The effect of alcohol consumption remains negative. The smoking coefficient is positive but insignificant. The effect of adjusted income $Income^*$ is somewhat large and statistically significant.

The estimates in Table 2 highlight the importance of using the economic model as a guide for the specification of the regression model. The estimated coefficients change dramatically after one includes all of the variables implied by the correctly specified theoretical model. The variation in estimated "effects" across the columns indicates that there are important unobserved health inputs.

The effects one can estimate due to unobserved health inputs, however, do not measure actual marginal effects. Our theoretical analysis indicates that the estimated effects are likely to provide lower bounds for the true effects. For dental cleaning, these negative lower bounds are not informative as we would expect good dental hygiene (more recent dental cleaning) to have a positive impact on health. Due to the statistical imprecision, we also learn little from this regression about the effect of smoking (packs per month) on self-reported health, unless one believes that smoking actually improves health. The lower bound on the detrimental effect of alcohol (binges per month), however, could be quite informative. The true marginal impact of alcohol on health is likely to be worse than the estimate found in column 4 of Table 2.

6.3 RLMS-HSE Analysis

A major problem with the above analysis is that the health outcome represents a lifetime of inputs, while we only look at recent ones. To examine changes in health as a function of recent health inputs, we use data from the Russian Longitudinal Monitoring Survey of Higher School of Economics (RLMS-HSE). We focus on the years from 2000 to 2005 in our examination of the model specification issues raised in the theoretical analysis. We use a less subjective measure of health than in the BRFSS by focusing on the change in a child's height over time as a function of her nutrition, smoking in the household, and a parent's attention to the child, as measured by the mother's time spent not at work. We restrict our analysis to children aged between 1 and 10, who live with mothers and are not suffering from any chronic diseases or medical conditions (which may have permanent effects on the child's health). Table 3 contains summary statistics for all of the variables used in this part of the analysis.

In Table 4 we estimate the change in a child's height from one year to the next as a function of the three health inputs of interest: mother's home time, caloric intake, and smoking in the household. We estimate all equations by two-step feasible GMM treating the three health inputs as well as (adjusted) household income as endogenous. We use prices of cigarettes, prices of different foods (milk, eggs, bread, etc), price per calorie for the food basket consumed by children (i.e. excluding alcohol), mother's value of time[13], and household income net of mother's contribution as instrumental variables. We also include demographic controls such as: dummy variables for each year of a child's age; the number of adult males and adult females; the number of children and teens in the household; mother's age and age squared; and dummies for the mother's educational attainment[14]. To avoid the most egre-

[13]To avoid a measurement error problem we compute the mother's value of time as the previous period's wage for those employed, or matched by region, year, and educational attainment to the 33rd percentile of the female wage distribution for those not employed.

[14]We use dummy variables for 8 years of schooling or less (secondary schooling or less), 10 years of schooling (high school graduate), community college but no college, some college, and a college degree.

gious reporting errors we drop observations with height increasing by more than 25cm or decreasing by more than 5cm in any 12 month period.

In specification 1 we do not control for prices of other consumption goods omitted from the regression nor do we include any income measure. It is the naive model that expresses health only as a function of observed inputs. Specification 2 adds total household income. Specification 3 adds adjusted household income net of spending on three health inputs included in the regression[15,16]. Specification 4 uses the theoretically correct income measure as well as prices of other consumption goods and the log of travel time to visit a pediatrician from each household's community as independent variables.

[15] As we argued in the theoretical model above one cannot include price of caloric intake common to children and adults since then the model will be non identified, we use price of caloric intake for basket of foods consumed by adults (i.e. including alcohol but excluding milk) to adjust household income.

[16] Adjustment with respect to mother's time at home is different from adjustment for other goods. If we define $\bar{I} = I_{w/o\ mom} + w_{mom} * 18 * 30$ the full monthly income of the household, where $I_{w/o\ mom}$ is household income without mother's contribution and w_{mom} mother's value of time. Then household income net of "spending" on mother's time at home L will be $\bar{I} - w_{mom}L = I_{w/o\ mom} + w_{mom} * (18 * 30 - L) = I$ where I is observed household income, including mother's wage income.

Moving from specifications 1, 2 or 3 to the more theoretically correct specification 4, we see important changes in the estimates of the three health production function "effects." The effect of mother's time spent at home increases by over 20 percent. The point estimate implies a third of a centimeter increase in height over the course of a year if the mother were to work part-time instead of full-time (100 versus 200 hours per month), and this larger "effect" actually represents a lower bound on the true marginal product of the time spent at home by the mother. The estimated impact of caloric intake also increases considerably. The preferred specification reveals that an additional 1000 calories per day results in at least a 0.65 centimeter growth in height over the course of a year. The true marginal product could be considerably higher, but because of missing inputs these data can only provide a lower bound for the effect. Once we use a more appropriately specified model, the estimated effect of smoking by adults in the household becomes negative. This estimated bound is not small. Additional 100 packs of cigarettes consumed per month in the household (about additional 3.3 packs per day) would lead to at least a half of a centimeter reduction in growth over the course of a year. The ad hoc specifications fail to uncover these important, large, and meaningful bounds for the magnitudes of the effects of health inputs on children's growth.

7 Summary

This paper demonstrates the power of economic theory to help a researcher specify empirical models of health production functions and interpret effects estimated using these correctly specified hybrid production functions. Provided observed and unobserved health inputs are not strongly complementary, the theoretical analysis reveals that the estimated effect of an observed productive input would actually be an estimate of a lower bound on the marginal product of the observed health input. For a "bad" observed input (e.g., smoking), the estimated impact would provide a lower bound on its true, marginal detrimental effect. In both situations, it is possible that the estimable effect of an observed health input can have the opposite sign of its true marginal impact. These bounds follow from theoretical ceteris paribus derivations; they do not depend upon any assumptions about endogeneity of inputs or the form of a statistical model.

If one has a priori knowledge about the direction of an input's true marginal effect, then an estimated effect might provide an estimate of an informative bound for the effect. But if one does not have this a priori knowledge, then the estimate will likely provide no useful information about the sign or the magnitude of the observable input's true marginal effect. If particular separability assumptions for the utility function are reasonable, however, then one can sign the true marginal impact of an input and obtain a more informative bound for its effect.

Our empirical analysis using a self-reported health measure and information on smoking, dental visits and binge drinking using BRFSS data generally support the implications of the theoretical analysis. Similarly the analysis of children's growth functions using data from the RLMS-HSE supports the use of the empirical models derived from economic theory. The sensitivity of the point estimates we obtain to the inclusion or the exclusion of the prices of the "unmeasured health inputs" and to the inclusion or exclusion of the expenditure allocated to the excluded inputs suggests that the types of theoretical issues we raise could have important substantive implications for health policy research. Do note that we have only used fairly simple linear specifications in our empirical analyses, and it is possible that better-specified nonlinear models might not exhibit such sensitivity to the exclusion or the inclusion of the variables indicated by economic theory.

Recognizing that we almost never observe all of the relevant inputs to a household production function has two key implications. First, economic theory provides explicit guidance about the types of variables one needs to incorporate in empirical analyses. Second, economic theory indicates that researchers should not interpret estimated effects as point estimates of the true marginal impacts of observed inputs. Under some plausible conditions, however, the estimated impacts can be interpreted as informative bounds on the actual marginal effects.

8 References

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

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, pp132-157.

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, No. 4 (December), pp. 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, No. 6, pp. 1263-98.

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), pages 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.

9 Appendix.

9.1 Derivation of the bias

In this section of the appendix we derive an expression for the bias of the estimated marginal effect of observed input 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. The consumer's problem (conditional on X) in this case can be written as:

$$\begin{gathered} \max_{C,Z} U(C,X,Z,F(X,Z)) \\ \text{s.t. } p_C C + p_Z Z = I^* = I - p_X X \end{gathered} \tag{20}$$

Optimality conditions for this problem are:

$$-\frac{p_Z}{p_C}U_C + (U_H F_Z + U_Z) = 0 \tag{21}$$

$$\begin{gathered} \Delta \equiv \tfrac{p_Z^2}{p_C^2}U_{CC} - 2\tfrac{p_Z}{p_C}U_{CZ} - 2\tfrac{p_Z}{p_C}U_{CH}F_Z + \\ +2U_{HZ}F_Z + U_H F_{ZZ} + U_{HH}F_Z^2 + U_{ZZ} \leq 0 \end{gathered} \tag{22}$$

Consider varying the observed input X by infinitesimal amount dX. 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 (21) we obtain:

$$
\begin{aligned}
&(-\tfrac{p_Z}{p_C}U_{CC}+U_{CH}F_Z+U_{CZ})dC+\\
&+(-\tfrac{p_Z}{p_C}U_{CZ}-\tfrac{p_Z}{p_C}U_{CH}F_Z+2U_{HZ}F_Z+U_{HH}F_Z^2+U_HF_{ZZ}+U_{ZZ})dZ+\\
&+(-\tfrac{p_Z}{p_C}U_{CX}-\tfrac{p_Z}{p_C}U_{CH}F_X+U_{HX}F_Z+U_{HH}F_ZF_X+U_HF_{ZX}+U_{ZX}+U_{ZH}F_X)dX=0
\end{aligned}
\tag{23}
$$

From the budget constraint dC can be expressed as a function of dZ:

$$dC=-\frac{p_Z}{p_C}dZ \tag{24}$$

Here we used the fact that combined expenditure on C and Z is held constant: ($I^*=I-p_XX=const$). Substituting dC from (24) into the equation above yields:

$$
\begin{aligned}
\Delta dZ=&\left(\tfrac{p_Z}{p_C}U_{CH}F_X-(U_{HH}F_ZF_X+U_HF_{ZX})\right)dX+\\
&+(\tfrac{p_Z}{p_C}U_{CX}-U_{HX}F_Z-U_{HZ}F_X-U_{ZX})dX
\end{aligned}
\tag{25}
$$

where Δ is the expression from second order condition above.

The price ratio $\frac{p_Z}{p_C}$ can be expressed from first order condition (21) as:

$$\frac{p_Z}{p_C} = \frac{U_H F_Z + U_Z}{U_C} \tag{26}$$

Substituting the expression for the price ratio in equation (25) above we obtain:

$$\begin{aligned}
\Delta dZ = \left(\tfrac{U_H}{U_C} U_{CH} F_X F_Z - (U_{HH} F_Z F_X + U_H F_{ZX})\right) dX + \\
+ (\tfrac{U_H F_Z}{U_C} U_{CX} - U_{HX} F_Z + \tfrac{U_Z}{U_C} U_{CH} F_X - U_{HZ} F_X + \tfrac{U_Z}{U_C} U_{CX} - U_{ZX}) dX = \\
= \Big[U_H F_Z F_X \left(\tfrac{U_{CH}}{U_C} - \tfrac{U_{HH}}{U_H} - \tfrac{F_{ZX}}{F_Z F_X} \right) + U_H F_Z \left(\tfrac{U_{CX}}{U_C} - \tfrac{U_{HX}}{U_H} \right) + \\
+ U_Z F_X \left(\tfrac{U_{CH}}{U_C} - \tfrac{U_{ZH}}{U_Z} \right) + U_Z \left(\tfrac{U_{CX}}{U_C} - \tfrac{U_{ZX}}{U_Z} \right) \Big] dX
\end{aligned} \tag{27}$$

Thus, the bias term $Bias = F_Z \left.\frac{dZ}{dX}\right|_{I - p_X X = const}$ can be expressed as the sum of the following four terms:

$$\begin{aligned}
Bias = Bias_1 + Bias_2 + Bias_3 + Bias_4 = \\
= \tfrac{U_H F_Z^2 F_X}{\Delta} \left(\tfrac{\partial}{\partial H} \left(\log \tfrac{U_C}{U_H} \right) - \tfrac{F_{ZX}}{F_Z F_X} \right) + \tfrac{U_H F_Z^2}{\Delta} \tfrac{\partial}{\partial X} \left(\log \tfrac{U_C}{U_H} \right) + \\
+ \tfrac{U_Z F_X F_Z}{\Delta} \tfrac{\partial}{\partial H} \left(\log \tfrac{U_C}{U_Z} \right) + \tfrac{U_Z F_Z}{\Delta} \tfrac{\partial}{\partial X} \left(\log \tfrac{U_C}{U_Z} \right)
\end{aligned} \tag{28}$$

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. Third term $Bias_3$ appears when Z has a direct impact on utility. And fourth term is present when both X and Z have direct impacts on utility.

9.2 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 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:

$$0 = dC = \frac{\partial C}{\partial p_C} dp_C + \frac{\partial C}{\partial I} dI \tag{29}$$

$$dH = \frac{\partial H}{\partial p_C} dp_C + \frac{\partial H}{\partial I} dI \tag{30}$$

Solve for dI from equation (29) and substitute this into (30)

$$\frac{dH}{dp_C} = \frac{\partial H}{\partial p_C} - \frac{\partial H}{\partial I}\frac{\frac{\partial C}{\partial p_C}}{\frac{\partial C}{\partial I}} \tag{31}$$

or, equivalently

$$\frac{dH}{dp_C} = \frac{H}{p_C}\left[\epsilon_{HC} - \eta_H \frac{\epsilon_{CC}}{\eta_C}\right] \tag{32}$$

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 (32) yields:

$$\frac{dH}{dp_C} = \frac{H}{p_C}\left[-\frac{s_C}{s_H} - \epsilon_{CC}\left(\frac{\eta_H}{\eta_C} + \frac{s_C}{s_H}\right)\right] \tag{33}$$

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

$$\frac{dH}{dp_C} = -\frac{H}{p_C s_H}\left[s_C + \frac{\epsilon_{CC}}{\eta_C}\right] = -\frac{H\epsilon^*_{CC}}{p_C s_H \eta_C} \tag{34}$$

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 we obtain:

$$\frac{\partial}{\partial H}\left(\log\frac{U_C}{U_H}\right) = \frac{d(p_C/p_H)}{dH}\frac{p_H}{p_C} = \frac{dp_C}{dHp_C} = \frac{\eta_C}{-\epsilon^*_{CC}}\frac{s_H}{H}. \tag{35}$$

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.

■

9.3 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:

$$\left.\frac{dH}{dX}\right|_{I^*=const,W=W^*} = \frac{\partial F}{\partial X} + \left.\frac{\partial F}{\partial Z}\frac{dZ}{dX}\right|_{I^*=const,W=W^*} \tag{36}$$

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:

$$\begin{aligned}&\max_{C,Z,W} U(C,W,F(X,Z))\\&\text{s.t. } p_C C + p_W W + p_Z Z = I^* \equiv I - p_X X\end{aligned} \tag{37}$$

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

$$\max_{Z,W} V(Z, W, F(X, Z); I^*) \equiv \max_{Z,W} U(I^* - p_W W - p_Z Z, W, F(X, Z)) \quad (38)$$

The first order conditions can be written as usual:

$$\begin{aligned} &V_W = 0 \\ &V_Z + V_H F_Z = 0 \end{aligned} \quad (39)$$

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

$$\begin{pmatrix} V_{WW} & V_{WZ} + V_{WH} F_Z \\ V_{WZ} + V_{WH} F_Z & V_{ZZ} + 2V_{HZ} F_Z + V_{HH} F_Z^2 + V_H F_{ZZ} \end{pmatrix} \leq 0 \quad (40)$$

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:

$$V_{WW} dW + (V_{WZ} + V_{WH} F_Z) dZ = -V_{WH} F_X dX \quad (41)$$

$$\begin{aligned}&(V_{ZZ}+2V_{ZH}F_Z+V_{HH}F_Z^2+V_HF_{ZZ})dZ+(V_{ZW}+V_{HW}F_Z)dW=\\&=-(V_{ZH}F_X+V_{HH}F_XF_Z+V_HF_{ZX})dX\end{aligned} \tag{42}$$

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 (40)).

Expressing dW from (41) and substituting it into (42):

$$dW=-\frac{V_{WH}F_XdX+(V_{WZ}+V_{WH}F_Z)dZ}{V_{WW}} \tag{43}$$

We finally obtain

$$F_Z\frac{dZ}{dX}=F_Z\frac{\frac{(V_{ZW}+V_{HW}F_Z)}{V_{WW}}V_{WH}F_X-(V_{ZH}F_X+V_{HH}F_XF_Z+V_HF_{ZX})}{\Delta_{22}-\frac{(V_{ZW}+V_{HW}F_Z)^2}{V_{WW}}} \tag{44}$$

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

$$\left.\frac{dH}{dX}\right|_{I^*=const,W=const}=\frac{\partial F}{\partial X}+\frac{\partial F}{\partial Z}\left.\frac{dZ}{dX}\right|_{I^*=const,W=const}, \tag{45}$$

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

$$F_Z \frac{dZ}{dX} = F_Z \frac{-(V_{ZH}F_X + V_{HH}F_XF_Z + V_HF_{ZX})}{\Delta_{22}} \tag{46}$$

In order to estimate relative magnitudes of the bias one would need to compare expressions (44) and (46).

In the case when we do not control for W the denominator in (44) is smaller in absolute value than the denominator in (46):

$$\Delta_{22} < \left(\Delta_{22} - \frac{(V_{ZW} + V_{HW}F_Z)^2}{V_{WW}}\right) < 0^{17} \tag{47}$$

This effect, as it works through the denominator, tends to amplify the bias in the case when we do not control for W.

In order to 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

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

in the bias:

$$B_2 \equiv \frac{(V_{ZW} + V_{HW} F_Z)}{V_{WW}} V_{WH} F_X F_Z \tag{48}$$

This term, however, has an indeterminate sign. To further analyze this term, it is useful to return to the original function U. Using definition (38) we find:

$$\begin{aligned} V_Z &= -p_Z U_C \\ V_H &= U_H \\ V_W &= -p_W U_C + U_W \end{aligned} \tag{49}$$

First order conditions above can then be written then as:

$$-p_W U_C + U_W = 0, -p_Z U_C + U_H F_Z = 0 \tag{50}$$

or

$$p_W = \frac{U_W}{U_C}, p_Z = \frac{U_H F_Z}{U_C} \tag{51}$$

Thus, we obtain:

$$\begin{aligned} & V_{ZW} = p_Z p_W U_{CC} - p_Z U_{CW} = \tfrac{U_H U_W F_Z}{U_C} \left(\tfrac{U_{CC}}{U_C} - \tfrac{U_{CW}}{U_W} \right) = \\ & = \tfrac{U_H U_W F_Z}{U_C} \tfrac{\partial}{\partial C} \left(\log \tfrac{U_C}{U_W} \right) \end{aligned} \tag{52}$$

$$V_{HW} = -p_W U_{CH} + U_{WH} = -U_W \left(\frac{U_{CH}}{U_C} - \frac{U_{WH}}{U_W} \right) = \\ = -U_W \frac{\partial}{\partial H} \left(\log \frac{U_C}{U_W} \right) \tag{53}$$

Using these one can write the bias term B_2 in equation (48) as:

$$B_2 = \frac{U_W^2 F_X F_Z^2}{V_{WW}} \frac{\partial}{\partial H} \left(\log \frac{U_C}{U_W} \right) \left[\frac{\partial}{\partial H} \left(\log \frac{U_C}{U_W} \right) - \frac{U_H}{U_C} \frac{\partial}{\partial C} \left(\log \frac{U_C}{U_W} \right) \right] \tag{54}$$

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 (47). In this situation 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 (48). The exact analogue of condition (54) 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.

10 Tables

Table 1. Summary Statistics for the BRFSS Sample.

Variable	N-obs	Mean	st. dev.	Min	Max
Subsample of males ages 25-55					
General Health	5534	-1.847	0.754	-4	-1
Dental Cleaning	5534	2.659	0.780	0	3
Alcohol	5534	0.828	2.773	0	30
Smoking	5534	0.204	0.576	0	2
Income (in $ 10,000)	5534	8.503	2.305	0.5	10
Income* (in $ 10,000)	5534	8.495	2.306	0.5	10
Age	5534	42.29	7.46	26	54
Education	5534	3.335	0.849	1	4
Price of dentist visit	5534	81.63	12.52	52.8	124.56
Cigarette Tax	5534	1.144	0.649	0.07	2.575
Beer price (6-pack)	5534	8.21	.477	6.89	10.18
Wine price	5534	6.982	1.020	4.82	10.95
Apartment Rent	5534	952	397	462	3475
Total Energy Costs	5534	189.35	44.43	113.53	328.81
Groceries	5534	102	9.196	86.3	166.5
Housing	5534	105	35.36	75.64	244.3
Utilities	5534	101.87	16.92	73.2	147.8
Price Index	5534	102.82	14.290	88.1	166.5
Health Care Price Index	5534	102.09	8.44	87.8	123.1

Table 2. Estimates for BRFSS Married Males Sample.

	(1)	(2)	(3)	(4)
Dental Cleaning	0.449	0.188	0.188	-0.028
	(0.091)	(0.106)	(0.106)	(0.121)
Alcohol	-0.001	-0.060	-0.059	-0.039
	(0.033)	(0.035)	(0.035)	(0.036)
Smoking	-0.088	0.197	0.199	0.087
	(0.150)	(0.159)	(0.159)	(0.165)
Age	-0.032	-0.039	-0.039	-0.044
	(0.016)	(0.015)	(0.015)	(0.015)
Age-squared/100	0.030	0.038	0.038	0.046
	(0.019)	(0.018)	(0.018)	(0.018)
Income		0.033		
		(0.008)		
$Income^*$			0.033	0.036
			(0.008)	(0.008)
Constant	-2.643	-2.136	-2.136	-2.306
	(0.416)	(0.415)	(0.415)	(0.439)
Observations	5,534	5,534	5,534	5,534
UnderIdentification stat:	80.50	60.80	58.60	56.12
P-value	0.0797	0.555	0.149	0.470
J statistic	84.28	73.60	73.61	47.58
P-value	0.0380	0.149	0.634	0.751

The dependent variable in all regressions is self-reported health status. The sample includes white "healthy" married males aged between 25 and 55 at the time of the survey. $Income^*$ (in $10,000) is household income net of spending on the three health inputs included in the regression. All equations are estimated by two-step feasible GMM with Dental cleaning, Alcohol consumption, Smoking and $Income^*$ (where present) treated as endogenous variables. Instruments are household imputed income, prices of dentist visits, local tax on cigarettes, prices of beer and wine, as well as price indices for health care, groceries, housing, energy and their second order interactions. Specification (4) includes price indices for health care, groceries, housing, utilities, and energy as independent variables. All regressions control for individual's age and age squared, dummies for educational attainment. Underidentification test is based on Kleibergen-Paap rk LM statistic. Robust standard errors are reported in parentheses.

Table 3. Summary Statistics for RLMS-HSE.

VARIABLES	mean	sd	min	max
Child-Household Specific Variables				
Child's change in height (cm/year)	7.842	5.070	-5	25
Mother time at home (100's hrs/month)	6.200	0.970	0.814	7.200
Caloric Intake (in 1000 per day)	1.522	0.563	0.246	4.933
Smoking in the HH (packs per month)	21.16	19.92	0	142.5
$Income^*$ (10K RUR)	0.542	0.527	-0.394	9.398
HH income w/o mother (10K RUR)	0.651	0.914	-0.429	15.08
Mother's value of time	9.618	7.550	0.244	87.25
Mother's Education	2.710	0.853	1	4
Mother's Age	30.03	5.895	17	52
# (adult) males in the HH	1.063	0.597	0	5
# (adult) females in the HH	1.344	0.624	1	5
# children in the HH	0.921	0.735	0	5
# teens in the HH	0.802	0.888	0	6
Prices				
Prices of pack of cigarettes	8.972	4.724	1.639	29.75
Price of calorie w/o milk	0.014	0.010	0.003	0.056
Price of calorie w/o Alcohol	0.014	0.010	0.003	0.052
Price of sugar (1Kg)	10.70	1.567	7.709	17.52
Price of kolbasa (1Kg)	40.60	11.42	21.07	81.97
Price of bread (1Kg)	8.071	2.502	2.727	22.31
Price of eggs (per 10)	11.79	2.683	4.942	24.79
Price of milk (per 1L)	6.352	2.163	2.120	17.67
Price of potatoes (per 1Kg)	3.732	1.207	0.988	7.438
Price of vodka (per 1L)	60.78	14.15	19.83	106.7
Price of utilities (per $1m^2$)	10.03	7.255	0.103	132.8
Pediatrician Availability	0.829	0.377	0	1
Log time to travel to pediatrician	0.591	1.319	0	5.011

Number of child×year observations is 3000. All nominal values are adjusted for inflation; the base year is 2000. The mother's value of time is computed as the mother's average hourly wage from the previous year if employed; if she did not work in the previous year we use the 33rd percentile of the region, year, and educational class specific wage distribution as the measure of her value of time. $Income^*$ is computed as full household income net of spending on three health inputs under consideration: mother's time at home, child's caloric intake and smoking by other household members.

Table 4. RLMS-HSE: Effect of Mother's Time at Home, HH Smoking and Child's Caloric Intake on Changes in Child's Height.

	(1)	(2)	(3)	(4)
Mother's time at home	0.287	0.301	0.297	0.349
(hrs per month /100)	(0.197)	(0.197)	(0.197)	(0.208)
Caloric Intake	0.422	0.581	0.580	0.647
(per day in 1000)	(0.483)	(0.504)	(0.503)	(0.533)
Smoking	0.016	0.019	0.018	-0.005
(packs per month)	(0.011)	(0.011)	(0.011)	(0.011)
Household Income		-0.242		
(in RUR10,000)		(0.220)		
$Income^*$			-0.255	-0.540
(in RUR10,000)			(0.234)	(0.394)
Observations	3,000	3,000	3,000	3,000
Underidentification stat:	104.7	90.98	91.33	78.72
P-value	0.155	0.670	0.441	0.577
J statistic	83.83	82.60	82.58	80.86
P-value	0.663	0.451	0.671	0.671

The dependent variable in all regressions is the annual change in a child's height. The sample includes all "healthy" children aged between 1 and 10 at survey date. $Income^*$ (adjusted for inflation in 10,000 Rubles) is full household monthly income net of spending on the three health inputs included in the regression. All specifications are estimated by two-step feasible GMM with the health inputs and $Income^*$ treated as endogenous variables. Instruments include prices of different foods (sugar, alcohol, milk, bread, meat, eggs etc), price of utilities, mother's value of time, household income net of mother's contribution, and travel time to a pediatrician. All regressions include household demographic controls (number of males, females, children and teens in the household), mother's age, age squared, dummies for mother's educational attainment, dummies for child's age. Specification (4) also includes as exogenous regressors prices of other goods (alcohol, utilities, average cost of calorie excluding milk) dummy for pediatrician services availability in the town, and travel time (in minutes) to the nearest pediatrician if none is present in town. Underidentification test is based on Kleibergen-Paap rk LM statistic. Standard errors (in parentheses) are clustered at the individual child level.