Econometric Game 2012, Case B:

How does maternal smoking during pregnancy affect infants' birthweight?

Group 4

April 19, 2012

1 Introduction

Low birthweight of infants is associated with adverse health related and economic outcomes. Black, Devereux, and Salvanes (2007) show that in the short run, newborns with low birthweights have higher mortality rates, and in the long run, those infants will later in life have lower earnings and education attainment.

There are many determinants of birthweight which range from genetic composition of parents, demographic and psychological factors, obstetric factors, and maternal behaviors such as cigarette smoking, alcohol consumption, caffeine intake, and narcotic addictions as explained in Kramer(000). Previous studies have shown that maternal smoking during pregnancy is associated with lower birthweight, but the extent to which this effect is causal needs further examination. The fact that maternal smoking may be correlated with other unobserved variables makes the econometric inference problem challenging. While reduced form evidence that does not adjust for endogeneity may serve a useful purpose in forecasting exercises, policy recommendations (such as excise tax changes) must be based on a causal model.

According to Abrevaya (2006), OLS estimates of average reduction in birthweight due to smoking are around 200–250 grams, but, if unobserved heterogeneity is correlated with smoking, the causal reduction in birthweight may be overestimated. His panel data results indicate that the correct effect should be in the range of 100–150 grams. To address causality, there have been efforts to employ instrumental variable (IV) strategies, such as WN and JS. (1999) who used cigarette taxes as instrument. Most of the previous IV studies find a higher estimate of reduction in birthweight than the OLS estimate, contradicting intuition, as well as large standard errors. Another approach would be to model the endogeneity issue as a time-invariant unobserved heterogeneity (i.e., fixed or random effect) and exploit variation in the time dimension of a panel data set (Abrevaya (2006)).

In this paper, we use panel data from 12360 women with observations from three consecutive births. First we perform model selection and estimate a linear fixed effects model for

birthweight. A potential issue, raised by Abrevaya (2006), is the possibility of *feedback effects*, i.e., that smoking behavior may be dependent on the outcome of previous births. Such an effect would call into question the standard strict exogeneity assumption. We present a model that helps gauge the extent of feedback and also suggests a panel IV approach to robust inference. As far as we know, we are the first in the literature to specifically address the feedback effect. Finally, we analyze the effects of smoking on different parts of the conditional distribution of birthweights; specifically, the conditional probability of giving birth to a child with a weight in certain intervals.

2 Data

The data is a panel of three consecutive births for a sample of 12,360 women. One of the main covariates of interest, the number of cigarettes per day that the woman reports having smoked during her pregnancy, is reported missing for 179 observations. We drop all women from the panel who do not have three reported measures of cigarette consumption (160 mothers), leaving a balanced panel of 12,200 mothers.

One critical issue to consider is how to capture cigarette consumption given both the presence of outliers and possible mismeasurement. As illustrated in Figure 1 (reported cigarettes per day for women conditional on cigarettes > 0), reported cigarette consumption that disproportionately clusters around multiples of 10, the number of cigarettes in half a pack. In addition, there are extreme outliers, with a maximum of 92.

3 Linear Panel Data Fixed Effects Model

The panel nature of the data set allows us to control for unobserved heterogenity at the mother's level. Let y_{ib} be the birthweight of birth b of woman i.[1] We postulate the following

[1] In our cross-sectional analysis we used 10-fold cross validation to select between a specification with log birthweight and one with birthweight in levels on the left-hand side. The latter came out on top, and so we continue with the additive specification here.

Figure 1: Reported Cigarettes per Day
(Conditional on Smoking)

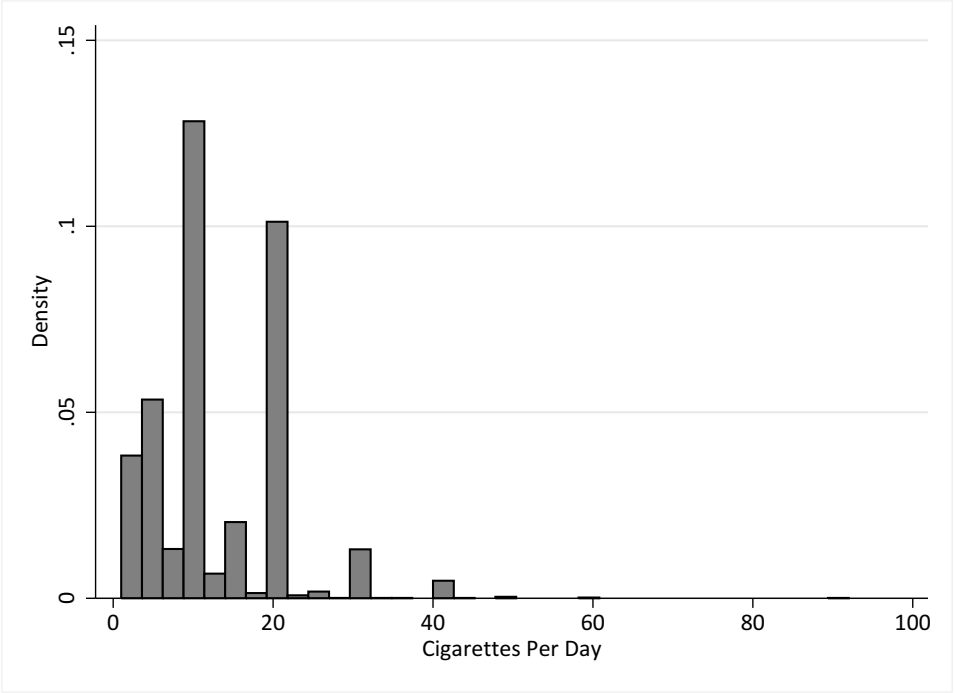

reduced linear panel-data model

$$y_{ib} = s_{ib}\beta + \gamma X_{ib} + c_i + u_{ib} \quad i = 1 \dots N, \quad b = 1, 2, 3 \tag{1}$$

where s_{ib} is the indicator for smoking during pregnancy, c_i is the (time-invariant) mother fixed effect, X_{ib} observed covariates at time t and u_{it} is a residual population term. Note that implicit in equation (1) lies the following assumption: after controlling for birth b covariates and birth b smoking status, there is no additional contribution from the previous (or future) birth covariates. Within the context of this model, it is possible to think about different ways of identifying the parameter of interest (β).

We will focus on two inference strategies. First, fixed effects estimation of β. We will discuss possible drawbacks of this inference procedure in the context of the feedback effect problem. Then, we will consider instrumental variable estimation of the parameter to account for such a problem.

3.1 Fixed effects estimation

Following the previous basic set-up, we first assume a linear model of the form

$$y_{ib} = \beta' X_{ib} + f(cigs_{ib}) + c_i + \epsilon_{ib},$$

where y_{ib} is the birthweight (in grams) of birth b of mother i, $f(cigs_{ib})$ is a function of the number of cigarettes per day reported consumed by the mother, c_i is a mother-specific, time invariant fixed effect, and X_{ib} are birth-specific variables potentially correlated with birthweight. Another potential outcome of interest would be gestation time, which could be analyzed in a duration model; we leave this for future work. For identification we impose the standard strict exogeneity condition

$$E[\varepsilon_{ib}|X_i, cigs_i] = 0, \quad b = 1, 2, 3,$$

where X_i and $cigs_i$ collect all regressors and smoking-related variables for mother i across all births.

If c_i is unobserved and correlated with cigarette consumption, cross-sectional methods may result in biased estimates of the coefficients of $f(cigs_{ib})$. Panel data allows for the estimation of individual specific fixed effects that control for all time-invariant components of c_i. However, this approach may not completely address the bias present if c_i is not time-invariant; we return to this point later in our discussion.

Our baseline specification for fixed effects is

$$y_{ib} = \beta' X_{ib} + \gamma s_{ib} + c_i + \epsilon_{ib}$$

where $f(cigs_{ib})$, the independent variable of interest, takes on the form of $s_{ib} = 1$ if $cigs_{ib} \geq 0$. X_{ib} includes child gender, mother age and age squared, and the five measures of prenatal care. Standard errors are estimated allowing for heteroskedasticity and autocorrela-

tion. See Table 1 attached at the end of the paper. Results from this fixed effects estimation are reported in Column 3, while results from the analogous OLS specification in which the data is treated as cross-sectional are reported in Column 1 for comparision[2]. The coefficient on smoking is estimated to be -165.00 with a standard error of 17.24 in the fixed effects model, which is much smaller in magnitude than the cross-sectional estimate of -274.3 (s.e. 10.37). This confirms the results of Abrevaya (2006), who finds similar magnitudes for OLS and fixed-effects regressions. The panel nature of the data also allows us to replace age controls with nonparametric controls for the time in years since last birth; we find that that this improves in-sample fit as evidenced by a reduction in the BIC but does not change the coefficient on smoking (Column 4).

There may be nonlinearities in relationship between cigarettes consumed and birthweight. To allow for this, we modify $f(cigs_{ib})$ by including two additional indicators for intensity of smoking. Let $m_{ib} = 1$ if the mother reports smoking at least 10 but less than 20 cigarettes per day (half a pack a day), and let $h_{ib} = 1$ if the mother reports smoking 20 or more cigarettes per day (a pack or more). We refer to these as moderate and heavy smokers, respectively. Column 5 reports coefficients from fixed effects estimates of

$$y_{ib} = \beta' X_{ib} + \gamma_1 s_{ib} + \gamma_2 m_{ib} + \gamma_3 h_{ib} + c_i + \epsilon_{ib}$$

Congruent with results from OLS estimates, intensity of smoking, above and beyond the indicator for smoking, continues to be an important predictor of birthweight. The coefficient on smoking is -122.53, which is significantly different from zero at 1% but smaller than the coefficient in the regression using only the binary smoking indicator. However, reporting moderate smoking predicts a decrease in expected birthweight of an additional 46.35 grams, which is significant at 10%, and reporting heavy smoking predicts a 88.85 gram decrease in expected birthweight which is significant at 1%.

[2]Note that cross sectional regressions include state of residence fixed effects and controls for time-invariant mother characteristics including age, race, marital status, and education.

In Column 6 we measure cigarette consumption ($f(cigs_{ib})$) using cigarettes per day in linear and squared terms in addition to the smoking dummy. The coefficients on the continuous measures are jointly significant at 1%. We compare the in sample and out of sample fit of models 4, 5, and 6 using the Bayesian information criterion and out of sample mean squared forecast error[3]. The model using the continuous cigarette measure has the lowest BIC and MSFE (269971.3 compared to 270210.5 for the model with only a smoking dummy and 269983.3 for the model with intensity controls). However, as discussed in the cross sectional analysis, the reported number of cigarettes per day is almost certainly subject to measurement error and there are large outliers in reported numbers of cigarettes that may impact model estimates. Thus, our preferred specification is the Column 5 fixed effects model with indicators for smoking intensity which has the second-lowest MSFE.

Comparing our preferred specification to those from the equivalent OLS regression in Column 2, we see that the coefficient on the smoking indicator is again significantly reduced. In addition, the impact of moderate smoking decreases from about -80, which was highly significant, to a marginally significant -46. Interestingly, however, the additional impact of heavy smoking on expected birthweight above and beyond any smoking is almost identical even when fixed effects are included.

In unreported results, we further allow for heterogeneity in the impact of smoking by adding interactions between the smoking indicator and time-invariant mother characteristics. When only a dummy for smoking is included as a main effect, the interactions between black and smoking and college graduate and smoking are statistically different from zero at 5%. However, when measures of smoking intensity are included, these interactions are neither individually nor jointly statistically significant. This suggests that heterogeneity in smoking effect by race (for example) is a result of differences in smoking intensity.

Taken as a whole, we find that the estimated impact of smoking on expected birthweight

[3]We compute the MSFE using "leave one group out" cross validation: we divide the mothers in to ten randomly-selected groups of equal size, estimate the model on nine groups, and predict outcomes for the omitted group. We repeat this for each group and compute the mean of the squared forecast errors.

is reduced by about 40% in fixed effects models compared to analogous cross-sectional estimates. This suggests that the cross-sectional estimates are subject to omitted variables bias from unobserved factors that impact both the conditional probability of smoking and infant birthweight.

3.2 IV estimation

3.2.1 Feedback effect

The feedback effect, as introduced by Abrevaya (2006), is the potential effect of previous birth outcomes on future smoking decisions. One may worry, for example, that a smoking mother whose first-born child is severely underweight may reconsider her smoking habits. One way to think about the feedback effect is that it arises through the dependence of the contemporaneous smoking decision (s_{ib}) on *past* births information.

$$s_{i2} = f(y_{i1}, s_{i1}, X_{i1}, X_{i2}, \eta_2) \tag{2}$$

$$s_{i3} = g(y_{i2}, s_{i2}, X_{i2}, X_{i1}, X_{i2}, X_{i3}, \eta_3) \tag{3}$$

The problem with such dependence, is that it conflicts with the *strict exogeneity* assumption, which is required to get consistency of the fixed effects estimator. More precisely, $\mathbb{E}[s_{12}u_{i1}]$ need not equal 0, as s_{i2} implicitily depends on u_{i1} through y_{i1}.

For the sake of concreteness, we ilustrate ideas within the context of a simple threshold model for the discrete choice problem of selecting smoking status.

$$s_{2i} = \begin{cases} 1 \quad if \quad \alpha_1 y_{1i} + \alpha_2 s_{i1} + \alpha_3' X_{i1} + \alpha_4' X_{i2} + \eta_2 > c^* \\ 0 \quad if \qquad\qquad\qquad \text{in other case} \end{cases} \tag{4}$$

Therefore, we can re-write

$$s_{i2} = \mathbf{1}[(\alpha_1\beta + \alpha_2)s_{i1} + (\alpha_1\gamma' + \alpha_3')X_{i1} + \alpha_1 c_i + \alpha_1 u_{i1} + \alpha_4' X_{i2} + \eta_2 > c^*]$$

Even under the strong assumption of exogeneity (or idenpendence) between u_{i1} and $(s_{i1}, \eta_2, X_{i1}, X_{i2})$, the decision s_{i2} depends on the first birth innovation. Therefore, $\mathbb{E}[s_{i2}u_{i1}] = 0$ need not be satisfied. The probit model provides a simple strategy to assess the relevance of the feedback effect in the data set. We use the discrete choice model in 4 to test the following null hypothesis: the previous period birthweight is equal to zero in the second birth period and in the third birth period. Rejection of the null hypothesis suggests that it is important to account for the feedback effect in our sample. In the second period we cannot reject the null hypothesis that coefficient on $y_{i1} = 0$. In the third period, we reject the null hypothesis that the effect on the lagged outcome variables y_{i1}, y_{i2} is zero at the .01 level.

To face the challenge posed by the feedback effect we propose the following strategy. Consider the following assumption. Let

$$\overline{u}_{i3} \equiv \begin{pmatrix} u_{i2} \\ u_{i3} \end{pmatrix}$$

Assumption 1. (IV estimation)

1. $\mathbb{E}[\overline{u}_{i3}s_1] = \mathbf{0}$
2. $\mathbb{E}[X_{ib} \otimes \overline{u}_{i3}] = \mathbf{0} \quad \text{for} \quad b = 2, 3$

Under these assumptions, the smoking status in birth $b = 1$, can be used as a valid instrument to estimate the parameter β. Let $\Delta_b a \equiv a_{ib} - a_{i(b-1)}$, for any random variable a. Note that the difference between third period and second period outcomes is given by:

$$\Delta y_{i3} = \beta \Delta s_{i3} + \gamma \Delta X_{i3} + \Delta u_{i3} \tag{5}$$

and, under part 2 of Assumption 1, the regressor ΔX_{i3} can be treated as exogenous. We report results for two IV specifications used to perform inference on β.

- **Model 1:** We use birth information from period 1 as an instrument for the change in smoking status between period 3 and period 2. The instruments are the smoking status at period 1 (s_{i1}), the birthweight outcome in period 1 (y_{i1}) and the levels of other first period covariates that vary over time, such as the measures of pre-natal care.

- **Model 2:** We drop y_{i1} from the set of regressors in Model 1.

The results are reported in the following table:

	(1)	(2)
	FD.birwt	FD.birwt
FD.smoker	-157.5	-256.3
	(286.74)	(288.08)
Robust F	2.349	2.474
J-Test	16.22	32.83
(p-value for J)	.26	.0030
N	12200	12200

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The first specification yields a point estimate smaller than the fixed effect estimator. The latter suggests that the feedback effect might actually pose an important challenge to exogeneity. There are three limitations associated to the first result:

- The confidence bands associated to Model 1 are large: $[-719.4, 404.5]$. In particular, they do not allow us to claim that a) the effect of smoking over birthweight is negative, and b) the effect is smaller than that captured by the fixed effects estimator. This is, of course, related to the strength of our set of instruments. As suggested by the first-stage this is potential source of concern.

- The test for overidentifying restrictions is rejected. This challenges the exogeneity assumption used to estimate β. One possible hypothesis that explains the violation of exogeneity concerns unobserved time varying heterogeneity; for example, drinking behavior. If drinking behavior is an important determinant of the birthweight, part of its effect will be captured in the residual term. Since contemporaneous drinking habits are presumably correlated with low weight outcomes in previous periods, the exogeneity assumption is likely to be violated.

- To address the exogeneity concern in the previous paragraph, we considered we dropped the birthweight outcome from the set of instruments. To our surprise, the test for overidentifying restrictions is no longer rejected. However, the point estimate goes back to the levels observed in the fixed effects model: a point estimate of -256.3 with confidence intervals of $[-820.98, 308.30]$.

4 Probabilities of High and Low Birthweights

In our cross-sectional analysis for Case A we used reduced-form probit regressions to assess the effect of smoking on the probability that the birthweight falls in certain intervals. This type of analysis has the potential to address possible non-linearities in the conditional distribution of birthweights, as it allows the researcher to focus on particular weight ranges and compare between them. This is an important exercise to perform since medical studies have shown that the health-related and economic costs of abnormal birth weight is largest at extreme birth weights.

In Section 3 we worked with a linear panel model of the form

$$y_{ib} = c_i + \gamma' X_{ib} + \varepsilon_{ib}, \tag{6}$$

with i indexing mothers and b births. Here c_i is the unobserved heterogeneity and X_{ib} are

the regressors (including a constant and the smoking dummy). Suppose that we strengthen the strict exogeneity assumption to

$$\varepsilon_{ib}|X_i \sim \mathcal{N}(0, \sigma^2),$$

where $X_i = (X_{i1}, \ldots, X_{i3})$. Furthermore, we make the correlated random effects (CRE) assumption that

$$c_i|X_i \sim \mathcal{N}(\psi + \xi'\bar{X}_i, \sigma_c^2), \tag{7}$$

where $\bar{X}_i = \sum_{b=1}^{3} X_{it}/3$. Then Wooldridge (2010, 2nd ed., Ch. 15.8.2) shows that

$$y_{ib} = \tilde{\gamma}'X_{ib} + \xi'\bar{X}_i + a_i + \varepsilon_{ib},$$

where $a_i|X_i \sim \mathcal{N}(0, \sigma_a^2)$.[4] For any cut-off birthweight A, this implies the probit model

$$\text{Prob}(y_{ib} \leq A|X_i) = \Phi(\delta'X_{ib} - \xi'\bar{X}_i),$$

where δ is obtained from $\tilde{\gamma}$ by subsuming A in the constant and flipping signs. This probability specification may readily be estimated in Stata using the command `probit` on pooled data. It is, however, imperative to cluster by mother as the disturbances may be correlated across births.[5]

We are of course also able to estimate right-tail probabilities consistent with the linear CRE specification, because

$$\text{Prob}(y_{ib} \geq A|X_i) = \Phi(-\delta'X_{ib} + \xi'\bar{X}_i).$$

[4] The parameter ψ has been subsumed in the constant, hence the change from γ to $\tilde{\gamma}$.

[5] A different approach would have been to impose the assumption that y_{i1}, y_{i2}, y_{i3} are i.i.d. conditional on X_i. Under this assumption, valid inference could be made with Stata's panel probit command `xtprobit`. While the added restriction may increase power, we choose to maintain the more robust clustered probit specification.

To assess whether the CRE assumption (7) is misspecified, we compared the estimated average partial effect (APE) of smoking from the above-mentioned probit model for Prob($y_{ib} \leq 2500$) with the APE from a fixed effects logit model with the same regressors. The latter model is only able to use 3238 of the mothers in the data, as the rest do not exhibit any variation in the incidence of low birthweights. The logit model's 99% confidence interval for the smoking dummy APE contains the confidence interval for the probit model, so we continue with the probit specification.

While our linear model (6) directly motivates CRE estimation of tail probability probit specifications, we seek a broader overview of the effects of smoking on the full conditional distribution of birthweights. In particular, we focus on the birth weight intervals

$$I_1 = [0, 2250), I_2 = [2250, 2500), I_3 = [2500, 2750), \ldots,$$
$$I_{11} = [4500, 4750), I_{12} = [4750, 5000), .$$

Note that we do not include a bin for the highest birth weights. There are only 84 observations in the dataset with birth weights greater than 5000 g, and so to aid the following graphical exposition we omit those observations for the present purposes. For each of the 12 bins we run a pooled (across mothers and births) CRE-type probit regression

$$\text{Prob}(y_{ib} \in I_n | X_i) = \Phi(\delta_n' X_{ib} + \xi_n' \bar{X}_i), \quad n = 1, \ldots, 12 \tag{8}$$

using left-hand side dummies $1_{\{y_i \in I_n\}}$. We then store estimates of the average partial effects (APEs)

$$APE_{mn} = E\left[\frac{\partial}{\partial X_{im}} \Phi(\delta_n' X_i + \xi_n' \bar{X}_i)\right], \quad m = 1, \ldots, M, \; n = 1, \ldots, 11.$$

Here M is the total number of covariates in X_i. The elements in X_i are similar to the baseline specification in Section 3: three smoking-related dummies, race, education dummies,

prenatal care dummies and the gender dummy.[6] As CRE estimation is hard to interpret for models with factor variables, we include the mother's age as well as its square instead of the indicator for years between consecutive births. The probit regressions are implemented in Stata using the commands `probit` (clustering by mother) and `margins`.

We emphasize that, unlike for tail probabilities, the probit model for *intervals* does not have a direct interpretation in terms of the linear model (6). However, we view the exercise as providing useful semi-reduced form evidence, which—if the CRE assumption (7) is approximately correct—is better able to correct for time-invariant unobserved heterogeneity.

Figure 2 displays the estimated APEs for the smoking dummy graphically. The horizontal axis lists the *upper* bound of the relevant bin. The dashed lines indicate *pointwise* (i.e., bin-by-bin) 99% heteroskedasticity robust (White) confidence bands, clustering by mother.[7] The results are strikingly similar to the cross-sectional plots we produced in Case A. In fact, the only material change is that the APEs shrink slightly toward 0 after controlling for unobserved heterogeneity. The broad shape of the graphs are the same as in the cross section. Maternal smoking has a significantly positive effect on the probability of giving birth to a child with low weight, except for very low birthweights. The effect on the probabilities of high birth weights are significant for relatively high birthweights (3500–4000 g) but not for more extreme weights. While weaker than the effect measured in the cross section, the estimated APEs are still economically substantial for intermediate weights, i.e., on the order of $\pm 3\%$ in the range 2750–4000 g.[8]

We confirm our surprising finding from the cross section that the absolute value of the

[6] In Case A we did the exercise separately for male and female babies. The results were almost identical, so here we just include a gender dummy for simplicity.

[7] Ideally, we would like to have reported *uniform* confidence bands that, under standard frequentist assumptions, would cover the entire graph (n, APE_{nm}) with 99% probability in repeated samples. For a finite number of bins, this could be accomplished by a Bonferroni correction, although the band would be very conservative for a modest number of bins. We are not aware of asymptotic theory providing tight uniform confidence bands in the present setting.

[8] As the standard errors are not uniform, we are technically not able to compare the magnitude of APEs across bins. One way to make formally sound comparative analysis would be to restrict attention to two bins of interest (corresponding to high and low weight, respectively) and test the joint hypothesis that both estimates equal zero against a one-sided alternative. The critical values for the two sub-hypotheses would have to be Bonferroni-adjusted.

Figure 2: CRE Probit Model, APE of Smoking

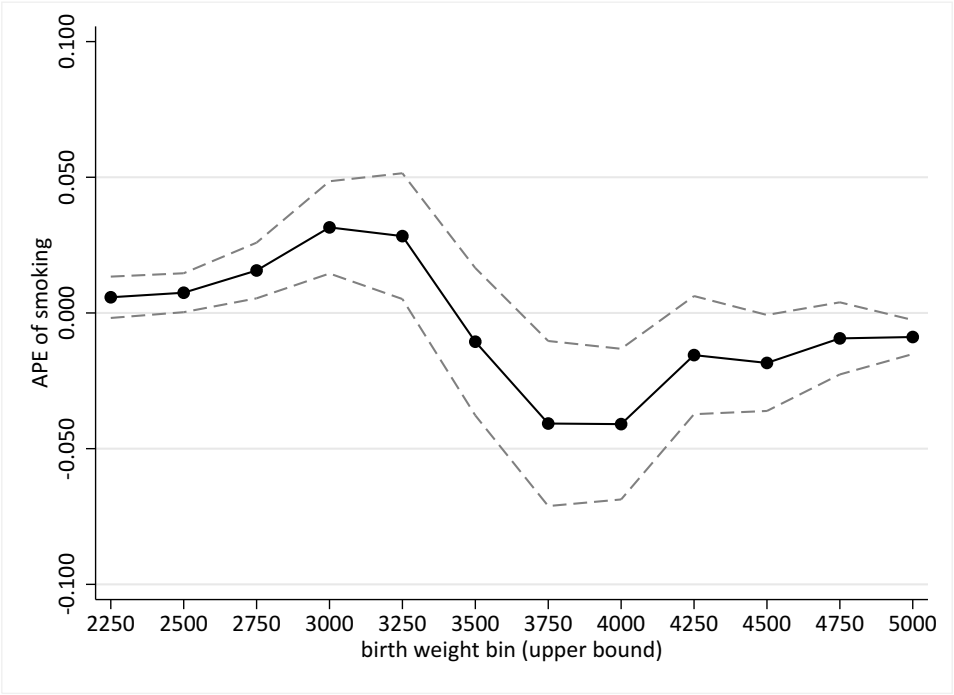

Figure 3: CRE Probit Model, APE of Heavy Smoking

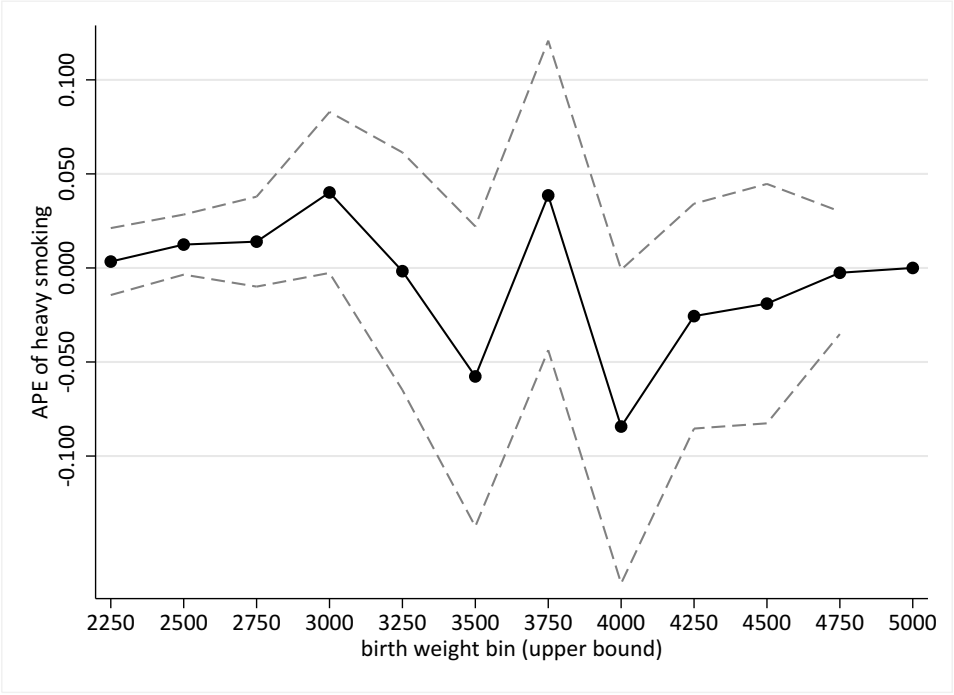

APE of the smoking dummy is small, statistically and economically, in the extreme bins (although the point estimates of the APEs have intuitive signs). It could be the case that only

heavy cigarette consumption substantially affects the tails of the birth weight distribution. Figure 3 graphically depict the APEs for the heavy smoking dummy (the results for the 12th bin $[4750, 5000)$ are not reported since there are no heavy smokers in this subsample). While some of the point estimates are large in magnitude, the standard error bands contain 0 for every bin. Thus, there is no statistical evidence that the *intensity* of smoking, measured in this way, has an additional effect on the relative probabilities of various birth weights.

5 Discussion

5.1 Limitations to Fixed Effects

The consistency of fixed effects estimates relies crucially on the assumption of strict exogeneity; formally, that

$$E[\epsilon_{ib}|X_i, s_i] = 0.$$

However, given that smoking is a choice made in a dynamic framework, this assumption may fail in two ways. The first is the feedback effect, which has been extensively addressed above. Another potential violation is a correlation between smoking choices and other unobserved behaviors that impact birth weight (such as alcohol or caffeine consumption, or exercise habits). If for example the first-period birthweight is low, the mother may quit drinking, which is tantamount to the second-birth error being correlated with first-birth outcomes. This means that even though our fixed effect estimate is of a lower magnitude than OLS estimate, we may still be overestimating the effects due to possible violation of the strict exogeneity condition.

5.2 Direct dependence on lagged outcomes

We have not explored dynamic specifications where lagged birth outcomes appear explicitly in the reduced-form equation. There are reasons to believe that birthweights have an au-

toregressive structure due to time-varying unobserved biological factors. If this is the case, the time-invariant unobserved heterogeneity will not be able to fully correct for the misspecification (although if biological determinants are largely genetic, this is less of a concern). In our panel IV model we experimented with adding lagged birthweights and/or smoking status explicitly as controls. We did not find any conclusive evidence that such a model extension was worthwhile, although this is certainly an interesting topic of future research.

5.3 Conclusions and future work

Integrating across our specifications, we have found evidence that cross-sectional OLS estimates are likely biased, and that some of this bias may be addressed using variation across births in a panel data setting. Our headline partial effect of smoking on birthweight is -160 in the fixed effects set-up, consistent with Abrevaya (woo6). We proposed a model to think about the potentially serious feedback effect and found evidence that the effect is present in our data set. A panel IV approach was unsuccessful in yielding sharp inference but we could not reject our basic model. Furthermore, experimentation with the instruments led to the tentative conclusion that a time-invariant heterogeneity term may not be sufficient to capture the endogeneity. Future work should try to uncover better instruments, such as cigarette tax changes. Our probit regressions confirmed our cross-sectional conclusion that smoking significantly affects the probability of high or low birthweights for moderately abnormal weights, but we did not find a significant effect for extreme birthweights. Since the latter are most relevant for cost/benefit analyses (and infant welfare), future work should focus on tail behavior. Perhaps a quantile regression approach adapted to panel data would deliver other interesting insights.

References

Abrevaya, J. (2006): "Estimating the effect of smoking on birth outcomes using a matched panel data approach," *Journal of Applied Econometrics*, 21, 489–519.

Black, S., P. Devereux, and K. Salvanes (2007): "From the cradle to the labor market? The effect of birth weight on adult outcomes," *Quarterly Journal of Economics*.

WN, E. and R. JS. (1999): "Can higher cigarette taxes improve bith outcomes?" *Journal of Public Economics*.

Table 1
Regressions of birth weight (grams) on smoking

	OLS		Fixed Effects			
	(1)	(2)	(3)	(4)	(5)	(6)
Smoker	-274.30***	-209.43***	-165.04***	-163.72***	-122.54***	-83.54**
	(10.37)	(22.10)	(17.24)	(17.21)	(22.84)	(26.94)
Moderate Smoker		-85.37**			-46.36	
		(27.67)			(25.22)	
Heavy Smoker		-97.34***			-88.85**	
		(24.12)			(28.47)	
Cigarettes Per Day						-8.53**
						(2.85)
Cigarettes Squared						0.09
						(0.07)
Male Child	134.81***	134.77***	136.64***	136.05***	136.17***	136.27***
	(5.44)	(5.41)	(5.18)	(5.16)	(5.16)	(5.16)
Prenatal Care	Yes	Yes	Yes	Yes	Yes	Yes
Marriage & Race	Yes	Yes	No	No	No	No
Education	Yes	Yes	No	No	No	No
Age, Age Squared	Yes	Yes	Yes	No	No	No
Dummies for Years Since Last Birth	No	No	No	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	No	No	No	No
Observations	36600	36600	36600	36600	36600	36600
Adjusted R^2	0.110	0.111	0.045	0.049	0.050	0.050
BIC	559598	559594	527489	527391	527391	527378

Regression coefficients are estimated using OLS in models (1) and (2), and using Fixed Effects in models (3), (4), (5) and (6). Heteroskedasticity-robust standard errors are reported in bracket and are clustered by state of residence in OLS models, and by mother in FE models. Asterisks denote statistical significance at the 0.1% (***), 1% (**), or 5% (*) level.