The Effect of Smoking on Obesity: Evidence from a Randomized Trial*

Charles Courtemanche†, Rusty Tchernis‡, and Benjamin Ukert§

April 26, 2016

Abstract

This paper aims to identify the causal effect of smoking on body mass index (BMI) using data from the Lung Health Study, a randomized trial of smoking cessation treatments. Since nicotine is a metabolic stimulant and appetite suppressant, quitting or reducing smoking could lead to weight gain. Using randomized treatment assignment to instrument for smoking, we estimate that quitting smoking leads to an average long-run weight gain of 1.8-1.9 BMI units, or 11-12 pounds at the average height. These results imply that the drop in smoking in recent decades explains 14% of the concurrent rise in obesity. Semi-parametric models provide evidence of a diminishing marginal effect of smoking on BMI, while subsample regressions show that the impact is largest for younger individuals, females, those with no college degree, and those with healthy baseline BMI levels.

JEL code I10, I12, C11

*We thank Katherine Flegal, Jason Fletcher, Sara Markowitz, Tom Mroz, and audiences at Georgia State University, Indiana University-Purdue University-Indianapolis, the University of Alabama at Birmingham, and the Southern Economic Association Annual Meeting for helpful feedback.

†Address for correspondence: Charles Courtemanche, Department of Economics, Andrew Young School of Policy Studies, Georgia State University, P.O. Box 3992, Atlanta, GA 30302-3992, e-mail: ccourtemanche@gsu.edu. Department of Economics, Andrew Young School of Policy Studies, Georgia State University and National Bureau of Economic Research.

‡Department of Economics, Andrew Young School of Policy Studies, Georgia State University, National Bureau of Economic Research, and Institute for the Study of Labor (IZA).

§Department of Economics, Andrew Young School of Policy Studies, Georgia State University.

1 Introduction

In the last 40 years obesity[1] rates have steadily increased in the United States (US), rising from 13% in the early 1960s to 35% in 2011-2012 (Flegal et al., 1998, Ogden et al., 2014). This rise in obesity has contributed significantly to increasing rates of diabetes, heart disease, and stroke (Mokdad et al. 2001, Manson et al. 1990, Rexrode et al. 1997), with Flegal et al. (2005) finding that obesity-related diseases lead to 112,000 deaths per year. Wang et al. (2011) project that by 2030 the number of obese adults in the US will grow by another 65 million. Cawley and Meyerhoefer (2012) estimate that obesity leads to $190 billion per year in medical expenses, while Wang et al. (2011) project that this number will increase by $48 to $66 billion by 2030.

During the same time frame, the percentage of adults who use tobacco in the US declined from 42% to 19% (National Center for Health Statistics, 2011). The 1964 Surgeon General's Report concluded that smoking leads to adverse health conditions such as lung cancer and heart disease and increases mortality risk (US Department of Health and Human Services, 1964). Subsequently, federal and state governments launched an aggressive tobacco control campaign featuring advertising restrictions, warning labels, information-spreading programs, cigarette taxes, and smoking bans in public places. Despite the success of these efforts in reducing smoking, tobacco is still responsible for one out of every five deaths in the United States and at least $130 billion per year in medical expenses (US Department of Health and Human Services, 2014).

The inverse trends in smoking and obesity raise the question of whether they are causally related, in which case tobacco control policies may had the unintended consequence of contributing to the rise in obesity. Quitting or reducing smoking could increase body weight since nicotine can act as an appetite suppressant and metabolic stimulant (Pinkowish, 1999). In standard economic models of body weight (e.g. Philipson and Posner, 1999), nicotine's

[1]Obesity is defined as having a body mass index (BMI) greater than 30, where BMI is equal to weight in kilograms divided by height in meters squared.

appetite-suppressing properties could be seen as decreasing the marginal utility of food consumption, leading to less eating and therefore lower body weight. Stimulating the metabolism would mean more calories burned holding physical activity constant, again reducing weight. On the other hand, smoking reduces lung capacity (Hedenstrom et al., 1986), which could lead to weight gain by increasing the marginal disutility from exercise.

A large public health literature documents that individuals tend to gain weight following smoking cessation. A meta-analysis of 15 studies found that individuals who quit smoking gain an average of four pounds more than a comparison group of continued smokers (U.S. Department of Health and Human Services, 1990). Another meta-analysis, which included a larger number of studies (62) but no comparison group, found that weight gain steadily increases in the year after smoking cessation, ultimately reaching 4.67 kg (10.3 lbs) (Aubin et al., 2012). Evidence regarding longer-run effects is mixed. Some studies have found that much of the weight gain after quitting smoking is temporary (Chen et al., 1993; Mizoue et al., 1998), but others conclude that the effect remains sizeable five to ten years after cessation (Flegal et al., 1995; Travier et al., 2012; Williamson et al., 1991). These associational estimates could be susceptible to bias from unobservable characteristics, such as time preference and level of interest in one's health, that likely influence both the probability of smoking cessation and weight trajectory. Additionally, studies that simply track changes in weight without utilizing a comparison group of continued smokers are susceptible to additional confounding from the tendency to gain weight with age.

The economics literature has attempted to move closer toward causality by examining the effects of plausibly exogenous sources of variation in economic factors that influence smoking on BMI. The results from this literature are mixed. Chou et al. (2004), Rashad et al. (2006), and Baum (2009) estimate positive relationships between cigarette costs and BMI. Since higher cigarette prices have been shown to reduce smoking, these results are consistent with reduced smoking leading to weight gain. However, Gruber and Frakes (2006), Courtemanche (2009), and Wehby and Courtemanche (2012) estimate the effect of cigarette

costs on BMI to actually be negative, while Nonnemaker et al. (2009) and Courtemanche et al. (forthcoming) find little evidence of an effect in either direction. The discrepancies in observed results in the literature hinge on methodological issues such as whether cigarette prices or tax rates are used as the measure of cigarette costs, whether time is modeled using a quadratic trend or time period dummies, and whether the difference between short-run and long-run effects is considered. A particularly controversial issue is whether cigarette costs can actually be considered exogenous. Cigarette prices may depend on the demand for cigarettes, while high cigarette taxes may be more politically palatable in states where a relatively small percentage of the population smokes. Fletcher (2014) considers a different tobacco-control policy – workplace smoking bans – and finds evidence that smoking cessation induced by these bans increases BMI.

To our knowledge, the only paper that uses a randomized intervention to estimate the causal effect of smoking on weight is Eisenberg and Quinn (2006; hereafter EQ). EQ use the Lung Health Study (LHS), which randomly assigned smokers to a comprehensive smoking cessation program and then tracked their health for five years. EQ do not actually use the LHS microdata, but instead take advantage of the fact that O'Hara et al. (1998) report differences between the treatment and control groups' average changes in weight and smoking status to compute a Wald instrumental variables (IV) estimate of the weight gain from quitting smoking. EQ find that quitting smoking leads to a very large average weight gain of 9.7 kg (21.4 pounds), about two to five times the magnitude typically found in the associational literature.

Despite their use of randomization, there are reasons to suspect that EQ's estimate is overstated. They use assignment into the smoking cessation program as an instrument for "sustained quitting", which is defined as being a medically verified quitter in all five LHS follow-up waves. The IV strategy therefore requires the strong assumption that the program only affects weight through its effect on sustained quitting. However, in addition to helping some smokers quit immediately and permanently, the program may help others quit smoking

in some but not all follow-up periods, and still others by reducing their number of cigarettes smoked per day even though they never quit entirely. To the extent that such partial quitting exists and influences BMI, EQ's estimated average effect of quitting smoking will be biased upwards.

We contribute to the literature on the effect of smoking on weight in several ways. First, we provide, in our view, the most reliable estimates to date of the average causal effect of quitting smoking on weight. We use the LHS microdata to exploit the randomized nature of the study while also constructing detailed smoking measures – such as cigarettes smoked per day and average carbon monoxide (CO) level over the entire five-year period of the study – that account for delayed or temporary quitting as well as smoking intensity. Our preferred estimates imply that quitting smoking leads to an average weight gain of 1.5-1.7 BMI units (10-11 pounds at the average height) at the end of the first year of the study period. The effect persists over time, reaching 1.7-1.9 BMI units (11-12 pounds) by the end of the fifth year – a magnitude that implies that the fall in smoking explains around 14% of the rise in BMI in recent decades. Our estimated effects are toward the high end of the range of results from the associational literature, but substantially smaller than EQ's estimate that uses randomization but relies on the "sustained quitting" measure.

Our paper also contributes by providing new information related to the heterogeneity of the effect across the smoking and weight distributions as well as by demographic characteristics. We estimate a semi-parametric instrumental variables model that allows the data to determine the functional form of the relationship between smoking and BMI. The results suggest a diminishing marginal effect, with additional smoking having little long-run impact on BMI beyond about a pack of cigarettes per day or a CO level of about 20 parts per million (ppm). We also conduct subsample analyses by age, gender, education, and baseline BMI and find that on average younger individuals, females, those with no college degree, and those with healthy baseline BMI levels gain the most weight in response to smoking cessation.

2 Data

This section provides a brief introduction to the LHS, with an emphasis on the information most relevant for our paper. O'Hara et al. (1993, 1998) provide a more detailed discussion of the LHS, and further information is also available online at https://www.clinicaltrials.gov.

The purpose of the LHS was to observe changes in the severity of chronic obstructive pulmonary disease (COPD) among smokers. The study consisted of 5887 smokers with initial ages between 35 and 59. Recruitment started in 1986 and ended in 1989. The clinical trial ended in 1994. To be eligible for selection, potential participants had to show signs of mild lung function impairment, have no history of certain medications, consume less than 25 drinks per week, and have no severe illnesses or chronic medical conditions. Each year all participants were extensively interviewed individually at a medical clinic near the residence of the participant (no more than 75 miles away from the participant's permanent residence). The data therefore consist of the baseline period (1989) plus five annual follow-up periods (1990 through 1994). Attrition was relatively low, as 5,297 individuals remained in the sample in the final wave. The attriters included 315 participants who died during the study period.

Participants were randomly assigned into three different groups: two treatment groups and one control group. Both treatment groups received a special intervention (SI) consisting of free nicotine gum, an intensive quit week, and frequent contact with support personnel with invitations to bring a spouse or relative to the meetings. The only difference between the two treatment groups is that, in addition to the SI, one group received an inhaled bronchodilator (SI-A) while the other received an inhaled placebo (SI-P). Most of the intensive intervention treatments were completed within the first 4 months of the study. The control group referred to as the usual care (UC) group received no intervention and members continued to use their own private sources for medical care.

The LHS collected information about weight, height, smoking behavior, family smoking habits, health status, and demographic characteristics. Weight and height were measured by

medical staff at the participants' clinic visits, so our BMI measure is not susceptible to the concern about measurement error that is common in the economics of obesity literature.[2] The data contain self-reported smoking information as well as CO test results. We consider three different measures of smoking: a dummy variable for whether the respondent currently smokes (clinically measured through the CO test), number of cigarettes typically smoked per day (self reported), and CO level in ppm. We also utilize the LHS' information on education (dummies for high school graduate, some college, and college graduate), gender (dummy for male), age (years), and marital status (dummy for married) as controls. Note that we do not control for race/ethnicity because 97% of LHS participants were white. Our sample is therefore not representative of the overall population of US smokers along this dimension.

Table 1 presents descriptive statistics for the three groups at the time of randomization. Average cigarette consumption was roughly 30 cigarettes per day, average CO level was about 26, and the average respondent was just slightly overweight. The summary statistics for all variables are very similar across the three groups, indicating the randomization was successful.

Figure 1 displays changes throughout the sample period in the average number of cigarettes smoked per day, objectively-verified smoking status, CO level, and BMI for each group. Sharp decreases in cigarette smoking, smoking status and CO level are evident for both treatment groups in the first year after the intervention. The decrease in smoking for the control group is much more moderate. Average BMI is trending upward for all three groups, but the two treatment groups experience much sharper increases in BMI than the control group in the first year. The graph therefore suggests both that the intervention was effective in reducing smoking and that smoking reduces BMI. We next use econometric methods to estimate the magnitude of these effects.

[2]See Courtemanche et al. (2015) for an overview of the challenges involved with using self-reported weight and height.

3 Econometric Analyses

Our econometric objectives are to identify 1) the average short- and long-run causal effects of quitting smoking on weight gain, 2) how the effect of smoking on weight changes across the smoking distribution, and 3) how the effect of smoking on weight varies by demographic characteristics and baseline BMI. We begin by using parametric regressions to answer the fist question and running falsification tests to evaluate the validity of the models. We then conduct semiparametric estimation allowing for a flexible relationship between smoking and weight to address the second question. Next, we answer the third question with subsample analyses. Finally, we address issues related to the generalizability of the results and show why our results differ from those of EQ.

3.1 Average Effects

3.1.1 Short Run

We begin by aiming to identify the average causal effect of quitting smoking on weight gain with a series of parametric regressions. Economists typically consider body weight to be a capital stock. Individuals start with an exogenous endowment of weight that changes over time due to depreciation as well as "investments" that take the form of caloric intake or expenditure. In the LHS, weight at the end of the first follow-up year can therefore be modeled as a function of weight at baseline and investments – such as smoking – in year one. This leads to the model

$$bmi_{i1} = \beta_0 + \beta_1 bmi_{i0} + \beta_2 S_{i1} + \beta_{3t}\mathbf{X}_i + \varepsilon_{i1} \tag{1}$$

where bmi_{i1} is individual i's BMI at the end of year 1, bmi_{i0} is BMI at the beginning of the study (year 0), S_{i1} is smoking in year 1, $\mathbf{X}_i$ is a vector of demographic controls that are assumed to be constant over time since they are only available for the baseline wave, and

ε_{i1} is period 1's error term. β_2, the coefficient of interest, gives the short-run association between smoking and weight. We estimate the OLS model given by (1) as well as an IV model that uses the randomized treatment assignment to instrument for S_i. The first stage of the IV model is given by

$$S_{i1} = \gamma_0 + \gamma_1 w_{i0} + \gamma_2 si_a_i + \gamma_3 si_p_i + \gamma_4 \mathbf{X}_i + \mu_{i1} \tag{2}$$

where si_a_i and si_p_i reflect whether the individual was assigned into the SI-A or SI-P treatment group, respectively. The second stage of the IV model is identical to (1) except it replaces S_{i1} with the predicted value generated by (2). In the IV model, β_2 can be interpreted as the short-run local average treatment effect (LATE) of intervention-induced changes in smoking on BMI. We estimate linear models in both stages due to their relative ease of interpretation, their ability to produce reliable average effects (e.g. Angrist and Pischke, 2009), and the inherent difficulties with non-linear IV estimation (e.g. Terza et al., 2008). We define S_{i1} three different ways: a dummy for smoking cessation, number of cigarettes smoked per day, and CO level. We next discuss these three smoking variables.

The first smoking measure is a dummy equal to one if and only if individual i was a medically validated non-smoker at the end of year one. We consider this to be a naive measure of smoking because it ignores variation in smoking intensity among smokers. This could lead to an overstatement of the average weight gain from quitting smoking estimated by IV models. When a quit dummy is used as the smoking measure, the IV estimator effectively scales the difference in BMI between the treatment and control groups by the difference in smoking cessation rates between the two groups. The validity of this estimator therefore hinges on the assumption that the randomized intervention only affected the BMIs of people who fully quit smoking. To the extent that the intervention also affected the BMIs of those who cut back on smoking but did not quit entirely, the difference in BMI will be scaled by too small a number and the resulting IV estimate will be too large. (This point is

central to our critique of EQ, and we will discuss it in more detail in Section 3.5.)

Our second smoking measure is therefore self-reported number of cigarettes smoked per day, with a value of zero assigned to those who reported quitting. This measure incorporates both reducing smoking and quitting entirely and therefore is not susceptible to the above criticism. In order to make the results using cigarettes per day comparable to $\widehat{\beta}_2$ from the regressions using the smoking cessation dummy, we need to compute an implied average weight gain from quitting smoking. Since the best way to do this is not immediately obvious, we consider several possibilities. First, we use the average weight that would be gained if all individuals in the sample switch from their baseline number of cigarettes to none. We do this by multiplying the coefficient estimate on the cigarettes smoked per day variable by each individual's number of cigarettes smoked at baseline, and then taking the average across all individuals. Formally, this means we compute $\left(\sum_{i=1}^{N}\beta_2 cigday_{i0}\right)/N$, where $cigday$ is cigarettes smoked per day and i indexes the N observations. Second, we use the same formula but average over only those individuals who actually quit smoking. Third, we only average over quitters from the treatment groups. Fourth, we compute the effect of quitting smoking from the sample mean baseline smoking level, i.e. $\beta_2\overline{cigday}_{i0}$. Fifth and sixth, we again compute $\beta_2\overline{cigday}_{i0}$ but using only quitters, and quitters from the treatment group, respectively. The estimated average weight gain from quitting smoking is similar using all six approaches, so we only report the results using the first method: averaging the predicted effects across all individuals. Results from the other approaches are available upon request.

Key limitation with cigarettes per day is its self-reported nature. At issue for the validity of our IV estimates is not whether cigarettes per day are reported with error, but whether this error is correlated with treatment status. It is not obvious that this is the case, but it is possible that, for instance, being assigned into the treatment group creates more pressure to report progress toward smoking cessation, leading to differentially large reporting error among the treatment group. Alternatively, perhaps reporting error simply rises with number of cigarettes smoked per day, in which case we would expect the amount of error in the follow-

up periods to be highest among the control group.

We therefore also utilize a third smoking variable that is both clinically measured and incorporates both the intensive and extensive margins of smoking: CO level from a test conducted during the follow-up interview. Using the CO regression estimates to compute the average weight gain from quitting smoking is somewhat more complicated than using cigarettes smoked per day since even non-smokers generally have a positive CO level. We therefore compute each individual's predicted effect of quitting smoking as the effect of switching from her baseline CO level to the mean CO level for non-smokers, rather than to a CO level of zero. For the mean CO level of non-smokers, we use Deveci et al.'s (2004) estimate of 3.61 ppm; this is similar to the mean CO level of verified non-smokers in the follow-up waves of the LHS. The average effect of quitting smoking on weight across the entire sample is therefore given by $\left(\sum_{i=1}^{N}\beta_2(CO_{i0}-3.61)\right)/N$ where CO is CO level in ppm. Note that CO levels are only available at baseline for 922 individuals, so our average effect is computed using only this portion of the sample (though our regressions still utilize the full sample). We doubt that this limitation is of consequence since reported numbers of cigarettes smoked per day at baseline are virtually identical for those with missing baseline CO levels and those with non-missing levels. We have also considered analogs of the other five approaches to computing the average effects of quitting smoking discussed above and verified that, as with cigarettes per day, the results are robust.

While using CO levels solves the probability of reporting error, it should be noted that it is not immune to all sources of measurement error. In particular, it only reflects smoking in the past couple of days. Therefore, for some people self-reported number of cigarettes smoked per day could actually be more indicative of typical smoking behavior than clinically measured CO. Consequently, we take an agnostic view about which measure is preferred and present the results for both alongside each other throughout the paper.[3]

[3]If we regress CO level in the first follow-up year on cigarettes smoked per day as well as the interaction of cigarettes with the two treatment dummies, the R-squared is 0.48. This suggests that, while cigarettes smoked per day and CO are highly correlated, they do convey different information. Additionally, the coefficients on the interaction terms are positive and significant, though small. In other words, measurement

3.1.2 Long Run

The above specifications estimate the short-run causal effect of smoking on BMI. We also aim to identify the long-run effect by asking how smoking across all five follow-up waves affects BMI at the end of the study (year five). Comparing the short- and long-run effects is important since evidence from the associational public health literature suggests that at least some of the weight gained after quitting smoking is temporary (Chen et al., 1993; Mizoue et al., 1998). Ideally, we would like to estimate

$$bmi_{i5} = \beta_0 + \beta_1 bmi_{i0} + \sum_{t=1}^{5} \beta_{2t} S_{it} + \beta_3 \mathbf{X}_i + \varepsilon_{i5} \tag{3}$$

where bmi_{i5} is individual i's BMI at the end of year 5 and S_{it} is smoking in year t. However, the need to utilize IV estimation prevents us from allowing separate coefficients for each of the five smoking variables, as this would require five instruments. In other words, in order to operationalize an IV model we need to compress the five years of smoking information into a single variable S_i. The easiest way to do this is to take a simple average across the five years:

$$S_i = \frac{S_{i1} + S_{i2} + S_{i3} + S_{i4} + S_{i5}}{5}. \tag{4}$$

However, this approach assumes that smoking in each of the five periods has the same effect on weight. To the extent that weight is a depreciating capital stock, we might expect smoking in more recent years to have a larger effect on BMI than smoking in more distant years. We therefore also estimate models defining S as a weighted rather than simple average of quit status in the five follow-up years:

$$S_i = \frac{S_{i1} + (1-\delta)S_{i2} + (1-\delta)^2 S_{i3} + (1-\delta)^3 S_{i4} + (1-\delta)^4 S_{i5}}{1 + (1-\delta) + (1-\delta)^2 + (1-\delta)^3 + (1-\delta)^4}. \tag{5}$$

Since we do not have a sufficient number of instruments to credibly estimate the depreciation

error does appear to be slightly correlated with treatment status. Both of these results underscore the importance of verifying that the results are similar using the two different measures.

rate δ, we simply try several plausible values: 0.05, 0.1, 0.15, 0.2, 0.25, and 0.3. In all our regressions, the coefficient estimate on the baseline BMI variable will barely be below one, so we consider it probable that the "true" value of δ is toward the low end of this range; i.e. there is little reason to consider values of δ above 0.3.

We estimate equation (3) using both OLS and IV, with si_a_i and si_p_i again serving as the instruments. We again use the three different measures of S_{it}: smoking cessation, cigarettes smoked per day, and CO level. For the cigarettes per day and CO regressions, we compute implied average effects of quitting smoking in the same manner as the short-run specifications.

3.1.3 Results

Table 2 reports the results of interest from the parametric regressions. Panel A presents the OLS and IV estimates of the effects of the different smoking measures on BMI. Panel B shows the estimated effects of the treatment dummies on the smoking variables from the first stage of the IV models, along with the F-statistic from a test of their joint significance. The first three columns show the effect of year 1 smoking on year 1 BMI (short-run effect), with the first column using the binary quitting variable as the smoking measure, the second using cigarettes smoked per day, and the third using CO. The last three columns present the effects of the simple averages of these three smoking measures across years 1-5 on BMI in year 5 (long-run effect). The results using the weighted averages, available in Appendix Table A1, are similar to those obtained using simple averages.

Coefficient estimates are shown in the table, with heteroskedasticity-robust standard errors in parentheses. The stars represent 0.1%, 1%, and 5% significance levels. For the regressions with the non-binary smoking measures cigarettes per day and CO, the implied average effects of quitting smoking on BMI are presented in brackets. In other words, the numbers in brackets from the cigarettes and CO regressions are comparable to the coefficient estimates from the quit status regressions. The row labeled "Hausman" gives the p-values

from Hausman tests of the consistency of the OLS estimator compared to IV. The sample sizes, provided in the row labeled "N", vary somewhat across specifications due to differing amounts of missing information. In unreported regressions (available upon request), we re-estimated the models using only observations with no missing smoking information and verified that any meaningful differences between the results cannot simply be attributed to the difference in samples.

The first column presents the short-run estimates using the quit dummy. The OLS regression estimates that quitting smoking increases BMI by 1.295 units, or 8.2 pounds at the US average height of 66.55 inches.[4] This is well within the range of estimates from the associational literature discussed in Section 1. The IV estimate is a larger 2.202 BMI units, and the Hausman test strongly rejects the consistency of OLS. This IV estimate equates to 13.9 pounds, which is larger than most estimates of the average short-run weight gain from quitting smoking from the associational literature.

The next two columns use the smoking measures that incorporate intensity: cigarettes per day and CO. In the IV specifications, we estimate that in the short run an additional cigarette smoked per day reduces BMI by 0.052 units while an additional ppm of CO reduces BMI by 0.077 units. The average effects of quitting smoking implied by these two regressions are 1.52 and 1.71 BMI units, which translate to 9.6 and 10.8 pounds at the mean height. These estimates are 31% and 22% smaller than the 13.9 pounds we obtained using the quit dummy. This is consistent with our prediction that neglecting to account for smoking intensity leads to an exaggerated IV estimate of the average weight gain from quitting smoking.

The last three columns turn to the long-run estimates. The key result is that the long-run effects are slightly stronger than the short-run effects. This is an important result, as the issue of whether the effect diminishes over time has been a point of contention in the associational literature, as discussed in Section 1. In the IV specification using average quit status, quitting for all five follow-up years is estimated to increase BMI by 2.646 units, or

[4]Average height is computed by taking a simple average of the male and female heights given by http://www.cdc.gov/nchs/fastats/body-measurements.htm.

16.7 pounds. An additional cigarette smoked per day over the five years reduces BMI by 0.065 units, while an additional ppm of average CO reduces BMI by 0.082 units. These latter two estimates imply average weight gains from quitting smoking of 1.91 and 1.81 units of BMI, or 12.0 and 11.4 pounds. As with the short-run estimates, these results suggest that incorporating smoking intensity is necessary to avoid overstating the magnitude of the weight gain from smoking cessation.

We can use these long-run results to estimate the percentage of the rise in BMI that can be attributed to falling smoking, under the admittedly strong assumption that the results generalize. This percentage is given by $\frac{dbmi}{dcigday}\frac{\Delta\overline{cigday}}{\Delta\overline{bmi}} * 100\%$. For $\frac{dbmi}{dcigday}$, we use the long-run IV estimate for cigarettes smoked per day: -0.065. $\Delta\overline{cigday}$ and $\Delta\overline{bmi}$ are the changes in the population means of cigarettes smoked per day and BMI among those at least eighteen years old. We compute these using the oldest and newest waves of the National Health and Nutrition Examination Survey (NHANES) that contain data on both smoking (self-reported) and BMI (medically measured): 1971-1974 (NHANES I) and 2011-2012. This period spans the entirety of the sharp rise in obesity, which did not begin until the late 1970s. During this time frame, average cigarettes smoked per day fell from 9.165 to 2.188, so $\Delta\overline{cigday} = 6.977$. Average BMI rose from 25.425 to 28.617, so $\Delta\overline{bmi} = 3.192$. Plugging in these numbers suggests that the drop in smoking explains 14.2% of the rise in BMI. We view this as a relatively substantial contribution to the trend. Courtemanche et al. (2015) examine the extent to which 27 different economic factors contributed to the rise in BMI, finding that the increased prevalence of big box grocers and restaurants explain 17% and 12%, respectively, while no other factors explain more than 4%.

Finally, the first-stage estimates in Panel B of Table 2 show that the treatment was effective in reducing smoking. In the short run, being assigned into the SI-A or SI-P groups increased the probability of quitting by 27-28 percentage points while decreasing cigarettes smoked per day by 11-12 and CO level by 8 ppm. In the long run, SI-A or SI-P assignment increased the fraction of the five follow-up years quit by 0.21 while decreasing average cigarettes

per day by 9 and average CO by 6-7 ppm. The treatment variables are all highly significant in the first stage and the F-statistics from the test of their joint significance are easily large enough to conclude that they are sufficiently strong instruments. Also noteworthy is the fact that there is essentially no difference in the coefficient estimates for the two treatment variables; in other words, the inhaled bronchodilator given to the SI-A group did not influence smoking. This also means that, though our IV model is technically overidentified, the instruments are not sufficiently distinct to make an overidentification test informative or to consider instrumenting for two endogenous variables.

3.1.4 Falsification Tests

We next conduct falsification tests to support our contention that the IV results using cigarettes per day and CO are more credible than those using the quit dummy. In our IV models, the identifying assumption is that the randomized treatment only influences BMI via the smoking variable. Our falsification tests evaluate this assumption by asking whether the instruments influence the BMIs of individuals who did not change their smoking habits during the sample period, according to each smoking measure. A significant association would provide evidence that the randomized intervention influenced BMI through pathways other than the particular smoking measure, invalidating the causal interpretation of the corresponding IV estimate. For the quit measure, we restrict the sample to those with values of 0, meaning those who smoked in the first follow-up wave for the short-run analysis and those who smoked in all five follow-up waves for the long-run analysis. For the continuous smoking measures, there are obviously very few individuals with literally no change in smoking across the sample period (i.e. the exact same number of cigarettes smoked per day/CO level in the baseline period as in the follow-up waves), so a judgment call is required as to what magnitude change in smoking should be considered "meaningful". We report results restricting the sample to those whose post-treatment level of smoking is within 25% of their baseline level; results using neighboring cutoffs are similar. Using these subsamples, we estimate the

reduced-form version of the short- and long-run IV models; i.e. we regress BMI on the two treatment variables plus the controls.

Ex ante, our prediction is that the binary quit measure will perform the worst in the falsification tests since it leaves people in the sample who did actually experience a meaningful change in smoking but did not quit entirely. Note that it is not obvious that the falsification tests will produce null results even for the smoking measures that incorporate intensity, though, since it is conceivable that the treatment could affect BMI through pathways other than smoking. For instance, perhaps being exposed to an intensive health-related intervention might increase some people's level of general health consciousness, which could lead to improved health behaviors along other dimensions besides smoking. The falsification tests are therefore important in assessing whether even our most conservative estimates of the effect of smoking on weight can be given a causal interpretation.

Table 3 reports the results. The left half of the table presents the results from the short-run falsification tests ("effects" of the treatment dummies on year 1 BMI for those with unchanged smoking status) while the right half shows the long-run results (year 5 BMI). For comparison purposes, the first column of each half of the table presents the reduced-form results for the full sample. The remaining three columns of each half include those with no meaningful changes in the quit, cigarettes per day, and CO variables, respectively. The sample sizes in the tests based on cigarettes per day are smaller than those using quit status simply because more individuals are excluded as the measure of smoking becomes more comprehensive. The sample sizes in the two columns using CO are very small because, as discussed previously, much of the sample is missing baseline CO information, preventing the calculation of the percentage change. The falsification tests using cigarettes per day are therefore much more highly powered – and consequently more informative – than those using CO.

The columns labeled "full sample" show that, in both the short and long run, the reduced-form effects of the two treatment variables on BMI are between 0.54 and 0.61 before excluding

any observations. Dropping those who quit smoking reduces the magnitude of these effects by about half, but significant effects of 0.21-0.28 remain. There is therefore clear evidence that the intervention affected BMI through a pathway besides quitting smoking, implying that the IV estimates using quitting smoking are too large. The falsification test results are much more favorable if we also exclude those with meaningful (>25%) changes in smoking intensity. In the two regressions that use cigarettes per day, the coefficient estimates for the treatment variables are small (between -0.014 and 0.091) and highly statistically insignificant. There is therefore no evidence that the exclusion restriction in the IV model is violated if *cigday* is used as the smoking measure. Excluding on the basis of changes in CO also leads to highly insignificant effects, with three being negative (the opposite direction of the full-sample relationship) and one positive. The estimates are imprecise due to the small sample size, so these results are not as compelling as those using *cigday*, but the lack of a clear pattern is at least somewhat reassuring. To summarize, the results in this section suggest that while the IV results using the quitting indicator are contaminated by an alternative effect, the results using both number of cigarettes and CO levels are more reliable.

3.2 Semi-Parametric Estimation

An issue with the parametric regressions for *cigday* and *CO* is that they assume that smoking intensity affects BMI linearly. This is a strong assumption, as it seems likely that there is either a non-linear dose-response effect of nicotine on metabolism/appetite or a non-linear effect of metabolism/appetite on weight-related behaviors. While it is not clear that this will bias estimates of the average weight gain from quitting smoking, such a restrictive functional form is likely to lead to systematically inappropriate predictions for at least some individuals. Moreover, given the complicated chain of biological and behavioral pathways through which smoking influences BMI, the nature of the non-linearity is not clear *ex ante*. In other words, it is not obvious that the non-linearity could be captured through simple approaches such as logarithmic or quadratic specifications. We therefore next estimate a semi-parametric

model that allows the data to determine the functional form of the relationship between smoking and BMI. Specifically, we implement Robinson's (1988) semi-parametric double residual estimator with local smoothing. This approach allows us to model the expectation of the dependent variable at every point on the distribution of the independent variable, thereby enabling the prediction of the weight gained (or lost) from switching from any level of smoking to any other level.[5]

Semi-parametric IV models can be estimated using a control function approach (Blundell and Powell, 2004; Lee, 2007). The first stage takes the same form as equation (2). The second stage differs from equation (1) in two ways. First, it does not specify the functional form for the smoking measure. Second, rather than using the predicted value of the smoking variable from the first-stage regression, the second stage includes the residual from the first stage as a regressor.[6] The second stage short-run regression can therefore be expressed as

$$bmi_{i1} = \beta_0 + \beta_1 bmi_{i0} + f(S_i) + \beta_2 \mathbf{X}_i + \beta_3 \widehat{\mu}_i + \varepsilon_i \tag{6}$$

where S is either *cigday* or CO and $\widehat{\mu}$ is the first-stage residual. The second stage long-run regression is similar but replaces bmi_{i1} with bmi_{i5} and S_i with the average smoking measures discussed previously.

The estimation was conducted using the Stata program "semipar" by Deparsy and Verardi (2012). The first step is to estimate E($bmi|S$), E($\mu|S$) and E($\mathbf{X}|S$), which are approximated by the predicted values $\widehat{bmi}$, $\widehat{\mu}$, and $\widehat{\mathbf{X}}$ by a kernel weighted local polynomial regression. The second step is to form the residuals $\widehat{\mu}_1 = bmi - \widehat{bmi}$, $\widehat{\mu}_2 = \mathbf{X} - \widehat{\mathbf{X}}$, $\widehat{\mu}_3 = \mu - \widehat{\mu}$. Then the coefficients $\widehat{\beta}_0$, $\widehat{\beta}_1$, $\widehat{\beta}_{\mathbf{2}}$, and $\widehat{\beta}_3$, representing the relationships between the independent variables and BMI, are estimated by regressing $\widehat{\mu}_1$ on $\widehat{\mu}_2$ and $\widehat{\mu}_3$. Thus, all parameters in equation (6) are identified except the relationship between cigarette consumption and BMI.

[5]For simplicity, we round smoking values to the nearest integer; e.g. if someone averaged 21.2 cigarettes per day over the five follow-up waves we round this to 21.

[6]For an overview of the control function approach to dealing with endogeneity, see Heckman (1979) and Heckman and Robb (1986).

The last step is, therefore, to identify this relationship with a non-parametric regression of cigarette consumption on the predicted BMI residual, $\hat{bmi}_{i1} - \widehat{\beta}_0 - \widehat{\beta}_1 bmi_{i0} - \widehat{\beta}_2 \mathbf{X}_i - \widehat{\beta}_3 \widehat{\mu}_i$. This relationship is estimated at every level of cigarette smoking, allowing independent marginal effects. The idea behind this strategy is to estimate the non-parametric cigarette function by the residual variation that is unrelated to the parametric independent variables.

We calculate the average effect of quitting smoking on BMI using the semi-parametric estimates as follows. When using the cigarettes smoked per day variable, we first calculate the change in predicted weight from switching from the number of cigarettes smoked at baseline to zero. We then take the average of these predicted changes across all individuals in the sample. The process for the CO variable is similar; the only difference is that we compute the predicted effect of switching to the average CO level for non-smokers of 3.61 ppm, as opposed to zero.

In semi-parametric estimation, the confidence interval becomes very wide at extreme values where there are very few observations. We therefore drop the top 1% of the smoking distribution, which means those who smoke more than 50 cigarettes per day on average across the five follow up years and those with average CO levels of over 50 ppm. We doubt that this restriction is consequential, since if we drop the same individuals in the parametric regressions the results (available upon request) remain similar.

Figures 2 and 3 present the short-run semi-parametric IV results for *cigday* and *CO*, respectively. The graphs display both the point estimates for each integer level of smoking and the 95% confidence intervals. Figure 2 shows that the short-run relationship between cigarettes smoked per day and BMI is highly nonlinear. Specifically, smoking has a diminishing marginal effect on BMI throughout most of the distribution, with the shape of the curve being approximately quadratic. Quitting smoking from levels of 10, 20, 30, and 40 cigarettes per day is predicted to lead to weight gains of 1.22, 1.58, 1.66, and 1.94 BMI units, respectively. Most of the effect of smoking on weight therefore appears to occur at levels below 20 cigarettes per day. Taken literally, this would suggest that heavy smokers could

cut back to a pack a day without fear of substantial weight gain. Figure 3 shows that the short-run effect of CO on BMI is less obviously non-linear than the effect of cigarettes per day. The curve is somewhat flat at very low levels of CO – specifically two to five ppm – but recall that even non-smokers often have non-zero CO so changes at such low levels probably do not reflect changes in smoking behavior. Starting at five ppm, the graph begins to take a quadratic shape, but unlike the graph for cigarettes per day we do not observe a complete leveling off until the far right tail of the distribution.

Figures 4 and 5 turn to the long-run results using simple averages of the smoking measures; the graphs using weighted averages are very similar and are available upon request. Figure 4 shows that the shape of the long-run relationship between cigarettes per day and BMI is roughly similar to the shape of the short-run relationship, as it is approximately quadratic and levels off at around 20 cigarettes per day. Figure 5 displays a similar pattern of results for CO level. CO has a diminishing marginal effect on BMI, and most of the weight gain from reduced CO comes at levels below about 20 ppm. The long-run relationship between CO and BMI therefore flattens out more quickly than the short-run relationship.[7]

The average effects of quitting smoking on BMI implied by these semi-parametric graphs are generally similar to those from the parametric specifications. Using cigarettes per day, the average effect of quitting is 1.67 BMI units in the short run and 1.93 in the long run, compared to 1.52 and 1.91 from the corresponding parametric regressions. For CO, the average effect is 1.80 in the short run and 1.99 in the long run, compared to the parametric regressions' estimates of 1.33 and 1.81. The results presented in this section suggest that the marginal effect of smoking on weight is likely to be modest for levels of smoking above 20 cigarettes a day, which would be impossible to detect using linear specifications.

[7]Note that there is some evidence that additional CO actually leads to higher BMI at the far right tail of the distribution: CO levels of around 47-50 ppm. However, this should be interpreted with caution as it is based on a very small number of individuals. Accordingly, the confidence intervals in this portion of the distribution are quite large.

3.3 Subsample Analyses

We next conduct subsample analyses to evaluate whether the effect of smoking on BMI differs by age, gender, education, or baseline BMI. Heterogeneous effects could occur because of differences in either the biological effects of nicotine on appetite or metabolism or the behavioral responses to these biological effects. Given the complicated nature of these relationships, we make no *ex ante* predictions about the patterns of heterogeneity. For age, we split the sample into three groups: those under 45, 45-54, and 55 and over at baseline. We use these splits because there are no individuals under 35 or over 64 in the LHS. For education, we consider subsamples of those with no college education, some college, and a four-year college degree or greater. There are not enough individuals with less than a high school degree or greater than a college degree to enable further stratification. For baseline BMI, we are interested in whether quitting smoking leads to larger weight gains among those who were already at risk of weight-related ailments prior to the intervention. We therefore split the sample into those with healthy weights at baseline (BMI$<$25) and those who were already overweight or obese (BMI$\geq$25).

Table 4 displays the results for cigarettes per day (Panel A) and CO (Panel B). For brevity, Table 4 contains only the results from long-run parametric IV regressions using simple averages of the smoking measures. Semi-parametric graphs for each subsample are available in Appendix Tables A1-A4. Short-run estimates and those using weighted rather than simple averages lead to broadly similar conclusions and are available upon request.

The results suggest that the effect of smoking on weight is strongest for younger individuals, women, those without a college degree, and those with healthy baseline BMIs. Quitting smoking leads to an average weight gain of 2.19-2.21 BMI units for those under 45, 1.88-1.98 for 45-54 year olds, and 1.33-1.45 for those 55 and older. One possible explanation is that the health consequences from obesity become more salient with age, so older individuals may have a stronger incentive than others to mitigate weight gain after smoking cessation. Again averaging over the four specifications, the average effect of quitting smoking on BMI

is 2.04-2.45 units for women compared to 1.56-1.59 for men. Stratifying by education, the average effects of smoking cessation are around 1.83-1.98 for those with no college education, 1.82-1.93 for those with some college but no degree, and 1.64-1.73 for those with a college degree. There is therefore some evidence of a small reduction in the effect of smoking on weight as education rises, with the largest gap being between those without a college degree and those with a degree. Perhaps education enables individuals to limit weight gain through an improved understanding of nutrition and exercise. Alternatively, education is correlated with income, and additional income may enable the purchase of healthier foods, gym memberships, or over-the-counter products that can help counteract weight gain. Finally, the average weight gain from quitting smoking is 1.98-2.05 BMI units for people with healthy baseline BMI levels compared to 1.7-1.82 for those who started the study overweight or obese. This result suggests that individuals who are at higher risk of health consequences from weight gain take more steps than others to limit the amount of weight gained after smoking cessation.

In all, though, perhaps the most striking results from Table 4 is that, while some heterogeneity appears to exist, the overall amount of heterogeneity is relatively small. Negative and highly significant effects of smoking on weight are evident for all subsamples. The smallest average effect of quitting smoking on BMI from any specification (55 and over, parametric, CO) is a still sizable 1.33. The lack of substantial heterogeneity in the effect within the sample provides perhaps some assurances that the results are generalizable outside the sample. The next section evaluates the generalizability issue in more detail.

3.4 External Validity

We next perform some checks related to external validity. One obvious concern about the generalizability of the results is that the LHS was conducted in the early 1990s, raising the question of the relevance for current policy debates. (With that said, many of the frequently cited associational estimates are from studies using data that are as old or older.) Another

concern related to generalizability is that the LHS' participants are not a random sample of smokers: participants had to desire to quit smoking, have mild (but not major) lung function impairment, and live within reasonable proximity of the locations for follow-up visits. As discussed in the Data section, the end result was a sample that was almost exclusively white (97%) and exclusively middle-aged (starting age 35-59, ending age 40-64).

We attempt to at least somewhat alleviate these concerns by conducting additional analyses with the National Health Interview Survey (NHIS), a large nationally representative survey conducted annually by the Centers for Disease Control and Prevention. The NHIS contains self-reported data on smoking, weight, and height, along with the same control variables used in our LHS analyses (except for baseline BMI, since the NHIS is not a panel). We use the NHIS to see if the association between cigarettes smoked per day and BMI varies along the dimensions of the generalizability issues: time period, race, and age. Obviously a causal analysis is not possible with the NHIS, but verifying that the association between smoking and weight is not particularly unique among the LHS population should provide at least some assurance that the causal effect is not likely to be unique either. We first estimate the association among the NHIS' best available analog to the LHS sample: white 35-64 year olds in 1990-1994 (the years of the five LHS follow-up waves). We then evaluate whether this association has changed over time by estimating the same model among 35-64 year old whites in the five most recent NHIS waves currently available: 2009-2013. Next, we examine the issue of lack of representativeness by race by returning to the 1990-1994 NHIS waves and restricting the sample to 35-64 year old non-whites. Finally, we estimate the model for whites of an age outside of the 35-64 range (i.e. 18-34 year olds combined with those 65+) in order to evaluate the implications of the lack of representativeness by age.[8]

Table 6 reports the results. The first column shows that, in the sample most comparable to the LHS, each additional cigarette smoked per day is associated with a reduction in

[8]The associations of the control variables with BMI are very different for the 18-34 year old age group and the 65+ age group. Therefore, in the regression combining 18-34 year olds with those 65 and older, we include as additional covariates the interactions of each control with an indicator for whether the individual is in the 18-34 portion of the sample or the 65+ portion.

BMI of 0.038 units. This implies an average weight gain from quitting smoking of 0.8 BMI units. The second column shows that the association between cigarettes smoked per day and BMI is stronger in the 2009-2013 sample than the 1990-1994 sample (-0.061 compared to -0.038), but the average effects of quitting smoking are nonetheless fairly similar (0.91 BMI units compared to 0.8) which is similar to the short-run OLS estimate from the LHS. This is because the average number of cigarettes smoked among smokers has dropped over the past two decades. In other words, β_2 may have grown over time but *cigday* has shrunk for the average smoker, leaving $\left(\sum_{i=1}^{N}\beta_2 cigday_{i0}\right)/N$ roughly constant. Next, the third column provides evidence that the association between smoking and BMI for non-whites is stronger than for whites, but the implied average effects of quitting smoking are similar. Again, this is because on average non-white smokers consume fewer cigarettes than white smokers. The final column shows that the association between smoking and BMI among those who are not between the ages of 35 and 64 is virtually identical to the association among those who are in this age range. The average effect of quitting smoking is, however, slightly smaller among the non-35-to-64 sample due to a lower number of cigarettes smoked among smokers.

In sum, though there is likely some heterogeneity across age, race, and time, these results provide at least some assurance that the lack of representativeness of the LHS is not driving our conclusions. Smoking is inversely associated with weight in all NHIS subsamples. The associations between cigarettes smoked per day and BMI all fall within a reasonably tight range of -0.038 to -0.061. The implied average effects of quitting smoking are all between 0.68 and 0.91 BMI units, which equate to 4.3 to 5.7 pounds. These magnitudes are within the range found in the associational literature and are well below the estimates from our LHS IV specifications. This underscores the importance of accounting for endogeneity when evaluating the relationship between smoking and weight.

3.5 Reconciling Our Results with Prior Literature

We close our empirical analysis by reconciling our results with those of EQ, who used previously published LHS summary statistics from O'Hara et al. (1998) to estimate a very large 21.4 lb average weight gain from smoking cessation. We first replicate EQ's results and then show that our finding of a considerably smaller effect can be attributed to our use of more comprehensive smoking measures.

We replicate EQ by computing a Wald IV estimate of the form

$$\hat{\beta}_{WALD} = \frac{\overline{bmi}_1 - \overline{bmi}_0}{\overline{quit}_1 - \overline{quit}_0} \tag{7}$$

where subscript one indicates the treatment group (combination of the SI-A and SI-P groups) and zero the control group (UC). $\overline{bmi}_1$ and $\overline{bmi}_0$ are average BMIs among the treatment and control groups, respectively, at the end of the study period (year 5). *quit* represents EQ's measure of quitting smoking, called "sustained quitting," which is a dummy variable equal to one if and only if the individual was a medically verified non-smoker in *all five* follow-up waves. This is a very stringent measure, as anyone who smokes any amount in any of the five follow-up years is classified as a non-quitter.

The validity of the Wald estimator hinges on the assumption that the intervention only affected the weight of individuals for whom $quit = 1$. To the extent that the intervention also affected the weight of any other individuals (i.e. those with $quit = 0$), the denominator will effectively be too small. The observed difference in average weight between the treatment and control groups will therefore be scaled by too small a number, and the estimated effect of quitting smoking on weight will consequently be overstated. We suspect that the Wald estimator's identifying assumption is violated since there are two types of individuals categorized by EQ as having $quit = 0$ whose smoking behavior (and therefore weight) likely responded to the intervention to at least some extent.

The first type consists of those who quit smoking for part but not all of the 5-year follow-

up period. If, for instance, someone quit smoking for the first two years, relapsed in year three, and then quit again for years four and five, this person is not classified as a quitter by EQ, but it seems likely that they would have gained almost as much weight as someone who quit for all five years. There are 1114 people in the treatment group who quit smoking in at least one follow up wave but were not sustained quitters. Therefore, not accounting for this group has the potential to substantially impact the results.

The second type consists of those who reduced smoking but did not quit entirely. Given the highly addictive nature of cigarettes, it seems likely that there are at least some people who were able to cut back on their cigarette intake as a result of the intervention but were unable to quit completely. Indeed, among those in the treatment group who never quit in any of the five follow-up waves, average cigarettes smoked per day still fell from 31 to 22. There is no reason to suspect that the biological pathways through which smoking affects weight occur only along the extensive margin of smoking, so people who cut back on smoking would likely experience at least some amount of weight gain. Additionally, some people may also be a blend of the two types; e.g. someone who responds to the intervention by gradually cutting back on smoking until successfully quitting at the end of the third year.

After replicating EQ's results using the "sustained quitting" variable, we then re-compute the Wald estimate using our more nuanced long-run smoking measures discussed earlier in Section 3.1.2. Our "average quitter" measure addresses the issue of people who quit in some but not all follow-up years. The average cigarettes per day and average CO variables also address the issue of cutting back but not quitting entirely.

Table 7 reports the results. The first column shows that, replicating EQ's Wald estimator, we obtain an average estimated weight gain from quitting smoking of 3.196 BMI units, or 20.13 lbs at the average height. This is very similar to the result obtained by EQ, differing slightly because EQ used weight as the dependent variable rather than BMI. (We are unable to directly use weight because the LHS microdata suppress height and weight and only provide BMI.) The second column shows that using simple average quitter rather than

sustained quitter reduces the average estimated weight gain from quitting smoking by about 17% to 2.655 BMI units. In the last two columns, we see that using the simple averages of cigarettes per day and CO attenuates this magnitude even further, to 1.84 and 1.58 BMI units, respectively. Ultimately, then, accounting for both temporary/delayed quitting and smoking intensity reduces the estimated average weight gain from smoking cessation by 42%-51% relative to using the naive sustained quitter measure. Since the Wald estimates using our preferred smoking measures from Table 7 are quite similar to those from our preferred long-run specifications in Table 2, we conclude that the difference between our results and those of EQ is due to the different smoking measures rather than our use of a covariate-adjusted regression model in Table 2. This is not surprising given the randomized design. Note, however, that the standard errors are lower in Table 2, so including covariates is still beneficial in that it improves the precision of the estimates.

4 Conclusion

This paper aimed to provide the most credible answers to date to several questions related to the relationship between smoking and weight. First, what is the average short-run causal effect of quitting smoking on body weight? Our preferred estimates suggest that this effect is around 1.5-1.7 BMI units, or 10-11 pounds at the average height. Second, does the weight gain from quitting smoking disappear over time? The answer appears to be no, as the weight gain actually becomes slightly larger in the long run. The long-run effect is around 11-12 pounds, which implies that the fall in smoking explains about 14% of the rise in obesity in recent decades. Third, how does the impact of smoking on weight vary across the smoking distribution? We find evidence of a diminishing marginal effect, with additional smoking having little long-run impact beyond about a pack of cigarettes per day or a CO level of 20 ppm. Finally, how does the effect of smoking on weight vary by age, gender, education, and baseline BMI? Our results suggest that, while quitting smoking leads to sizeable weight gain

for all subsamples, the impacts are largest for younger individuals, females, those with no college degree, and those with healthy baseline BMI levels.

Our estimated average effects of quitting smoking on weight fall within the range of estimates from the associational public health literature, albiet toward the high end of the range. It is clear, though, that our estimates are markedly smaller than those of EQ despite the fact that they utilized the same randomized intervention. The fact that our results are closer to the associational estimates than to those of EQ illustrates a broader methodological point about the dangers of using IV estimation uncritically even when the instrument is randomized or as-good-as-randomized. The randomization merely ensures the validity of the estimated reduced-form relationship between the instrument and outcome. Obtaining a reliable second-stage estimate requires the assumption that the endogenous variable is the only pathway through which the randomized instrument affects the outcome. This can be a difficult assumption to satisfy. As our paper shows, even if conceptually there is only one pathway through which the intervention can plausibly impact the outcome, careful measurement of that pathway is critical.

Our results also have interesting implications for the economics literature on tobacco control policies. As discussed in the introduction, the literature on the effect of cigarette costs (prices or taxes) on BMI reaches conflicting conclusions, with several studies suggesting the effect is either very small or negative – implying that quitting or reducing smoking actually leads to weight *loss*. Is it possible that the causal effect of an aggressive smoking cessation program is to increase BMI while the casual effect of higher cigarette costs is either zero or negative? This seems conceivable for two reasons.

First, the LATE from a price-induced reduction in smoking may differ from the LATEs from smoking ban-induced or aggressive smoking cessation program-induced reductions in smoking. Different smokers could be affected by these different types of interventions; for instance, it seems reasonable to think marginal smokers would be the ones to respond to cigarette price increases whereas those with strong addictions would be the ones to volunteer

for a comprehensive program. Perhaps those with strong addictions are relatively more likely to "quit at all costs", even if it means gaining a substantial amount of weight.

Second, perhaps cigarette prices/taxes affect BMI through pathways besides smoking behavior. In other words, people who quit smoking in response to higher cigarette prices may gain weight, but this could be counteracted by weight losses among those whose cigarette consumption is unchanged – a large share of the population given the price inelasticity of cigarettes (Chaloupka and Warner, 2000). Smokers who do not reduce their consumption when prices rise experience potentially sizeable negative income effects, which could lead to weight loss by reducing overall food consumption or frequency of eating out at restaurants. Moreover, cigarette taxes generate revenue for the state, which can be used to provide funding for nutrition education or health-related programs such as Medicaid. These, in turn, could reduce the BMIs of even non-smokers.

To close, we should emphasize that our results should not be interpreted as suggesting that individuals should be reluctant to quit smoking out of fear of gaining weight. The large body of epidemiologic evidence that smoking is bad for health implies that any increase in obesity-related ailments after quitting smoking is far outweighed by the health improvements along other dimensions. Instead, our findings should be interpreted as a call for further investigation into medical and policy interventions that can limit the weight gain from smoking cessation, and potentially for expanded insurance coverage for medical interventions that prove effective. If this side effect of an otherwise healthy decision could be eliminated, it stands to reason that the health gains from tobacco control efforts would become even greater.

References

[1] Angrist, J., and Pischke, J.S., 2009. *Mostly harmless econometrics: an empiricist's companion*, Princeton University Press: Print.

[2] Aubin, H., Farley, A., Lycett, D., Lahmek, P., and Aveyard, P., 2012. Weight gain in smokers after quitting cigarettes: meta-analysis. *BMJ*, 345: e4439.

[3] Baum, C., 2009. The effects of cigarette costs on bmi and obesity. *Health Economics*, 18(1): 3-19.

[4] Blundell, R., and Powell, J., 2004. Endogeneity in semiparametric binary response models. *Review of Economic Studies*, 71: 655-679.

[5] Cawley, J. and Meyerhoefer, C., 2012. The medical care costs of obesity: an instrumental variables approach. *Journal of Health Economics*, 31: 219-230.

[6] Chaloupka, F. and Warner, K., 2000. The economics of smoking. Culyer, A. and Newhouse, J. (Eds.) Chapter 29 in *Handbook of Health Economics*, Amsterdam: North-Holland, 1539-1627.

[7] Chen, Y., Horne, S., and Dosman, J., 1993. The influence of smoking cessation on body weight may be temporary. *American Journal of Public Health*, 83(9): 1330-1332.

[8] Chou, S , Grossman, M. and Saffer, H., 2004. An economic analysis of adult obesity: results from the behavioral risk factor surveillance system. *Journal of Health Economics*, 23(3): 565–87.

[9] Courtemanche, C., 2009. Rising cigarette prices and rising obesity: coincidence or unintended consequence? *Journal of Health Economics*, 28(4): 781–98.

[10] Courtemanche, C., Pinkston, J., Ruhm, C., and Wehby, G., *forthcoming*. Can changing economic factors explain the rise in obesity? *Southern Economic Journal*.

[11] Courtemanche, C., Pinkston, J., and Stewart, J., 2015 Adjusting body mass for measurement error with invalid validation data. *Economics and Human Biology*, 19: 275-293

[12] Eisenberg, D., and Quinn, B. C., 2006. Estimating the effect of smoking cessation on weight gain: an instrumental variable approach. *Health Services Research*, 41(6): 2255–2266.

[13] Finkelstein, E. A., Trogdon, J. G., Cohen, J. W., and Dietz,W., 2009. Annual medical spending attributable to obesity: payer-and service-specific estimates. *Health Affairs*, 28(5): 822-831.

[14] Flegal, K., Troiano, R., Pamuk, E., Kuczmarski, R., and Campbell, S., 1995. The influence of smoking cessation on the prevalence of overweight in the United States. *New England Journal of Medicine*, 333(18): 1165-1170.

[15] Flegal, K., Carroll, M., Kuczmarski, R. and Johnson, C., 1998. Overweight and obesity in the united states: prevalence and trends, 1960-1994. *International Journal of Obesity*, 22: 39-47.

[16] Flegal, K., Graubard, B., Williamson, D. and Gail, M., 2005. Excess deaths associated with underweight, overweight, and obesity. *Journal of the American Medical Association*, 293: 1861-1867.

[17] Fletcher, J., 2014. The effects of smoking cessation on weight gain: new evidence using workplace smoking bans. *Forum for Health Economics and Policy*, 17(2): 105-129.

[18] Gregory, C. A., and Ruhm, C. J., 2011. Where does the wage penalty bite?, Grossman, M and Mocan, N., (Eds.) in *Economic Aspects of Obesity*, University of Chicago Press: 315-347.

[19] Gruber, J. and Frakes, M. 2006. Does falling smoking lead to rising obesity? *Journal of Health Economics*, 25(2): 183-197.

[20] Heckman, J., 1979. Sample selection bias as a specification error. *Econometrica*, 47(1): 153-161.

[21] Heckman, J., and Robb, R., 1986. Alternative methods for solving the problem of selection bias in evaluating the impact of treatments on outcomes. *Drawing Inferences from Self-selected Samples*, 63-107.

[22] Hedenstrom, P., Malmberg, H., and Fridriksson, H. V., 1986. Reference values for lung function tests in men: regression equations with smoking variables. *Upsala Journal of Medical Science*, 91: 299-310.

[23] Lee. S., 2007. Endogeneity in quantile regression models: A control function approach. *Journal of Econometrics* 141: 1131-1158.

[24] Manson, J. E., Colditz, G. A., Stampfer, M. J., Willett, W. C., Rosner, B., Monson, R. R., Speizer, F. E., and Hennekens, C. H., 1990. A prospective study of obesity and risk of coronary heart disease in women. *New England Journal of Medicine*, 322(13): 882-889.

[25] Mizoue, T., Ueda, R., Tokui, N., Hino, Y., and Yoshimura, T., 1998. Body mass decrease after initial gain following smoking cessation. *International Journal of Epidemiology*, 27(6): 984-989.

[26] Mokdad, A. H., Bowman, B. A., Ford, E. S., Vinicor, F., Marks, J. S., and Koplan, J. P., 2001. The continuing epidemics of obesity and diabetes in the united states. *Journal of the American Medical Association*, 286(10): 1195-1200.

[27] U.S. Department of Health and Human Services, 2012. Health, united states, 2011: with special feature on socioeconomic status and health. U.S. Hyattsville, MD: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention. National Center for Health Statistics.

[28] Nonnemaker, J., Finkelstein, E., Engelen, E., Hoerger, T., and Farrelly, M., 2009. Have efforts to reduce smoking really contributed to the obesity epidemic? *Economic Inquiry*, 47(2): 366-376.

[29] Ogden, C.L., Carroll, M.D., Kit B.K., and Flegal, K. M., 2014. Prevalence of childhood and adult obesity in the united states, 2011-2012, *Journal of the American Medical Association*, 311(8): 806-814.

[30] Philipson, T. and Posner, R., 1999. The long-run growth in obesity as a function of tehnological change. National Bureau of Economic Research Working Paper No. 7423.

[31] U.S. Department of Health and Human Services, 1964. Smoking and health: report of the advisory committee of the surgeon general of the public health service. Atlanta, GA: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, Coordinating Center for Health Promotion, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health.

[32] U.S. Department of Health and Human Services, 1990. The health benefits of smoking cessation: A report of the surgeon general. Atlanta, GA: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, Coordinating Center for Health Promotion, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health.

[33] U.S. Department of Health and Human Services, 2014. The health consequences of smoking—50 years of progress: A report of the surgeon general. Atlanta, GA: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, Coordinating Center for Health Promotion, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health.

[34] Ogden, C. and Carroll, M., 2010. Prevalence of overweight, obesity, and extreme obesity among adults: United states, trends 1960–1962 through 2007–2008. *Division of Health and Nutrition Examination Surveys*.

[35] Ogden, C., Carroll, M., Kit, B., and Flegal, K., 2012. Prevalence of obesity in the united states, 2009–2010. *NCHS data brief*, no 82.

[36] O'Hara, P., Connett, J., Lee, W., Nides, M., Murray, R., and Wise, R., 1998. Early and late weight gain following smoking cessation in the Lung Health Study. *American Journal of Epidemiology*, 148(9): 821-830.

[37] O'Hara, P., Grill, J. Rigdon, M. A., Connett, J. E., Lauger, G. A.,and Johnston,J. J., 1993. Design and results of the initial intervention program for the lung health study. *Preventive Medicine*, 22(3): 304-315.

[38] Pinkowish, M., 1999. Hand in glove: smoking cessation and weight gain. *Patient Care*, 33(2): 134.

[39] Rashad, I., Grossman, M., and Chou, S., 2006. The super size of america: An economic estimation of body mass index and obesity in adults. *Eastern Economic Journal*, 32(1): 133-148.

[40] Rexrode, K. M., Hennekens, C. H., Willett, W. C., Colditz, G. A., Stampfer, M. J., Rich-Edwards, J. W., Speizer, F. E., and . Manson, J. E., 1997. A prospective study of body mass index, weight change, and risk of stroke in women. *Journal of the American Medical Association*, 277(19): 1539-1545.

[41] Robinson, P., 1988. Root-n-consistent semiparametric regression. *Econometrica*, 56(4): 931-954.

[42] Terza, J., Basu, A., and Rathouz, P., 2008. Two-Stage Residual Inclusion Estimation: Addressing Endogeneity in Health Econometric Modeling. *Journal of Health Economics*, 27(3): 531-543.

[43] Travier, N., Agudo, A., May, A., et al., 2012. Longitudinal changes in weight in relation to smoking cessation in participants of the EPIC-PANACEA study. *Preventive Medicine*, 54: 183-192.

[44] Verardi, V. and Debarsy, N., 2012. Robinson's square root of N consistent semiparametric regression estimator in Stata. *The Stata Journal*, 12(4): 726–735.

[45] Wang,Y. C., McPherson, K., Marsh, T., Gortmaker, S. L., and Brown, M., 2011. Health and economic burden of the projected obesity trends in the USA and the UK. *Lancet*, 378(9793): 815-825.

[46] Wehby, G, and Courtemanche, C., 2012. The heterogeneity of the cigarette price effect on body mass index. *Journal of Health Economics*, 31(5): 719-729.

[47] Williamson, D. F., Madans, J., Anda, R. F., Kleinman, J. C., Giovino, G. A., and Byers, T. , 1991. Smoking cessation and severity of weight gain in a national cohort. *New England Journal of Medicine*, 324(11): 739-745.

Table 1 – Summary Statistics
(Means, with Standard Deviations in Parentheses)

Variable	SI-A	SI-P	UC
Age	48.41	48.55	48.43
	(6.84)	(6.83)	(6.80)
Cigarettes per Day	29.59	29.49	29.53
	(14.08)	(13.60)	(14.11)
Carbon Monoxide	25.97	26.70	25.98
	(13.47)	(12.67)	(12.66)
BMI	25.42	25.67	25.55
	(3.91)	(3.92)	(3.92)
High School Degree	0.29	0.30	0.29
	(0.45)	(0.46)	(0.45)
Some College	0.34	0.34	0.35
	(0.47)	(0.47)	(0.47)
College Degree	0.22	0.22	0.24
	(0.41)	(0.42)	(0.42)
Male	0.60	0.64	0.64
	(0.49)	(0.48)	(0.48)
Married	0.65	0.67	0.65
	(0.47)	(0.47)	(0.47)
Observations	1961	1962	1964

Table 2 – Parametric Regression Results

	Short Run (BMI Year 1)			Long Run (BMI Year 5)		
	Quit	Cigarettes	CO	Average Quit	Average Cigarettes	Average CO
Panel A: Effects of Smoking Measures on BMI						
OLS	1.295***	-0.031***	-0.032***	2.299***	-0.052***	-0.060***
	(0.047)	(0.001)	(0.002)	(0.077)	(0.002)	(0.003)
		[0.91]	[0.72]		[1.52]	[1.33]
IV	2.202***	-0.052***	-0.077***	2.646***	-0.065***	-0.082***
	(0.148)	(0.004)	(0.006)	(0.285)	(0.007)	(0.011)
		[1.52]	[1.71]		[1.91]	[1.81]
Hausman	0.000	0.000	0.000	0.211	0.047	0.030
Panel B: Effects of Treatment Dummies on Smoking Measures (First Stage of IV)						
SI-A	0.277***	-11.952***	-8.025***	0.206***	-9.092***	-6.634***
	(0.013)	(0.489)	(0.464)	(0.012)	(0.431)	(0.416)
SI-P	0.272***	-11.390***	-7.728***	0.211***	-8.874***	-6.380***
	(0.013)	(0.495)	(0.472)	(0.012)	(0.437)	(0.417)
F Statistic	320.6	371.1	190.3	227.8	284.3	163.6
N	5345	5344	5274	5446	4966	4517

Notes: Heteroskedasticity-robust standard errors are in parentheses. ***,** and * indicate significance at the 0.1, 1, and 5 percent levels. For the non-binary smoking measures, the implied average effect of quitting smoking is in brackets. The controls for education, gender, marital status, age, and baseline BMI are included in all regressions.

Table 3 – Falsification Test Results

	Short Run (BMI Year 1)				Long Run (BMI Year 5)			
	Full Sample	Quit=0	<25% Change in Cigs.	<25% Change in CO	Full Sample	Average Quit=0	<25% Change in Avg. Cigs.	<25% Change in Avg. CO
SI-A	0.596***	0.278***	0.091	-0.253	0.541***	0.210*	0.086	-0.368
	(0.048)	(0.051)	(0.130)	(0.365)	(0.075)	(0.082)	(0.121)	(0.342)
SI-P	0.614***	0.214***	0.049	-0.127	0.561***	0.252**	-0.014	0.204
	(0.049)	(0.052)	(0.127)	(0.378)	(0.074)	(0.084)	(0.118)	(0.367)
N	5345	3812	1703	231	5446	2958	1379	211

See notes for Table 2.

Table 4 – Subsample Results for Age and Gender
(IV with Year 5 BMI Only)

	Age			Gender	
	<45	45 to 54	≥55	Women	Men
Average Cigarettes	-0.074***	-0.066***	-0.052***	-0.091***	-0.051***
	(0.014)	(0.010)	(0.015)	(0.015)	(0.008)
	[2.21]	[1.98]	[1.45]	[2.45]	[1.56]
N	1514	2275	1177	1894	3072
Average CO	-0.093***	-0.083***	-0.069***	-0.101***	-0.069***
	(0.021)	(0.015)	(0.021)	(0.020)	(0.012)
	[2.19]	[1.88]	[1.33]	[2.04]	[1.59]
N	1374	2065	1078	1714	2803

See notes from Table 2.

Table 5 – Subsample Results for Education and Baseline BMI
(IV with Year 5 BMI Only)

	Education			Baseline BMI	
	No College	Some College	College Graduate	< 25	≥ 25
Average Cigarettes	-0.068***	-0.065***	-0.059***	-0.072***	-0.061***
	(0.011)	(0.014)	(0.013)	(0.010)	(0.011)
	[1.98]	[1.93]	[1.73]	[2.05]	[1.82]
N	2082	1722	1162	2369	2597
Average CO	-0.088***	-0.080***	-0.071***	-0.090***	-0.077***
	(0.016)	(0.021)	[0.019]	(0.014)	(0.016)
	[1.83]	[1.82]	[1.64]	[1.98]	[1.70]
N	1895	1555	1067	2161	2356

See notes from Table 2.

Table 6 – Comparison of Associations Between Smoking and BMI in Different NHIS Samples

(OLS Results Only)

	1990-1994 White Age 35-64	2009-2013 White Age 35-64	1990-1994 Non-White Age 35-64	1990-1994 White Age Not 35-64
Cigarettes	-0.038***	-0.061***	-0.059***	-0.040***
	(0.002)	(0.0 05)	(0.007)	(0.002)
	[0.80]	[0.91]	[0.79]	[0.68]
N	51253	53136	10114	60424

Notes: The controls for education, gender, marital status, and age are included in all regressions. NHIS sampling weights are used. See other notes from Table 2.

Table 7 – Reconciling Our Results with those of EQ

(IV with Year 5 BMI Only)

	Sustained Quit (EQ)	Average Quit	Average Cigarettes	Average CO
Smoking Variable	3.196***	2.655***	-0.063***	-0.072***
	(0.736)	(0.605)	(0.015)	(0.021)
			[1.84]	[1.58]
N	5446	5446	4966	4517

Notes: No control variables are included. See other notes for Table 2.

Figure 1: Changes Over Time in BMI and Smoking for Treatment and Control Groups

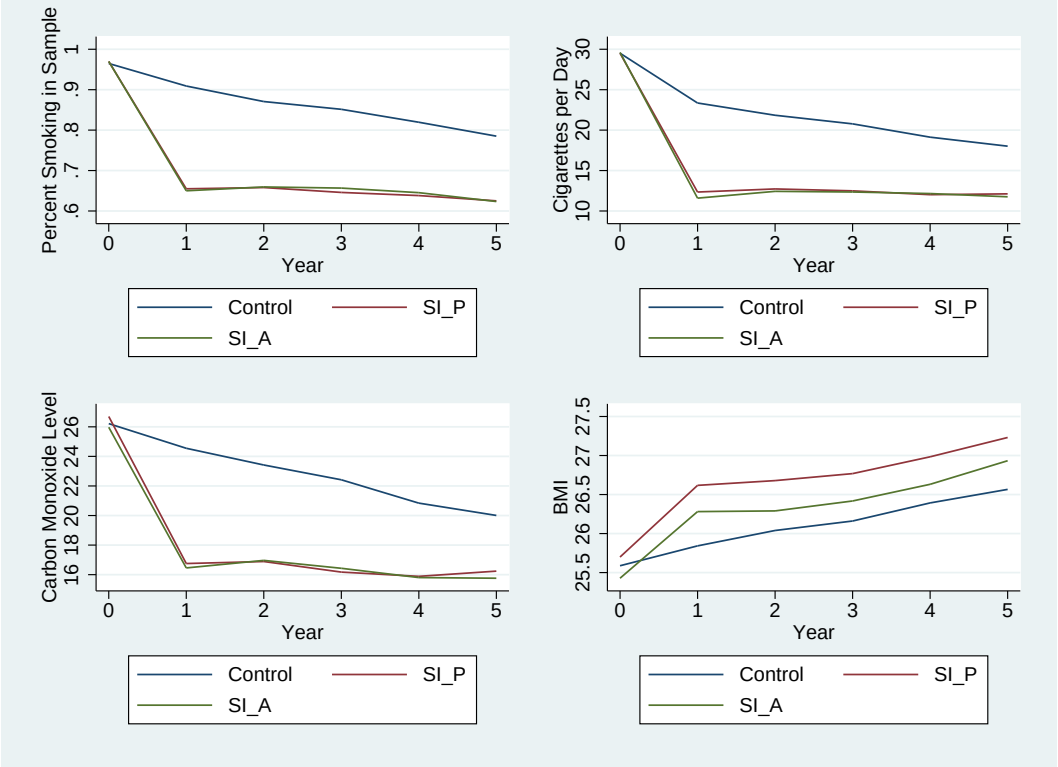


Figure 2: Estimated Short-Run Effect of Cigarettes Per Day on BMI from Semi-Parametric Model

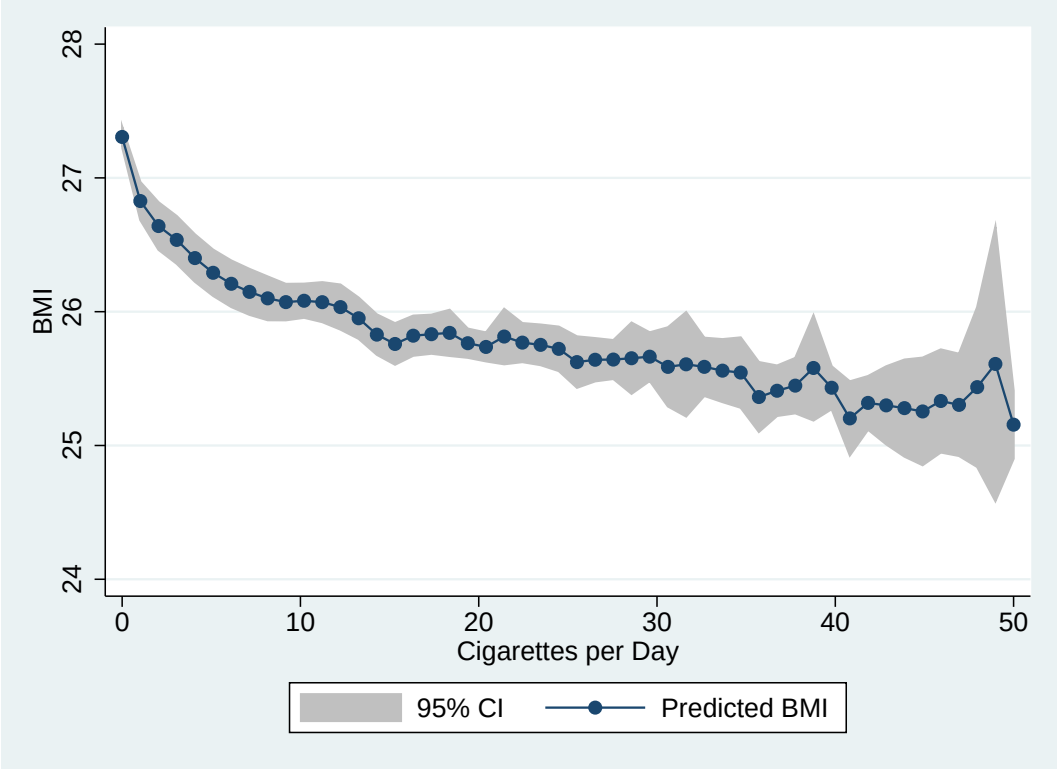


Figure 3: Estimated Short-Run Effect of CO Level on BMI from Semi-Parametric Model

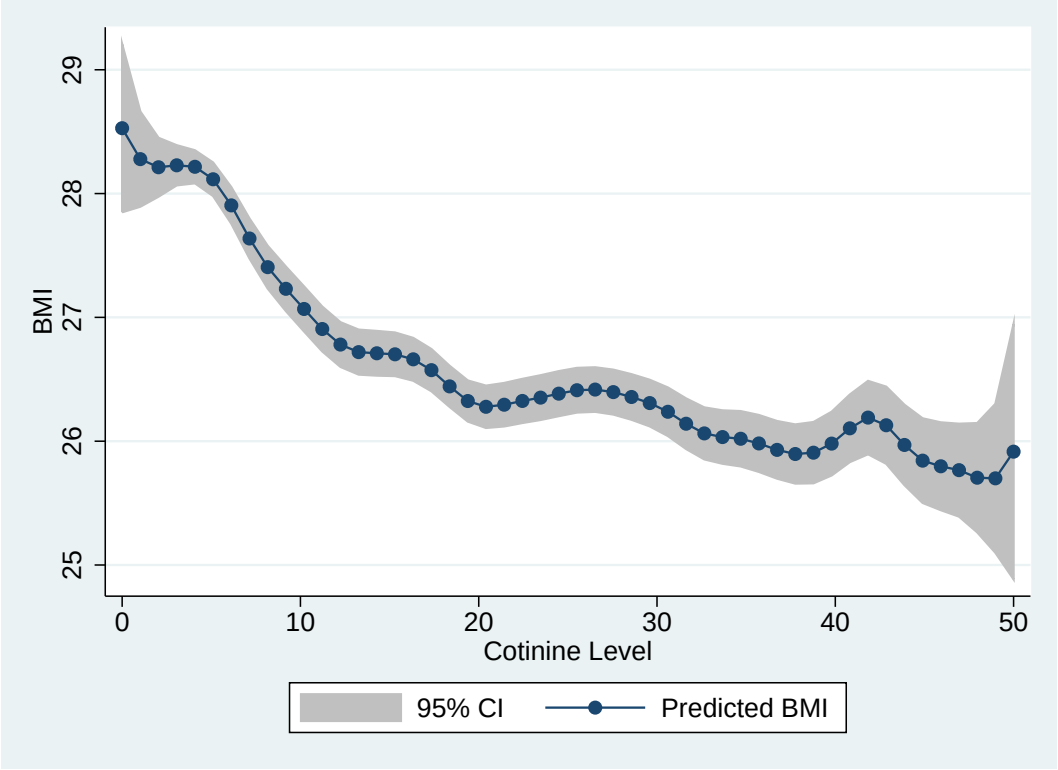


Figure 4: Estimated Long-Run Effect of Simple Average Cigarettes Per Day on BMI from Semi-Parametric Model

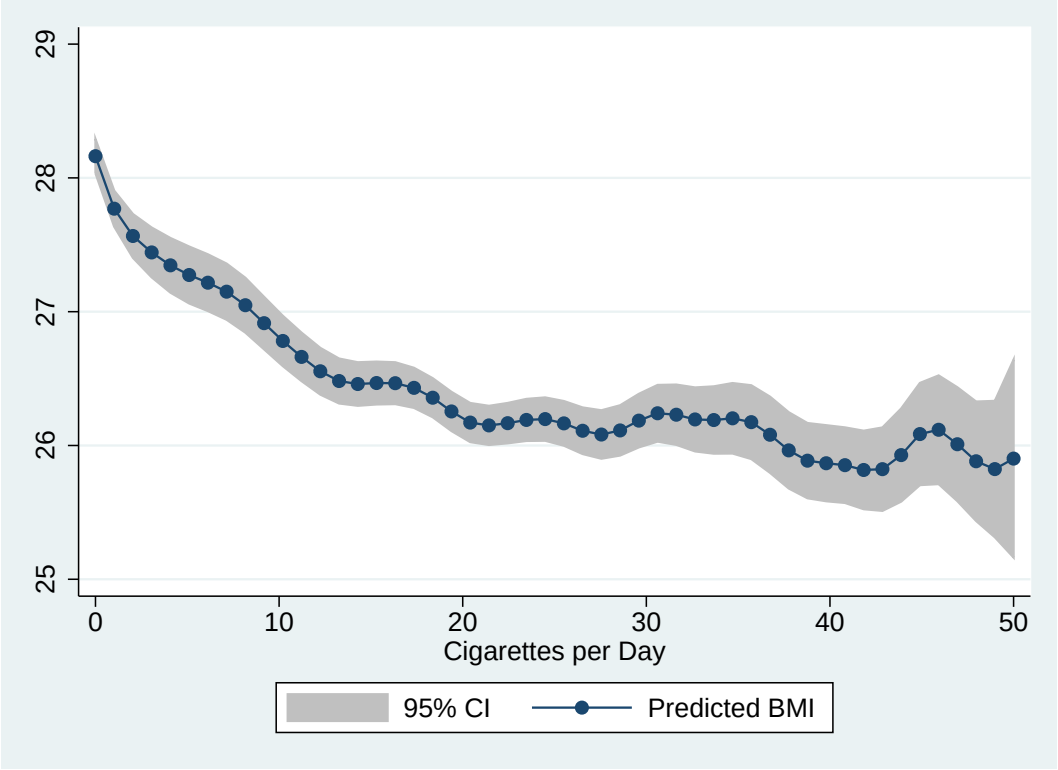
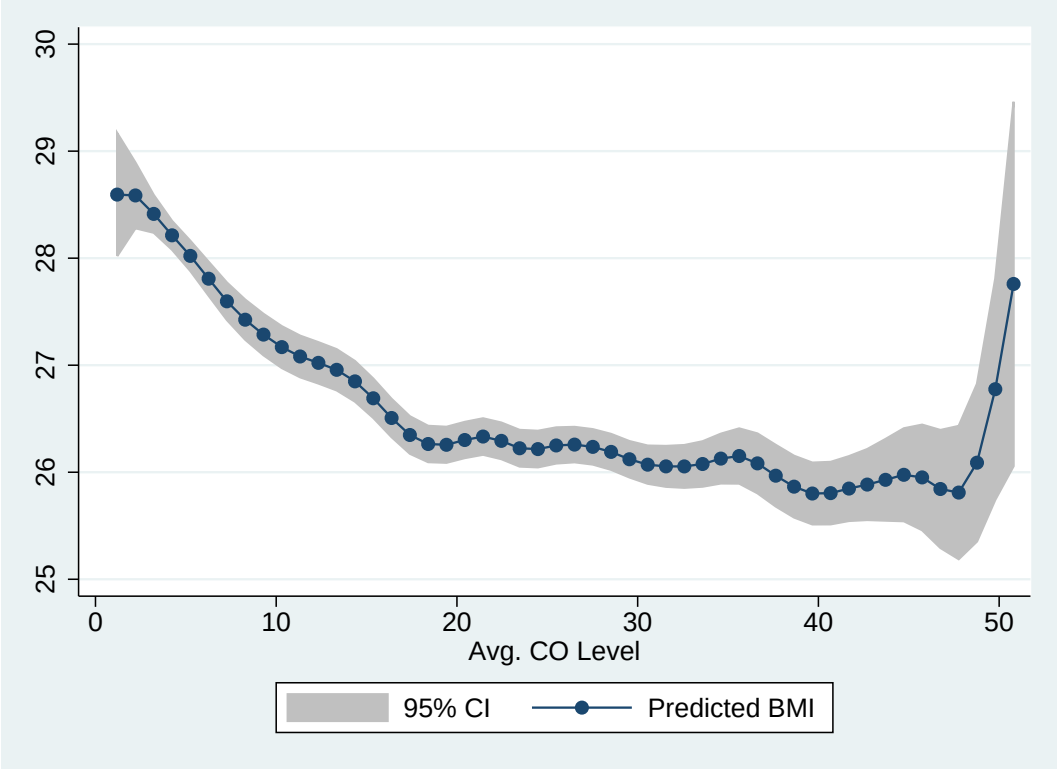


Figure 5: Estimated Long-Run Effect of Simple Average CO on BMI from Semi-Parametric Model



Appendix Table 1 – Sensitivity of Estimates to Different BMI Depreciation Rates

(IV With Year 5 BMI Only)

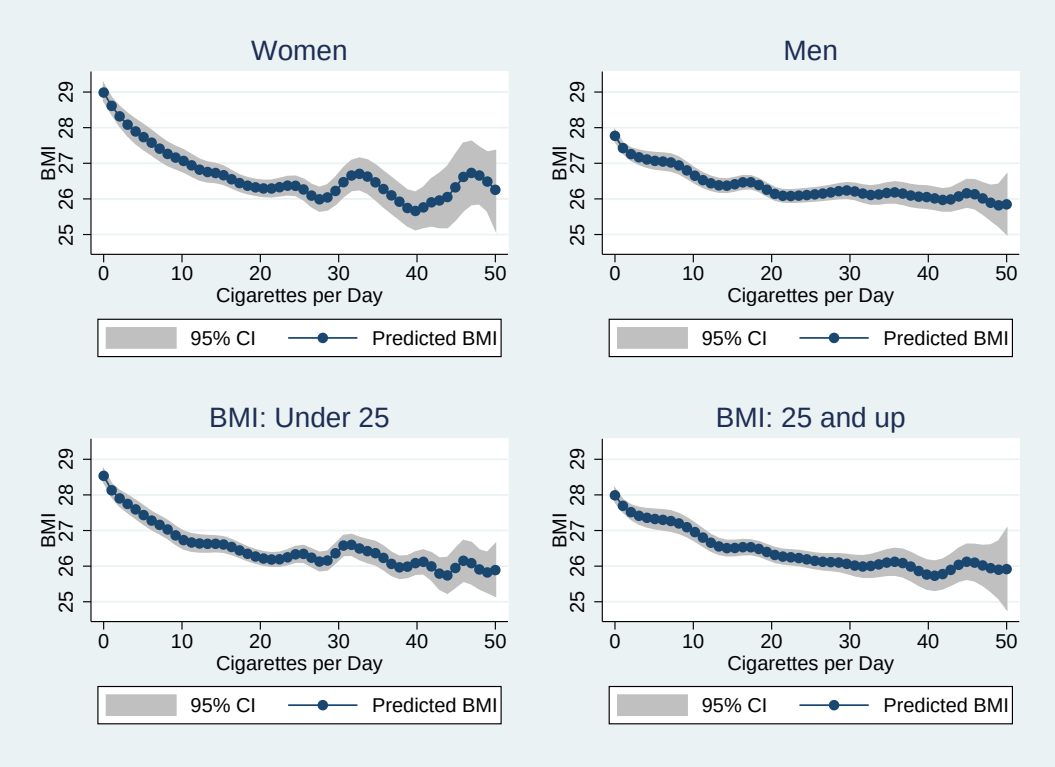
	Long Run (BMI Year 5)		
	Average Quit	Average Cigarettes	Average CO
$\delta = 0$	2.646***	-0.065***	-0.082***
	(0.285)	(0.007)	(0.011)
		[1.91]	[1.81]
$\delta = 0.05$	2.675***	-0.066***	-0.083***
	(0.288)	(0.007)	(0.011)
		[1.94]	[1.83]
$\delta = 0.1$	2.706***	-0.067***	-0.084***
	(0.291)	(0.007)	(0.011)
		[1.97]	[1.86]
$\delta = 0.15$	2.738***	-0.068***	-0.086***
	(0.294)	(0.008)	(0.011)
		[2.00]	[1.89]
$\delta = 0.2$	2.772***	-0.069***	-0.088***
	(0.297)	(0.008)	(0.011)
		[2.03]	[1.93]
$\delta = 0.25$	2.808***	-0.070***	-0.089***
	(0.300)	(0.008)	(0.011)
		[2.07]	[1.96]
$\delta = 0.30$	2.844***	-0.072***	-0.091***
	(0.304)	(0.008)	(0.012)
		[2.10]	[2.00]
N	5446	4966	4517

See notes for Table 2.

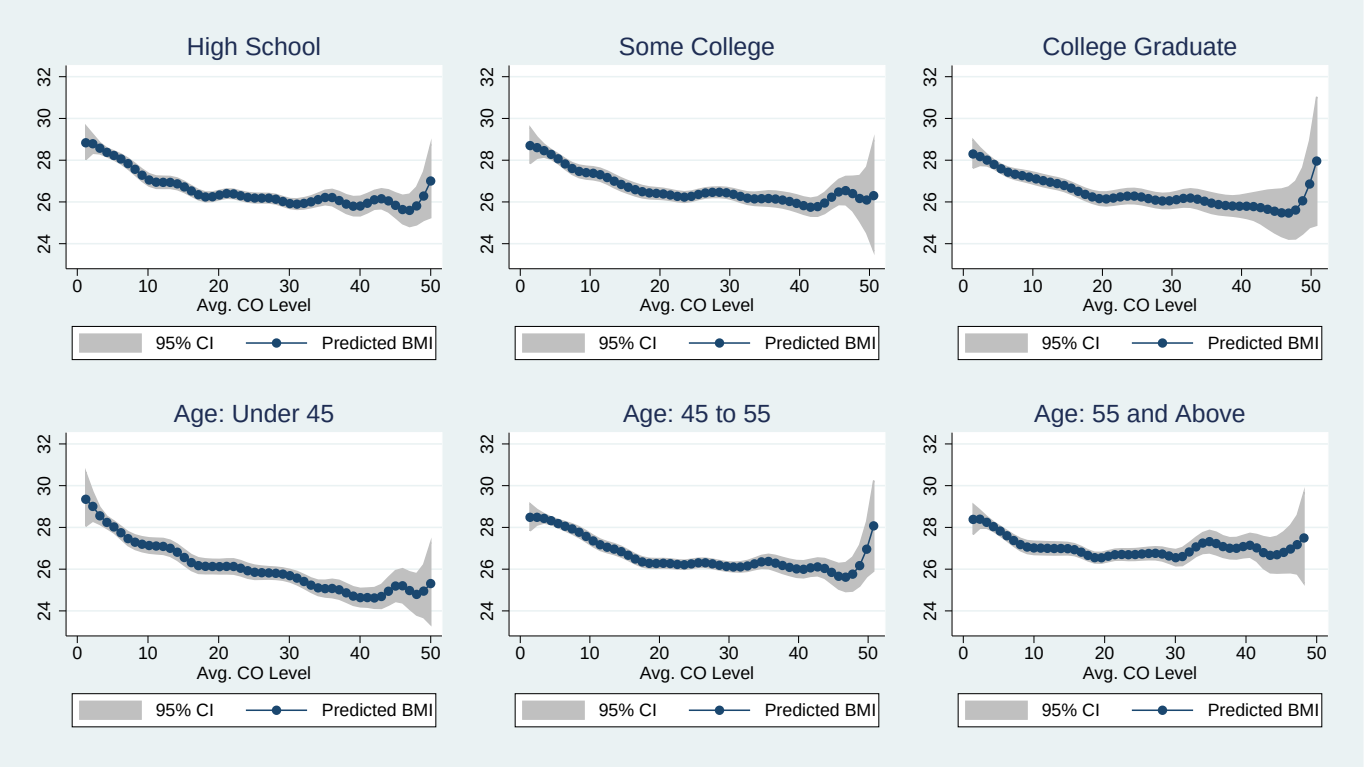
Appendix Figure A1: Semi-Parametric Graphs of Long-Run Effects of Simple Average Cigarettes per Day on BMI for Education and Age Subsamples



Appendix Figure A2: Semi-Parametric Graphs of Long-Run Effects of Simple Average Cigarettes per Day on BMI for Gender and Baseline BMI Subsamples



Appendix Figure A3: Semi-Parametric Graphs of Long-Run Effects of Simple Average CO on BMI for Education and Age Subsamples



Appendix Figure A4: Semi-Parametric Graphs of Long-Run Effects of Simple Average CO on BMI for Gender and Baseline BMI Subsamples

