Fiscal Externalities and Underinvestment in Early-Life Human Capital: Optimal Policy Instruments for a Developing Country

Nicholas Lawson*

Université du Québec à Montréal

Dean Spears†

University of Texas at Austin and Indian Statistical Institute – Delhi

February 26, 2019

PRELIMINARY – PLEASE DO NOT CITE

Abstract

Early-life health is an important form of human capital in developing countries. Health shapes the physical and cognitive development of children and influences their future productivity, and given that about one-tenth of women in the developing world are clinically underweight, maternal nutrition is a particularly important input. This paper considers policy instruments that could be used to improve maternal nutrition in the presence of three distortions that generate underinvestment. First, early-life health has fiscal externalities: healthier babies become more productive adults, who pay more in taxes; second, parents may be uninformed and not recognize the importance of maternal nutrition; and finally, low women's status may imply a preference against adequately feeding young mothers. We present a simple model of maternal nutrition investments, and solve for an optimal subsidy; however, in a developing country with an inefficient fiscal system, such a subsidy could be prohibitively costly. We therefore calibrate our model to data from the India Human Development Survey-II, to evaluate a richer set of policy instruments in a setting with high maternal malnutrition. We find that a non-linear policy instrument combining a modest subsidy with a large award for reaching a basic level of investment could generate significant welfare gains.

Keywords: early-life health, maternal nutrition, fiscal externalities, information constraints, women's status, India

*lawson.nicholas@uqam.ca.

This research was supported by the FRQSC Research Support for New Academics program (Soutien à la recherche pour la relève professorale, grant 2018-NP-204783). We thank Ndéye Séne Mbaye and Marc-André Geraldo-Demers for excellent research assistance. Any errors or omissions are the responsibility of the authors.

†dspears@utexas.edu.

1 Introduction

Early-life health is an important part of human capital in developing countries (Currie and Vogl, 2013; Deaton, 2013). Investment in early-life human capital begins even *in utero*, before birth (Alderman and Behrman, 2006). A baby's birth weight shapes its economic productivity in adulthood,[1] and poor maternal nutrition (meaning, pre-pregnancy body mass and weight gain in pregnancy) is among the most important factors that leads to poor early-life growth. About one-tenth of women in the developing world are underweight at the time that they become pregnant, and in the case of India, which we study in our calibrated numerical analysis, that percentage is considerably higher at about one-third (Coffey and Spears, 2019).[2] As a result, policies that could improve human capital by addressing any underinvestment in maternal nutrition could have very large benefits.

In the classic economic model of the household, parents choose optimal investments in children's human capital (Behrman, 1997; Strauss and Thomas, 2007). However, especially in developing countries, parents may underinvest in maternal nutrition. This could be true for several reasons. One reason is that early-life health has fiscal externalities. Fiscal externalities occur when one person's behaviour has implications for others in the economy because of effects on the government's budget constraint (Buchanan, 1966; Lawson, 2017a,b). In this case, healthier babies become stronger, smarter and more productive adults, who make more money and pay more in taxes to the government (as well as other potential spillover benefits). No prior paper has modeled the optimal policy instrument to correct this fiscal externality. Another reason is that parents, especially in developing countries, may be uninformed and not recognize the importance of maternal nutrition. Finally, in some societies, women's social status may imply a preference against adequately feeding young mothers; this is particularly well-documented in India (Das Gupta, 1995).

Everyone in the economy has an interest in the workers and taxpayers who babies grow up to be. Because families do not internalize this consequence of their decisions, parents

[1] In part because of data availability, much of this evidence is from developed countries (Behrman and Rosenzweig, 2004; Black et al., 2007; Currie and Rossin-Slater, 2015; Almond et al., 2017); however, theory and evidence point to even larger effects in developing countries (Spears, 2012; Currie and Vogl, 2013).

[2] We use the clinical definition of underweight, which means having a body mass index below 18.5. See Coffey (2015) for more details, and Drèze and Sen (2013) for more background on the Indian context.

will tend to underinvest in their children. In particular, women may not gain enough weight before or during pregnancy; in the Indian context, Coffey (2015) estimates that, in the average full-term pregnancy, women gain only 7 kg of body weight, despite being likely to begin pregnancy underweight. Therefore, policy-makers need an instrument to correct this important distortion. The contribution of this paper is to model and calibrate such an optimal policy instrument. Although other papers have quantified the importance of early-life human capital, our paper is the first to solve the core public-economics problem for early-life nutrition policy: an optimal policy design problem that recognizes the fact that governments must influence maternal nutrition through the endogenous choices of households.

The paper proceeds in stages, to explore different deviations from optimal human capital investment. We begin with an unrealistic baseline, intended to highlight the fact that policy must work through households' economic choices. In this baseline we ask: what investment in maternal nutrition would the social planner choose, if it could simply choose directly to maximize the utilitarian outcome? Of course, this is unrealistic because the household is making an endogenous decision to invest in their children, and the state cannot directly provide something to the person (the infant who will become a worker) whose human capital it wants to invest in; instead, the state has to work through the household, using a subsidy for example, and raising the funds for the subsidy will necessarily create distortions elsewhere in the economy, which creates a welfare tradeoff.

In the main models of the paper, we solve this tradeoff for the optimal subsidy to maternal nutrition. In addition to the central distortion that we study (namely, fiscal externalities), we further enrich the model to include two additional one-parameter distortions: information constraints, and women's status. In the informational constraint, we hypothesize that parents do not realize the importance of maternal nutrition. This would be consistent with a large body of literature in economic demography that describes important gaps in households' understanding of the production functions of child health (Preston and Haines, 1991; Coffey and Spears, 2017). In the women's status distortion, we allow the model to incorporate a preference against feeding young mothers, motivated by the possibility that families may have a hierarchical structure by age and gender.

These additional distortions imply an even larger underinvestment in maternal nutrition, requiring a quantitatively large early-life nutrition subsidy to correct it. This is a key result of our paper: fiscal distortions caused by poor early-life health prove to be large, even in a model of optimal, costly taxation. However, in further analysis we note that in a second-best world with distortionary taxation, this yields inefficiently high taxes, which suggests that perhaps a type of non-linear policy instrument could be considered: an incentive to reach a certain threshold of investment, such as the final weight of the mother at birth, or the weight gain during pregnancy.

In the second core result of our paper, we ask a more complicated optimal-policy question: what is the optimal maternal nutrition budget, and what is the optimal way to split it into a per-unit nutrition subsidy and an incentive which consists of a threshold and a prize amount? In a homogeneous population, this question of the division of the budget is trivial: policy should always use the prize (that is, the prize weakly dominates a subsidy). However, in a heterogeneous population – where different women have different pre-pregnancy weights and effective costs of calories – policy might need both instruments. For example, one possible complication is that richer women might find it easier to meet the threshold, so there are regressivity implications of using the nutritional prize alone. We therefore focus on India – a country in which maternal malnutrition is a particularly severe problem – and examine the optimal policy question using a version of our model which we calibrate to data from the India Human Development Survey-II. We find that the optimal combined policy incorporates a modest per-unit subsidy and a large award (on the order of $160 US) for meeting a basic level of maternal nutrition. The total spending on optimal maternal nutrition policy is considerable, at about 0.68% of India's GDP, and would generate net welfare improvements that are equivalent to 2.4% to 3.2% of current mean consumption for households with pregnant mothers.

Ultimately, these results highlight that a more effective prize could be constructed with better information, meaning better public monitoring of nutrition. In our final result, we ask by how much we could improve policy if we could observe birth weight before and at the end of pregnancy, to condition on weight gain during pregnancy. This would be a much more intensive policy, because it would require routine monitoring of the weight of non-pregnant women (rather than waiting to only weigh women when they become pregnant). Our results indicate a possibility for additional welfare improvements from such a policy, though the policies we test only raise welfare by about an additional 0.1% in India relative to the simple combined policy; it is likely that a different policy design could be found that would better use this added information. So one way to understand these results is as a quantification of the social value of investments in the state capacity required to develop such a health and nutrition surveillance system.

The rest of the paper proceeds as follows. Section 2 presents our simple model of investments in children, along with the social planner's allocation and the second-best policy in the presence of fiscal externalities; we subsequently extend the model to consider distortions from information and women's status, while section 3 presents the analysis of the optimal combined policy in a calibrated version of the model. Section 4 concludes the paper.

2 Simple Model of Investments in Children

We now present our baseline model of investments in children. We start with a simple model of a unitary household with one child, in which the parent must decide how much of their income to allocate to their own consumption, and how much to allocate to investments in the child's human capital, which from now on we will assume take the form of investments in maternal nutrition. We allow for a fiscal externality, and we will solve the model for the social planner's allocation, and for the second-best allocation when the government is limited to using a per-unit subsidy on maternal nutrition investments. For the moment, we abstract from information constraints or problems of inefficiently low women's status (that is, a non-unitary household model), but we will return to these issues in section 2.1.

Our model has two periods, and consists of a mass of unitary parents who each have one child. To simplify the algebra, we assume for now that the parents are ex-ante identical, but in the calibration of our model we will permit heterogeneity in wages and information/preferences. Thus, the representative parent has a wage of w in period 1, and chooses their labour supply L subject to a disutility term $e(L)$, receiving pre-tax income of wL. The parent faces a proportional tax rate t and receives a lump-sum transfer d, and also chooses how to allocate their after-tax income $wL(1-t)+d$ between their own consumption x and maternal nutrition investments c (where the letter c can be thought of as representing the word *child*). The government provides a proportional subsidy at rate b on investments in maternal nutrition, which means that the household budget constraint is given by:

$$wL(1-t) + d = x + (1-b)c.$$

Maternal nutrition investments influence the productivity, and thus the before-tax income, of the child when they become an adult in period 2; denoting that productivity as p, we have $p = f(c) + \alpha$, where $f(c)$ is a non-linear function that is concave in c, and α is a random variable encompassing all other inputs that determine productivity, which has mean zero and is uncorrelated with c. The government also operates an income tax in period 2, at a flat rate τ on child income p, which is used to pay for public good G.

For simplicity, we assume that the government cannot borrow or save between periods, and thus must balance its budget in each period. In period 1, the government budget constraint is:

$$wLt = bc + d + q$$

where q represents other expenditures of the government aside from the maternal nutrition subsidy and lump-sum transfer. In the second period, for simplicity, we assume that the tax rate τ is fixed and G adjusts to balance the budget; because the mean of α is zero, the second-period revenue of the government is simply $\tau f(c)$:

$$G = \tau f(c)$$

The parent gets utility from their own consumption x and disutility from labour supply L, as well as utility from their child's expected consumption $z = (1-\tau)f(c) + G$, according to a utility function $U(x, z, L)$. Assuming that it is the child's expected consumption that enters the utility function, rather than taking the expected parental utility over the distribution of α, considerably simplifies the algebra.[3] To simplify, we assume separability in utility, so that $U = u(x) + v(z) - e(L)$, where u and v are concave functions and e is a convex function. For the rest of this section of the paper, we simplify further by assuming that $u(x)$ is linear, which rules out income effects; in the simulations in section 3, we will allow for diminishing marginal utility of income.

[3]Although we do not make the parent's utility a function of risk in the second period, we assume that τ exists for redistributive purposes and that it is politically infeasible to remove it.

We take as our social welfare function the expected utility of the parent:

$$
\begin{aligned}
W &= x + v(z) - e(L) \\
&= wL(1-t) + d - (1-b)c + v(f(c)(1-\tau) + G) - e(L). \qquad (1)
\end{aligned}
$$

To begin to understand the implications of the model, consider the first-best allocation: if a social planner could choose any c directly, they would choose the c that maximizes (1) when the government budget constraints are substituted in:

$$W = wL - c - q + v(f(c)) - e(L)$$

and therefore $v'(z)f'(c) = 1$ defines the optimal c, which is unique given the concavity of v and f.

However, in the real world, a government cannot impose their preferred value of c, though they can influence that choice with public policy. We allow the government to choose the subsidy rate b on investments in maternal nutrition c, and we use the principle of sufficient statistics presented in Chetty (2009) to consider the effect of such a subsidy on welfare. In particular, the welfare effect of a marginal change in b can be written as the direct effect $\frac{\partial W}{\partial b}$ plus the effects caused by changes in t and G to balance the government budgets:

$$\frac{dW}{db} = \frac{\partial W}{\partial b} + \frac{\partial W}{\partial t}\frac{dt}{db} + \frac{\partial W}{\partial G}\frac{dG}{db}. \qquad (2)$$

The sufficient statistics approach makes use of envelope conditions on individual choices: the effects of b on c and L are relevant for welfare only through their effects on t and G, because c and L are chosen by the parent in a privately-optimal way, so $\frac{\partial W}{\partial c} = 0$ and $\frac{\partial W}{\partial L} = 0$.

The various derivatives can be solved as follows:

$$
\begin{aligned}
\frac{\partial W}{\partial b} &= c \\
\frac{\partial W}{\partial t} &= -wL \\
\frac{\partial W}{\partial G} &= v'(z) \\
\frac{dt}{db} &= \frac{1}{wL}\frac{c(1+\varepsilon_{cb})}{1 - \frac{t}{1-t}\varepsilon^t}
\end{aligned}
$$

$$\frac{dG}{db} = \frac{\tau f'(c)c}{b}\varepsilon_{cb}$$

where ε_{cb} is the elasticity of c with respect to b, and $\varepsilon^t \equiv \frac{(1-t)}{L}\frac{dL}{d(1-t)}$ is the elasticity of taxable income. Therefore, we have:

$$\frac{dW}{db} = c - \frac{c(1+\varepsilon_{cb})}{1 - \frac{t}{1-t}\varepsilon^t} + v'(z)\frac{\tau f'(c)c}{b}\varepsilon_{cb}. \tag{3}$$

Next, we can use the parent's first-order condition with respect to c: $v'(z)f'(c)(1-\tau) = (1-b)$. This allows us to replace $v'(z)f'(c)$ in the final term of (3):

$$\frac{dW}{db} = c\left[1 - \frac{1+\varepsilon_{cb}}{1 - \frac{t}{1-t}\varepsilon^t} + \frac{\tau(1-b)}{b(1-\tau)}\varepsilon_{cb}\right]. \tag{4}$$

This expression implicitly defines the optimal value of b as the one at which the welfare derivative $\frac{dW}{db} = 0$. As a baseline, we can consider what would happen if there was no income tax distortion in period 1 – that is, if L was exogenously fixed, so that $\varepsilon^t = 0$. Then $\frac{dW}{db}$ would simplify to:

$$\frac{dW}{db} = \frac{c\varepsilon_{cb}}{b(1-\tau)}(\tau - b)$$

which leads to the simple conclusion that the optimal b is equal to τ: the subsidy on maternal nutrition is set to perfectly offset the second-period marginal tax rate τ. At such an allocation, the parent's first-order condition with respect to c is:

$$v'(z)f'(c) = \frac{1-b}{1-\tau} = 1$$

and so the first-best allocation would be achieved by $b = \tau$.

Of course, the real world is characterized by tax distortions, and developing countries often have particularly inefficient fiscal systems; India is known to be particularly inefficient at raising and distributing revenues, with one study (Programme Evaluation Organisation, 2005) finding that "for one rupee worth of income transfer to the poor, the [Government of India] spends Rs.3.65". In the presence of fiscal distortions, ε^t may well be large, implying a smaller $\frac{dW}{db}$ and a lower optimal value of the maternal nutrition subsidy b. Later, this will motivate us to consider other policy instruments that may be less costly.

2.1 Extension of Model to Information and Women's Status Distortions

The model presented above highlighted the role for a subsidy to investments in maternal nutrition in the presence of a fiscal externality from future taxes: parents may underinvest in their child's productive capabilities, including insufficient attention to maternal nutrition, because they are aware that their child will not fully benefit from such investments. However, this is just one of many potential reasons for underinvestments in children, and here we will focus on two factors that could have significant negative effects on maternal nutrition: information constraints and low women's status. Simply put, parents may not realize the importance of maternal nutrition for the future productivity of their child, and even if they do realize it, women's preferences and needs may be underappreciated if households are hierarchical and paternalistic. In the rest of section 2, we will consider how these additional distortions alter the optimal subsidy problem, and we will demonstrate that their welfare consequences in the context of our model are identical conditional on a parameter that measures underinvestment in maternal nutrition.

2.1.1 Imperfect Information

First, we consider the possibility that the parent does not fully understand the child human capital production process; in particular, we assume that the parent believes that the function giving expected child productivity as a function of c is $\hat{f}(c)$, where we further assume that $\hat{f}'(c) = \beta f'(c)$ with $\beta < 1$ so that the parent underestimates the slope of the function.[4] The optimal solution to such a problem could be an information intervention aimed at improving the parent's comprehension of the importance of maternal nutrition, but we assume that such an intervention would be prohibitively costly.

Therefore, we consider the effect of imperfect information on the welfare derivative for the optimal subsidy b. The welfare function is exactly the same as (1) above, but the parent maximizes a slightly different function:

$$\hat{W} = wL(1-t) + d - (1-b)c + v(\hat{f}(c)(1-\tau) + G) - e(L)$$

[4] Presumably the parent also overestimates the level of the function at $c = 0$ – that is, $\hat{f}(0) > f(0)$ – so that their expected child human capital level will be correct at some intermediate value of c.

and as a result they will set:

$$\frac{\partial \hat{W}}{\partial c} = -(1-b) + \hat{f}'(c)(1-\tau)v'(\hat{z}) = 0 \tag{5}$$

where $\hat{z} = \hat{f}(c)(1-\tau) + G$.

Because $\frac{\partial W}{\partial c}$ is no longer equal to zero (although it remains true that $\frac{\partial W}{\partial L} = 0$), the welfare derivative is now:

$$\frac{dW}{db} = \frac{\partial W}{\partial b} + \frac{\partial W}{\partial t}\frac{dt}{db} + \frac{\partial W}{\partial G}\frac{dG}{db} + \frac{\partial W}{\partial c}\frac{dc}{db}.$$

The partial derivatives with respect to b, t and G are unchanged from the baseline model, as are the equations for $\frac{dt}{db}$ and $\frac{dG}{db}$; meanwhile, the partial derivative for c is:

$$\frac{\partial W}{\partial c} = -(1-b) + f'(c)(1-\tau)v'(z)$$

which means that our welfare derivative becomes:

$$\frac{dW}{db} = c - \frac{c(1+\varepsilon_{cb})}{1 - \frac{t}{1-t}\varepsilon^t} + v'(z)\frac{\tau f'(c)c}{b}\varepsilon_{cb} + (f'(c)(1-\tau)v'(z) - (1-b))\frac{c}{b}\varepsilon_{cb}.$$

However, to replace $v'(z)$, we only have our equation (5) for $\frac{\partial \hat{W}}{\partial c}$, which is written in terms of $v'(\hat{z})$. To keep the notation simple, we define a new parameter ρ such that $v'(z) = \rho\beta v'(\hat{z})$.[5] As long as $\hat{z}$ is not too much smaller than z, ρ must be greater than 1: if the parent underestimates the increase in z they can achieve by raising c, they will tend to underinvest in c, leading to a low value of z and a high value of $v'(z)$.[6] Indeed, if the parent correctly estimate the level of $f(c)$ (but not the slope), then $\hat{z} = z$ and $\rho = \frac{1}{\beta}$. Therefore, we can write our welfare derivative as follows:

$$\frac{dW}{db} = c\left[1 - \frac{1+\varepsilon_{cb}}{1 - \frac{t}{1-t}\varepsilon^t} + \frac{\rho\tau(1-b)}{b(1-\tau)}\varepsilon_{cb} + \frac{(\rho-1)(1-b)}{b}\varepsilon_{cb}\right]$$

which can be simplified to:

$$\frac{dW}{db} = c\left[1 - \frac{1+\varepsilon_{cb}}{1 - \frac{t}{1-t}\varepsilon^t} + \frac{(\rho+\tau-1)(1-b)}{b(1-\tau)}\varepsilon_{cb}\right]. \tag{6}$$

[5] ρ will generally be a function of c, as the value needed to satisfy the equality could vary as $z(c)$ and $\hat{z}(c)$ vary with c, but we will ignore this dependence to keep the notation simple.

[6] If $\hat{z}$ is significantly smaller than z, because the parent not only underestimates the slope of $f(c)$, but also the level, the parent might choose to make a large investment in c because they believe that $v'(\hat{z})$; in that case, ρ could be smaller than 1.

If $\rho > 1$, so that parents do indeed underinvest in maternal nutrition as a result of the information constraint, then this welfare derivative is larger (more positive) than (4) from the baseline model, conditional on the values of the elasticities. Therefore, unless the elasticities are significantly different in this version of the model, the optimal subsidy b will tend to be larger in the presence of informational constraints, which is intuitive given that the parent has a supplementary reason to underinvest in maternal nutrition: they don't recognize the importance of such investments.

Once again, in a simple world in which first-period taxation is not distortionary, so that $\varepsilon^t = 0$, this welfare derivative would give us a simple condition for the optimal subsidy: $b = \frac{\rho+\tau-1}{\rho}$, which increases with ρ. In the real world, where raising and distributing revenues is costly and distortionary, paying for a large subsidy to raise b to correct both a fiscal externality and an information constraint is likely to create significant distortions. For example, if we assume $\tau = 0.15$ and $\varepsilon^t = 0.595$ as in the calibrated model in section 3, and use $\varepsilon_{cb} = 0.129$ and $c = 0.0222$ at a baseline value of $b = 0.1$ (as is also true for the mean value of c in the calibrated model), then the optimal subsidy with $\rho = 1.5$ is significant at $b = 0.283$. In the absence of distortionary taxation, the optimal subsidy would be even larger at $b = 0.433$, but the taxes to pay for such a subsidy would be inefficiently high, and this once again raises the question of whether an alternative policy could be more efficient at improving maternal nutrition.

2.1.2 Inefficiently Low Women's Status

Before moving on to a examination of the options for policy design beyond a simple subsidy, we now consider a third possibility for underinvestment in maternal nutrition: even if the parental unit was not deterred from investing by future taxes, and even if they knew how important it was, they simply might not choose to do it. More precisely, the family's decision-maker might choose not to make such investments, if the family was not unitary as supposed up to this point, but rather consisted of multiple agents who face a conflict over the distribution of resources.

To model this scenario, we now extend our model to a non-unitary household, where our parental unit is divided into two halves: a father and a mother. We assume that the father and mother have different preferences, such that the father places a relatively higher weight on consumption and labour supply relative to utility from z:

$$U_f = (1+\delta)x + (1-\delta)v(z) - (1+\delta)e(L)$$

$$U_m = (1-\delta)x + (1+\delta)v(z) - (1-\delta)e(L)$$

where subscripts f and m are used to represent the father and the mother respectively, and $\delta > 0$ parametrizes their preferences for x and L relative to $v(z)$.[7] To simplify the problem, we assume that the father and mother supply labour jointly at a common wage, so that there is no conflict over *who* will perform the required labour; if the model was generalized to allow for separate labour supplies L_f and L_m at separate wages w_f and w_m, the results that follow would be unchanged except that the welfare derivative would have an extra term of ambiguous sign reflecting the welfare consequences of distortions to relative labour supplies.

If we define the social welfare function as the average of U_f and U_m, it is identical to the baseline welfare function in (1):

$$W = \frac{U_f + U_m}{2} = wL(1-t) + d - (1-b)c + v(f(c)(1-\tau) + G) - e(L).$$

However, we assume that the mother has inefficiently low status, and thus the decision on c and L is not taken to maximize W, but rather to maximize the following:

$$\begin{aligned}\hat{W} &= \frac{(1+\epsilon)U_f + (1-\epsilon)U_m}{2} \\ &= (1+\delta\epsilon)(wL(1-t) + d - (1-b)c) + (1-\delta\epsilon)v(f(c)(1-\tau) + G) - (1+\delta\epsilon)e(L).\end{aligned}$$

The welfare derivative takes the same general form as in the information constraint case, because c is not chosen efficiently:

$$\frac{dW}{db} = \frac{\partial W}{\partial b} + \frac{\partial W}{\partial t}\frac{dt}{db} + \frac{\partial W}{\partial G}\frac{dG}{db} + \frac{\partial W}{\partial c}\frac{dc}{db}.$$

Given that the social welfare function is unchanged, the partial derivatives are the same as in the version of the model with imperfect information, as are the derivatives of the budget

[7] If the δ terms are not placed on the disutility from labour supply, the welfare derivative will include an extra term reflecting the fact that the status distortion drives up labour supply, so that raising b is more beneficial than in the baseline model because it raises taxes which counteract this distortion. Essentially, in that case, taxes are less distortionary than usual, and we have chosen to abstract from this effect.

constraints $\frac{dt}{db}$ and $\frac{dG}{db}$. The only thing that has changed is the family's first order condition for c that is used to replace $v'(z)f'(c)$ in $\frac{\partial W}{\partial G}$ and $\frac{\partial W}{\partial c}$:

$$\frac{\partial \hat{W}}{\partial c} = -(1+\delta\epsilon)(1-b) + (1-\delta\epsilon)f'(c)(1-\tau)v'(z) = 0$$

which gives us $v'(z)f'(c) = \frac{(1+\delta\epsilon)(1-b)}{(1-\delta\epsilon)(1-\tau)}$. Defining $\sigma \equiv \frac{1+\delta\epsilon}{1-\delta\epsilon} > 1$, our welfare derivative becomes:

$$\frac{dW}{db} = c\left[1 - \frac{1+\varepsilon_{cb}}{1-\frac{t}{1-t}\varepsilon^t} + \frac{(\sigma+\tau-1)(1-b)}{b(1-\tau)}\varepsilon_{cb}\right].$$

This expression for $\frac{dW}{db}$ is identical to (6) if $\sigma = \rho$. In other words, conditional on the elasticities, the forms of imperfect information and low women's status that we consider in this model have exactly the same welfare implications: they lead to inefficiently low investments in maternal nutrition c, and thus they add a positive term to the welfare derivative measuring the added beneficial impact of a maternal nutrition subsidy b which can correct this distortion as well as the fiscal externality. Therefore, in what follows, we will collapse these two models into one, in which a single parameter will be used to represent these two additional sources of underinvestment in maternal nutrition.

3 Optimal Combined Policy

The previous section presented the welfare implications of a maternal nutrition subsidy b that can correct the fiscal externality as well as an information constraint and/or an inefficiently-low mother's status. However, as already mentioned, in a developing country with inefficient fiscal infrastructure, raising the necessary revenues for a maternal nutrition subsidy and then distributing those resources to the targetted population could be excessively costly. This raises the question of whether we could design an alternative policy that would be more efficient.

As an example, let us consider an incentive to reach a certain threshold of investment: suppose that the government gives a prize α to all mothers whose weight passes some threshold value. Our policy question now comes in two parts: what is the optimal budget, and what is the optimal way to split it into a per-unit subsidy and an incentive which consists of a threshold and a prize amount?

In the homogenous population we have been considering to this point, the answer to the optimal division of the budget into subsidy and prize is trivial: a threshold prize dominates a per-unit subsidy. The prize can always replicate the effect of the optimal per-unit subsidy by giving an award equal to bc at threshold c, where c is the optimal value under the subsidy; and given concavity of $v(z)$ and non-convexity of $u(x)$, reducing the award marginally from bc would still leave the threshold c preferable to receiving nothing, so the government can get away with paying less than bc (and raising distortionary taxes less) to achieve investment c.

It is possible to do an alternative sufficient statistics analysis in the case of a threshold and an award α; if we define $\alpha(c)$ as the level of award needed to make the representative parent choose c in the baseline model, the welfare derivative with respect to c becomes:

$$\frac{dW}{dc} = \alpha'(c) - 1 + f'(c)v'(z) - \frac{\alpha'(c)}{1 - \frac{t}{1-t}\varepsilon^t}$$

where $\alpha(c)$ is defined by:

$$\alpha'(c) = 1 - f'(c)\left[\tau v'(f(c_0)(1-\tau) + G) - v'(z)\right]$$

and where c_0 is the alternative value of c that the parent would choose if they decided not to seek the reward and received no subsidy. The same expression for $\frac{dW}{dc}$ applies in the case of imperfect information and low women's status, though $\alpha(c)$ would be defined differently.

The expression above is less simple to use than the earlier equations for the optimal subsidy, as it depends on functional forms of $v(z)$ and $f(c)$, since the optimization decision of the individual is no longer marginal in nature. However, a larger issue is that, in the presence of a heterogeneous population as in the real world, our conclusion about the superiority of a threshold prize no longer holds. We may prefer to use two instruments – a per-unit subsidy and a prize – so as to target different sections of the population; and a variety of other policy instruments are available, including subsidies and prizes that are wealth- or income-contingent (if we can reliably measure wealth or income), or policies that are conditional on weight gain during pregnancy rather than final weight, if that is observable.

In any case, answering the optimal policy question in the presence of a heterogenous population is not possible with the sufficient statistics method; the intuitions provided by that model have helped us to understand the implications of factors such as imperfect information and low women's status, but to go further, we need a calibrated numerical model that can be simulated to perform policy experiments. The following subsection presents the calibration of our model to the Indian context.

3.1 Calibration of Model

In the rest of the paper, we focus on the specific case of India. As already noted, about one-third of women in India are underweight at the time that they become pregnant, and deficits in children's human capital in India are particularly profound and well-studied (Coffey and Spears, 2018). India is also quantitatively important when considering maternal nutrition for the simple reason that about one-fifth of all births now occur in India. Therefore, India is a setting in which policies that address underinvestment in maternal nutrition could have very large benefits.

We calibrate our model using data from the India Human Development Survey-II (Desai and Vanneman, 2018), which is a nationally representative survey of 42152 households across India in 2011-12. For our purposes, the IHDS-II has the advantage of containing data on both income and the weight of potential mothers, as well as variables which can be used as proxies for each household's information and women's status. However, the data on the weight of women is imperfect, as we have only one observation per woman,[8] and very few observations are identified as coming from currently pregnant mothers. Other datasets have better data on the weight of pregnant women in particular, but given the importance for our welfare analysis of calibrating the model across the income distribution, we consider that the advantages of the IHDS-II outweigh the disadvantages.

To begin with, we need to specify functional forms for the utility function and the child investment function $f(c)$. For the former, we assume that $u(x) = \frac{x^{1-R}}{1-R}$ is CRRA with declining-marginal-utility parameter R, whereas $v(z) = \eta \log(z)$ and $e(L) = \frac{1}{\gamma}L^{\gamma}$. The child investment function is a bit more complicated, as in the model it transforms monetary investments into expected child productivity, whereas we have data on the weight of the mother, which is an intermediate input into child productivity that is itself a function of monetary investments. To be clear, if we define weight of the mother as k, we have $f(c) \equiv$

[8] Technically, there are two variables for a respondent's weight, AP8 and AP9, but these are simply two measurements taken at the same time.

$g(k(c))$, and we need to calibrate both $g(k)$ and $k(c)$.

We start by defining $k(c)$ in reverse: what is the cost c for k units of weight for the mother? An FAO website on the energy requirements of pregnancy (Food and Agriculture Organization of the United Nations, 2004) estimates that pregnancy with a gestational weight gain of 12 kg imposes a total energy cost of about 77000 kilocalories; however, Coffey (2015) finds that the average Indian women gains only 7 kg in pregnancy, and that there appears to be little variation in this average across starting weight levels. Therefore, we assume that the "standard" weight gain of $k = 7$ kg requires 44917 kilocalories, or 164 per day over 9 months; by normalizing k in this way, we treat k as the weight of the mother above a basic pre-pregnancy baseline, including weight gain before and during pregnancy. Meanwhile, the standard Mifflin-St Jeor (Mifflin et al., 1990) equation suggests that gaining an extra kg (outside of pregnancy) requires a sustained increase of 10 kilocalories per day; if we assume that pregnant mothers need to consume these added calories over 45 months (5 times the duration of pregnancy), this implies that the slope of calories with respect to mother's weight is 50 per pregnancy-day, and our equation is $kcal = -186 + 50k$. Meanwhile, the cost of calories is presumably convex in the quantity consumed; we use results from Subramanian and Deaton (1996) that indicate that the price per 1000 kilocalories varies from 0.88 rupees in the bottom decile of the population to 1.5 in the top decile. Given net weights in our data (assumed to represent the gain in pregnancy plus added weight above a baseline) that range from 7 to 14.61 kg across quartiles, and assuming that price is linear in mother's weight, we have that the price per 1000 kilocalories is $0.3095 + 0.0815k$ per day in 1983 rupees; combining these two equations, and using our estimate that average household consumption was 21.1 rupees per day in 1983,[9] we have the following equation for cost as a percentage of current income:

$$c(k) = \frac{(-186 + 50k)(0.3095 + 0.0815k)}{21100} = \frac{-2.7283 + 0.015k + 0.1931k^2}{1000} \tag{7}$$

which can be inverted to give us $k(c)$:

$$k(c) = \frac{-0.015 + \sqrt{0.015^2 + 0.7724(2.7283 + 1000c)}}{0.3862}.$$

[9] Subramanian and Deaton (1996) find a mean of 115 rupees of expenditures per month per capita in their sample, and the 1981 Census estimates that average household size was 5.5; dividing by 30 days per month, this implies 21.1 rupees per household per day.

For the functional form of $g(k)$, we assume $g(k) = g_0 + g_1 \hat{k}^{g_2}$, where $\hat{k} = k - k_0$ for a "baseline" weight k_0 that is given by $k(0) = 3.72$. We use estimates from several papers to pin down the parameters of this function. Girard and Olude (2012) find that 0.45 kg of mother's weight is associated with 0.105 kg of the child's birth weight,[10] and Alderman and Behrman (2006) find that raising a newborn out of low-birth-weight status (they consider about a 1 kg gain) is associated with a 7.5% increase in earnings as well as other monetary benefits that are 76% as large in present value, implying a 13.2% overall increase in p. As a result, we conclude that 1 kg of mother's weight raises $E(p)$ by 0.0308 at $k = 7$. Meanwhile, Strauss (2000) estimates that a baby that is 1 kg less than "normal" earns 7.77% less as an adult; we take 14.5 kg to be a "normal" mother's weight gain,[11] and so this implies a 7.77% drop in $E(p)$ for $k = 10.21$ relative to $k = 14.5$. We also assume that $E(p) = 1$ at $k = 7$, and this gives us 3 equations for g_0, g_1, and g_2:

$$g_0 + g_1(7 - k_0)^{g_2} = 1$$

$$g_1 g_2 (7 - k_0)^{g_2 - 1} = 0.0308$$
$$\frac{g_0 + g_1(14.5 - k_0)^{g_2}}{g_0 + g_1(10.21 - k_0)^{g_2}} = 0.0777$$

and these equations can be solved for $g_0 = 0.8080$, $g_1 = 0.1027$, and $g_2 = 0.5262$.

We then allow the wage distribution $\phi(w)$ to be parameterized as the exponential function of a cubic spline with nodes at 0.25, 0.75, 1.5, 4, and 10 (the mean of the income distribution will be targetted to 1). We allow for a fraction $\theta(w)$ of households to suffer from information constraints and/or low women's status (though we model it explicitly as an information constraint); this fraction varies across the wage distribution according to a cubic spline with the same nodes as the wage distribution, though we constrain the fraction that is underinformed to decline with wages. We assume that underinformed agents act as if the function linking child productivity to investments is given by $\beta f(c) + (1 - \beta) f(\bar{c})$, where $\bar{c}$ is the estimated mean of c in our data and where we calibrate β to match the IHDS-II

[10] Girard and Olude (2012) evaluate the effect of nutrition education and counselling on both outcomes, effectively allowing us to interpret the education and counselling as an instrumental variable for the effect of mother's weight on birth weight of the child.

[11] The standard recommendations are 11.5-16 kg for a women of healthy BMI, and 12.5-18 kg for an underweight woman. Given that a significant fraction of young Indian women are underweight, we take the average of these 4 quantities as a rough estimate of a "normal" baseline.

moments. Therefore, to keep the number of parameters manageable, we assume that there are only two types of families in our model with respect to information/women's status: underinformed, and fully informed (for whom $\beta = 1$).

This gives us the full set of parameters listed in Table 1: those in panel B will be chosen to match a set of moments in the IHDS-II, whereas those in panel A are selected to match standard estimates from the literature. Specifically, the lump-sum transfer d is set to 0.074 to match a report in the Indian Express (2017) of total social service spending of 7.4% of GDP, while q is set so that $d + q = 0.133$ to correspond with total government spending representing 13.3% of GDP in 2016-17 (Trading Economics, 2018). We assume modest growth of the relative size of government in setting the 2nd-period tax rate $\tau = 0.15$. Finally, we set γ to correspond to the value of the elasticity of taxable income, and given the limited evidence on the value of such an elasticity in developing countries, we use the finding in Piketty et al. (2014) that the 95th percentile of elasticities for the top 1 percent in OECD countries is 0.595; in the context of the inefficient fiscal infrastructure in India, this seems likely to be a lower bound. In a homogeneous-agent version of our model, we can show that:

$$\gamma = \frac{1 + \varepsilon^t}{\varepsilon^t} - \frac{w(1-t)R(1+\varepsilon^t)}{x\varepsilon^t}$$

and so we use this equation in our calibration to choose γ such that $\varepsilon^t = 0.595$. We find that, in the final calibrated model, the income-weighted average of the right-hand side gives us $\gamma = 1.6152$.

The moments used to set the parameters in panel B of Table 1 all come from the IHDS-II, and can be found in Table 2. We only keep observations representing women of ages 16 to 55 (which are the ages found in the much smaller pregnant sub-sample) and without missing values for a series of variables used in the regressions described below,[12] and focus on two variables: an adjusted per capita household income, which we calculate as household income divided by the square root of the number of members of the household, and the woman's weight as measured by variable AP9 with the level normalized so as to have $k = 7$ in the bottom quartile.[13] We drop observations with average weight or income below the 1st

[12]These variables are the SN2 series of variables, the some of the MM series of variables, the weight variable AP9, the activity status variable RO7, the caste and religion variable GROUPS, and education EDUC7.

[13]Since very few observations are identified as being currently pregnant, we assume that AP9 measures

Table 1: Parameters for Calibration

Name	Description	Value
	Panel A: Selected to Match Literature	
g_0	level parameter of child productivity function	0.8080
g_1	slope parameter of child productivity function	0.1027
g_2	curvature parameter of child productivity function	0.5262
d	lump-sum transfer	0.074
q	other government expenditures	0.059
τ	2nd-period tax rate	0.15
γ	exponent on labour supply in $e(L)$	1.6152
	Panel B: Parameters for Matching of Moments	
$\phi(w)$	wage distribution	$*$
$\theta(w)$	percentage uninformed	$*$
β	information of underinformed households	$*$
R	coefficient of relative risk-aversion	$*$
η	welfare weight on $\log(z)$	$*$

Notes: Panel A presents parameters for which values are selected from or to match estimates from the literature. Panel B presents parameters whose values will be calibrated in a moment-matching procedure, as indicated by the use of a $*$ in place of a numerical value.

percentile or above the 99th percentile, then divide this sample into quartiles of adjusted per capita income, and use a total of 16 quantities: in each quartile of the income distribution, we use the mean and standard deviation of income (normalized to an overall mean income of 1) and the mean and adjusted standard deviation of maternal nutrition investments. Since our model cannot capture all the sources of variation in mother's weight, we run a regression of weight on a quadratic in per capita income and a series of variables that measure information or women's status,[14] and use the fitted values to calculate the standard deviation of mother's weight. Then, since we observe women's weight and not actual investments, we back out estimated investments from our function for $c(k)$ in (7), and use the delta method to calculate

the weight before pregnancy, and that each woman will add an expected 7 kg of weight during pregnancy. The average pre-pregnancy weight in the lowest quantile thus corresponds to $k = 0$, the minimum theoretical level of maternal nutrition investments.

[14]These include the SN2 series of variables on acquaintances with various professionals, the HB series of variables about health and pregnancy beliefs, some of the MM series of variables on the use of TV, radio, newspaper, computers, and mobile phones, variables for primary activity status, education, and caste and religion, and a variable we constructed that indicates whether a woman is head of the household. The HB variables contain a large amount of missing values, so rather than drop missing observations, we recode missing values as zeros and add dummies to our regressions for missing values of each variable.

standard deviations for c.

Table 2: Moments from IHDS-II by Quartile

Description	Q1	Q2	Q3	Q4
mean of income	0.2484	0.5175	0.9059	2.3282
standard deviation of income	0.0875	0.0808	0.1605	1.1302
mean of maternal nutrition investment	0.0068	0.0125	0.0203	0.0387
standard deviation of maternal nutrition investment	0.0085	0.0110	0.0144	0.0220

Notes: This table presents the set of moments used for the calibration.

The calibration procedure is as follows. We start at a baseline vector of parameter values, and we define a grid of 89 points on the wage distribution; at each point we solve for the utility-maximizing c and L for both informed and uninformed households, given parameters and values of t and G. We then perform a cubic spline of c and L over a much finer wage grid ranging from 0.1 to 10 at intervals of 0.00001, and we solve for $f(c)$ and incomes on this finer grid, as well as the budget-balancing values of t and G, and then iterate the whole procedure until t and G converge. This allows us to calculate the moments in our model that correspond to the quantities in Table 2. We then iterate over parameter values using fminsearch in Matlab, looking for the vector of parameter values that minimizes the sum of weighted squared differences between the moments in our model and in the data, where each moment is weighted by the inverse of the mean of that category of moments across quartiles in the data.

This finally gives us the calibration results presented in Table 3. The values for $\phi(w)$ and $\theta(w)$ are the values at the nodes of the cubic splines, which are located at 0.25, 0.75, 1.5, 4, and 10. Since a distribution must sum to 1, one of the values of $\phi(w)$ can be normalized, and so we chose $\exp(-12.5)$ for the first node. The results are fairly intuitive and sensible; the low value of $\beta = 0.138$ indicates that underinformed households – who are numerous at the bottom of the distribution – have very limited information about the value of maternal nutrition. Figure 1 presents the wage distribution and the values of $\theta(w)$ over the entire distribution of wages, while Figures 2, 3, and 4 show the baseline distributions of c, k and $E(p)$. Meanwhile, Table 4 presents the values of the moments by quartile in the calibrated model, and a comparison to Table 2 indicates that the model fits the data very well.

Table 3: Calibrated Values of Model Parameters

Name	Description	Estimate
$\phi(w)$	wage distribution	$\exp(-\{12.5, 11.45, 12.95, 19.96, 398.27\})$
$\theta(w)$	percentage uninformed	$\{0.776, 0.448, 0.276, 0.095, 0.004\}$
β	information of underinformed parents	0.118
R	coefficient of relative risk-aversion	0.441
η	welfare weight on $\log(z)$	0.398

Notes: This table presents the calibrated values of the moments from panel B of Table 1.

Figure 1: Wage Distribution and Percentage Uninformed

(a) Wage Distribution (b) Percentage Uninformed $\theta(w)$

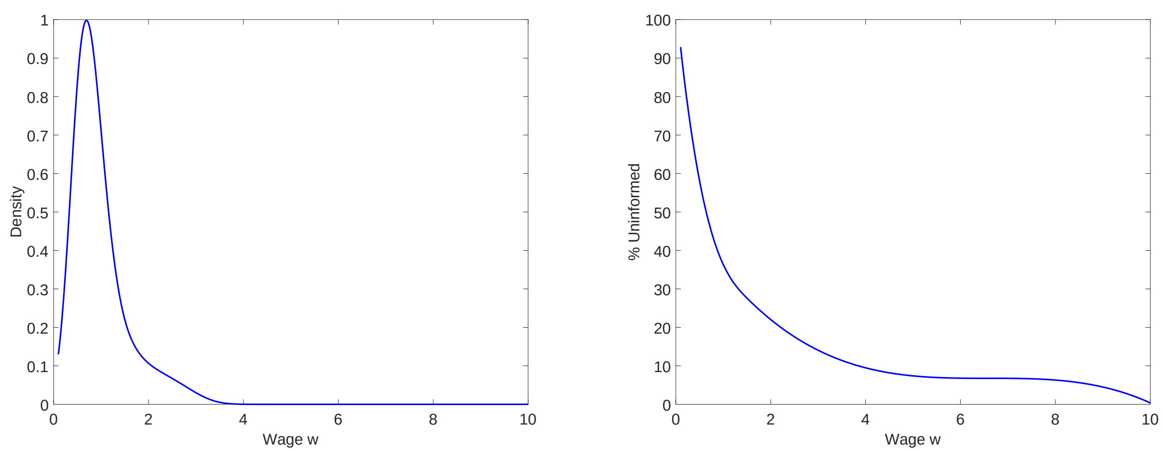

Table 4: Moments from Calibrated Model by Quartile

Description	Q1	Q2	Q3	Q4
mean of income	0.2119	0.5156	0.8935	2.3440
standard deviation of income	0.0994	0.0871	0.1458	1.1292
mean of maternal nutrition investment	0.0065	0.0135	0.0205	0.0383
standard deviation of maternal nutrition investment	0.0079	0.0114	0.0141	0.0223

Notes: This table presents the set of moments calculated from the calibrated model.

3.2 Optimal Policy in Heterogeneous Population

With the calibration of the numerical model provided in the previous subsection, we can now consider optimal policy. To begin with, we should note that, for simplicity and to correspond with the model, we focus on subsidies and awards based on maternal nutrition investments c rather than mother's weight k. This choice is without loss of generality for the award, as

Figure 2: Baseline Values of c Across Wage Distribution

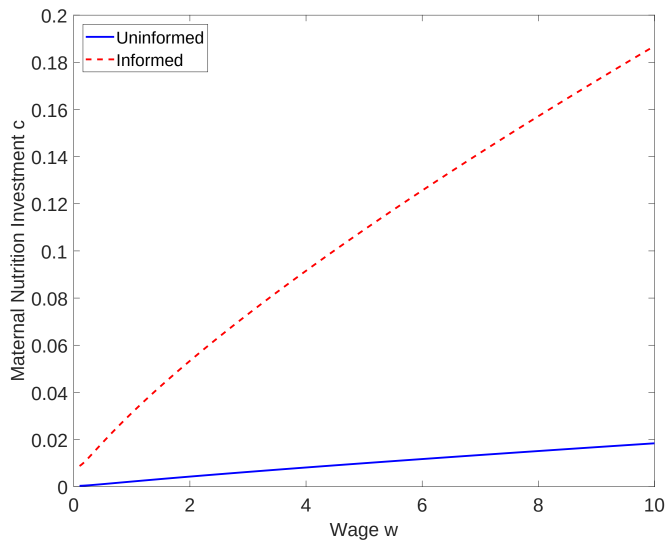

Figure 3: Baseline Values of k Across Wage Distribution

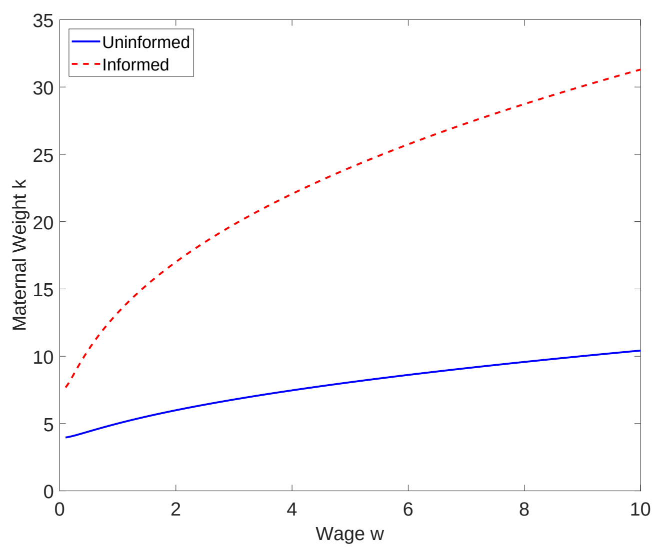

it is binary and our function linking c to k is deterministic and monotone, but it means that our linear subsidy in c would correspond to a slightly convex subsidy in k.

All optimal policy results can be found in Table 5. We consider two overall scenarios: in the first, we assume that the lump-sum transfer d (as well as q) is exogenously fixed, and solve for optimal maternal nutrition policy conditional on d. In the second, we allow for the possibility that redistribution is underprovided by the Indian state – which our calibrated

Figure 4: Baseline Values of $E(p)$ Across Wage Distribution

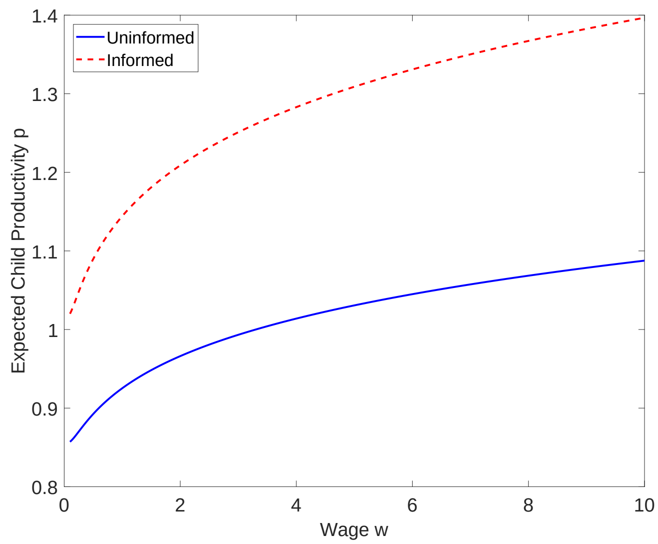

model indicates to be true – and increase d to the value that would be optimal in the absence of any maternal nutrition decision, which is $d = 0.1418$; we then solve for optimal maternal nutrition policy conditional on that value of d. As will be observed, the optimal policy is less generous in the second scenario, because maternal nutrition policy no longer is required for pure redistribution and because of the reduced fiscal space, but the quantitative size of the subsidies and awards remains large.

We begin by considering the optimal subsidy b on maternal nutrition investments c if no award is available. As can be seen in the table, the percentage subsidy is large: $b = 0.4213$ if d is fixed, and $b = 0.3700$ if $d = 0.1418$. Relative to the baseline averages, the average mother's weight k increases by about 2.4 to 3 kg, and future productivity increases by about 5 to 6%. The welfare gain from such a policy in present-value terms is equivalent to a 0.42% to 0.54% increase in mean consumption, which is moderate in size: given that the average value of c was about 0.02 at baseline, these welfare improvements are equivalent to a range of 20-30% of baseline spending on maternal nutrition.

As noted above, this simple subsidy is unlikely to be the optimal policy if more complex policy instruments are available, as it requires a significant increase in taxes for each unit of subsidy, and because of the regressivity implications of strongly subsidizing an activity that is particularly undertaken by wealthier families. Therefore, we next consider the optimal

Table 5: Numerical Results for Optimal Policy

Policy Design	Optimal Value(s)	Welf. Gain	$E(k)$	$E(p)$
Panel A: Fixed $d = 0.074$				
baseline	–	–	9.638	1.0418
maternal nutrition subsidy b	0.4213	0.538%	12.608	1.1036
combined policy $\{b, c_a, \alpha\}$	$\{0.1140, 0.0283, 0.0702\}$	3.183%	13.558	1.1491
combined weight-gain policy (version 1)	$\{0.1359, 0.0207, 0.0722\}$	3.129%	14.551	1.1642
combined weight-gain policy (version 2)	$\{0.0000, 0.0207, 0.0750\}$	3.129%	14.560	1.1643
combined weight-gain policy (version 3)	$\{0.2345, 0.0286, 0.0710\}$	3.283%	14.103	1.1583
Panel B: Optimized $d = 0.1418$				
baseline	–	–	9.467	1.0388
maternal nutrition subsidy b	0.3700	0.422%	11.848	1.0899
combined policy $\{b, c_a, \alpha\}$	$\{0.1102, 0.0258, 0.0179\}$	2.432%	13.194	1.1421
combined weight-gain policy (version 1)	$\{0.1821, 0.0202, 0.0124\}$	2.427%	14.327	1.1604
combined weight-gain policy (version 2)	$\{0.0000, 0.0203, 0.0161\}$	2.427%	14.338	1.1606
combined weight-gain policy (version 3)	$\{0.2378, 0.0269, 0.0154\}$	2.554%	13.893	1.1544

Notes: This table presents the optimal policies in our calibrated model; panel A presents the case in which the lump-sum transfer d is left unchanged, while panel B presents the case in which d is optimized in a model without a maternal nutrition decision. The third column presents the welfare gain as an equivalent percentage increase in consumption from the baseline average, where the baseline features $d = 0.074$ and $d = 0.1418$ respectively. The fourth and fifth columns present the average values of k and p at the optimum. The 3 versions of the combined weight-gain policy are described in the text.

combined policy when both a per-unit subsidy and a prize conditional on a threshold are available. We find that the per-unit subsidy is much lower, at about 0.11, but the award is quite economically significant, reaching $\alpha = 0.0702$ for fixed d and 0.0179 for optimized d. The threshold for the award is quite low, less than $c = 0.03$, so that it is accessible to everyone (even poor families with limited information about maternal nutrition); given that the marginal value of c is decreasing with the size of the investment, solving underinvestment at the bottom of the distribution is most important. The welfare gain of such a combined policy is quite large, equivalent to an increase in consumption of 2.4 to 3.2 percentage points; this amounts to a welfare gain of well over 100% of baseline spending on maternal nutrition.

Another way to understand the quantitative significance of the optimal combined policy is to consider the magnitude of the policy instruments: in panel A, mean c at the optimum is 0.0336, so the subsidy of $b = 0.1140$ amounts to spending of 0.0038, and combined with the award of $\alpha = 0.0702$, this means government spending on maternal nutrition policy that equals 7.4% of average baseline income, or the same amount as the lump-sum transfer d. This amount is nearly 4 times the baseline average spending on maternal nutrition by households. Given the average annual household income of 143811 rupees in our IHDS-II data, and given that we have converted expenditures into a percentage of expenditures during the period of pregnancy, the award alone corresponds to a transfer of 7572 rupees (or $163 in 2011 US dollars, given an average exchange rate of 46.59 in that year) to each pregnant mother that reaches a threshold of investment that corresponds to 12.63 kg of mother's weight above the baseline pre-pregnancy weight of the bottom quartile.[15]

Alternatively, considering that the average household size in our IHDS-II sample is 5.626, and that the 2011 Indian Census found a crude birth rate of 21.8 per 1000 population, there is a 12.26% chance of any given household having a baby in any given year; this implies that the overall government spending amounts to an investment in maternal nutrition policy of $0.074 \times 0.75 \times 0.1226 = 0.68\%$ of Indian GDP. Either of these quantities gives us the same conclusion: the optimal maternal nutrition policy in India is quantitatively significant.[16]

Finally, we consider a policy that requires more intense monitoring: what if we could observe birth weight before and at the end of pregnancy, to condition transfers on weight gain during pregnancy? It is not immediately clear how to implement such a policy in our model, as our model does not distinguish between weight gain before and during pregnancy. However, building on the finding in Coffey (2015) that the average weight gain during pregnancy in India varies little across starting weight levels, we model this as a subsidy and/or award based on the value of c relative to its baseline, laissez-faire value, or the value of c that each parent would choose absent any subsidy or award. We call this value c_0, and we consider 3 different versions of the policy: version 1 conditions both the subsidy and the award on $\Delta c \equiv c - c_0$, whereas version 2 only conditions the award on Δc while paying the subsidy on the full value of c, and version 3 does the opposite (award on c, subsidy on Δc). From the results in Table 5, it is clear that only version 3 offers the possibility of additional welfare gains; as noted above, raising maternal nutrition at the bottom of the distribution is more

[15]In the IHDS-II data, the actual value of this baseline pre-pregnancy weight is 47.01 kg in the bottom quartile, so the threshold corresponds to a maternal weight of 59.64 kg at birth.

[16]Even in the case in which d is increased to 0.1418, the optimal combined subsidy and award amounts to 2.14% of average baseline income. In that case, our results would indicate that basic redistribution should be roughly doubled *and* that a further 0.2% of GDP should be invested in maternal nutrition policy; the quantitative conclusions of the paper remain significant in that case.

important than at the top, and version 3 keeps the total subsidy payments manageable by not paying enormous amounts to the wealthiest families (who tend to have healthier mothers), allowing for generous per-unit subsidies on weight gain above 7 kg that help to support an ambitious weight target for uninformed low-income families. However, the welfare gains of such a policy relative to the basic combined policy are fairly small, at 0.1% to 0.12%. Using the intuition of these results, however, we expect that more complicated policies could be designed that could more profitably use the additional information; for example, a policy that features a concave subsidy and an award, allowing the government to aggressively target their resources on the least-fortunate mothers.

4 Conclusion

A large literature has documented the empirical importance of early-life health, especially in developing countries, and especially in cases where maternal nutrition is poor. This situation creates fiscal externalities: all taxpayers have an interest in the next generation receiving the human capital inputs necessary to be productive earners. These externalities can be exacerbated when parents are either unaware of the importance of investing in maternal nutrition, or uninterested in making sufficient investments due to low women's status in hierarchical and paternalistic households. This paper has contributed a calibrated model of the government's optimal policy design problem that takes into consideration the pragmatic constraint that households make their own economic choices of how much of their resources to devote to maternal nutrition. We find that the optimal policy corrective to this distortion could be quantitatively large. As a result, investments in state capacity to conduct better-informed nutrition policy — such as through a health and nutrition surveillance system — could generate valuable policy tools.

References

Alderman, H. and J. R. Behrman (2006): "Reducing the Incidence of Low Birth Weight in Low-Income Countries Has Substantial Economic Benefits," *World Bank Research Observer*, 21, 25–48.

Almond, D., J. Currie, and V. Duque (2017): "Childhood Circumstances and Adult Outcomes: Act II," Working Paper no. 23017, NBER.

Behrman, J. R. (1997): "Intrahousehold Distribution and the Family," in *Handbook of Population and Family Economics*, ed. by M. R. Rosenzweig and O. Stark, vol. 1A, Elsevier, 125–187.

Behrman, J. R. and M. R. Rosenzweig (2004): "Returns to Birthweight," *Review of Economics and Statistics*, 86, 586–601.

Black, S. E., P. J. Devereux, and K. G. Salvanes (2007): "From the Cradle to the Labor Market? The Effect of Birth Weight on Adult Outcomes," *Quarterly Journal of Economics*, 122, 409–439.

Buchanan, J. M. (1966): "Externality in Tax Response," *Southern Economic Journal*, 33, 35–42.

Chetty, R. (2009): "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 1, 451–488.

Coffey, D. (2015): "Pre-Pregnancy Body Mass and Weight Gain During Pregnancy in India and Sub-Saharan Africa," *Proceedings of the National Academy of Sciences*, 112, 3302–3307.

Coffey, D. and D. Spears (2017): *Where India Goes: Abandoned Toilets, Stunted Development and the Costs of Caste*, Harper Collins.

——— (2018): "Child Height in India: Facts and Interpretations from the NFHS-4, 2015-16," *Economic & Political Weekly*, 53, 87.

——— (2019): "Neonatal death in India: The effect of birth order in a context of maternal undernutrition," working paper, UT-Austin.

Currie, J. and M. Rossin-Slater (2015): "Early-Life Origins of Life-Cycle Well-Being: Research and Policy Implications," *Journal of Policy Analysis and Management*, 34, 208–242.

Currie, J. and T. Vogl (2013): "Early-Life Health and Adult Circumstance in Developing Countries," *Annual Review of Economics*, 5, 1–36.

Das Gupta, M. (1995): "Life Course Perspectives on Women's Autonomy and Health Outcomes," *American Anthropologist*, 97, 481–491.

Deaton, A. (2013): *The Great Escape: Health, Wealth, and the Origins of Inequality*, Princeton University Press.

Desai, S. and R. Vanneman (2018): "India Human Development Survey-II (IHDS-II), 2011-12," Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2018-08-08. `https://doi.org/10.3886/ICPSR36151.v6`.

DRÈZE, J. AND A. SEN (2013): *An Uncertain Glory: India and its Contradictions*, Princeton University Press.

FOOD AND AGRICULTURE ORGANIZATION OF THE UNITED NATIONS (2004): "Human Energy Requirements: Report of a Joint FAO/WHO/UNU Expert Consultation," http://www.fao.org/docrep/007/y5686e/y5686e0a.htm, chapter 6: Energy Requirements of Pregnancy, viewed on February 7, 2019.

GIRARD, A. W. AND O. OLUDE (2012): "Nutrition Education and Counselling Provided during Pregnancy: Effects on Maternal, Neonatal and Child Health Outcomes," *Paediatric and Perinatal Epidemiology*, 26, 191–204.

LAWSON, N. (2017a): "Fiscal Externalities and Optimal Unemployment Insurance," *American Economic Journal: Economic Policy*, 9, 281–312.

——— (2017b): "Liquidity Constraints, Fiscal Externalities, and Optimal Tuition Subsidies," *American Economic Journal: Economic Policy*, 9, 313–343.

MIFFLIN, M. D., S. T. ST JEOR, L. A. HILL, B. J. SCOTT, S. A. DAUGHERTY, AND Y. O. KOH (1990): "A New Predictive Equation for Resting Energy Expenditure in Healthy Individuals," *American Journal of Clinical Nutrition*, 51, 241–247.

PIKETTY, T., E. SAEZ, AND S. STANTCHEVA (2014): "Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities," *American Economic Journal: Economic Policy*, 6, 230–271.

PRESTON, S. H. AND M. R. HAINES (1991): *Fatal Years: Child Mortality in Late Nineteenth-Century America*, Princeton University Press.

PROGRAMME EVALUATION ORGANISATION (2005): "Performance Evaluation of Targeted Public Distribution System (TPDS)," http://planningcommission.nic.in/reports/peoreport/peo/peo_tpds.pdf, Planning Commission, Government of India, viewed on January 30, 2019.

PTI (August 11, 2017): "Need more social spending on health, education: Survey," https://indianexpress.com/article/india/need-more-social-spending-on-health-education-survey-4792481/, Indian Express.

SPEARS, D. (2012): "Height and cognitive achievement among Indian children," *Economics & Human Biology*, 10, 210–219.

STRAUSS, J. AND D. THOMAS (2007): "Health over the Life Course," in *Handbook of Development Economics*, ed. by T. P. Schultz and J. A. Strauss, vol. 4, Elsevier, 3375–3474.

STRAUSS, R. S. (2000): "Adult Functional Outcome for Those Born Small for Gestational Age: Twenty-six-Year Follow-up of the 1970 British Birth Cohort," *Journal of the American Medical Association*, 283, 625–632.

Subramanian, S. and A. Deaton (1996): "The Demand for Food and Calories," *Journal of Political Economy*, 104, 133–162.

Trading Economics (2018): "India Central Government Total Expenditure to GDP," `https://tradingeconomics.com/india/government-spending-to-gdp`, viewed on February 7, 2019.