Behavioral Feedback: Do Individual Choices Influence Scientific Results?*

Emily Oster, Brown University and NBER

March, 2019

Abstract

Consider a case in which a new research finding links a health behavior with good health outcomes. A possible consequence is additional take-up of this behavior among individuals who engage in other positive health behaviors. If this occurs, then later analysis of observational data may be biased by the change in selection. Even sampling-driven false positive results may be confirmed in subsequent work. This paper asks whether these dynamic biases occur, whether they are quantitatively important in empirical settings, and whether standard selection-on-observables adjustments are sufficient to address them. Using data from vitamin supplementation and diet behaviors I show that selection responds endogenously to health recommendations. When behaviors are more recommended, they are more frequently adopted by individuals who engage in other positive health behaviors (not smoking, exercising) and who are better educated and richer. Further, the relationship between these behaviors and health outcomes changes over time. When behaviors are more recommended they are more strongly associated with positive health outcomes, including survival, weight and heart health. The effects are large and adjustment for selection on observables is insufficient to address the bias. This suggests research findings themselves may endogenously bias observational studies.

1 Introduction

Health recommendations – about the best diet, the optimal amount of exercise, vitamins supplements, and other behaviors – change over time. These changes are often in response to new

*I am grateful for comments to Isaiah Andrews, Amy Finkelstein, Matthew Gentkzow, Florian Gunsilius, Emir Kamenica, Ilyana Kuziemko, Matthew Notowidigdo, Jesse Shapiro, Andrei Shleifer, Heidi Williams and participants in a seminar at the US Census. I am grateful to Valeria Zurla, Marco Petterson, Claire Hug, Julian DeGeorgia, Sofia LaPorta, Cathy Yue Bai, James Okun and Geoffrey Kocks for outstanding research assistance. This paper benefits from the Nurses Health Study data, funded with NIH grant UM1 CA186107.

evidence from the medical or public health literature about links between these behaviors and health outcomes.

At the same time, adherence to health recommendations varies across people. Positive health behaviors tend to cluster – people who exercise are also less likely to smoke, for example – and they also correlate with education and income (e.g. Berrigan et al, 2003; Friel, Newell and Kelleher, 2005; Finke and Huston, 2003; Kirkpatrick et al, 2012; Cutler and Lleras-Muney, 2010; Cutler, Lleras-Muney and Vogel, 2008; Goldman and Smith, 2002). These correlations make it a challenge to learn about links between behaviors and health outcomes in observational data where there is a possibility of omitted variable bias. This issue is well-known (e.g. Greenland et al, 1999; Vandenbroucke et al, 2007).

This paper connects these two facts and argues that in the presence of differential *response* to new health advice, bias in estimates of the impact of health behaviors on heath outcomes may be dynamic and endogenously respond to changes in recommendations. This adds to the challenge of learning from observational data. In addition, it means that sampling-driven false positive (or negative) results may be self-reinforcing.

To be concrete, consider a hypothetical case in which researchers are evaluating the relationship between pineapple and cardiovascular health. Imagine that although the true effect is zero, due to sampling variability they find in their study a positive impact of pineapple on heart attacks. Imagine this leads to positive pineapple-related news coverage, or even a change in official guidelines about the role of pineapple consumption in preserving heart health. In response to this, some people would increase their consumption of pineapple. It is likely that these would be the people who are most concerned about their health. But this group is also likely to be engaged in other heart-healthy behavior like exercise, not smoking and eating other fruits and vegetables. As a result of this differential adoption of the recommendation, later studies of the pineapple-heart health relationship may see a larger link between pineapple and health, even though the initial effect was a statistical accident, since a bias has now been created by changes in selection.[1]

The goal of this paper is to explore whether these dynamic responses occur in health settings and whether they are quantitatively important.

I begin, in Section 2, by developing a formal version of the pineapple intuition above. I present a utility model in which individuals differ in their overall valuation of health and they can choose

[1]This discussion, and indeed this paper overall, presumes that the actual size of the causal effects is the same in each period. This seems reasonable since these are intended as biological relationships, unlikely to change substantially within a population on a year to year time frame.

to invest in costly health behaviors. I show in this model that when the (perceived) health value of a behavior increases as a result of a change in recommendation, its covariance with other positive health behaviors and proxies for health valuation also increase. A consequence is that the relationship between this target health behavior and health outcomes will vary with recommendations. This covariance between recommendations and bias will occur even when the analyses include observable controls, if these controls are incomplete. I note that under this model true treatment effects of zero will be especially unstable.

The theory suggests the possibility of these dynamic biases. The bulk of the paper focuses on empirical evidence about the response to changes in recommendations. The model suggests two key tests. First, does the composition of who engages in a behavior change when it is more (or less) recommended? Second, do these change affect later empirical links between health behavior and health outcomes?

Beyond looking for the presence of these dynamics, I ask if they matter. How large are the effects and to what extent are they mitigated by inclusion of controls? If changes in selection have only minimal impacts on the relationship between health behavior and health outcomes, or these changes are fully addressed by standard controls, then patterns of changing selection may be of academic interest but may not be important in influencing conclusions about health behaviors.

I evaluate two contexts: vitamin supplementation and dietary patterns. In both cases I made use of several datasets, described in more detail in Section 3, including the National Health and Nutrition Examination Survey ("NHANES"), the Nielsen HomeScan panel ("HomeScan") and the Nurse Health Study ("NHS").

Health recommendations for vitamins D and E have changed over time. This is the most stark in the case of vitamin E, where a pair of studies in the early 1990s suggested they could prevent cancer and then a meta-analysis in 2004 suggested they actually *increase* mortality.

Overall consumption of supplements responds to recommendations. I show that the composition of those who take the supplements also changes over time. For example, in the period before the two positive vitamin E studies, those who take vitamin E are an 0.7 percentage points less likely to smoke. During the 1993 to 2004 period, when vitamin E was thought to lower mortality, those who take vitamin E are 4 percentage points less likely to smoke. After 2004, when the supplement is less recommended, this falls again to 1.6 percentage points. Putting this together, I argue that these patterns are consistnet with larger changes in vitamin E consumption among those who do not smoke. Similar patterns show up in education. Prior to the positive health news about vitamin

E in the early 1990s, college-educated individuals in the NHANES were 4.6 percentage points more likely to consume this supplement than those with less than a high school education. By the late 1990s, after the recommendation changed, this increased to 13 percentage points. After 2004, when bad news came out, this dropped again to 4.3 percentage points.

Similar patterns appear for smoking, exercise, diet quality, household income and education, across both Vitamin D and E, and in all three datasets considered.

The changing selection impacts the relationship between these behaviors and outcomes. In the Nurse Health Study I estimate links between vitamin E and short-term mortality over time. In the period before 1993 taking vitamin E is associated with an insignificant 10 percent reduction in two-year death risk. After the positive health recommendation, this jumps to a highly significant 20 to 30 percent reduction in the mortality risk. After the negative vitamin E news in 2004, later surveys show the effect of supplementation on mortality is again around 10 percent and not significant. These changes occur with and without controls for smoking behavior. Again, this suggests the selection into vitamin E consumption must be larger among those with otherwise lower mortality risk.

I find similar changes in published work. For both vitamins, papers which use data collected *after* the behavior is more recommended are more likely to find a significant impact on cancer prevention. These dynamics are apparent even though the published papers typically adjust for a wide variety of important controls. This suggests that adjusting regressions for observed covariates may not address the dynamic bias I identify.

The second example in the paper focuses on three dietary patterns: sugar consumption, saturated fat consumption and consumption of a Mediterranean-style diet, each of which have seen changing recommendations over time. Specifically, sugar and fat have become less recommended over time, and the Mediterranean diet more recommended.

As in the case of vitamins, changes in these dietary patterns in response to changing recommendations are larger for some groups than others. Those who decrease their sugar consumption over time are more likely to exercise, less likely to smoke, have higher income and more education. Similar patterns show up for fat and the Mediterranean diet (opposite in the latter case as it becomes more recommended over time).

These changes in selection are reflected in large changes in the apparent link between these dietary patterns and health outcomes over time. When these behaviors are more recommended, they are more strongly associated with lower BMI and better heart health. These effects are large.

In the case of sugar, in the earliest period of the data higher sugar consumption is actually associated with a *lower* BMI. By the latest period of the data it is strongly associated with higher BMI. This suggests that the changes in selection are large enough to change not just the magnitude but the sign of the relationship.

I use regression evidence to show formally that the relationship between dietary choices and socioeconomic status or other health behaviors moves in tandem with the relationship between dietary choices and health outcomes.

As in the vitamin case, including standard observable controls – demographics and controls for other health behaviors – does not eliminate these dynamic biases. Even if we consider regressions with comprehensive controls included, large, visible, changes in the links between behavior and outcomes persist.

Overall, the evidence in these empirical sections suggests that these dynamics are qualitatively important for treatment effect estimates and that controlling for observed covariates is, at least in these settings, insufficient to address these dynamic biases. These findings bode especially poorly for our ability to learn about null effects. A true effect of zero will be unstable, as false positives and false negatives will be self-reinforcing. Once a significant result has been obtained in some setting, biases may be enduring until better data (for example, from a randomized trial) is available.

A natural question, which I take up briefly in the conclusion, is whether there are solutions to these issues. Randomized controlled trials are the gold standard for causal inference and, where feasible, would dramatically improve the evidence in these settings. But: they are expensive, difficult to run and in some of these settings probably close to impossible to implement. Better research designs using non-randomized data, which have not gotten the attention in the public health literature that they have in economics (Angrist and Pischke, 2010), could significantly improve the conclusions here. This could include using food-specific taxes, or discontinuities in vitamin recommendations across age groups. Finally, when these are infeasible it may be possible to use changes in selection patterns directly to evaluate robustness (i.e. Heckman, 1978; Altonji et al, 2005; Oster, 2018).

A primary contribution of this paper is to extend work on the limits of observational data. Many authors have noted that observational evidence in health settings often appears biased and may be contradicted by randomized trials (Autier et al, 2014; Maki et al, 2014; Brownlee et al, 2010). The results here suggest that the observational findings themselves may contribute to the creation of bias. Further, they cast doubt on the ability of standard adjustment-for-observables

techniques to address these dynamic biases.

The paper also contributes to a large literature in economics on the relationship between socioeconomic status and adherence to health recommendations (e.g. Berrigan et al, 2003; Friel, Newell and Kelleher, 2005; Finke and Huston, 2003; Kirkpatrick et al, 2012; Cutler and Lleras-Muney, 2010; Cutler, Lleras-Muney and Vogel, 2008; Goldman and Smith, 2002; Kowalski, 2018) and on consumer response to health information (e.g. Cutler, 2004; Chern et al, 1995; Brown and Schrader, 1990; Chang and Just, 2007; Roosen et al, 2009; Kinnucan et al, 1997; Ippolito and Mathios, 1995).

2 Theoretical Framework

In this section I briefly formalize the intuition described in the introduction. The second subsection discusses the research process and empirical implications of these dynamics.

2.1 Model of Behavior

2.1.1 Setup

I consider a set of individuals who have the option to undertake a set of health behaviors from a vector $\Lambda = (\Lambda_1, ..., \Lambda_n)$. Assume each behavior Λ_j is binary, i.e. $\Lambda_j \in \{0,1\}^n$, with a value of 1 indicating undertaking the behavior. The assumption that behaviors are binary is taken for simplicity of exposition. All results would hold if $\Lambda \in [0,1]^n$ instead.

Health behavior j has a perceived health value $\kappa_j \geq 0$. Individual i has a health benefit function $U_i = \alpha_i \sum_{j=1}^{n} \kappa_j \Lambda_j$. This utility varies across individuals in α_i. We will define individual i as having a higher health value than individual j if $\alpha_i > \alpha_j$. These health values are drawn *iid* from some arbitrary non-degenerate distribution on $\mathbb{R}_+$ with positive density everywhere, so $E[\alpha_i] > 0$.

The assumption of a linear form in the health value (i.e. using $\sum_{j=1}^{n} \kappa_j \Lambda_j$) introduces weak substitutability of different behaviors; the main results developed here would strengthen if the health behaviors were complements. The assumption on α_i rules out some distributions but allows, for example, various Gaussian distribution options.

Each behavior also has a cost, which is specific to individual i and denoted $c_{i,j}$ for behavior j. These costs are drawn *iid* from a normal distribution with mean $\overline{c_j} > 0$ and variance σ_j^2. This allows for heterogeneity in costs across individuals and average differences across behaviors, but assumes these costs are independent of other characteristics of individuals. The distribution of $c_{i,j}$ is drawn independently of the α_i values.

Each individual chooses their optimal set of behaviors, trading off their utility value of health against the cost. We can write the problem for individual i as:

$$max_{\Lambda} \alpha_i \sum_{j=1}^{n} \kappa_j \Lambda_j - \sum_{j=1}^{n} c_{i,j} \Lambda_j$$

Note that $c_{i,j}$ may be zero or negative so individuals may engage in some of these behaviors even if they do not confer health benefits. The individual adopts j if the health benefit of this behavior exceeds its cost: $\alpha_i \kappa_j \geq c_{i,j}$.

Under this model, individuals with a higher health value will undertake more health behaviors on average than those with a lower health value. They will also be more likely to engage in any particular health behavior with a positive health value.

Individuals realize some positive health outcome (e.g. low cholesterol, healthy weight) which is a function of these health behaviors. I assume this outcome is a linear function of health behaviors and write

$$Y_i = \eta + \sum_{\Lambda_j \in \Lambda} (\vartheta_j \Lambda_{i,j}) + \epsilon_i$$

where the coefficients ϑ_j represent the true impact of each behavior Λ_j on the health outcome. Note that $\vartheta_j = 0$ would imply a behavior j does not matter for health outcome Y. Assume that $E[\epsilon_i | \{c_{i,j}\}, \alpha_i] = 0$.

This structure for outcomes assumes that there is no treatment effect heterogeneity, an assumption which will be important for the results later. In these contexts, where the effects I posit are biological, this assumption may be more appropriate than in some other settings.

2.1.2 Change in Value of Behavior

This paper is primarily concerned with the dynamics that occur when there is a change in the (perceived or actual) value of a behavior. Here, I will develop the simplest case in which a behavior moves from having no perceived health value to having a positive value. In the case where the health value is initially positive these results may still hold (especially if the initial health value is small) but they maybe not. In this sense, the result here is intended as a possibility result to develop intuition, which will be tested in the bulk of the paper. Note that all proofs for the results below appear in Appendix B.

Timing Consider a behavior Λ_j which is an element of Λ. In period $t = 0$ behavior Λ_j has a value $\kappa_j = 0$. Note that people may still engage in the behavior at baseline, for example if their cost of undertaking it is negative.

Between time $t = 0$ and $t = 1$ there is a (potentially misleading) signal about behavior Λ_j which leads people to update their beliefs, such that in period $t = 1$ the belief is $\kappa_j > 0$.

I will be concerned with the change, between $t = 0$ and $t = 1$, in the relationship between (i) behavior Λ_j and other behaviors; and (ii) behavior Λ_j and health outcomes.

Behavior Selection Dynamics Let $\Lambda_{j\prime}$ be a behavior with $\kappa_{j\prime}$ constant and positive in both $t = 0$ and $t = 1$. That is, this is a behavior which is understood to have health benefits. Recall from the setup above that this behavior is more likely to be undertaken by individuals with a high health value. The first result relates this behavior to Λ_j in period $t = 0$ relative to $t = 1$.

Proposition 1 *Given behaviors Λ_j and $\Lambda_{j\prime}$defined as above, $Cov_{t=1}(\Lambda_j, \Lambda_{j\prime}) > Cov_{t=0}(\Lambda_j, \Lambda_{j\prime}) = 0$.*

This proposition shows that the relationship between the behavior of interest and the other positive health behaviors will strengthen after the change in recommendation. This result is immediate in this simple case where there is no perceived health value of Λ_j in the baseline period, since there is no positive covariance between the behaviors at $t = 0$.

This first proposition links behavior Λ_j to other health behaviors. In addition, we can consider links to other covariates. Specifically, assume that we are able to observe a variable Z, which is positively related to the health value α_i. This is intended to capture a variable like education or income. Similar logic leads to the following proposition.

Proposition 2 *For any random variable Z such that it is independent of $c_{i,j}$ and $\mathbb{E}[Z|\alpha_i]$ is increasing in α_i, $Cov_{t=1}(\Lambda_j, Z) > Cov_{t=0}(\Lambda_j, Z) = 0$.*

This indicates that we expect the relationship between behavior Λ_j and the covariate Z to strengthen after the change in recommendation.

Disease-Behavior Dynamics These above results relate directly to changes in selection. I turn now to the implications for the estimated relationship between behavior Λ_A and health outcomes.

Proposition 3 *Let Ω be a strict subset of Λ, which excludes behavior Λ_j and at least one other behavior p for which $\vartheta_p > 0$. Then, as long as κ_j and ϑ_j are not too large, we can derive the following results. Precise conditions are given in Appendix B.*

(A) $Cov_{t=1}(\Lambda_j, Y) > Cov_{t=0}(\Lambda_j, Y)$

(B) $Cov_{t=1}(\Lambda_j, Y|\Omega) > Cov_{t=0}(\Lambda_j, Y|\Omega)$.

This says that as the behavior becomes more recommended, and thus the selection on the behavior changes, the estimated effect of the behavior on health outcomes will change. This will be true even if researchers observe and adjust for some of the confounding variables, as long as they do not observe all of them. Note if all elements of Λ were observed and controlled for then it would be possible to estimate the true effect of Λ_j on Y in all periods and these effects would not vary over time.

It is important to note that the results in this section are sensitive to the assumptions (detailed above) on distributions of c and α, among other things. As a result, it may be best to view these as possibility results. The purpose of this discussion is simply to make clear that we *could* see these types of dynamics; the empirical work will focus on whether we do.

2.2 Research Process and Data Implications

The preceding subsection derives conditions under which we would see the dynamics described in the simple example in the introduction. Here, I consider overlaying a research process over these results, and ask how we can look for these patterns in the data.

I consider the research process evaluating the effect of behavior Λ_j on outcome Y. It will be helpful to pull out the particular behavior Λ_j in the estimating equation for Y and rewrite the equation for Y as

$$Y_i = \eta + \beta\Lambda_{j,i} + \mathring{\vartheta}\mathring{\Lambda}_i + \epsilon_i$$

where $\mathring{\vartheta}$ and $\mathring{\Lambda}_i$ represent the coefficient and behavior vectors with behavior Λ_j removed.

Assume the research process is as follows. In each year, researchers draw a sample of individuals and collect data on behavior Λ_j, outcome Y and a set of other variables Θ. This vector Θ may include some elements of $\mathring{\Lambda}$, along with elements of Z (other demographics, etc). Following the data collection, they estimate the effect of Λ_A on Y using the feasible equation below, where we

have introduced subscripts t to indicate the estimation is specific to a year. Note that β also has a t subscript to indicate that it may not be equal to the true β.

$$Y_{it} = \alpha_t + \beta_t \Lambda_{j,it} + \varsigma_t \Theta_{it} + \epsilon_{it}$$

Typically, researchers then report β_t.

Based on the model above, there are some dynamics we might see in the data.

First, propositions 1 and 2 highlight the dynamics of the relationship between Λ_j, other health behaviors and elements of Z. Second, Proposition 3 suggests estimated effects β_t may vary over time with the selection patterns, and this may be true even if the outcome regressions include controls, if these controsl are incomplete.

I use the data described below to look for these patterns in empirical settings.

It is important to note that if we observe all the elements of $\mathring{\Lambda}$, or if the observed controls are sufficient to fully explain $\mathring{\Lambda}$, then these latter dynamics will not occur. In that case, we expect the β_t coefficients to be the same in each period, and equal to β. In this sense, observing that they are different provides an (indirect) test for whether the included controls fully capture the omitted factors.

Stability Before moving on to the empirical results, it is worth noting that under this model and research process a true treatment effect of zero will be fragile and difficult to sustain relative to a true positive or negative treatment effect.

To see why, consider the case where $\beta = 0$, and imagine that initially selection is minimal so in the population the estimated treatment effect would also be zero. Due to sampling variability, with repeated sampling, 5% of the time the data will yield a significant positive or negative impact. When this happens, and if people do not recognize the changes in selection, the endogenous behavioral reaction will reinforce that finding. In later periods the estimated treatment effect in the population will be biased.

In contrast, a true treatment effect that is positive (or, conversely, negative) will be self-reinforcing. We expect a move away from a positive effect (for example) *only if* the sampling variability in the data generates a significant negative effect despite the true significant positive effect. This will happen strictly less than 2.5% of the time.

3 Data

I develop the empirical evidence for these patterns in two contexts: vitamin supplementation and dietary choices. In the case of vitamins the target behaviors - the Λ_j above - are vitamin D and vitamin E supplementation. For diet, I consider sugar consumption (measured as a share of carbohydrates), saturated fat (as a share of total fat) and an index of adherence to a Mediterranean diet.[2]

In addition to these target health behaviors, there are three key data elements: additional health behavior measures ($\Lambda_{j\prime}$ above), covariates which are correlated with health value (Z above) and a set of health outcomes (Y above). Below, I detail the key variables and the datasets used.

3.1 Key Variables

Health Behavior ($\Lambda_{j\prime}$) The health behaviors in $\Lambda_{j\prime}$ should be those which are widely considered to be positive health behaviors. I use exercise, not smoking cigarettes and, in the case of vitamins, a metric of overall diet quality.[3] For graphing purposes, I will sometimes show results using a summary measure based on principal components analysis.[4]

Proxies for Health Value (Z) I proxy for health value with socioeconomic status - education and household income.

Health Outcomes (Y) For vitamins, the health outcomes considered are mortality and cancer incidence. For diet, BMI and cardiovascular health.

3.2 Data Sources

3.2.1 NHANES

The National Health and Nutrition Examination Survey (NHANES) is a nationally representative survey which has been run, in some form, since the 1960s. In this project, I use data from the

[2]For both sugar and fat I focus on the shares of carbohydrates and total fat here rather than the share of calories since the recommendations for total fat and carbohydrates move over time, and this share focus better isolates "sugar" and "saturated fat" as the key target components.

[3]I do not look at vitamin-taking behavior as another metric when I consider diet, since vitamins are only sometimes recommended during this period.

[4]In NHANES, this health behavior factor is based on exercise, cigarettes and (for vitamins) diet quality. In HomeScan it is based on cigarettes and diet quality. For the Nurse Health Study, it is based only on cigarettes.

NHANES III (1988 through 1994) and from the continuous NHANES (beginning in 1999/2000 through 2012/2013).

Information on vitamin supplementation is obtained from the vitamin supplement modules. I focus on individual vitamin supplements - that is, is someone taking a single-ingredient Vitamin D or E supplement. Information on diet is generated from the daily dietary recalls in the study. I generate a Mediterranean diet score as described in Trichopoulou et al (2003).[5]

I extract data on education, income and other demographics from the demographic survey portion of the NHANES. The NHANES also provides a measure of exercise, which I standardize within year, as well as information on smoking behavior. In addition, I use as a measure of diet quality vegetable consumption relative to the median. The advice to consume vegetables has stayed fairly constant over time, while overall advice has varied, so the goal in using this metric is to capture something which is always an indicator of a "good" diet.

To study health outcomes, I extract information on cardiovascular health and BMI. I construct an index of heart health based on blood pressure and cholesterol. All of the health measures are collected objectively, based on weighing and measuring the individual, taking blood samples and blood pressure.

3.2.2 Nielsen HomeScan Data

The Nielsen HomeScan panel tracks consumer purchases using at-home scanner technology. Households that are part of the panel are asked to scan their purchases after all shopping trips. The Nielsen data records the UPC of items purchased. Einav, Leibtag and Nevo (2010) validate the reliability of the HomeScan panel. I use Nielsen data available through the Kilts Center at the University of Chicago Booth School of Business. These data span 2004 through 2016.

These data will be used to look at selection in vitamin purchases over time. They do not contain information on health, and variation in nutrient data coverage over time makes it difficult to analyze diet. However, it is possible to look at vitamins, in particular by generating variables indicating whether the household purchased each vitamin supplement during each year. Information on household education and income can then be used to analyze selection. In addition, I incorporate data on cigarette purchases and an expenditure-based measure of diet quality as measures of other health behaviors.

[5]The score assigns a value of 0 or 1 in nine dietary elements, where a value of 1 is given if someone is either above the median in a good food category (vegetables, fish, etc) or below the median in a detrimental food category (dairy, meat, sugar).

3.2.3 Nurses Health Study

Finally, I use data from a panel of nurses, the Nurses Health Study (NHS), to analyze selection patterns in Vitamin E. The NHS recruited a cohort of approximately 120,000 female nurses in 1976 and conducted by-mail surveys of the cohort every two years. The study is ongoing and response rates are very good. Information is available on dietary patterns, some behavioral characteristics and outcomes, including mortality. The cohort is described in more detail in Colditz et al (1997).

I use data from 1984 to 2006 on vitamin E supplementation, smoking behavior and mortality.

3.2.4 Publications

In the analysis of vitamins I also draw information from published work on the relationship between vitamin supplements and cancer. I locate publications in two ways. First, I scrape Pubmed for "Vitamin X and cancer" and extract relevant studies, limiting to studies in journals in the top 20% in terms of impact factor. Second, I extract lists of publications from meta-analyses of these relationships. I focus on observational studies and exclude RCTs.

For each original study I then extract information on the treatment (either vitamin D or vitamin E supplementation), the outcome (a type of cancer), the years of data covered in the study, the population characteristics and, importantly, whether the study indicated a significant reduction in cancer with vitamin supplementation. I focus on significance rather than magnitude because given the varying approaches across studies, and the varying types of cancer, it is difficult to compare magnitudes. In all, the resulting dataset includes 82 studies of vitamin D supplementation, and 83 studies of vitamin E supplementation.

I then allocate studies to time periods based on the timing of the data used in the analysis. Some studies are allocated in parts to different time periods. For example, if a study includes data from 2005 through 2010, I assign a 50% weight to the period before 2008, and a 50% weight to the period after. I then summarize the (weighted) share of significant cancer reductions across time periods. I residualize results with respect to the type of cancer studied (this does not affect the findings).

Publications lag data, and as a result it is not feasible to look at results from the most recent time periods. In particular, in the case of Vitamin D we do not have data past 2010, and in the case of Vitamin E, we do not have sufficient data past 2004.[6]

[6]In addition to the general issue that publications lag data, there appears to be a drop in publications after negative news comes out about both Vitamin E and Vitamin D. This is an interesting dynamic, although outside the

4 Results: Vitamins

Recommendations about vitamin D and E supplementation have varied over time. I identify timing of significant information events for each vitamin supplement. These events include changes in government recommendations, the advice of national organizations and major research findings. Data from Google Trends are used to validate these events where possible.

In the case of vitamin D, in 2007 there were a number of pieces of positive news. This included a widely cited New England Journal of Medicine (NEJM) article on the benefits of vitamin D (Holick, 2007), coverage in the NY Times and other outlets (e.g. Nagourney, 2007) and a corresponding spike in Google searches. On the other side, around 2012 there was a push back on much of this. An Institute of Medicine report (Rosen et al, 2012) suggested the purported benefits of vitamin D were overblown, and media coverage in the same period reinforced this (e.g. Bakalar, 2012a; Kolata, 2012; Balakar, 2012b).

Vitamin E has gone in and out of fashion even more dramatically. In 1993, a pair of widely covered studies in NEJM cited large benefits of vitamin E in preventing heart disease in men and women (Rimm et al, 1993; Stampfer et al, 1993). In 2004, however, a similarly widely cited report suggested vitamin E supplementation could actually *increase* mortality (Miller et al, 2005).[7] This latter fact was based on randomized controlled trial data, and therefore not affected by the selection biases which had grown over time.

Trends in vitamin consumption over time are shown in Figure 1. For vitamin D these data come from the NHANES and HomeScan data; for vitamin E I also include data from the Nurse's Health Study. There is clear evidence that purchase and consumption of these products vary around the events identified above. Vitamin D consumption rises in both the NHANES and HomeScan after 2007, and we can see evidence of declines - especially in the longer time scale in HomeScan - after 2012. For vitamin E we observe both a sharp increase in consumption after 1993 and a decline after 2004.

4.1 Changes in Behavioral Selection

The first implication of the theory is that the group who consumes vitamins *after* they are more recommended will be different on other dimensions than those who consume before. Specifically, they will be disproportionately positively selected in terms of other health behaviors and socioeconomic

scope of this paper.

[7] This study was released in 2004 but the publication in print is January, 2005

status.

To explore this, I construct yearly correlations between vitamin consumption and other variables (health behaviors, socioeconomic status measures). These correlations are adjusted for age and gender; they are extracted from a regression of each selection variable on an indicator for taking the vitamin supplement, along with age and gender controls, and household controls in the case of the HomeScan data.

As an initial illustration of the results, Figure 2 shows the relationship between vitamin E consumption and health behaviors in the NHANES, HomeScan and Nurse Health Study. The top row shows correlations with smoking in all three datasets; the bottom shows the relationship with exercise in the NHANES and diet quality in the NHANES and HomeScan.

In these graphs, we observe the correlations changing when recommendations change. For example, consider Panel A, with evidence on smoking in the NHANES data. In the NHANES-III data, collected before the positive news about Vitamin E came out, smokers are 0.7 percentage point less likely to take a vitamin E supplement. In the period between 1993 and 2004, when vitamin E is more recommended, this jumps to 4 percentage points. Following the bad news about vitamin E in 2004, the relationship shrinks again.

The time frame in the HomeScan is more limited, but we see here the change after 2004, when the correlation goes from more to less negative.

Panel C shows a similar pattern in the Nurse Health Study. In this case the data is from a panel of individuals, increasing the precision of the estimates and allowing a more direct interpretation around adoption. The relationship between smoking and vitamin E changes after 1993, an again (in the opposite direction) after 2004. Given that this is a panel, it implies that those who newly adopt vitamin E after 1993 must be less likely to be smokers, and those who dis-adopt after 2004 must also be less likely.

Panel D shows the same pattern in the NHANES for exercise, and Panels E and F for diet quality. Overall, in this example, we see clearly that those who adopt this vitamin supplement when it is more recommended are also more likely to be engaging in other positive health behaviors.

Table 1 summarizes the full set of correlations in the data. This adds vitamin D to the analysis, and also correlates with socioeconomic status. For each vitamin, I divide years into three groups - before the positive recommendation, during the period of positive recommendation, and after that period – and report the average correlation in each period. The year ranges differ for the two vitamins, but in each case the central column of the panel - which is italicized - shows the "positive

recommendation" period.

The patterns in this table echo the illustrative results in Figure 2. In the periods in which each vitamin is more recommended the consumers are less likely to smoke, more likely to exercise, eat a better quality diet and are richer and better educated. This appears across multiple datasets and for both vitamins. The differences over time are, for the most part, significant, suggesting this is not just sampling variability. One exception to the overall pattern is when studying vitamin D in the NHANES we do not observe a decline in the relationship after 2012; this may be due to this cell being based on only one post-2012 year. Notably, in the HomeScan data, which covers a longer period after 2012, we do observe a clear reduction in the relationship.

As a final summary, I look at the composite metrics of health behavior and socioeconomic status, and illustrate their correlation with vitamin consumption over time. Figure 3 shows these results. Consistent with Table 1, for both other health behaviors and socioeconomic status, the correlations move with the recommendations.

It is worth saying that these changes are, in many cases, large. For example, in the late 1980s there is virtually no relationship between education and vitamin D consumption. However, by 2011 increasing one education category increases the chance of consuming vitamin D supplements by 4 percentage points, or about 25%. The relationship between education and vitamin E shows similarly large changes.

4.2 Outcome-Behavior Links

The evidence above shows movements in the selection of who engages in these health behaviors around changes in the recommendations. Here, I turn to whether these changes are reflected in changes in the relationship between vitamin supplements and health *outcomes.*

I begin by showing evidence on the link between vitamin E and mortality in the Nurses Health Study. Echoing the construction of the correlations above, I generate a time series of estimates of the impact of vitamin E supplements on two-year-survival in the NHS data. The outcome in this regression is a dummy for surviving two years after the survey (that is, until the next survey round) and the coefficient of interest is on vitamin E supplementation. The regression includes comprehensive age controls, and I divide the resulting coefficients by the average death rate over that period, to reflect the fact that cohort is aging. Magnitudes are interpreted as the percent decrease in mortality rate.

The results are shown in in Figure 4. The solid dots show the estimated correlations with no

controls other than age.[8] The effects move significantly with the recommendations. In the early period, taking vitamin E is associated with an insignificant 10 percent reduction in death risk. After 1993, this jumps to a highly significant 20 to 30 percent reduction in the mortality risk. In 2006, the first survey after the 2004 release of evidence undermining the value of vitamin E, the effect of supplementation on mortality is again smaller (around 10 percent) and not significant.[9]

Section 2 notes that if we can fully adjust for selection then these patterns may arise in uncontrolled regressions, but will be addressed by adding controls. In most settings, the key variables to control for are basic demographics (i.e. education, income). In this case, the construction of the data already effectively holds these constant, as all participants in the study are nurses. The differences in mortality effects over time may already be muted relative to what we would see if the data had a more heterogenous group. However, in the results above we observe changes in the relationship between smoking and vitamin consumption. Smoking is an important contributor to mortality, so it is reasonable to ask whether including a control for smoking in the regression dramatically changes the results.

The hollow markers in Figure 4 show these effects with an adjustment for smoking behavior. The basic patterns remain and, indeed, the magnitude of the difference between pre-1993 and post-1993 is almost unchanged by this control. It is notable that the change from pre-2004 to post-2004 is muted in this series relative to the series without controls, suggesting that the changes in selection on smoking behavior is driving some (but by no means all) of the dyamics. In combination with the fact that we are already holding many demographic variables constant here, this suggests these dynamics are not eliminated by standard observable controls.

As a second piece of evidence, I evaluate published results on the relationships between vitamins and cancer. Using the dataset of published work described in Section 3 I show the change in significance of published results over time in Figure 5.

The changes in research findings follow changes in recommendations. For Vitamin D, published results with data post-2007 are much more likely to find significant relationships between Vitamin D supplementation and cancer. For Vitamin E, a similar pattern occurs for results with data post-1993. Again, these changes are sizable. There is a thirty percentage point increase in the chance of a significant result between pre- and post-2007 period. Although the data coverage from

[8]Note that this cohort is entirely women and they are all nurses, so effectively these regression also control for gender and occupational status.

[9]The decrease in death rate seem to occur in the 1992 survey, initially; this survey takes place in part over 1993 due to attempts to contact everyone, which may explain this pattern.

publications ends before the negative news about vitamins in both cases we would expect studies using later data to be less frequently significant in both cases.

The publications from which these results are drawn typically control extensively for demographics and other variables. This provides further evidence that such controls do not fully address these dynamic biases.

The evidence suggests the changes in selection detailed above are quantitatively important for empirical conclusions about the relationship between these behaviors and outcomes.

It is important to acknowledge here the central alternative explanation for these findings, which is that the actual effect of these supplements varies over time. This is difficult to test directly, but I believe it is unlikely to be the explanation for what we see in the data. It is difficult to imagine what non-selection based changes over time would explain the large and abrupt swings in, for example, the mortality-vitamin relationship. At the same time, the selection patterns offer an obvious explanation.

5 Results: Diet

I move now to the second example of dietary patterns. I analyze the three dietary patterns: sugar intake, saturated fat intake and adoption of a Mediterranean diet.

Recommended levels of sugar and fat intake have fallen over time. In the case of sugar, the revision of the US Dietary guidelines in 2000 mark the first mention of avoiding sugar for health reasons (Krauss et al, 2000). Following this, in 2011/2012 there was a surge of popular media interest in the dangers of sugar (e.g. Taubes, 2011; Chang, 2012; 60 Minutes "Sugar").

Saturated fat was first restricted in the US dietary guidelines in 1990, with a suggested limit of 10% of calories (Peterkin, 1990). In 2005 this limit was lowered to 7% (Thompson and Veneman, 2005).

The Mediterranean diet has been, conversely, increasingly recommended over time, largely due to positive findings from major research studies. The first of these was in 2004, when two articles in the Journal of the American Medical Association (JAMA) showed positive health benefits from the diet (Knoops et al, 2004; Esposito et al, 2004). In the 2009 period additional studies argued for cognitive benefits (Feart et al, 2009). Finally, and most notably, a large randomized trial released in 2013 showed significant cardiovascular benefits (Estruch et al, 2013).

Figure 6 shows variation in diet over time in the NHANES data. The changes in diet line up

with the changes in recommendations. Sugar declines over time, saturated fat declines over time and the Mediterranean diet spikes, especially in the last years in the data. This suggests people are responding to these highlighted events.

Unlike in the case of vitamins these trends in recommendation are secular, so we cannot look for the reversals in selection that were apparent above, although the direction of changes we expect are different for sugar and fat versus the Mediterranean diet.

5.1 Changes in Behavioral Selection

I construct yearly correlations between diet and other health behaviors or socioeconomic status as I did for vitamins, here using the NHANES data.

Table 2 shows the summary correlations between dietary choice and behaviors or socioeconomic status over time; this echoes Table 1 on vitamins. For both sugar and fat the later periods of the data correspond to more caution about excessive consumption; for the Mediterranean diet, the later periods correspond to more evidence in favor of the diet. In most cases, the correlations between diet and other variables move with the recommendations. The changes over time are significant.

Consider, as an example, sugar and exercise. In the earliest period of the data, individuals who report exercising more eat *more* sugar than those who exercise less - the correlation is positive. In the later periods of the data the correlation is negative, and this negative correlation grows over time. For sugar we also see this sign reversal for the correlations with education and income. This suggests that the change in selection is large enough not just to change the size of the relationship but also the direction.

The same patterns show up for sugar and smoking (without the sign reversal), and for fat and the Mediterranean diet. Most notable for fat is the large changes in relationships between saturated fat consumption and other behaviors and socioeconomic status between the pre-2005 and post-2005 period. For the Mediterranean diet the relationships between consuming this diet type and other health outcomes and socioeconomic status are always positive. But they are growing considerably over time. One exception is smoking; those who consume a more Mediterranean-style diet are always smoking less and the magnitude of this relationship does not grow significantly over time.

Figure 7 echoes Figure 3, showing the full evolution of correlations between diet and health or socioeconomic status indices over time. These graphs echo the results in Table 2: correlations respond to the recommendations in most cases.

5.2 Outcome-Behavior Links

I turn next to the implications of these changes in selection for the relationship between diet and health outcomes. The outcomes I consider here are weight (measured by BMI) and an index of heart health including metrics of blood pressure and cholesterol. In both cases, I generate correlations over time between the outcome and the diet measures. Having created these, I analyze them in two ways. First, it is possible to look at graphs similar to those shown above, and simply look at whether the correlations move with the recommendations. Second, I can look statistically at whether the relationship between the diet behavior and the outcomes moves together with the relationship between the selection variable and the diet behavior.

Beginning graphically, with BMI as the key example, Figure 8 shows the evolution of the BMI correlations with diet over time. The solid, filled-in series, shows the coefficients with no controls other than age and gender. Although there is more noise in these than in the selection figures above, we do see movement in these relationships around the changes in recommendations. Again, in some cases these effects are large enough to change the sign of the relationship. For sugar, in the earlier periods greater sugar consumption is actually associated with a lower BMI; in the later period, it is strongly associated with higher BMI.

The hollow series in Figure 7 show these coefficients from regressions which include a standard set of demographic controls common in this literature (education, income, race, marital status) and the controls for health behaviors (exercise and smoking). Including these controls represents a natural approach to combating the selection issues here. In the language of the theory, if we are able to include controls which fully address differences in behavior choices (either by including all elements of $\mathring{\Lambda}$ or by including controls that fully capture these differences) then the estimated treatment effects with controls will not vary over time, even if the selection in behavior does vary.

In practice, although the inclusion of controls does shrink the difference between the coefficients over time, the magnitude of the change is small. We still see the same movements in coefficient size corresponding to the change in recommendations.

Table 3 provides regression evidence on the co-variance between these diet-health gradients and the diet-selection coefficients generated in Section 5.1. We can focus first on Panel A of the Table, where the diet-outcome regressions are run only with age and gender controls. In all four cases - the two outcomes and the two different selection measures - there is a significant relationship between the two series.

Panel B re-estimates these relationships but where the diet-outcome regression includes the full set of demographic and health controls. The coefficients shrink slightly, but remain highly significant.

A natural question is why we do not also look at evidence from publications, similar to the case of vitamins. In fact, much less has been published on these relationships in the public health literature than on vitamin and mortality links, and the changes in recommendations are largely more recent. Given that publications lag data, there was not sufficient publication sample size to explore this.

Overall, this analysis suggests that the selection changes here are large enough to have significant effects on what a researcher would conclude about the relationship between these treatments and outcomes. In the earliest period of the data, someone running observational regressions with these data would conclude that sugar, if anything, slightly lowers weight. In the later period, the effects are huge and positive. And this is true even if the researcher was careful to include demographics and other health behavior controls; the selection on *unobservables* is simply too important to ignore.

6 Discussion and Conclusion

In this paper I analyze the role of health behavior change in driving biases in estimates of the impact of health behaviors on outcomes. I outline a simple data generating process in which changes in health recommendations differentially change health behaviors for different groups and show that these changes may influence estimated relationships between behavior and health over time. Using data on vitamin supplementation and diet I demonstrate that these dynamics occur in data. The degree of selection in behaviors varies over time, and the relationship between behavior and health also varies with these changes in selection. These dynamics are quantitatively important.

The results add to existing caution in interpreting observational results in settings like these. The problem of omitted variable bias is well known, but these results suggest such bias may be dynamic and, indeed, may *respond* to research findings. This suggests that awareness of the changes in recommendations over time should inform discussions about the plausible degree of bias in estimates. The findings are a reminder that these research results are not produced in a vacuum and, indeed, the underlying economics of health behavior are crucial for understanding the scale of the bias. One take-away is that current approaches in these literature – typically, using standard selection on observables adjustments – are very unlikely to yield causal results.

A natural question is how we can improve our understanding of causality here. One clear answer is more use of randomized controlled trials (RCTs). Indeed, in several of the settings here, our best evidence does come from RCTs. In the case of Vitamin E, fore example, the research that urged caution on consumption was a meta-analysis of randomized controlled trials. By this point in time the observational data was quite biased in favor of vitamin E, but randomized data is not subject to these biases.

However, expanding the scope of RCTs is not a panacea. These studies are expensive, it is often difficult to really impact subjects behavior around diet and it can be especially challenging to evaluate long-term outcomes like development of cancer or mortality. A growing literature (see e.g. Angrist et al, 2017) seeks to combine RCTs with observational estimates of the kind developed here. These approaches would be challenged, however, by the endogenously changing bias demonstrated above. This leaves us seeking solutions beyond expansion of randomized trials.

Angrist and Pischke (2010) argue that among the central advances in empirical economics in the past decades has been an improvement in non-experimental research design, which serves to improve our ability to get causal estimates of treatment effects of interest. These advances – more use of difference in difference, sharp and fuzzy RD, event studies – have been slower to spread into the type of public health problems addressed here. In light of the findings above it seems even more crucial to consider the possibility of improved research designs in these settings. Policies like sugar taxes or vegetable purchase subsidies could potentially be utilized, or discontinuities in health recommendations across age or other characteristics.

A third option, more akin to a robustness check and feasible even without a new research design, would be to *use* the dynamic selection to adjust the estimates. One approach in this spirit would rely on a Heckman selection framework (e.g. Heckman, 1978) again using multiple time periods to pin down causal parameters. A second would rely on an assumption about selection on observed and unobserved factors, as in Altonji et al (2005) and Oster (2018). In either case, it would be feasible to combine these assumptions with multiple periods of data to ask what causal effect is consistent with the combination of observed treatment effects and changes in selection. These type of approaches clearly fall short of the ideal of RCTs or improved research designs, but could force a recognition of the dynamic selection issues that would be valuable in these discussions.

I provide two final notes on extensions here. First, although this paper focuses in particular on individual health behaviors, the dynamics here may be present in other settings (parental behaviors, for example) where individual choices vary over time. Second, the logic in this paper may relate to

a broader literature - in economics and elsewhere - on technology adoption (e.g. Griliches, 1957; Skinner and Staiger, 2005). That literature often discusses an S-curve in adoption, with initially slow adoption, then a period of faster adoption, and then an asymptote to full adoption. An open question in many settings is why some users adopt first - one possibility is that the driver of earlier adoption is general knowledge or selection (e.g. Strang and Soule 1998); a second is that the users with the largest benefits adopt first (e.g. Hall and Kahn, 2003). The first of these explanations is parallel to the claims made here. If this first explanation is correct, the link between outcomes and technology adoption should grow initially as the technology is first adopted. If, on the other hand, initial adoption is driven by larger benefits, we should observe the link between technology and outcomes is everywhere decreasing as adoption moves forward. This may suggest a way to test between these theories in these settings.

References

Altonji, Joseph G, Todd E Elder, and Christopher R Taber, "Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools," *Journal of political economy*, 2005, *113* (1), 151–184.

Angrist, Joshua D and Jörn-Steffen Pischke, "The credibility revolution in empirical economics: How better research design is taking the con out of econometrics," *Journal of economic perspectives*, 2010, *24* (2), 3–30.

___ , Peter D Hull, Parag A Pathak, and Christopher R Walters, "Leveraging lotteries for school value-added: Testing and estimation," *The Quarterly Journal of Economics*, 2017, *132* (2), 871–919.

Autier, Philippe, Mathieu Boniol, Cécile Pizot, and Patrick Mullie, "Vitamin D status and ill health: a systematic review," *The lancet Diabetes & endocrinology*, 2014, *2* (1), 76–89.

Bakalar, Nicholas, "Risks: Too Much Vitamin D May Backfire," *New York Times*, Jan 17, 2012.

___ , "Regimens: Study Questions Benefit of Extra Vitamin D," *New York Times*, October 23, 2012.

Bao, Ying, Jiali Han, Frank B Hu, Edward L Giovannucci, Meir J Stampfer, Walter C Willett, and Charles S Fuchs, "Association of nut consumption with total and cause-specific mortality," *New England Journal of Medicine*, 2013, *369* (21), 2001–2011.

Berrigan, David, Kevin Dodd, Richard P Troiano, Susan M Krebs-Smith, and Rachel Ballard Barbash, "Patterns of health behavior in US adults," *Preventive medicine*, 2003, *36* (5), 615–623.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of political Economy*, 1992, *100* (5), 992–1026.

Brown, Deborah J and Lee F Schrader, "Cholesterol information and shell egg consumption," *American Journal of Agricultural Economics*, 1990, *72* (3), 548–555.

Brownlee, Iain A, Carmel Moore, Mark Chatfield, David P Richardson, Peter Ashby, Sharron A Kuznesof, Susan A Jebb, and Chris J Seal, "Markers of cardiovascular risk are not changed by increased whole-grain intake: the WHOLEheart study, a randomised, controlled dietary intervention," *British Journal of Nutrition*, 2010, *104* (1), 125–134.

Chang, Hung-Hao and David R Just, "Health Information Availability and the Consumption of Eggs: Are Consumers Bayesians?," *Journal of Agricultural and Resource Economics*, 2007, pp. 77–92.

Chang, Kenneth, "What we Eat: Doubts by the Teaspoonful," *New York Times*, June 12, 2012.

Chern, Wen S, Edna T Loehman, and Steven T Yen, "Information, health risk beliefs, and the demand for fats and oils," *The Review of Economics and Statistics*, 1995, pp. 555–564.

Colditz, Graham A, JOANN E MANSON, and SUSAN E HANKINSON, "The Nurses' Health Study: 20-year contribution to the understanding of health among women," *Journal of Women's Health*, 1997, *6* (1), 49–62.

Cutler, David M, "Behavioral health interventions: what works and why," *Critical perspectives on racial and ethnic differences in health in late life*, 2004, *643*, 674.

___ **, Adriana Lleras-Muney, and Tom Vogl**, "Socioeconomic status and health: dimensions and mechanisms," Technical Report, National Bureau of Economic Research 2008.

___ **and** ___ , "Understanding differences in health behaviors by education," *Journal of health economics*, 2010, *29* (1), 1–28.

Einav, Liran, Ephraim Leibtag, and Aviv Nevo, "Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?," *Quantitative Marketing and Economics*, 2010.

Esposito, Katherine, Raffaele Marfella, Miryam Ciotola, Carmen Di Palo, Francesco Giugliano, Giovanni Giugliano, Massimo D'armiento, Francesco D'andrea, and Dario Giugliano, "Effect of a Mediterranean-style diet on endothelial dysfunction and markers of vascular inflammation in the metabolic syndrome: a randomized trial," *JAMA*, 2004, *292* (12), 1440–1446.

Estruch, Ramón, Emilio Ros, Jordi Salas-Salvadó, Maria-Isabel Covas, Dolores Corella, Fernando Arós, Enrique Gómez-Gracia, Valentina Ruiz-Gutiérrez, Miquel Fiol, José Lapetra et al., "Primary prevention of cardiovascular disease with a Mediterranean diet," *New England Journal of Medicine*, 2013, *368* (14), 1279–1290.

Feart, Catherine, Cécilia Samieri, Virginie Rondeau, Hélène Amieva, Florence Portet, Jean-François Dartigues, Nikolaos Scarmeas, and Pascale Barberger-Gateau, "Adherence to a Mediterranean diet, cognitive decline, and risk of dementia," *JAMA*, 2009, *302* (6), 638–648.

Finke, Michael S and Sandra J Huston, "Factors affecting the probability of choosing a risky diet," *Journal of Family and Economic Issues*, 2003, *24* (3), 291–303.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams, "Place-Based Drivers of Mortality: Evidence from Migration," *Stanford University Working Paper*, 2018.

Friel, Sharon, John Newell, and Cecily Kelleher, "Who eats four or more servings of fruit and vegetables per day? Multivariate classification tree analysis of data from the 1998 Survey of Lifestyle, Attitudes and Nutrition in the Republic of Ireland," *Public health nutrition*, 2005, *8* (2), 159–169.

Goldman, Dana P and James P Smith, "Can patient self-management help explain the SES health gradient?," *Proceedings of the National Academy of Sciences*, 2002, *99* (16), 10929–10934.

Greenland, Sander, James M Robins, and Judea Pearl, "Confounding and collapsibility in causal inference," *Statistical science*, 1999, pp. 29–46.

Hall, Bronwyn H and Beethika Khan, "Adoption of new technology," Technical Report, National bureau of economic research 2003.

HECKMAN, JAMES J, "DUMMY ENDOGENOUS VARIABLES IN A SIMULTANEOUS EQUATION SYSTEM1," *Econometrica*, 1978, *46* (4), 931–959.

Holick, Michael F, "Vitamin D deficiency," *New England Journal of Medicine*, 2007, *357* (3), 266–281.

Ippolito, Pauline M and Alan D Mathios, "Information and advertising: The case of fat consumption in the United States," *The American Economic Review*, 1995, *85* (2), 91–95.

Kinnucan, Henry W, Hui Xiao, Chung-Jen Hsia, and John D Jackson, "Effects of health information and generic advertising on US meat demand," *American Journal of Agricultural Economics*, 1997, *79* (1), 13–23.

Kirkpatrick, Sharon I, Kevin W Dodd, Jill Reedy, and Susan M Krebs-Smith, "Income and race/ethnicity are associated with adherence to food-based dietary guidance among US adults and children," *Journal of the Academy of Nutrition and Dietetics*, 2012, *112* (5), 624–635.

Knoops, Kim TB, Lisette CPGM de Groot, Daan Kromhout, Anne-Elisabeth Perrin, Olga Moreiras-Varela, Alessandro Menotti, and Wija A Van Staveren, "Mediterranean diet, lifestyle factors, and 10-year mortality in elderly European men and women: the HALE project," *JAMA*, 2004, *292* (12), 1433–1439.

Kolata, Gina, "Healthy Women Advised Not to Take Calcium and Vitamin D to Prevent Fractures," *New York Times*, Jun 13, 2012.

Kowalski, Amanda E, "Behavior within a Clinical Trial and Implications for Mammography Guidelines," Technical Report, National Bureau of Economic Research 2018.

Krauss, Ronald M, Robert H Eckel, Barbara Howard, Lawrence J Appel, Stephen R Daniels, Richard J Deckelbaum, John W Erdman Jr, Penny Kris-Etherton, Ira J Goldberg, Theodore A Kotchen et al., "AHA Dietary Guidelines: revision 2000: A statement for healthcare professionals from the Nutrition Committee of the American Heart Association," *Circulation*, 2000, *102* (18), 2284–2299.

Loftfield, Erikka, Marilyn C Cornelis, Neil Caporaso, Kai Yu, Rashmi Sinha, and Neal Freedman, "Association of Coffee Drinking With Mortality by Genetic Variation in Caffeine Metabolism: Findings From the UK Biobank," *JAMA internal medicine*, 2018, *178* (8), 1086–1097.

Maki, Kevin C, Joanne L Slavin, Tia M Rains, and Penny M Kris-Etherton, "Limitations of observational evidence: implications for evidence-based dietary recommendations," *Advances in nutrition*, 2014, *5* (1), 7–15.

Miller, Edgar R, Roberto Pastor-Barriuso, Darshan Dalal, Rudolph A Riemersma, Lawrence J Appel, and Eliseo Guallar, "Meta-analysis: high-dosage vitamin E supplementation may increase all-cause mortality," *Annals of internal medicine*, 2005, *142* (1), 37–46.

Nagourney, Eric, "Shortage of Vitamin D May Weaken the Elderly," *New York Times*, May 1, 2007.

Oster, Emily, "Unobservable selection and coefficient stability: Theory and evidence," *Journal of Business & Economic Statistics*, 2018, pp. 1–18.

Pei, Zhuan, Jörn-Steffen Pischke, and Hannes Schwandt, "Poorly measured confounders are more useful on the left than on the right," *Journal of Business & Economic Statistics*, 2018, (just-accepted), 1–34.

Peterkin, BB, "Dietary guidelines for Americans," *Journal of the American Dietetic Association*, 1990, *90* (12), 1725–1727.

Rimm, Eric B, Meir J Stampfer, Alberto Ascherio, Edward Giovannucci, Graham A Colditz, and Walter C Willett, "Vitamin E consumption and the risk of coronary heart disease in men," *New England Journal of Medicine*, 1993, *328* (20), 1450–1456.

Roosen, Jutta, Stéphan Marette, Sandrine Blanchemanche, and Philippe Verger, "Does health information matter for modifying consumption? A field experiment measuring the impact of risk information on fish consumption," *Review of Agricultural Economics*, 2009, *31* (1), 2–20.

Rosen, Clifford J, Steven A Abrams, John F Aloia, Patsy M Brannon, Steven K Clinton, Ramon A Durazo-Arvizu, J Christopher Gallagher, Richard L Gallo, Glenville Jones, Christopher S Kovacs et al., "IOM committee members respond to Endocrine Society vitamin D guideline," *The Journal of Clinical Endocrinology & Metabolism*, 2012, *97* (4), 1146–1152.

Stampfer, Meir J, Charles H Hennekens, JoAnn E Manson, Graham A Colditz, Bernard Rosner, and Walter C Willett, "Vitamin E consumption and the risk of coronary disease in women," *New England Journal of Medicine*, 1993, *328* (20), 1444–1449.

Strang, David and Sarah A Soule, "Diffusion in organizations and social movements: From hybrid corn to poison pills," *Annual review of sociology*, 1998, *24* (1), 265–290.

Sugar, *60 Minutes*, *2012.*

Taubes, Gary, "Is Sugar Toxic?," *New York Times*, April 17, 2011.

Thompson, Tommy G and Ann M Veneman, "Dietary guidelines for Americans 2005," *United States Department of Health and Human Services and United States Department of Agriculture*, 2005.

Thorisson, Hermann, "Coupling methods in probability theory," *Scandinavian journal of statistics*, 1995, pp. 159–182.

Trichopoulou, Antonia, Tina Costacou, Christina Bamia, and Dimitrios Trichopoulos, "Adherence to a Mediterranean diet and survival in a Greek population," *New England Journal of Medicine*, 2003, *348* (26), 2599–2608.

Vandenbroucke, Jan P, Erik Von Elm, Douglas G Altman, Peter C Gøtzsche, Cynthia D Mulrow, Stuart J Pocock, Charles Poole, James J Schlesselman, Matthias Egger, Strobe Initiative et al., "Strengthening the Reporting of Observational Studies in Epidemiology (STROBE): explanation and elaboration," *PLoS medicine*, 2007, *4* (10), e297.

Yang, Quanhe, Zefeng Zhang, Edward W Gregg, W Dana Flanders, Robert Merritt, and Frank B Hu, "Added sugar intake and cardiovascular diseases mortality among US adults," *JAMA internal medicine*, 2014, *174* (4), 516–524.

Figure 1: **Vitamin Consumption Levels over Time**

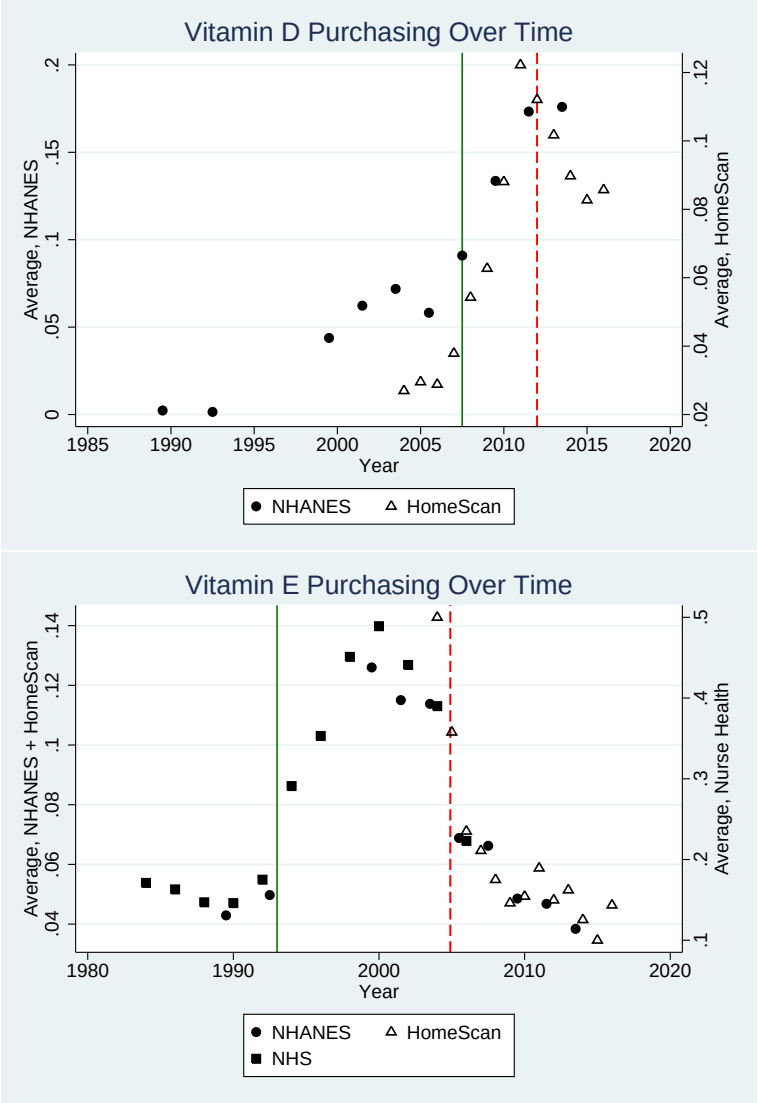

Notes: These graphs show evolution of vitamin consumption or purchasing over time in the NHANES, HomeScan and Nurse Health Study. HomeScan data is based on purchase behavior and NHANES and NHS are based on reported consumption. Events are marked with vertical lines; details of the events appear in Appendix Table A1.

Figure 2: **Vitamin E Consumption and Health Behaviors**

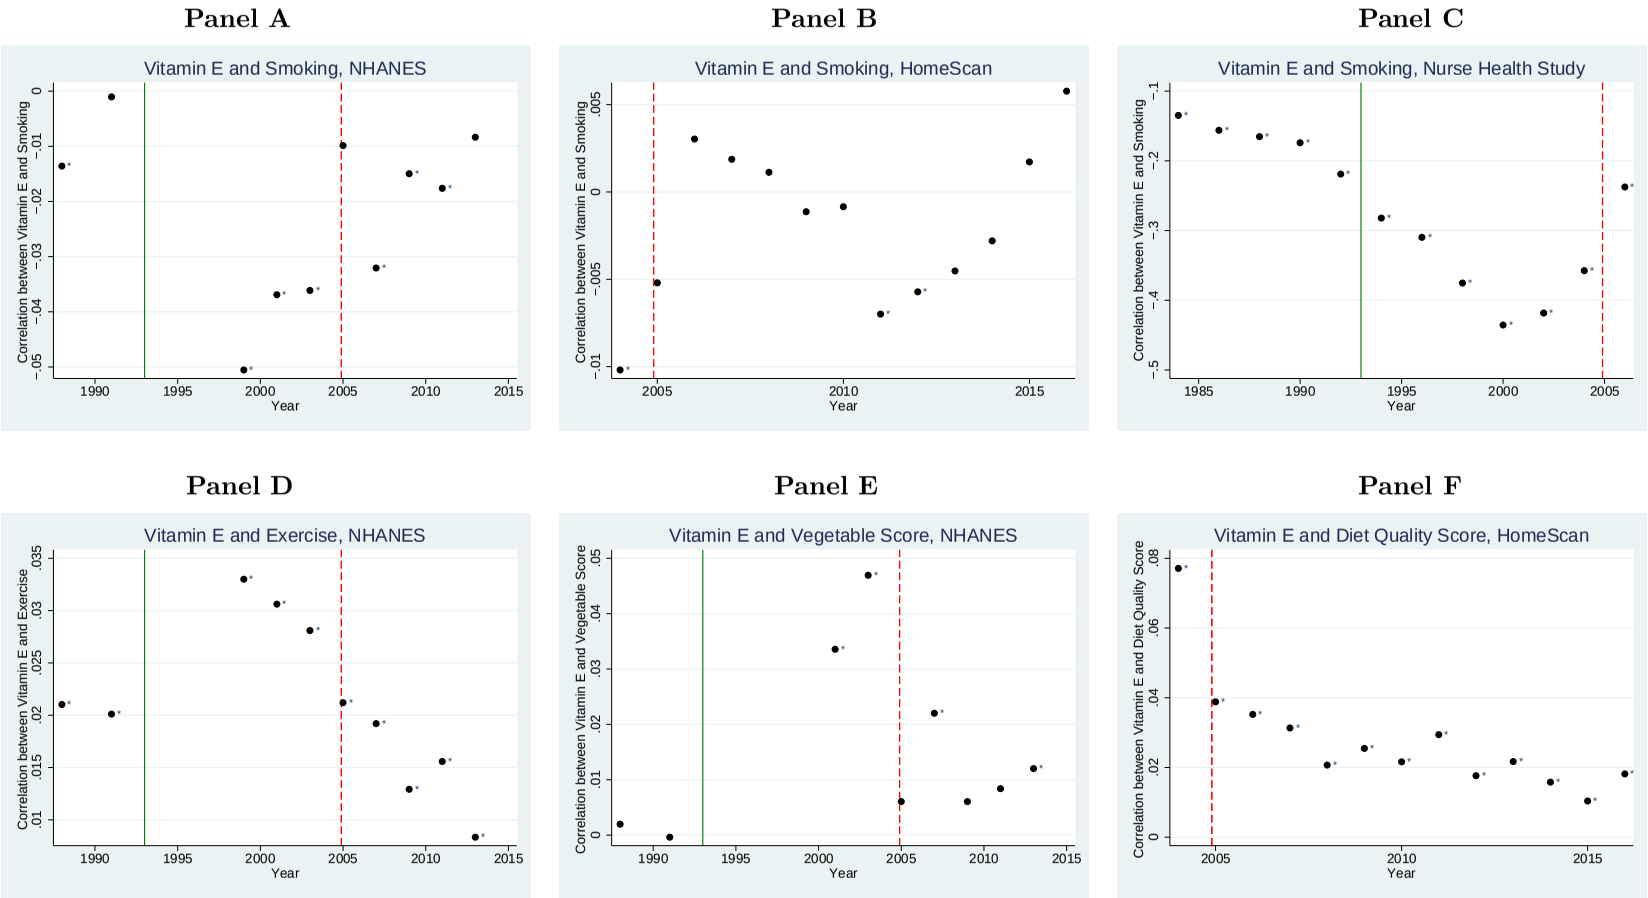

Notes: These figures show the correlation between Vitamin E consumption, various health behaviors and education over time. Each point comes from a regression of an indicator for vitamin E consumption on the health behavior or education variable; also included in this regression are age and (when relevant) gender controls. Panels A and B consider the relationship between vitamin E and smoking in the NHANES (Panel A) and the Nurse Health Study (Panel B). In the NHS this effect is reported as a share of smokers; in the NHANES it is a percentage point difference. Panel C looks at the relationship between Vitamin E and exercise in the NHANES, and Panel D at the relationship between vitamin E and exercise. Events are marked with vertical lines; details of the events appear in Appendix Table A1. Solid green lines indicate the release of positive news about vitamin E; dotted red lines indicate the release of negative news. ∗indicates significance at the 5% level.

Figure 3: **Vitamin Consumption, Socioeconomic Status and Health Behaviors**

Panel A

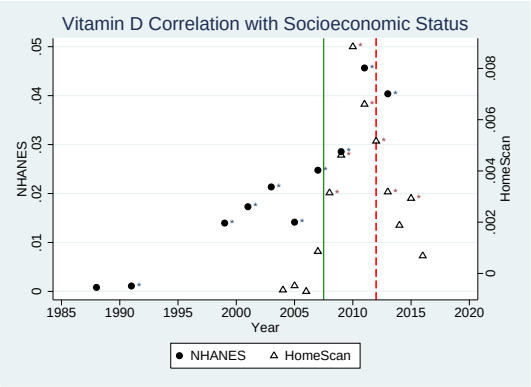

Panel B

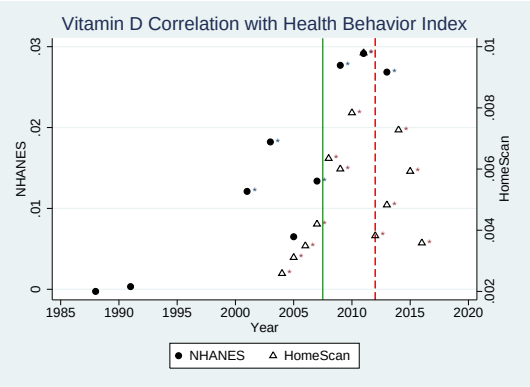

Panel C

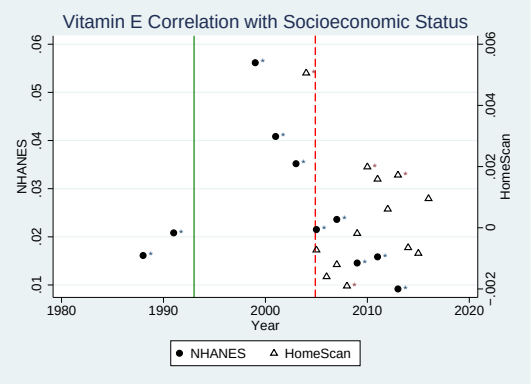

Panel D

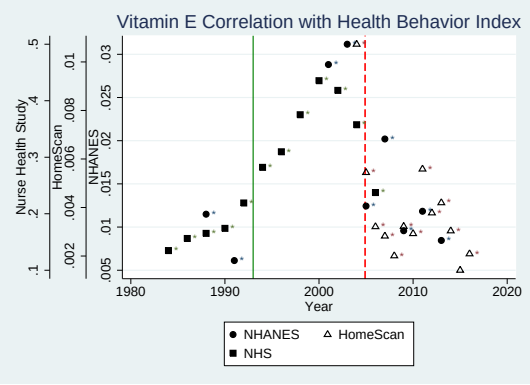

Notes: These figures show changes in the correlation between vitamin consumption and health behaviors or socioeconomic status over time. Socioeconomic status is an index created as the first principal component of education and income. The health behavior index is created based on the set of health behavior available in each dataset (for NHANES: first principal component of exercise, smoking and diet quality; for HomeScan: first principal component of smoking, diet quality; for NHS: smoking). Panels A and B look at vitamin D, Panels C and D look at Vitamin E. Events are marked with vertical lines; details of the events appear in Appendix Table A1. Solid green lines indicate the release of positive news; dotted red lines indicate the release of negative news. ∗indicates significance at the 5% level.

Figure 4: **Vitamin E and Mortality: Nurse Health Study**

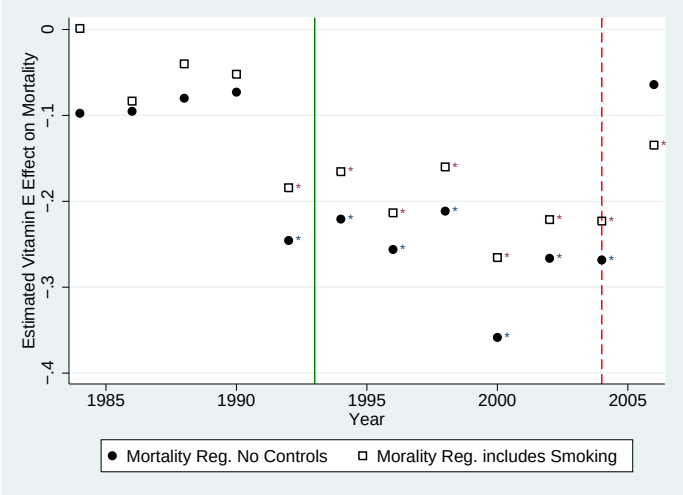

Notes: This figure shows the estimated impact of vitamin E consumption on two-year mortality in the Nurse Health Study data over time. Regressions in each period include age fixed effects. Gradients are scaled as a share of the average death rates so can be interpreted as a percent decrease in death rate as a result of reported vitamin E consumption. Events are marked with vertical lines; details of the events appear in Appendix Table A1. Solid green lines indicate the release of positive news about vitamin E; dotted red lines indicate the release of negative news. ∗indicates significance at the 5% level.

Figure 5: **Evidence from Publications on Vitamins**

Panel A

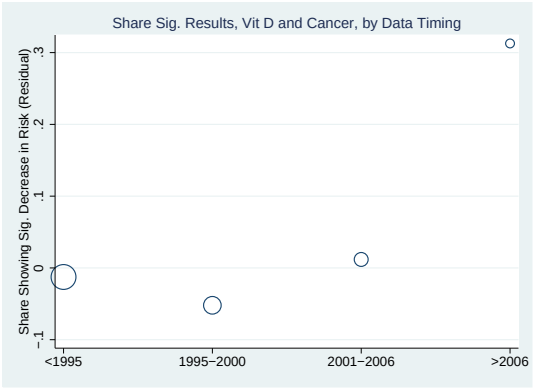

Panel B

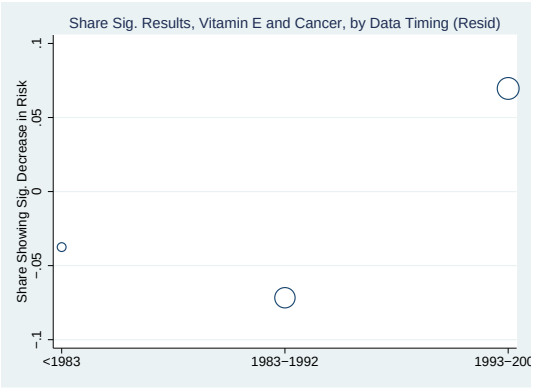

Notes: These figures show the share of significant vitamin-cancer relationships in published work using data from each period. Publications are identified from Pubmed searches and from published meta-analyses. The outcome (significant negative relationship between vitamin supplementation and cancer) is residualized with respect to the type of cancer. Studies with data which overlaps the time periods is assigned a partial weight in each time period. Panel A focuses on Vitamin D, Panel B on Vitamin E. The markers are scaled to represent the number of studies in each bin.

Figure 6: **Changes in Diet Behaviors over Time**

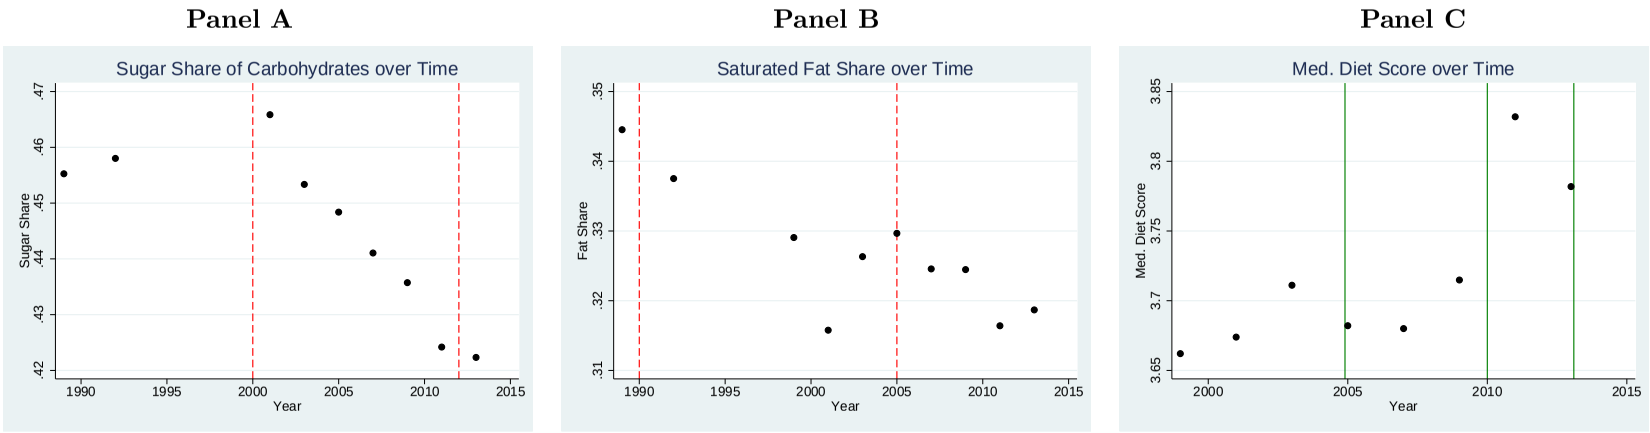

Notes: These graphs show evolution of dietary patterns over time in the NHANES data. Panel A shows changes in the sugar share of carbohydrates; Panel B shows changes in Saturated fat as a share of total fat; Panel C shows changes in the Mediterranean diet score. This last metric is calculated as described in Trichopoulou et al (2003). Events are marked with vertical lines; details of the events appear in Appendix Table A1. Solid green lines indicate the release of positive news about vitamin E; dotted red lines indicate the release of negative news.

Figure 7: **Dietary Patterns, Socioeconomic Status and Health Behaviors**

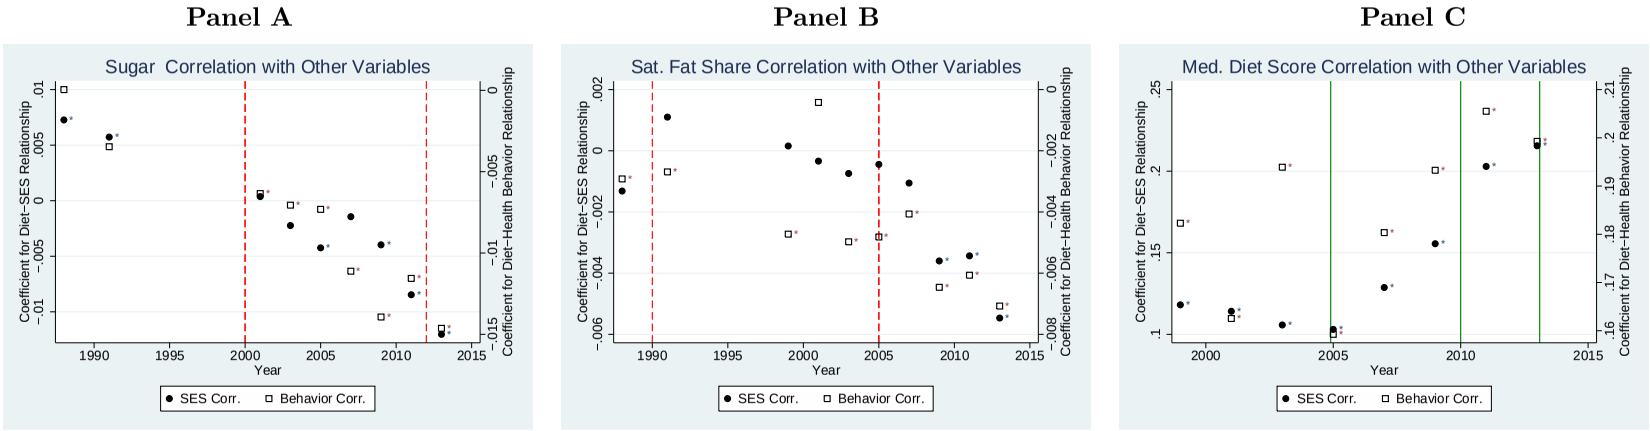

Notes: These figures show changes in the correlation between diet and health behaviors or socioeconomic status over time. Socioeconomic status is an index created as the first principal component of education and income. The health behavior index is the first principal component of exercise and smoking. Panels A looks at sugar, Panels B at saturated fat and Panel Cat the Mediterranean diet. Events are marked with vertical lines; details of the events appear in Appendix Table A1. Solid green lines indicate the release of positive news about vitamin E; dotted red lines indicate the release of negative news. ∗indicates significance at the 5% level.

Figure 8: **Diet and BMI Relationship Over Time**

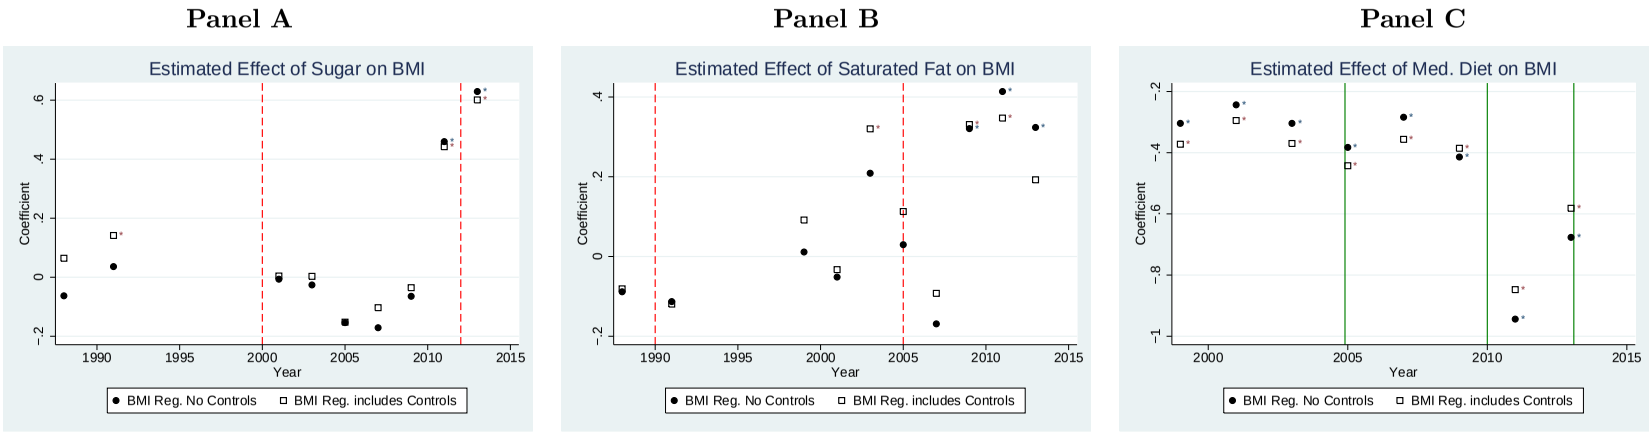

Notes: These graphs show the evolution of the correlation between BMI and diet over time. The filled in circles are coefficients from regressions of BMI on diet including only age and gender controls. The hollow circles come from regressions which also include demographic and health behavior controls. Events are marked with vertical lines; details of the events appear in Appendix Table A1. Solid green lines indicate the release of positive news about vitamin E; dotted red lines indicate the release of negative news. ∗indicates significance at the 5% level.

Table 1: **Correlation between Vitamin Consumption and Health Behaviors/Socioeconomic Status**

	Panel A: Vitamin D			**Panel B: Vitamin E**		
	Before 2007	*2007 - 2012*	After 2012	Before 1993	*1993 - 2004*	After 2004
Smoking						
NHANES (1991-2013)	-0.014‡,‡‡	*-0.042*	-0.047*	-0.007‡	*-0.041*‡	-0.016
HomeScan (2004-2015)	-0.004‡,‡‡	*-0.014*‡	-0.010		*-0.010**,‡	0.002
NHS (1984-2006)				-0.153‡,‡‡	*-0.368*‡	-0.237
Exercise						
NHANES(1991-2013)	0.003‡,‡‡	*0.019*	0.024*	0.020‡,‡‡	*0.031*‡	0.015
Higher quality diet						
NHANES(1991-2013)	0.005‡,‡‡	*0.026*	0.026*	0.0008‡,‡‡	*0.043*‡	0.011
HomeScan (2004-2015)	0.020‡	*0.038*‡	0.018		*0.077**,‡	0.018
Education						
NHANES(1991-2013)	0.011‡,‡‡	*0.030*	0.031*	0.015‡	*0.038*‡	0.015
HomeScan (2004-2015)	0.0003‡	*0.006*‡	0.002		*0.0049**,‡	-0.0002
Income						
NHANES(1991-2013)	0.007‡,‡‡	*0.010*‡	0.015*	0.007‡	*0.014*‡	0.005
HomeScan (2004-2015)	-0.0003‡,‡‡	*0.0009*‡	0.0003		*0.0008**,‡	-0.0003

Notes: This table shows correlations between vitamin consumption, other health behaviors and socioeconomic status over time. To generate these, in the NHANES I use the full dataset to regress vitamin consumption on the behavior or socioeconomic status variable, interacted with period. I include controls for the period, as well as for age, age square and gender (all interacted with the period to allow for flexibility). The procedure is the same in the HomeScan and NHS data, but the HomeScan controls for age, household size and household competition, and the NHS only for age. Smoking results in the NHS are reported as a share reduction in the smoking rate relative to those who do not take Vitamin E. The periods are divided based on the events detailed in Appendix Table A1. * indicates a single year of data used in this cell; ‡ significantly different from next period at 5% level; ‡‡significantly different from two periods later at 5% level.

Table 2: **Correlation between Dietary Choices and Health Behaviors/Socioeconomic Status**

	Panel A: Sugar Share of Carbohydrates			**Panel B: Saturated Fat (Share of Fat)**			**Panel C: Med. Diet Score**		
	Before 2000	2000-2012	After 2012	Before 1990	1990-2005	After 2005	Before 2005	2005-2010	After 2010
Smoking	0.0104‡,‡‡	0.0262‡	0.0444*	0.0062*,‡‡	0.0097	0.0120	-0.443	-0.406	-0.421
Exercise	0.0020‡,‡‡	-0.0053	-0.0052*	-0.0018*,‡‡	-0.0008‡	-0.0044	0.094‡‡	0.115	0.151
Education	0.0077‡,‡‡	0.0001	-0.0041*	-0.0008*	-0.0008	-0.0016	0.070‡,‡‡	0.097‡	0.134
Income	0.0018‡,‡‡	-0.0026‡	-0.0059*	-0.0006*	-0.001	-0.0014	0.050‡‡	0.048‡	0.083

Notes: This table shows correlations between dietary choices, other health behaviors and socioeconomic status over time. All data come from the NHANES, 1988 through 2013. To generate these, in the NHANES I use the full dataset to regress vitamin consumption on the behavior or socioeconomic status variable, interacted with period. I include controls for the period, as well as for age, age square and gender (all interacted with the period to allow for flexibility). The periods are divided based on the events detailed in Appendix Table A1. * indicates a single year of data used in this cell; ‡ significantly different from next period at 5% level; ‡‡significantly different from two periods later at 5% level.

Table 3: **Co-movement: Behavior-Outcome Gradient and Selection Gradients**

Panel A: Outcome Regression With No Controls				
Outcome:	BMI Gradient	Heart Health Gradient	BMI Gradient	Heart Health Gradient
SES Gradient [n=27]	-5.84***	0.617**		
	(1.06)	(0.161)		
Health Behavior Gradient [n=27]			-5.08***	0.627**
			(1.51)	(0.195)
Panel B: Outcome Regression With Controls				
Outcome:	BMI Gradient	Heart Health Gradient	BMI Gradient	Heart Health Gradient
SES Gradient [n=27]	-4.07***	0.328**		
	(1.08)	(0.145)		
Health Behavior Gradient [n=27]			-3.90***	0.31*
			(1.34)	(0.172)
Diet Behavior FE	YES	YES	YES	YES

Notes: This table shows evidence on the co-movements between diet-selection correlations and diet-health outcome correlations. The left hand side in each regression is a regression coefficient from a regression of health outcome on diet choice. The right hand side is the regression coefficient from a regression of dietary choices on either a composite socioeconomic status measure or a composite health behavior measure. In all cases the selection regressions include only age and gender controls. In Panel A the health outcome-behavior regressions also adjust only for age and gender. In Panel B these also include controls for race, marital status, education, income, exercise and smoking behavior. $*$indicates significance at the 10% level, $**$indicates significance at the 5% level, $***$ indicates significance at the 1% level.

Appendix A: Figures and Tables

Table A1: **Information Events**

Behavior	***Event 1***	***Event 2***	***Event 3***
Vitamin D Supplementation	**2007, Positive:** Several studies, NEJM summary piece, NY Times coverage suggest Vitamin D good for health (cancer, fractures, etc). Corresponding growth in Google Trends.	**2011/2012, Negative:** IOM report suggests Vitamin D overblown, corresponding summary articles, coverage in NY Times. Additional studies in 2012 with similar findings. Google Trends stagnation.	
Vitamin E Supplementation	**1993, Positive:** Two studies in NEJM report reduction in heart disease for both men and women with use of Vitamin E supplements	**2004, Negative:** Widely covered meta-analsis of Vitamin E shows high doses increase mortality. Large Google trends spike.	
Sugar in Diet	**2000, Negative:** First explicit mention in US Dietary guidelines of avoidance of added sugars.	**2011/2012, Negative:** Extensive media coverage of health costs of sugar; "toxic sugar" in NY Times and 60 Minutes Segment.	
Saturated Fat	**1990, Negative:** First explicit restriction on saturated fat share in US dietary guidelines (<10%)	**2005, Negative:** Further restrict saturated fat to <7% for people with heart disease.	
Mediterranean Diet	**2004, Positive:** Two JAMA articles show health benefits of Mediterranean diet. Google trends spike.	**2009/2010, Positive:** Series of articles on role of Mediterranean diet in addressing cognitive decline. Google trends spike.	**2013, Positive:** Large randomized trial shows mortality reductions from Mediterranean Diet. Heavily covered in media. Very large Google trend spike.

Notes: This table shows the information events identified for each outcome. Events were identified by searching for well-cited publications, media coverage and Google search spikes.

Appendix B: Proofs

Proof of Proposition 1 (Baseline Case) At time $t = 0$, if $\kappa_j = 0$, then whether subject i undertakes Λ_j is solely determined by $c_{i,j}$, which is independent of $\Lambda_{j\prime}$. Thus, $Cov_{t=0}(\Lambda_j, \Lambda_{j\prime}) = 0$.

We then note

$$Cov_{t=1}(\Lambda_j, \Lambda_{j\prime}) = \mathbb{E}1[\alpha_i\kappa_j \geq c_{i,j}]1[\alpha_i\kappa_{j\prime} \geq c_{i,j\prime}] - \mathbb{E}1[\alpha_i\kappa_j \geq c_{i,j}]\mathbb{E}1[\alpha_i\kappa_{j\prime} \geq c_{i,j\prime}]$$

where $1[\cdot]$ is an indicator function.

Given that the costs are normally distributed and independent, denoting by Φ the cdf of the standard normal distribution, and using the law of iterated expectations, we obtain:

$$Cov_{t=1}(\Lambda_j, \Lambda_{j\prime}) = \mathbb{E}\Phi(\frac{\alpha_i\kappa_j - c_j}{\sigma_j})\Phi(\frac{\alpha_i\kappa_{j\prime} - c_{j\prime}}{\sigma_{j\prime}}) - \mathbb{E}\Phi(\frac{\alpha_i\kappa_j - c_j}{\sigma_j})\mathbb{E}\Phi(\frac{\alpha_i\kappa_{j\prime} - c_{j\prime}}{\sigma_{j\prime}})$$

where all expectations are taken with respect to health value α_i.

Note, that on this step we used the fact that both behaviors are independent conditionally on α_i, which is implied by the linear form.

The right hand side of the inequality has the form of $\mathbb{E}f(\alpha_i)g(\alpha_i) - \mathbb{E}f(\alpha_i)\mathbb{E}g(\alpha_i)$, where f, g are strictly increasing (given that $\kappa_j, \kappa_{j\prime} > 0$) bounded functions.

Given that α_i is not degenerate and has non-zero density everywhere by assumption, by the covariance inequality (Thorisson, 1995) this value is positive. Hence, $Cov_{t=1}(\Lambda_j, \Lambda_{j\prime}) > 0$.

Proof of Proposition 2 (Covariance with Other Variables) We assume that Z is independent of $c_{i,j}$. Since at time t we have $\kappa_j = 0$, then Λ_j only depends on $c_{i,j}$. Hence, $Cov_{t=0}(\Lambda_j, Z) = 0$. Now we will show that $Cov_{t=1}(\Lambda_j, Z) > 0$.

Similarly to the proof of Proposition 1, to establish the positive covariance between the variables at time $t = 1$ we use the law of iterated expectations and the fact that conditionally on α_i, Λ_j and Z are independent.

$$Cov_{t=1}(\Lambda_j, Z) = \mathbb{E}(\Phi(\frac{\alpha_i\kappa_j - c_j}{\sigma_j})\mathbb{E}[Z|\alpha_i]) - \mathbb{E}\Phi(\frac{\alpha_i\kappa_j - c_j}{\sigma_j})\mathbb{E}(\mathbb{E}[Z|\alpha_i]).$$

The assumption of increasing $\mathbb{E}[Z|\alpha_i]$ and the covariance inequality yield the result.

Proof of Proposition 3 (Disease-Behavior Dynamics)

(A) We can write the covariance thus:

$$Cov_{t=0}(\Lambda_j, Y) = Cov_{t=0}(\Lambda_j, \mu + \sum_{\Lambda_r \in \Lambda} \vartheta_r\Lambda_r + \epsilon_i) = \sum_{\Lambda_r \in \Lambda} \vartheta_r Cov_{t=0}(\Lambda_j, \Lambda_r).$$

Case 1 If $\vartheta_j = 0$,then from Proposition 1 it follows that $Cov_{t=0}(\Lambda_j, Y) = 0$. Analogously, from Proposition 1 it also follows that $Cov_{t=1}(\Lambda_j, Y) > 0$. Hence, in this case, **(A)** is established.

Case 2 If $\vartheta_j \neq 0$ then) and

$$Cov_{t=0}(\Lambda_j, Y) = \vartheta_j Var_{t=0}(\Lambda_j)$$
$$Cov_{t=1}(\Lambda_j, Y) = \vartheta_j Var_{t=1}(\Lambda_j) + \sum_{\Lambda_r \in \Lambda/\Lambda_j} \vartheta_r Cov_{t=1}(\Lambda_j, \Lambda_r)$$

As a result, $Cov_{t=1}(\Lambda_j, Y) > Cov_{t=0}(\Lambda_j, Y)$ if and only if

$$\vartheta_j Var_{t=1}(\Lambda_j) + \sum_{\Lambda_r \in \Lambda/\Lambda_j} \vartheta_r Cov_{t=1}(\Lambda_j, \Lambda_r) > \vartheta_j Var_{t=0}(\Lambda_j)$$

Proposition 1 establishes that $\sum_{\Lambda_r \in \Lambda/\Lambda_j} \vartheta_r Cov_{t=1}(\Lambda_j, \Lambda_r) > 0$ but does not tell us how $Var_{t=1}(\Lambda_j)$ compares to $Var_{t=0}(\Lambda_j)$.
At time $t = 0$ we have $\kappa_j = 0$. Thus, $Var_{t=0}(\Lambda_j) = \mathbb{E}1[c_{i,j} \leq 0] - (\mathbb{E}1[c_{i,j} \leq 0])^2 = \Phi(\frac{-c_j}{\sigma_j}) - \Phi^2(\frac{-c_j}{\sigma_j})$.
At $t = 1$, $\kappa_j > 0$. Hence, $Var_{t=1}(\Lambda_j) = \mathbb{E}1[c_{i,j} \leq \alpha_i \kappa_j] - (\mathbb{E}1[c_{i,j} \leq \alpha_i \kappa_j])^2 = \mathbb{E}\Phi(\frac{\alpha_i \kappa_j - c_j}{\sigma_j}) - (\mathbb{E}\Phi(\frac{\alpha_i \kappa_j - c_j}{\sigma_j}))^2$.
It is possible that this variance is lower in $t = 1$ than in $t = 1$ if κ_j is very large at time $t = 1$;that is, if almost everyone adopts it, then $Var_{t=1}(\Lambda_A) \approx 0$. However, this will not happen as long as κ_j is relatively small.

(B)

Case 1 Assume $\vartheta_j = 0$. Recall Ω is defined as subset of Λ which excludes at least one behavior Λ_p for which $\vartheta_p \neq 0$. We will show the proof under the assumption that a single behavior is excluded from Ω;the result would strengthen with more behaviors excluded.
At time $t = 0$ we can write

$$Cov_{t=0}(\Lambda_j, Y|\Omega) = \sum_{\Lambda_r \in \Lambda} \vartheta_r Cov_{t=0}(\Lambda_j, \Lambda_r|\Omega).$$

Note that since $\kappa_j = 0$ at $t = 0$ we have Λ_j independent of $\Lambda_r$for any r, even conditioning on Ω. Hence, $Cov_{t=0}(\Lambda_j, Y|\Omega) = 0$.
At time $t = 1$, the proof of Proposition 1 shows that $Cov_{t=1}(\Lambda_j, \Lambda_r) > 0$. For $\Lambda_r \in \Omega$, we have $Cov_{t=1}(\Lambda_j, \Lambda_r|\Omega) = 0$. However, given that behavior Λ_p is not included in Ω we have $Cov_{t=1}(\Lambda_j, \Lambda_p|\Omega) > 0$ and, as a result, $Cov_{t=0}(\Lambda_j, Y|\Omega) > 0$.

Case 2 Assume $\vartheta_j > 0$. Combining the logic in **(A)** above with that in case 1 here, we can see the inequality holds if

$$\vartheta_j Var_{t=1}(\Lambda_j|\Omega) + \vartheta_p Cov_{t=1}(\Lambda_j, \Lambda_p) > \vartheta_j Var_{t=0}(\Lambda_j|\Omega)$$

As above, this will hold as long as κ_j is not very large at time $t = 1$.