Together We Tip the Scale: The Spatial Concentration of Obesity[*]

Panos Margaris[†] Johanna Wallenius[‡]

March 4, 2026

Abstract

In this paper, we develop a life cycle model that features food consumption, exercise, and deviation from the reference BMI to rationalize spatial concentration and the educational gradient in obesity in the US. Agents face ex ante heterogeneity with respect to the area of residence and the strenuousness of occupation, and choose general consumption, food consumption, savings, and exercise. BMI is determined by calorific balance and affects health and the distribution of medical spending, the probability of survival, and the level of utility. We find that the reference BMI plays an important role in the spatial and educational gradients observed in obesity, and moving an individual from a low to a high obesity region increases average weight by 4.6 pounds. We introduce a GLP-1 treatment policy targeting individuals with the highest food preferences and find substantial direct effects on BMI, with total reductions ranging from 1.3 to 4.8 pounds depending on the region and targeting intensity. Importantly, we document meaningful general equilibrium effects through endogenous adjustment of reference BMI, with spillover effects accounting for 25% to 48% of the direct treatment effect.

JEL Classification: I12, I14, D15, E21

Keywords: Obesity, BMI, Networks, Time Use, Life Cycle

[*]We thank conference and seminar participants at McGill, SSE, EEA-ESEM 2025, World Congress 2025 and Workshop on the Economic Challenges of Demographic Inequalities. Wallenius gratefully acknowledges financial support from the Knut and Alice Wallenberg Foundation, the Ragnar Söderberg Foundation and the Torsten Söderberg Foundation. The computations were enabled by resources provided by Compute Canada.

[†]Department of Economics, Concordia University, panos.margaris@concordia.ca

[‡]Department of Economics, Stockholm School of Economics johanna.wallenius@hhs.se

1 Introduction

The prevalence of obesity in the US is high, with more than 30% of American adults classified as obese (2018, NHIS). When the definition is broadened to include overweight individuals, the share rises to approximately 66%. Both figures are up markedly from 9% and 38%, respectively, in the mid-1970s.[1] While average BMI is high throughout the US and in all demographic groups, there are regional and socioeconomic differences. Obesity rates are higher in the South than in the West, and higher among less educated individuals than their more educated counterparts.[2] In this paper, we seek to understand the socioeconomic and spatial variation in BMI in the US. To do so, we develop a heterogeneous agent, life cycle model, where individuals work in occupations with different strenuousness, choose food consumption and exercise, and face uncertainty with respect to health, medical spending, and health insurance coverage.

Using data from the Centers of Disease Control (CDC), Current Population Survey (CPS), and National Health Interview Survey (NHIS), we document three stylized facts about BMI and obesity: (i) average BMI and the prevalence of obesity have been steadily increasing over the last several decades, (ii) high school graduates are more likely to be overweight/obese compared to college graduates, and (iii) obesity is spatially concentrated.

The spatial concentration of obesity can partially be explained by different socioeconomic environments across areas of the US. We perform spatial analysis at the county level, controlling for observable characteristics, state fixed effects, and the spatial correlation of residuals. We find that obesity is indeed spatially correlated, and that differences in education, income, racial composition, and the labor market cannot fully explain the spatial concentration. We use this to motivate including region-specific reference BMI in the utility function of agents. We use the spatial analysis to create what we term "hot spots" and "cold spots" of obesity. Throughout the paper, we refer to "hot spots" and "cold spots" as areas in the US that have higher or lower than average obesity prevalence.

[1] The most common method to measure and identify obesity is the Body Mass Index (BMI), which is measured as kg/m^2. Individuals are classified as obese if their BMI is greater than or equal to 30. BMI of 25 or above is considered overweight.

[2] 32.5% of individuals are classified as obese in the South, compared to 27.7% in the West (NHIS, 2018). 34.5% of high school graduates are classified as obese, compared to 28% of those with at least some years of college education (NHIS, 2018).

To get at potential drivers of high BMI, we also document that work has become increasingly sedentary and less strenuous over time, and that calorie consumption has moderately increased. However, we show that the reduction in the strenuousness of occupations conceals an important composition effect. Occupations have become on average more sedentary because of the increase in the share of college graduates who are more likely to be employed in sedentary occupations. Moreover, college occupations have become on average less strenuous over time. In contrast, occupations of high school graduates, who are more likely to be obese, have become more strenuous over time due to job polarization (Autor, Katz and Kearney, 2006). Moreover, even though the average strenuousness of occupations has decreased in all areas of the US, occupations in the "hot spots" of obesity are more strenuous than the occupations in the "cold spots" of obesity.

Furthermore, we find that, even though both high school and college graduates are more likely to be obese in the "hot spots" relative to the "cold spots" of obesity and high school graduates are more likely to be obese compared to college graduates within each area, the difference in obesity prevalence is an artifact of the Body Mass Index (BMI) cut-off for the classification of obesity. The difference between the average BMI across areas is significantly smaller than the difference in obesity prevalence suggests.[3] The exception are college graduates in "cold spots", who have significantly lower BMI compared to college graduates in "hot spots" and high school graduates in general.

We develop a life cycle model to disentangle the drivers of the spatial and socio-economic variation in BMI. In the model, agents make decisions regarding food consumption and exercise in a realistic environment where BMI is determined by calorific balance, which depends on the lifestyle choices of the agents, and their occupation, which affects strenuousness during working hours. In addition, to rationalize the obesity patterns, we introduce the concept of reference BMI that is area specific. This takes the form of conspicuous consumption with loss aversion in networks, similar to Bramoullé and Ghiglino (2022). Agents derive disutility from deviating from the healthy BMI and the reference BMI, which is the average BMI of agents re-

[3]The obesity prevalence of high school graduates is 44% in "hot spots" and 38.2% in "cold spots", and the obesity prevalence of college graduates is 35.8% in "hot spots" and 25.1% in "cold spots" of obesity. High school graduates have an average BMI of 29.8 in "hot spots" and 29.2 in "cold spots", while college graduates have an average BMI of 29 in "hot spots" and 27.6 in "cold spots". A one unit change in BMI reflects 6.9 pounds in weight for the average male in the US.

siding in the same area. In our model, overweight individuals face objective costs in terms of worse health, higher medical spending and higher mortality, and subjective costs from having a higher BMI compared to their peers and deviating from the healthy BMI.

We find that the reference BMI and the disutility from exercise are the most important factors in accounting for the spatial pattern to BMI. If the obesity hot spot was assigned the reference BMI of the obesity cold spot, average BMI in the hot spot would decline by 0.7 units. If the hot spot was assigned the cold spot's disutility of exercise, average BMI in the hot spot would decline by 0.3 units. To put this in perspective, the baseline difference in average BMI across the two regions is 1.2, and 1 unit of BMI corresponds to just shy of 7 pounds for the average adult male in the US. The most important factors in accounting for the education gradient to BMI are the relative weight on healthy vs. reference BMI and the disutility from exercise. If high school graduates placed the same weight on healthy BMI relative to reference BMI as college graduates, the average BMI of high school graduates would decline by 0.5 units. If high school graduates were assigned the disutility of exercise of college graduates, average BMI of high school graduates would decline by 0.5 units. For context, the benchmark difference in BMI over education is 1.5.

While our model is designed with cross-sectional comparisons in mind, we are also able to use it to learn about factors that may have contributed to the rise in obesity over time. We find that changes in initial BMI, reference BMI, occupational MET, and education shares together account for between 60% and 65% of the rise in average BMI since the 1970s. Of these, reference BMI is by far the largest contributing factor, accounting for close to half of the increase. Yet, this leaves a meaningful portion unaccounted for by the model. We can, nevertheless, draw two particularly interesting conclusions from our temporal analysis. First, contrary to what has been posited in the literature, we can rule out MET as an important driver of rising obesity. We can also rule out food prices. Second, from our analysis, it is clear that regardless of what all factors are driving the rise in obesity, the amplification effect from social norms – the BMI reference point – is large.

Having shown that reference BMI plays an important role in rationalizing the patterns of BMI over space and time, it is natural to ask what implications it has for policies designed to mitigate the rise in obesity. We introduce a GLP-1 treatment policy with three targeting scenarios:

treating individuals in the top 10%, 20%, or 30% of the food preference distribution by reducing their preferences for food to those of the median. While in the preceding exercises the BMI reference point was varied exogenously, in the GLP-1 policy exercise we allow the BMI reference point to adjust endogenously. Under the most aggressive scenario (treating the top 30%), the direct policy effect on average BMI ranges from 0.43 units (3.0 pounds) in the cold spot to 0.87 units (6.1 pounds) in the hot spot. Crucially, allowing reference BMI to adjust endogenously generates substantial spillover effects, adding between 0.14 and 0.42 BMI units to the total policy impact. These externality effects represent between 25% and 48% of the direct treatment effect, depending on region and targeting intensity. The results demonstrate that the social multiplier effect from reference BMI adjustment is not merely a theoretical construct but has quantitatively important implications for policy evaluation. Ignoring these general equilibrium effects would lead to a nontrivial underestimate of the aggregate impact of pharmacological weight-loss interventions. Moreover, the heterogeneity in externality effects across regions—with the hot spot exhibiting the largest spillovers—underscores that the effectiveness of obesity interventions depends critically on the local distribution of preferences and baseline reference BMI.

There is previous research studying obesity from different angles, notably the increase over time, the socioeconomic gradient, and the spatial concentration of obesity. In particular, the literature on the rapid increase in obesity focuses on the: (i) increase in calorie intake (Cutler, Glaeser and Shapiro, 2003), (ii) the reduction in physical activity (Griffith, Lluberas and Lührmann, 2016) or (iii) a combination of the two (Lakdawalla, Philipson and Bhattacharya, 2005). The empirical literature has studied the relationship between socioeconomic status and obesity, in particular education (see among others Baum and Ruhm (2009)) and the spread of obesity through social networks (Christakis and Fowler, 2007; Cohen-Cole and Fletcher, 2008). Our results about the importance of the reference point in rationalizing spatial patterns of BMI is broadly consistent with the empirical work of Datar and Nicosia (2018) who use a natural experiment to determine the effect of reference BMI on weight. They study the assignment of military families to different bases around the US and find that moving to a county with a higher obesity rate has a significant effect on one's own BMI. To the best of our knowledge, we are among the first to study the salient trends of BMI in a unified, structural framework with conspicuous consumption and loss aversion. This framework is essential to rationalize the observed patterns of BMI on a disaggregated level.

The rest of the paper is organized as follows. Section 2 presents the stylized facts. Section 3 introduces the model. Section 4 describes the parameterization of the model, while Section 5 discusses the fit of the model. Section 6 discusses the quantitative results of the decomposition and policy exercises, while Section 7 concludes.

2 Stylized Facts

In this section, we present stylized facts with respect to BMI and obesity, and some key factors that influence BMI, namely the strenuousness of occupations, time use, calorie consumption, and food prices. We document patterns across regions and over education, but also pervasive changes over time. These salient observations motivate our analysis and are crucial in disciplining the model.

2.1 Data and Methodology

We use county-level data on obesity, inactivity, and socioeconomic characteristics from the Centers of Disease Control and Prevention (CDC) and the US Department of Agriculture (USDA). We use individual-level data on medical spending and health insurance from the Medical Expenditure Survey (MEPS), time use data from the American Time Use Survey (ATUS) and the American Heritage Time Use Study (AHTUS), income and strenuousness of occupation data from the American Community Service (ACS), BMI and food expenditure data from the National Health Interview Survey (NHIS), ATUS Eating and Health Module and Current Population Survey Food Security Supplement (CPS-FSS).

To study the spatial dimension of obesity, we want to sort geographic regions into clusters of high and low obesity. The first step is to show that there is spatial correlation in obesity. We do this using a spatial regression and Moran's I statistic. We classify obesity clusters as "hot spots" and "cold spots". The analysis is conducted at the county level. We aggregate up to the state level by imposing that if 75% of counties in a state are classified as hot spots, then the state is classified as a hot spot.[4] Using this classification, we estimate the average BMI

[4] The states that are classified as hot spots are: Alabama, Arkansas, Indiana, Kentucky, Louisiana,Mississippi, Missouri, Oklahoma, South Carolina, Tennessee, and West Virginia. The states that are classified as cold spots are: California, Colorado, Connecticut, District of Columbia, Idaho, Massachusetts, Minnesota, Montana, Nevada, New Hampshire, New Jersey, New Mexico, New York, Oregon, Rhode Island, Utah, Vermont, and Wyoming. The

and obesity prevalence by cluster, education, and age and the strenuousness of occupations by cluster and education. Some datasets, such as the NHIS, do not contain information on the state of residence. There we use geographical regions instead of the estimated clusters.

To study the education gradient to BMI, we group individuals into two education categories, which we term high school and college. Throughout, we use high school to refer to individuals with at most a high school diploma and college to refer to individuals with at least some college.

2.2 BMI and Obesity

A BMI of around 22 is considered optimal from a general health perspective.[5] According to standard classifications, a BMI of 25 or higher is categorized as overweight, while a BMI of 30 or above is classified as obese. Average BMI in the US is high, close to 28, and 30% of the population is currently classified as obese. Less educated individuals have on average higher BMI and higher incidence of obesity than their more educated counterparts. Average BMI and the obesity rate have increased steadily over the last 50 years, while the education gap has remained constant over this time period. See Figure 1 for details.

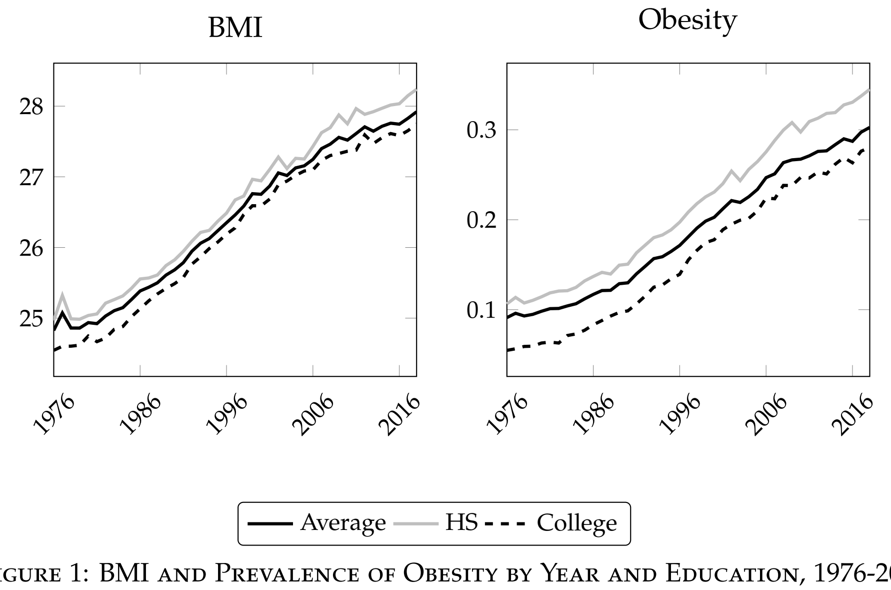

FIGURE 1: BMI AND PREVALENCE OF OBESITY BY YEAR AND EDUCATION, 1976-2018

Notes: Average BMI and obesity prevalence by education between 1976 and 2018. Individuals are classified as obese if $BMI \geq 30$.

Source: NHIS (1976-2018).

remaining states are classified as neutral.

[5] All-cause mortality is lowest at BMI between 20.0 and 24.9 (de Gonzalez et al., 2010).

We observe a striking concentration of obesity in specific areas of the US, especially the South (see Figure 2).

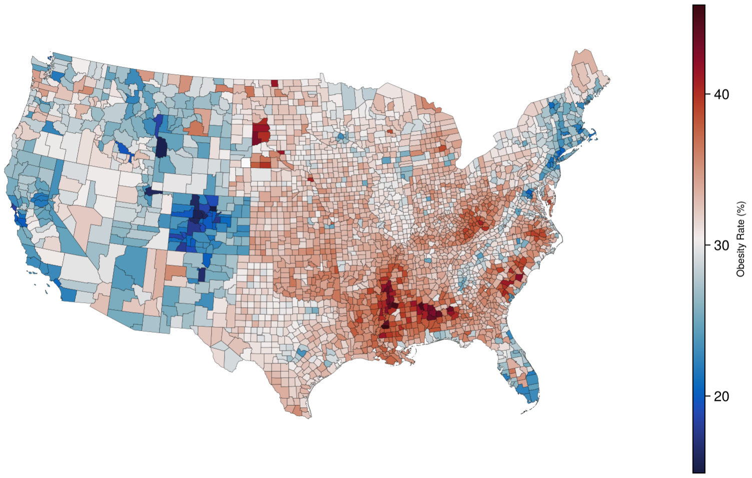

Figure 2: Prevalence of Obesity by US County

Notes: Obesity prevalence by county in the US. Individuals are classified as obese if $BMI \geq 30$.
Source: CDC Behavioral Risk Factor Surveillance System (BRFSS).

Average BMI and obesity rates are high and rising throughout the US. Regional differences persist over time. The exception is the Northeast, which exhibits a slower increase in BMI over time relative to the rest of the country. See Figure 3 for details.

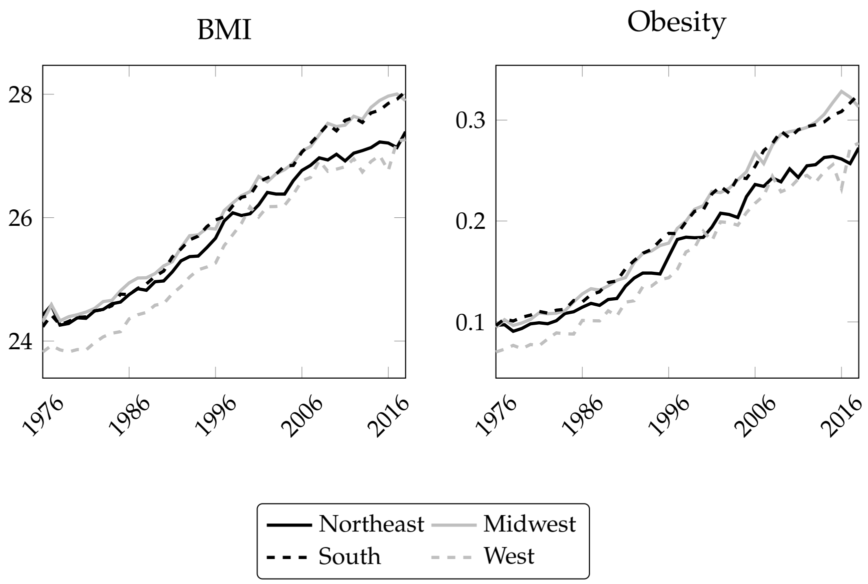

Figure 3: BMI and Prevalence of Obesity by Year and Region, 1976-2018

Notes: Average BMI and prevalence of obesity by year and region. Individuals are classified as obese if $BMI \geq 30$. The data are cohort adjusted.
Source: NHIS (1976-2018).

Zooming in on the educational gradient in BMI by region, we document that: (i) college graduates persistently have a lower BMI compared to high school graduates across all regions, (ii) both high school and college graduates have higher BMI in the South and Midwest compared to individuals with similar educational attainment in the rest of the country, and (iii) all education types across the US face a similar rising trend in BMI, with the exception of college graduates in the Northeast and West where the rise is somewhat less steep (see Figure 4).

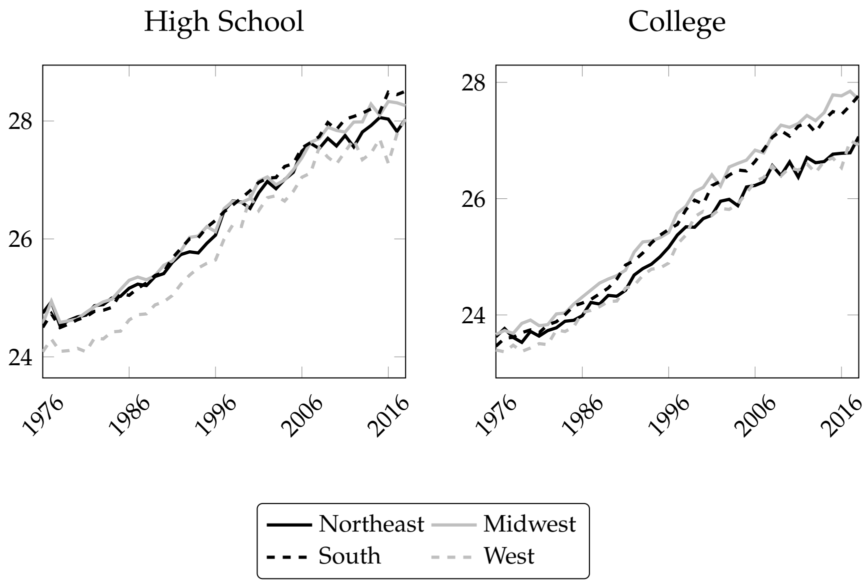

Figure 4: BMI by Year, Region, And Education, 1976-2018

Notes: Average BMI by year, education, and region. The data are cohort adjusted.
Source: NHIS (1976-2018).

2.3 Strenuousness of Occupation

The strenuousness of occupations potentially plays a crucial role in BMI trends, since working individuals spend around a quarter of their time working. We use data on the strenuousness of occupations estimated in Tudor-Locke et al. (2011), which assigns a Metabolically Equivalent Task (MET) value to each occupation. MET measures how much more strenuous an occupation is relative to being at rest. The total calories consumed per working hour is the product of MET and the Basal Metabolic Rate (BMR), which depends of the individual's gender, age, height and weight. As is common in the literature, the MET of each occupation is assumed to remain constant over time and the fall in average MET is caused by a shift to more sedentary occupations. The composition of occupations has become less strenuous over time in all areas of the US, but there is a significant gap across regions due to different distributions of industries and occupations. See Figure 5 for details.

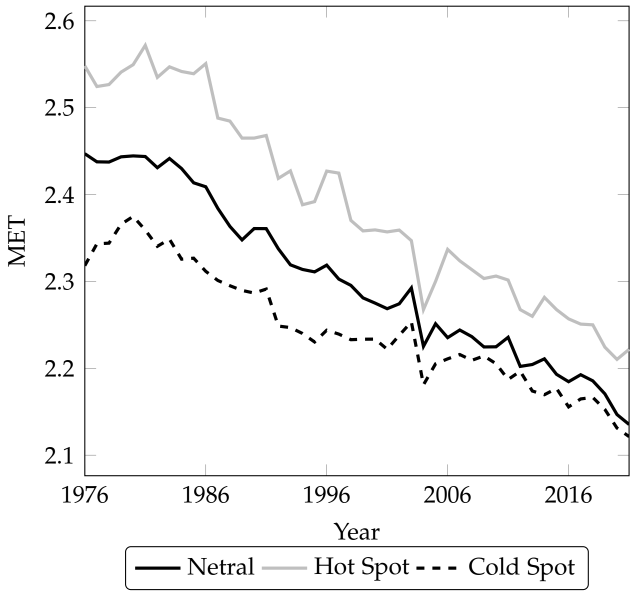

FIGURE 5: MET BY YEAR AND CLUSTER

Notes: Average MET by cluster and year. We use MET data by occupation from Tudor-Locke (2011) and merge it with the CPS occupational classification. Since each occupation has an assigned MET that does not change over time, the fall in the average MET is caused by changes in the distribution of occupations.

Source: CPS (1976-2021) and Tudor-Locke (2011).

Decomposing the average MET by education, we observe a striking pattern over time. High school graduates have moved to slightly more strenuous occupations, in contrast to college graduates who have shifted to occupations with clearly lower MET over time (see Figure 6).

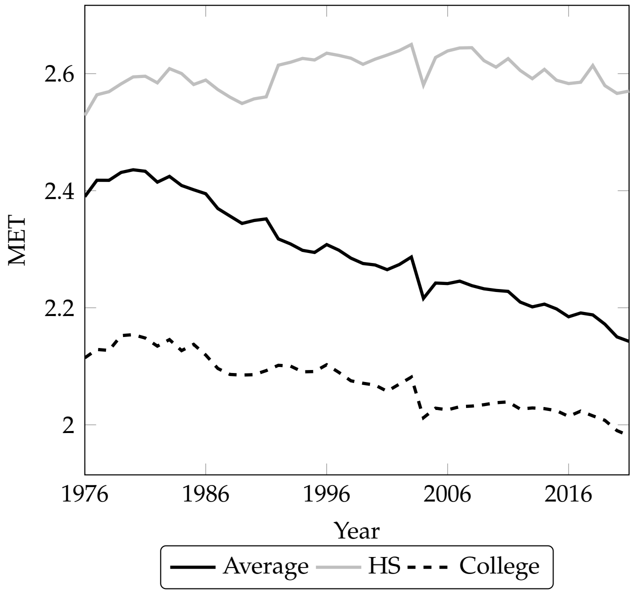

Figure 6: MET by Year and Education

Notes: Average MET by year and education.
Source: CPS (1976-2021) and Tudor-Locke (2011).

In sum, the observed fall in average MET in all regions of the US is driven by: (i) college graduates working in less strenuous occupations, and (ii) the increasing share of college graduates over the last decades. Further, the gap in educational attainment across regions can partially account for the gap in the average MET across said regions. See Table 1 for college shares by region and time period.

Table 1: College Education by Cluster, 1976 & 2021

Cluster	1976	2018
Hot	0.19	0.33
Neutral	0.23	0.39
Cold	0.25	0.37

Notes: Fraction of individuals with at least a year of college education by cluster in 1976 and 2018.
Source: NHIS (1976-2018).

2.4 Time Use

Physical activity is an important determinant of BMI and can partly account for the education gradient in BMI. As seen from Figure 7, college graduates are more likely to engage in any type of exercise, and spend on average more time exercising compared to high school graduates.

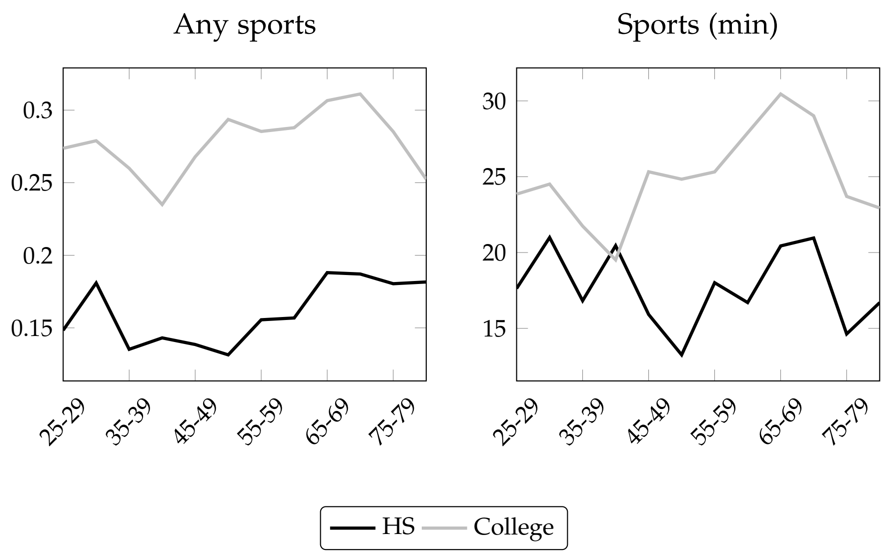

Figure 7: Sports and Exercise by Education and Age

Notes: Share of individuals reporting any type of exercise and average time spent exercising by education and age.
Source: ATUS (2003 - 2018),

This affects the average MET during non-working hours (Figure 8). Moreover, college graduates engage in more strenuous activities during their non-working hours, even excluding sports. This is a combination of engaging in more strenuous activities related to home production, and the time spent on these activities. Even though these differences may appear small, the total time spent in these activities has a substantial effect on calories burned[6]

[6]For the average 40-year-old male of healthy weight, working 8 hours per day, a difference of 0.1 in MET during no-working hours means a difference of 112 Kcal, or a 30 min brisk walk per day.

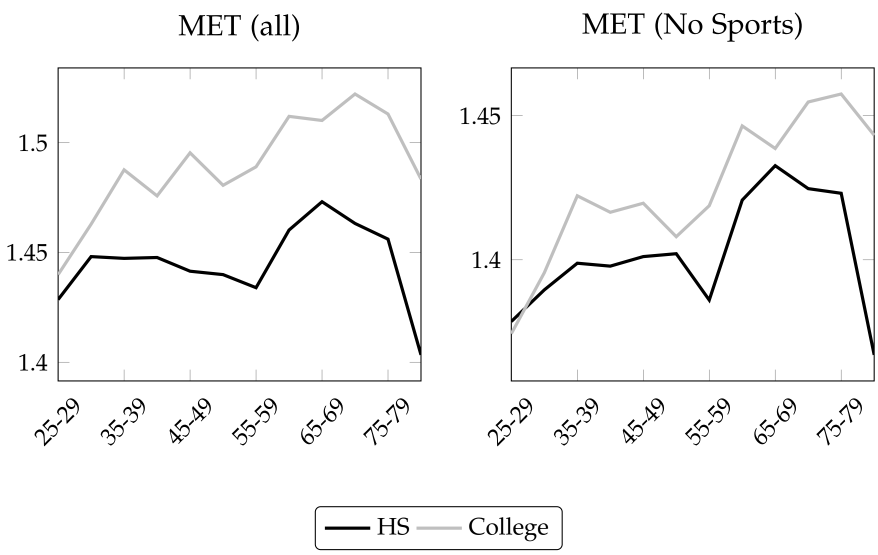

Figure 8: MET During Leisure by Education and Age

Notes: Average MET during leisure including and excluding sports. The average MET is estimated as the weighted average of the MET based on time spent on each activity.
Source: ATUS (2003 - 2018).

2.5 Calories

All survey data suffer from measurement error, but the precise measurement of calorie consumption is especially challenging. Individuals tend to under-report calorie consumption, and the misreporting increases with BMI. Archer, Hand and Blair (2013) use NHANES data and find that the reported calorie consumption is physiologically implausible given the reported weight. Our own estimation using USDA data suggests the same. In this paper, we circumvent this issue by estimating the implied calorie consumption of male workers in the US, given their weight, height, age, and the strenuousness of their occupation and leisure time.

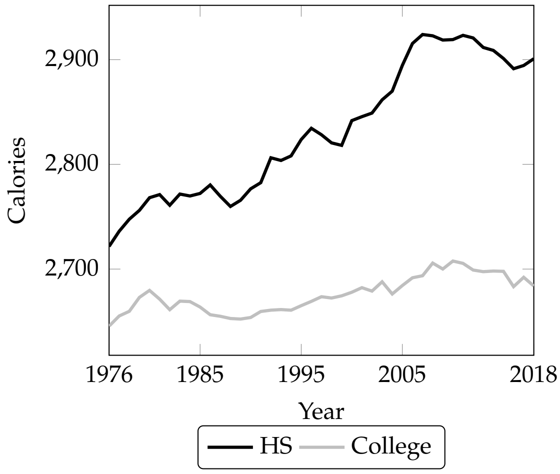

Figure 9: Calorie Consumption by Year and Education

Notes: We estimate the calorie consumption which is required for individuals of a certain age, height, weight, occupation, and strenuousness of activities during leisure to maintain their current weight. We merge ATUS and CPS data and estimate the average strenuousness of non-working time, and the strenuousness of occupations. BMI data are available as part of the Eating and Health Module in the ATUS.

Source: CPS (1976-2018), ATUS (2003 - 2018), Tudor-Locke et al. (2011).

We find that college graduates experience only a moderate increase in calorie consumption over the last five decades (see Figure 9). Rather, a significant part of the rise in their average BMI can be explained by a reduction in the strenuousness of occupation. In contrast, high school graduates increased their calorie consumption substantially over the time period in question. This is necessary to rationalize the increase in average BMI given the shift to more strenuous occupations during the same time period.

2.6 Food Prices

Prices are an important driver for consumer demand and might play a role in explaining the patterns of obesity. However, we document that the relative price of food does not have a clear trend over the time period under study (see Figures 10 and 11). More generally, empirical results suggest that food prices (Anderson and Matsa, 2011) and access to food (Allcott et al., 2019) do not contribute to the rise in obesity.

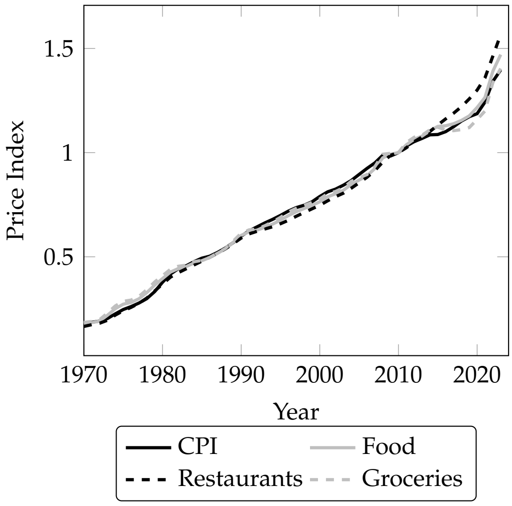

Figure 10: Food and Consumer Price Indices, 1976 - 2021

Notes: Consumer Price Index (CPI) for all items and food at home, normalized to 100 in 1976.
Source: FRED Data, Consumer Price Index (1976-2021).

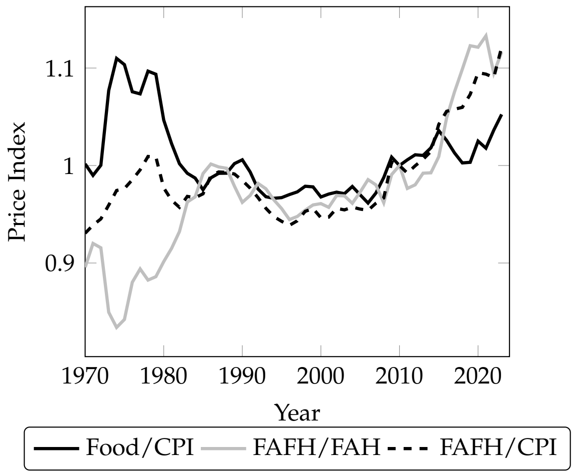

Figure 11: Relative Price of Food, 1976 - 2021

Notes: Relative price of food at home computed as the ratio of the food CPI to the all-items CPI. T
Source: FRED Data, Consumer Price Index (1976-2021).

3 Model

In order to rationalize the differences in BMI across clusters and education, we develop a life cycle model where heterogeneous agents make decisions regarding food consumption and exercise in a realistic environment where BMI is determined by calorific balance, which depends on the lifestyle choices of the agents, and their occupation which affects the strenuousness during working hours. Agents are ex-ante heterogeneous over education, occupation, and cluster, and face uncertainty with respect to health, survival, medical spending and health insurance coverage status. Agents make decisions with respect to general consumption, food consumption, savings, and exercise.

3.1 Preferences and Demographics

Consider an economy populated by overlapping generations of agents. Each period a continuum of agents, who live for at most J periods, is born. At age $j = 0$, each agent "inherits" a level of education, i, a cluster, k and an occupation o. The type of the agent $\{i, k, o\}$ determines the strenuousness of the occupation and the life cycle profile of productivity.

At the beginning of every period, agents observe their BMI BMI_j, assets a_j, and health h_j and whether they receive a health insurance offer from their employer. If they receive a health insurance offer, they enroll in the health insurance plan and pay the premium. If they do not, they remain uninsured. After age 65 all agents are enrolled in Medicare. Agents then observe the health and medical spending shocks z_j^h and z_j^m and make decisions regarding general consumption c_j, food consumption f_j, exercise n_j, and savings a_{j+1}.

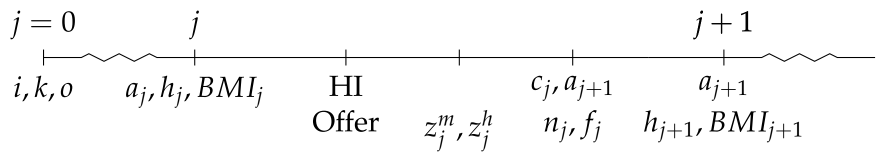

Agents face the following lifetime utility:

$$U = \sum_{j=1}^{J} \beta^{j-1} \left[\prod_{m=2}^{j} P_{m-1}(h_m) \right] u_j(c_j, f_j, n_j, BMI_j) \tag{1}$$

where β is the discount factor, $P_{m-1}(h_m)$ is the probability of survival from age $m-1$ to age m given health h_m. The period utility function is:

$$\begin{aligned} u_j(c_j, f_j, n_j, BMI_j) = {} & b + \alpha_1 \ln c_j + \alpha_2 \frac{f_j^{1-\gamma}}{1-\gamma} - \alpha_{3,i,k} \frac{n_j^{1+\nu_e}}{1+\nu_e} \\ & - \alpha_4 \left(\omega_i (BMI_j - BMI_k)^2 + (1-\omega_i) \left(BMI_j - 22\right)^2 \right) \cdot \mathbb{1}_{\{BMI_j > BMI_k\}} \\ & - \alpha_4 (1-\omega_i) \left(BMI_j - 22\right)^2 \cdot \mathbb{1}_{\{BMI_j \leq BMI_k\}} \end{aligned} \tag{2}$$

$BMI_h = 22$ denotes the "healthy" BMI and BMI_k is the reference BMI in cluster k. We choose this utility function to capture the salient features of lifestyle choices. Individuals enjoy general consumption and food consumption, with food preferences exhibiting CRRA utility with coefficient γ. The disutility from exercise has an inverse Frisch elasticity of ν_e and varies by education and cluster through $\alpha_{3,i,k}$, reflecting systematic differences in the cost of physical activity across education groups and geographic areas.

Crucially, agents exhibit loss aversion with respect to BMI: they face disutility when their BMI exceeds the reference BMI of their cluster, but not when it falls below it. When BMI exceeds the reference point, agents are penalized for deviations from both the reference BMI and the healthy BMI, with education-specific weights ω_i determining the relative importance of the reference BMI. When BMI is below the reference point, agents are only penalized for deviations from the healthy BMI. The constant b follows Hall and Jones (2007).

3.2 BMI

BMI is determined by calorific balance, namely the balance between calories-in and calories-out. Calories-in are determined by food consumption and calories-out are determined by current weight, age, exercise, and strenuousness of occupation. We assume that the strenuousness of occupation is exogenous and is constant throughout the life cycle. In particular, the weight and age of the agent determines the Basal Metabolic Rate (BMR)[7] which is the amount of calories individuals burn at rest. The strenuousness of occupation and exercise is measured in Metabolically Equivalent Task (MET) units which determine how much more strenuous an activity is relative to being at rest. The total calories-out are determined by:

$$Calories_{out} = BMR * \left(MET_{occupation} \times \ell + MET_{exercise} \times n + MET_{leisure} \times (1 - \ell - n)\right) \quad (3)$$

The weight change of the agent is determined by[8]:

$$\Delta W = (f - Calories_{out}) \div 7700 \quad (4)$$

This in turn determines the BMI of the agent:

[7] We use the Mifflin-St Jeor Equation where BMR is determined as $10 \times weight + 6.25 \times height - 5 \times age + 5$. In our model height is fixed by education to reflect the fact that college graduates are on average taller and ceteris paribus have a higher BMR.

[8] The number of excess calories that an individual needs to consume in order to gain one kg of weight is 7700.

$$BMI_j = \frac{W_j}{H_i^2} \tag{5}$$

where W_j is the weight of the agent and H_i is the average height of agents with education i.

3.3 Health

Agents face uncertainty with respect to health and medical spending. Health has two states, $h_j \in \{g, b\}$ where g denotes good health and b bad health. The transition probabilities between states depend on education, cluster, age, current health and BMI:

$$\pi_{j,i,k} = Pr\left(h_{j+1} = \psi \mid h_j = \varphi, j, i, k, BMI_j\right), \quad \psi, \varphi \in \{g, b\} \tag{6}$$

where $\pi_{j,i,k}$ is the probability of transitioning from health state φ to health state ψ given age, education, cluster, and BMI.

3.4 Medical Spending and Health Insurance

Agents face uncertainty with respect to out-of-pocket medical spending μ_j. Until age 65 agents are either uninsured or insured through their employer with a probability that depends on age and education. If agents face medical spending that does not allow them to achieve a minimum level of consumption, they are automatically enrolled in Medicaid. At age 65 all agents are enrolled in Medicare.

The distribution of medical spending shocks is determined by the health state, age, and education of the agent as in Margaris and Wallenius (2023).

The total out-of-pocket medical spending of the agent is determined by:

$$\mu_j = \left(1 - q\left(ins_j\right)\right) z_j^m + pr_j\left(ins_j\right) \tag{7}$$

where ins_j is the insurance status of the agent, $q\left(ins_j\right)$ is the coinsurance rate offered by the insurance, z_j^m is the medical spending shock, and $pr_j\left(ins_j\right)$ is the premium paid by the agent. The insurance status of the agent depends on age and education:

$$ins_j = \begin{cases} Medicare & if \quad j \geq J_R \quad , \\ Medicaid & if \quad c < \bar{c} \quad , \\ GHI, & w/prob. \quad \zeta\left(i,j\right) \quad , \\ uninsured & \end{cases} \tag{8}$$

3.5 Budget Constraints

The agent receives income from working, savings, social security, and government transfers, if necessary to guarantee a minimum level of consumption. The agent allocates disposable income between general consumption, food consumption, out-of-pocket medical spending, and savings:

$$\tilde{y}_j = \left(1 + \tau^c\right) c_j + p_j f + \mu_j + a_j \tag{9}$$

where τ^c is the tax rate on general consumption, and p_j is the price of food per calorie. $\tilde{y}_j$ is the disposable income of the agent:

$$\tilde{y}_j = \left(1 - \tau\left(y_j\right)\right) e\left(j,i,k,o\right) + \left(1 + r\right) a_{j-1} + ss_j + Tr_j \tag{10}$$

where $\tau\left(y_j\right)$ is the progressive income tax rate, $e\left(j,i,k,o\right)$ is the earnings of the agent as a function of age, education, cluster and occupation, r is the interest rate, ss_j is the social security benefit, and Tr_j is the government transfer.

3.6 Government

The government collects taxes via a progressive labor income tax, a tax on general consumption, and the Medicare premium. The government also pays social security benefits, covers a fraction of medical spending via Medicare and provides government transfers that guarantee a minimum level of consumption. We use the income tax function as in Guner, Kaygusuz and Ventura (2014):

$$\tau\left(y_j\right) = 1 - \tau_0 y_j^{\varphi} \tag{11}$$

4 Parameterization

We follow a two-step procedure for the parameterization of the model. In the first stage, we estimate as many parameters as possible directly from the data, and set some parameters based on commonly adopted values in the literature. In the second stage, we calibrate the vector of remaining parameters, in order to minimize the distance between moments generated from the model simulations and their data counterparts. Here we provide an overview of the estimations and calibration; please see the Appendix for further details.

4.1 Estimated Parameters

The parameters/processes estimated from the data are: (i) the education-specific distribution of occupational MET values, (ii) the education and cluster-specific average MET during non-working hours, (iii) the initial distribution of BMI by cluster and education, (iv) the initial distribution of health status and the transition probabilities by cluster and education, (v) the distribution of medical spending shocks by education, (vi) the probability of receiving a group health insurance offer by education and age, (vii) the labor income process, and (viii) the relative food price by education and cluster.

4.1.1 MET

We estimate the distribution of MET values of occupations and the average MET during non-working hours by education and cluster. We estimate the median and 75th percentile of occupational MET by merging CPS data with the MET data by occupation estimated in Tudor-Locke et al. (2011). Due to the change in the classification of occupations, we use Machine Learning to re-classify occupations prior to 2003 based on the distribution of occupations between 2004 and 2018. We use a Random Forest Classifier and the classification is based on race, marital status, industry, age, gender, education, hourly wage, and average hours worked per week.

High school types are more likely to work in strenuous occupations than college types. Occupations are more strenuous in the obesity hot spots than the cold spots. See Table 2 for a breakdown of the strenuousness of jobs across regions and education types.

Table 2: Occupational Strenuousness by Cluster and Education

Education	Cluster	Sedentary (%)	Moderate (%)	Strenuous (%)
High School	Hot Spot	32.3	43.4	24.3
	Neutral	34.0	43.3	22.7
	Cold Spot	34.1	43.1	22.8
College	Hot Spot	55.1	37.9	7.0
	Neutral	57.7	36.6	5.7
	Cold Spot	58.5	36.1	5.4

Notes: Distribution of occupational strenuousness categories by education and geographic cluster. Occupations are classified as sedentary (low MET), moderate, or strenuous (high MET) based on their physical intensity.
Source: CPS-ATUS (2003-2018) and Tudor-Locke et al. (2011).

MET during non-working hours is estimated using data from the ATUS and MET values of each reported activity from Tudor-Locke et al. (2009) weighted by its duration. We estimate the MET for exercise by cluster and education, and the MET for non-exercise activities by cluster and education. We find that college graduates have higher MET during non-working hours compared to high school graduates, and that the gap is marginally larger for exercise-related activities (see Table 3).

Table 3: MET of Sports by Cluster and Education

Cluster	Average	High School	College
Hot	1.455	1.438	1.472
Neutral	1.456	1.433	1.479
Cold	1.466	1.443	1.489

Notes: Average leisure sports MET (Metabolic Equivalent of Task) by cluster and education. Higher values indicate more intense physical activity during leisure. A MET of 1.0 corresponds to being at rest.
Source: CPS and ATUS (2003-2018).

4.1.2 BMI

There is substantial variation in average BMI across education and cluster at age 25, which can have substantial life cycle effects. High school graduates enter the economy with a higher BMI compared to college graduates, and the "hot spot" of obesity has a higher average BMI even at age 25 (see Figure 12).

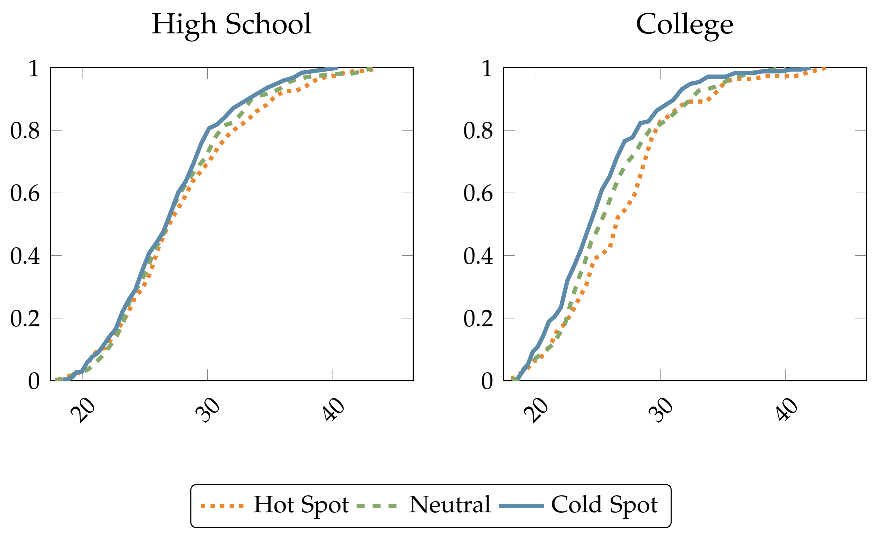

FIGURE 12: INITIAL DISTRIBUTION OF BMI BY EDUCATION

Notes: Initial distribution of BMI by education and cluster at age 25. *Source:* CPS-ATUS (2003-2018).

4.1.3 Health and Medical Spending

In our model, health has two states, good and bad. The initial distribution of health and the transition probabilities are estimated using PSID data. At age 25, agents enter with a BMI drawn from the initial distribution of BMI by education and cluster, which determines the initial level of health. The age-specific transition probabilities are estimated separately by health state, cluster, and education using a Probit model.

Medical spending shocks are estimated using MEPS data by health, education, and age as in Margaris and Wallenius (2023). In a nutshell, we construct three bins for medical expenditures, below the 50th, between the 50th and the 95th, and above the 95th percentile. Each individual is assigned into these bins according to their education, current level of health, and age.

4.1.4 Labor Income

We estimate the labor income process using pooled CPS–ATUS data separately by educational attainment.

For each education group, we estimate the following baseline specification:

$$\ln w_i = X_i\beta + Z_i\gamma + \varepsilon_i, \tag{12}$$

where w_i denotes real labor income, X_i includes measures of MET, BMI-based weight class indicators, cluster, and age polynomials, and Z_i includes additional demographic controls such as marital status, race, and sex.

Individuals are assigned to weight classes based on their position in the BMI distribution within their education group and cluster. For each cluster-education cell, we compute the 25th, 50th, and 75th percentiles of BMI. Individuals are then classified into four categories according to their relative position in this distribution: (i) BMI below the 25th percentile, (ii) between the 25th and 50th percentiles, (iii) between the 50th and 75th percentiles, and (iv) above the 75th percentile.

This procedure captures relative body mass within each local and educational reference group rather than absolute BMI thresholds.

To isolate the structural components of labor income associated with physical attributes and regional characteristics, we implement a two-step residualization. First, we estimate the baseline regression including the full set of controls. We then set the coefficients on MET, weight class, and regional indicators to zero and compute predicted log income excluding these components.

In the second step, the resulting "clean" income measure is regressed on age and age squared. This yields age profiles that have correct levels while preserving the estimated effects of MET, cluster, and weight class which are included in the model.

4.1.5 Food Prices

We estimate the price per calorie by cluster and education using data from the USDA and CPS. First, we estimate the consumption of calories that is consistent with current BMI, occupation, and age by cluster and education using data from the CPS (Table 4). Second, we estimate the real equivalized food expenditure from the USDA by cluster and education. In our model, prices per calorie are different across clusters and education to reflect: (i) the differences in cost of living between areas in the US, and (ii) the different eating habits of high school and college graduates (Table 5).[9]

[9] In our model, we consider a single composite food consumption good and abstract from food at home and food away from home. Moreover, due to peer effects, high school and college graduates might make different choices with respect to the prices for food away from home (e.g., eating fast food vs. eating at restaurants). The difference in the average price per calorie by education reflects these two effects.

Table 4: Calories By Cluster And Education

Cluster	Average	High School	College
Neutral	3150	3275	2924
Hot	3023	3164	2846
Cold	2986	3138	2816

Notes: Average calorie consumption by cluster and education. The calorie consumption is estimated as the amount of calories that individuals of a certain age, height, weight, occupation, and strenuousness of activities during leisure need to consume in order to maintain their current weight.
Source: CPS and ATUS (2003-2018)

Table 5: Relative Food Price By Cluster And Education

Cluster	High School	College
Hot Spot	0.65	0.87
Neutral	0.59	0.95
Cold Spot	0.63	1

Notes: Relative price of food per calorie by cluster and education. The price per calorie is estimated as the ratio of real equivalized food expenditure to the estimated calorie consumption by cluster and education.
Source: Authors' calculation using USDA and CPS (2003-2018)

4.2 Exogenous Parameters

Here we present the parameters that are set exogenously to values that are commonly found in the literature.

4.2.1 Preferences and Demographics

A model period is 5 years. Agents enter the economy at age 25 ($j = 1$) and survive to at most age 95 ($j = 14$). Population growth is set to 1% annually. The discount factor β is set to $1.01^5 \approx 1.051$ to account for the fact that agents value the future appropriately given the observed savings behavior. The effect of bad health on mortality (θ) is set to 0.2, meaning that bad health increases mortality risk by 20%.[10] The weight of the utility of consumption is normalized to $\alpha_1 = 1.0$.

Table 6: Demographic Parameters

Parameter	Description	Value
J	Maximum Number of Periods	14 (95 years old)
J_R	Retirement Period	8 (65 years old)
n	Population Growth	0.05 (1% annually)

[10]This is a conservative estimate based on existing literature (DeSalvo et al., 2006).

4.2.2 Government

We use the estimates of Guner, Kaygusuz and Ventura (2014) for the parameters of the income tax function, to reflect the progressivity of the tax function in the US. The consumption floor is set to \$5,000 annually (Ozkan, n.d.).

TABLE 7: Tax Parameters

Parameter	Description	Value
τ_0	Tax level	0.902
ϕ	Tax progressivity	0.036
τ^c	Consumption tax	0.05

Source: Guner et al. (2014)

Since it is computationally infeasible to track the lifetime earnings of agents in order to calculate social security benefits, we follow a stylized representation of actual US system. We estimate the average life-time labor income by cluster and education and calculate the average benefits as in Zhao (2017).

TABLE 8: Social Security Benefits

Average Lifetime Earning	Marginal Replacement Rate
$y_i \in [0, 0.2\bar{y})$	90%
$y_i \in [0.2, 1.25\bar{y})$	33%
$y_i \in [1.25, 2.46\bar{y})$	15%
$y_i \in [2.46, \infty)$	0%

Notes: We estimate the average lifetime income by cluster and education and apply the marginal replacement rates to estimate the average level of pension.
Source: Zhao (2017).

4.3 Calibration

The remaining parameters are calibrated internally to match key moments from the data. We employ a simulated method of moments approach to minimize the distance between model-generated moments and their empirical counterparts.

4.3.1 Calibrated Parameters

The disutility from exercise $\alpha_{3,i,k}$ varies by education and cluster, making it the only calibrated parameter in the model with this two-dimensional heterogeneity. These parameters capture systematic differences in the cost of exercise across regions and education groups, reflecting factors such as access to recreational facilities, time constraints, and cultural attitudes toward physical activity. Allowing the disutility of exercise to vary along both dimensions is essential for the model to match the joint distribution of exercise participation and intensity across clusters and education groups.

The weight of the utility of food α_2 is age-specific and varies across preference types. The age profile is approximated by an non-linear process. This parsimonious functional form proxies for the gradual strengthening of food preferences with age, capturing in reduced form the role of habit formation without the computational burden of tracking an additional endogenous state variable. We additionally allow for 10 food preference types that scale the age profile, also governed by a non-linear process. The resulting α_2 is a $J \times 10$ matrix, where each entry is the product of the age-specific and type-specific components. Importantly, the levels of α_2 are not education or cluster specific—they are common to all agents of a given age and preference type.

The distribution of food preference types across the population is, however, cluster and education specific. We construct 10 distinct food preference types based on the deciles of the national BMI distribution at age 25. Agents draw their initial BMI from the distribution specific to their education and cluster, and the initial BMI draw determines the agent's food preference type through their position in the *national* BMI distribution. For example, an agent whose initial BMI places them in the top 10% of the national distribution is assigned the highest food preference type, regardless of their education or cluster. Conversely, an agent in the bottom 10% nationally receives the lowest food preference type. Because high school graduates and agents in obesity hot spots tend to have higher initial BMI on average, they are more likely to be assigned higher food preference types. This mechanism generates endogenous heterogeneity in the composition of food preference types across education and cluster groups, ensuring that observed differences in BMI partly reflect systematic differences in food preferences that are disciplined by the data, while keeping the preference parameters themselves invariant across education and cluster.

The weight on BMI deviation penalties α_4 governs the intensity of disutility agents experience when their BMI deviates from the reference and healthy benchmarks. The education-specific weights ω determine the relative importance of reference BMI versus healthy BMI in the utility function. These parameters capture the degree to which individuals benchmark their weight to social norms versus medical guidelines, and we allow this benchmarking behavior to differ by education.

The coefficient of relative risk aversion for food consumption γ governs the curvature of the utility of food, while the inverse of the Frisch elasticity of exercise ν_e determines the responsiveness of exercise choices to changes in the marginal utility of leisure.

The constant in the utility function b is calibrated to improve the fit of the model to BMI by cluster and education. Since b enters the flow utility additively, it raises the total utility of being alive. Wealthier individuals, who have lower marginal utility of consumption, implicitly place a higher value on continued life, giving them a stronger incentive to invest in health through lower BMI. The implied value of a statistical life year (VSLY) is 4.2 times average per capita consumption, which is consistent with the lower end of estimates in the literature.[11]

4.3.2 Calibration Targets

The calibrated parameters are chosen to match the following moments: (i) average BMI by cluster and education for working-age individuals, (ii) age profiles of BMI within each education-cluster cell, (iii) relative frequency of exercise participation by cluster and education, (iv) average time spent exercising conditional on participation, by cluster and education, and (v) life expectancy differentials between high school and college graduates.

The calibration procedure ensures that the model replicates the observed patterns of BMI, exercise, and mortality across demographic groups, while allowing the model to speak to the underlying structural drivers of these patterns through the estimated preference parameters.

5 Model Fit

Our analysis focuses on working-age male adults for several important reasons. First, the available data become increasingly sparse at older ages, making reliable estimation difficult. Second, there is a strong cohort effect in BMI patterns, whereby different birth cohorts exhibit distinct trajectories that are not fully captured by age and period effects alone. Third,

[11] The implied VSLY is \$185,000. Glover et al. (2023) use the following equation to determine b within the model: $\bar{u}(\cdot) + b = \frac{\partial u}{\partial c} \times 6.25\bar{c}$, where $\bar{u}(\cdot)$ denotes the average flow utility excluding b, $\frac{\partial u}{\partial c}$ is the marginal utility of consumption, and $\bar{c}$ is average annual per capita consumption.

survival bias becomes increasingly severe at older ages, as individuals with higher BMI face elevated mortality risk—an issue that is only partially addressed by our mortality model. Fourth, and perhaps most importantly, our model does not account for the reverse causal relationship between health and BMI that becomes more prominent late in life. Specifically, serious health shocks may lead to reductions in BMI through involuntary lifestyle changes or the physiological effects of disease, a mechanism that operates in the opposite direction from the obesity-health relationship we model. Finally, focusing on working-age adults allows us to cleanly identify the effects of occupational strenuousness (MET) on BMI, as this mechanism operates primarily during the working years and becomes less relevant after retirement. For these reasons, we concentrate our analysis on ages where the primary mechanisms we study—food consumption, exercise, occupational MET, and the role of reference BMI—are most cleanly identified.

The model does quite well in matching the targeted moments. High school graduates have a higher BMI compared to college graduates in all cluster, and college graduates in the cold spot have substantially lower BMI compared to college graduates in the other clusters (Table 9). This is despite the fact that most of the heterogeneity across clusters and education contributes to a pattern of obesity that is opposite to the observed data. In particular: (i) high school graduates spend more calories during work due to higher occupational MET, and (ii) high school graduates are less likely to be insured during their working years, facing on average higher out-of-pocket medical spending due to obesity.

Table 9: BMI by Cluster and Education

	High School		College		Average	
	Simulation	Data	Simulation	Data	Simulation	Data
Hot Spot	29.64	29.31	28.63	28.11	29.33	28.94
Neutral	29.23	29.11	27.61	27.45	28.59	28.45
Cold Spot	28.48	28.48	26.75	26.95	27.73	27.81

Notes: Simulation results and data for average BMI by cluster and education for working-age males.
Source: Simulation results and CPS-ATUS (2010-2018).

The effect of income on BMI is ex-ante ambiguous. On the one hand, life expectancy is endogenous which means that, ceteris paribus, college graduates should value their BMI, and by extension health, higher than high school graduates (Hall and Jones, 2007). College graduates can achieve a higher level of consumption and a loss of per-period utility caused by higher mortality has a greater effect compared to high school graduates. On the other hand, higher income relaxes the budget constraint and individuals can achieve a higher level of both general and food consumption.

The disutility from exercise is allowed to vary by education and cluster, as discussed in the parameterization. This is necessary to match the substantial differences in exercise participation and intensity observed in the data across both dimensions.

Table 10: Exercise by Cluster and Education

	High School		College	
	Simulation	Data	Simulation	Data
Hot Spot	0.66	0.61	0.8	0.79
Neutral	0.71	0.6	0.97	0.93
Cold Spot	0.75	0.73	1.0	1.0

Notes: Simulation results and data for exercise by cluster and education for working-age males. We normalize the time spent exercising for college graduates in the cold spot to one.
Source: Simulation results and ATUS (2010-2018).

As shown in Table 10, the model matches the observed gradient in exercise across both education groups and clusters reasonably well. College graduates exercise more than high school graduates, and exercise participation increases as we move from the hot spot to the cold spot, consistent with the data.

Figures 13 and 14 illustrate the model's ability to match the age profiles of BMI observed in the data. The model successfully captures the key empirical patterns across both education groups and geographic clusters. Figure 13 shows that the model replicates the educational gradient in BMI, with high school graduates exhibiting higher average BMI than college graduates throughout the life cycle. The model also captures the hump-shaped age profile of BMI, with both groups experiencing increases in BMI through middle age followed by modest declines in later years. Figure 14 demonstrates that the model fits the data well across education, cluster, and age. The simulated BMI profiles closely track the data for both high school and college graduates in both the obesity hot spot and cold spot, capturing both the level differences across regions and the distinct life cycle dynamics within each education-cluster cell. The close fit across these multiple dimensions provides confidence that the calibrated model successfully captures the key mechanisms governing BMI determination and validates our use of the model for counterfactual policy analysis.

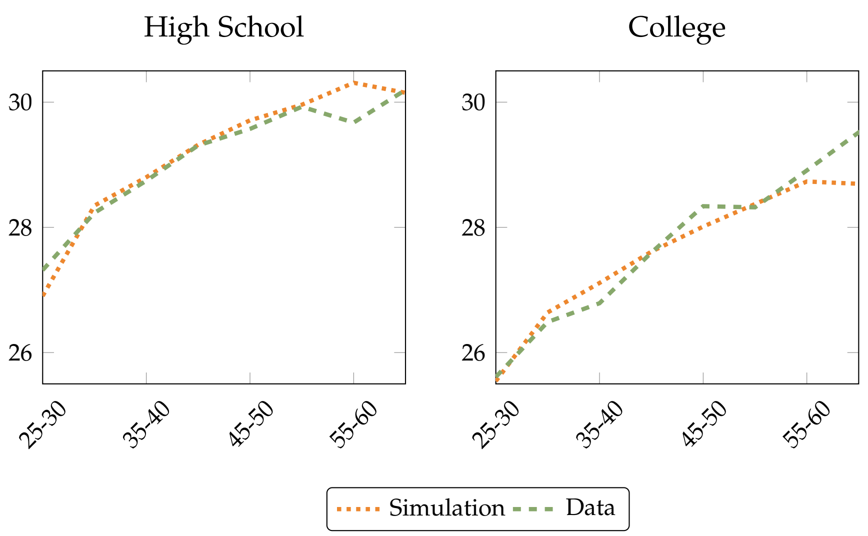

Figure 13: BMI by Education and Age, Simulation vs. Data

Notes: Comparison of simulated and observed BMI profiles by education and age.
Source: Simulation results and CPS-ATUS (2010-2018).

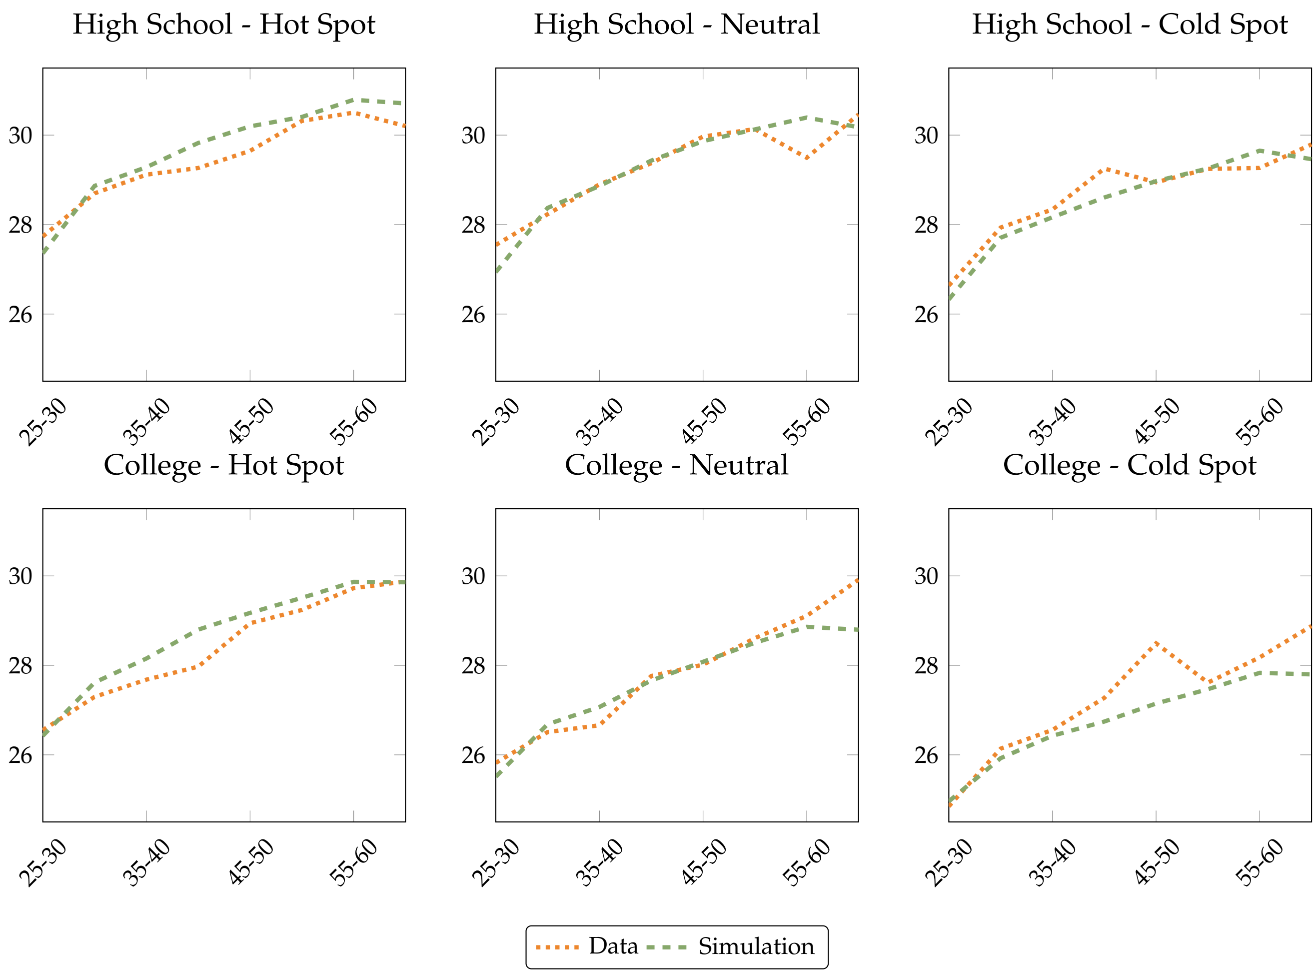

FIGURE 14: BMI BY EDUCATION, CLUSTER, AND AGE, SIMULATION VS. DATA

Notes: Comparison of simulated and observed BMI profiles by education, geographic cluster, and age.
Source: Simulation results and CPS-ATUS (2010-2018).

6 Counterfactuals

Having parameterized the model, we want to use it to understand what drives the spatial and socioeconomic variation in BMI. While our model is designed with cross-sectional comparisons in mind, we can also use the model to learn about the factors that have contributed to the rise in obesity over the last several decades. Lastly, we also use our framework to study the effect of new weight-loss treatments.

6.1 Spatial Decomposition

To study the determinants of the variation in BMI across regions of the US, we conduct an exercise where we assign the parameter values specific to the obesity hot spot to the cold spot, and vice versa. In particular, we in turn vary the following region-specific parameters governing: the BMI reference point, the disutility of exercise, the occupational MET, food prices, labor income, initial BMI, and the share of college graduates. The results are summarized in Table 11.

TABLE 11: Cluster Decomposition

Counterfactual	Hot Spot Avg	Cold Spot Avg
Reference BMI	-0.71	0.67
Disutility Exercise	-0.29	0.39
Education CDF	-0.13	0.22
Initial BMI	-0.22	0.12
MET	0.07	-0.06
Food Prices	-0.03	0.04
Labor Income	0.01	-0.02

Notes: Change in average BMI when assigning the cold spot's parameter values to the hot spot (column 1) and vice versa (column 2). Each row represents a counterfactual where a single parameter is varied while holding all other parameters at their benchmark values. Reference BMI is held fixed across exercises.
Source: Simulation results.

We find that there are four main drivers of the regional differences in BMI: reference BMI, disutility from exercise, initial BMI, and educational composition. Reference BMI is particularly important in rationalizing the spatial patterns to BMI. When the hot spot is assigned a lower BMI reference point corresponding to that of the cold spot, the average BMI in the hot spot declines by 0.71. To put this in perspective, recall that the difference in average BMI between the hot spot and cold spot is roughly 1.2 in the benchmark economy and that 1 unit of BMI corresponds to just below 7 pounds for the average adult man in the US. When the hot spot is assigned the cold spot's lower disutility of exercise, average BMI in the hot spot goes down by 0.29 units. Initial BMI also plays a meaningful role, with average BMI in the hot spot declining by 0.22 when assigned the cold spot's initial BMI distribution. The education composition also contributes to the regional differences. When the hot spot is assigned the educational composition of the cold spot, average BMI in the hot spot declines by 0.13. Other factors either play a negligble role or operate in the opposing direction, such as occupational MET. Notably, the higher occupational MET in the high obesity cluster actually puts downward pressure on BMI in that region. The comparable results from assigning the hot spot's values to the cold spot closely mirror these effects.

In sum, we find that allowing individuals' preferences regarding their own weight to depend on the average BMI of individuals in the region is crucial for reconciling the spatial variation in BMI. To validate this finding, we can compare our results to the empirical literature. Perhaps the cleanest comparison is to Datar and Nicosia (2018), which studies the assignment of military families to different bases around the US. They estimate that a 1 percentage point increase in reference-group obesity raises individual BMI by 0.08 points, implying a predicted effect of 0.96 BMI points for the observed 12 percentage point difference between regions. Our model therefore delivers effects that are of a comparable order of magnitude, but more conservative than these reduced-form estimates.

6.2 Education Decomposition

To study the determinants of the education gradient in BMI, we run a counterfactual exercise where we assign the parameter values specific to the high school type to the college type, and vice versa. In other words, we vary the following education-specific parameters one at a time: the relative weight on healthy vs. reference BMI, the disutility of exercise, the occupational MET, food prices, labor income, the probability of a GHI offer, medical spending, and initial BMI. The results are summarized in Table 12.

Again, the narrative is the same whether one looks at assigning the high school type the values of the college type or the reverse. The quantitatively most important factors in accounting for the education gradient to BMI are the relative weight on healthy vs. reference BMI and the disutility from exercise. When the high school type puts an equally high weight on healthy BMI relative to reference BMI as the college type, the average BMI of high school graduates declines by 0.54. When the high school type is assigned a lower disutility of exercise corresponding to that of the college type, average BMI for high school graduates declines by 0.49. For reference, the benchmark difference in BMI over education is 1.5. Differences in food prices and labor income also play non-negligible roles in rationalizing the BMI patterns over education, with the average BMI of high school graduates declining by 0.12 and 0.07, respectively, when faced with the food prices and labor income of college graduates. Differences in initial BMI and medical spending play only minor roles. Recall that high school graduates are employed in more strenuous occupations than college graduates. Thus, as was the case with the regional decomposition, occupational MET does not help rationalize the patterns in the data. Rather, this mechanism goes in the wrong direction.

Our results suggest that, while benchmarking one's weight to that of people in close geographical proximity is essential for rationalizing the spatial patterns of BMI, preference heterogeneity is critical in accounting for the socioeconomic patterns. Part of that preference heterogeneity, however, relates to the weight that is assigned to reference BMI. Moreover, one should bear in mind that the education composition affects reference BMI.

TABLE 12: Education Decomposition

Counterfactual	High School	College
Relative Weight	-0.54	0.67
Disutility Exercise	-0.49	0.75
Food Prices	-0.12	0.07
Labor Income	-0.07	0.03
Initial BMI	-0.02	0.0
MET	0.7	-0.61
GHI Offer	0.01	0.0
Medical Spending	0.01	0.01

Notes: Change in average BMI when assigning the college type's parameter values to the high school type (column 1) and vice versa (column 2). Each row represents a counterfactual where a single parameter is varied while holding all other parameters at their benchmark values. Reference BMI is held fixed across exercises.
Source: Simulation results.

6.3 Time Decomposition

Average BMI has increased substantially since the 1970s across all clusters, rising from 25.0, 25.1, and 24.6 in the hot spot, neutral, and cold spot, respectively, to 29.3, 28.6, and 27.7 in the current steady state (Table 13). Our framework can also be used to learn about factors that may have contributed to this rise over the last five decades. However, one should keep in mind that our model was designed for the express purpose of studying cross-sectional variation in BMI. As such, it omits some factors that are likely to have played a role in the rise in obesity over time.

We conduct an exercise where we change initial BMI, reference BMI, occupational MET, and the education shares to reflect the values for the 1970s. We present results where each of these parameters is varied one at a time as well as when all of these changes are implemented simultaneously. The results are summarized in Table 14.

When all four changes are implemented, the model accounts for between 60% and 65% of the increase in average BMI since the 1970s, depending on the region. When we break down the effects one by one, we see that reference BMI is by far the single largest factor, accounting for

46-50% of the increase. Initial BMI also plays an important role, accounting for 15-21% of the increase in the average BMI. The literature has stressed the role of MET, arguing that the fact that jobs have become more sedentary over time has contributed to the obesity epidemic. We find that this plays only a small role, since, as documented in the stylized facts, high school graduates are in fact not employed in more sedentary occupations than in the 1970s. The fact that the population has become more educated has served to slightly curb the increase in BMI over time.

Given that the considered mechanisms account for roughly 60-65% of the increase in average BMI since the 1970s, there remains a meaningful portion of the increase that is unaccounted for. Clearly, something other than initial BMI, reference BMI, occupational MET, and education shares has contributed to the rise in average BMI. The aforementioned mechanisms that we consider are all ones that are relatively straightforward to pin down empirically. One can hypothesize that the preference for food (beyond what we pick up through changes in initial BMI) or the weight placed on healthy BMI relative to reference BMI might also have changed. These hypothetical changes are, however, very hard to discipline. Additionally, forces outside the model could well be at play. For example, there is evidence that the calorie-density of food has changed over time. Duffey and Popkin (2011) document that energy density, portion size, and the number of eating occasions have all increased since 1977. Moreover, Mancino, Todd and Lin (2009) find that the increase in calorie share from food away from home has contributed to the rise in calorie consumption over time. In our model, agents choose calories and do not distinguish between food at home and food away from home, calorie density, or portion size.

Nevertheless, we can draw a few very interesting conclusions, both directly from the data and viewing the world through the lens of our model. First, we are able to rule out a few potential mechanisms. In particular, and contrary to what the previous literature has posited, the decline in the strenuousness of occupations plays only a minor role in accounting for the time trends in BMI. We also document that relative food prices display no marked trend, ruling it out as a potential driver. Second, we show that the amplification effect from reference BMI is sizable. It remains a bit of a mystery why BMI has risen so much over the last several decades, but whether the rise in BMI was precipitated by a shift in preferences, or the rise in the calorie-density of food, or something else, it is clear from our analysis that social norms in the form of benchmarking to people in close proximity has magnified this effect substantially.

Table 13: 70's Counterfactuals

Counterfactual	Hot Spot	Neutral	Cold Spot
Data	25.04	25.07	24.64
Benchmark	29.33	28.59	27.73
All	26.54	26.31	25.89
Initial BMI	28.44	27.91	27.27
Reference BMI	27.35	26.84	26.24
MET	29.19	28.51	27.71
Education	29.34	28.7	27.78

Notes: Average BMI across clusters under different counterfactuals. Data shows observed average BMI in the 1970s. Benchmark shows the current steady state. Each counterfactual row shows the average BMI when a single parameter is set to its 1970s value, holding all other parameters at their current benchmark values. All implements all four counterfactuals simultaneously. Reference BMI is held fixed across exercises.
Source: Simulation results.

Table 14: Fraction of the 1970s–Today BMI Increase Explained by Counterfactuals

Counterfactual	Hot Spot	Neutral	Cold Spot
All	0.65	0.65	0.60
Initial BMI	0.21	0.19	0.15
Reference BMI	0.46	0.50	0.48
MET	0.03	0.02	0.01
Education	-0.00	-0.03	-0.02

Notes: Fraction of the increase in average BMI between the 1970s and today that is explained by each considered mechanism. Values represent (Benchmark BMI - Counterfactual BMI) / (Benchmark BMI - 1970s Data). Positive values indicate the mechanism contributes to the observed BMI increase; negative values indicate the mechanism offsets the increase. Reference BMI is held fixed across exercises.
Source: Simulation results.

6.4 GLP-1 Treatment Policy Experiment

Having shown that reference BMI plays an important role in rationalizing the patterns of BMI over space and time, it is natural to ask what implications it has for policies designed to mitigate the obesity epidemic. Next, we consider the exercise of introducing GLP-1 receptor agonist treatment. In the preceding exercises, the social norm, i.e., the BMI reference point, was varied exogenously. In other words, average BMI did not correspond to a fixed point of the social interaction mechanism. As a result, changes in individual behavior did not feed back into reference-group body weight. In contrast, under the GLP-1 policy experiment, we allow the BMI reference point to adjust endogenously. This permits us to assess whether large-scale weight loss among treated individuals affects untreated individuals through changes in reference BMI.

In our framework, we consider a policy scenario that targets individuals with the highest food preferences, mimicking the appetite-suppressing effect of GLP-1 treatment. Specifically, we reduce the food preferences of individuals in the top 10%, 20%, or 30% of the preference distribution to those of the median type. Because in our model food preferences are assigned based on initial BMI at age 25, treatment exposure varies systematically by region and education group. As shown in Table 15, the fraction of individuals treated ranges from 2.8% (college graduates in cold spots in the 10% scenario) to 41.9% (high school graduates in hot spots in the 30% scenario). In order to assess the magnitude of the spillover effect, we run two variants of each policy scenario: holding reference BMI fixed (direct policy effect only), and allowing reference BMI to adjust endogenously (which generates additional externality effects). The results are presented in Table 16.

Table 15: Fraction of Individuals Treated by GLP-1 Policy

	Top 10%		Top 20%		Top 30%	
Cluster	HS	College	HS	College	HS	College
Hot Spot	13.9	11.1	29.4	17.5	41.9	40.6
Neutral	10.1	6.8	22.5	16.7	35.9	27.3
Cold Spot	9.5	2.8	19.1	11.8	36.0	17.4

Notes: Fraction of individuals in each education-cluster group with food preferences in the top 10%, 20%, and 30% of the national distribution, respectively, who receive GLP-1 treatment.
Source: Simulation results.

Table 16: GLP-1 Policy Effects by Cluster

	Top 10% Treated			Top 20% Treated			Top 30% Treated		
Cluster	Policy	Ext.	Total	Policy	Ext.	Total	Policy	Ext.	Total
Hot Spot	-0.344	-0.159	-0.504	-0.593	-0.293	-0.886	-0.869	-0.417	-1.287
Neutral	-0.192	-0.024	-0.216	-0.365	-0.112	-0.477	-0.520	-0.144	-0.664
Cold Spot	-0.136	-0.054	-0.190	-0.280	-0.174	-0.454	-0.428	-0.252	-0.681

Notes: Policy effect = Direct effect with fixed reference BMI. Externality effect = Additional effect from endogenous reference BMI adjustment. Total effect = Combined policy and externality effects.
Source: Simulation results.

The results show substantial and heterogeneous effects across the three policy scenarios. Under the most aggressive scenario – treating the top 30% of the food preference distribution – the direct policy effect on average BMI ranges from a reduction of 0.43 in the cold spot to

0.87 in the hot spot. The externality effects from endogenous reference BMI adjustment add between 0.14 and 0.42 BMI points, yielding total effects ranging from 0.68 to 1.29 across clusters. Under more modest targeting – top 10% – the effects are proportionally smaller but still meaningful, with total reductions ranging from 0.19 to 0.50 BMI points. Recall that 1 unit of BMI corresponds to just shy of 7 pounds for the average American adult male.

The variation in spillovers reveals important heterogeneity. The hot spot exhibits the largest externality effect under the 30% scenario (0.42), reflecting the fact that this region has both the highest treatment rates (42% for high school graduates) and the largest initial BMI disparities. In contrast, the cold spot shows more modest externalities (0.25), consistent with lower treatment penetration, especially among college graduates (17%). This amplification through social norms demonstrates that spillover effects are not uniform but depend on the local distribution of preferences and the baseline reference BMI.

Quantitatively, these results indicate that pharmacological weight-loss interventions generate sizable direct effects and meaningful general equilibrium effects through social interactions. The externality effects represent between 25% and 48% of the direct policy effect, depending on the region and targeting intensity. This implies that ignoring endogenous adjustment of norms leads to a nontrivial understatement of the aggregate impact of treatment. The magnitude of externalities scales roughly proportionally with the fraction of the population treated, suggesting that broader treatment coverage amplifies social multiplier effects. At the same time, even under the most aggressive treatment scenario, regional disparities in BMI persist, suggesting that structural differences in preferences and environments continue to matter beyond the direct effects of treatment.

7 Conclusions

This paper investigates the pronounced spatial and educational gradients in obesity across the US. Despite the widespread rise in obesity over recent decades, substantial heterogeneity persists: obesity rates are higher in the South than in the West, and higher among high school graduates than college graduates. To understand these patterns, we develop and estimate a heterogeneous-agent life cycle model in which individuals make decisions regarding food consumption and exercise, face realistic health and medical spending risk, and derive disutility from deviating both from a healthy BMI and from the reference BMI of their geographic area. This reference BMI mechanism, modeled as conspicuous consumption with loss aversion, captures the idea that individuals benchmark their weight against local social norms.

Our quantitative analysis yields several key findings. First, reference BMI emerges as the dominant factor explaining spatial variation in obesity. Assigning the obesity hot spot the reference BMI of the cold spot reduces average BMI in the hot spot by 0.7 units—accounting for more than half of the observed 1.2-unit regional gap. The disutility from exercise also contributes meaningfully, explaining an additional 0.3 units. Second, the educational gradient in BMI is primarily driven by differences in preferences: the relative weight high school graduates place on reference BMI versus healthy BMI, and their higher disutility from exercise. Each of these mechanisms accounts for approximately 0.5 BMI units of the observed 1.5-unit education gap. Third, our temporal decomposition reveals that changes in reference BMI account for nearly half of the 60-65% of the rise in obesity since the 1970s that our model can explain. Importantly, we find that occupational strenuousness and food prices—factors often emphasized in the literature—do not play an important role in explaining rising obesity. The increasing strenuousness of high school occupations due to job polarization, if anything, works against the observed obesity trends.

The central role of reference BMI has important implications for policy. We simulate a GLP-1 treatment policy that targets individuals with the highest food preferences, allowing reference BMI to adjust endogenously in response to treatment. Under the most aggressive scenario (treating 30% of the distribution), direct effects reduce average BMI by 0.43 to 0.87 units across regions. Crucially, the endogenous adjustment of reference BMI generates substantial spillover effects, adding 0.14 to 0.42 BMI units—representing 25% to 48% of the direct effect. These social multiplier effects are quantitatively large and vary systematically across regions, with the obesity hot spot exhibiting the largest spillovers. The results suggest that standard policy evaluations ignoring general equilibrium effects through social norms may significantly underestimate the aggregate impact of weight-loss interventions.

Our analysis abstracts from several potentially important mechanisms. We do not distinguish between food consumed at home versus away from home, nor do we model the quality or nutritional content of food. Given evidence that the rise in dining out and changes in food quality may have contributed to rising obesity, incorporating these dimensions represents a valuable direction for future work. Additionally, our model does not account for the reverse causal relationship between health shocks and BMI that becomes more prominent at older ages, nor does it consider the role of psychological factors beyond reference BMI that may influence eating and exercise behaviors. Finally, while our framework captures regional variation through reference BMI, a richer spatial model that accounts for migration decisions and the dynamic evolution of local norms could provide additional insights into the geographic concentration of obesity. Despite these limitations, our findings demonstrate that social norms, operating through reference BMI, are a quantitatively important determinant of obesity patterns and a critical channel through which health policies generate spillover effects.

References

Allcott, H., Diamond, R., Dubé, J.P., Handbury, J. et al. 2019. "Food Deserts and the Causes of Nutritional Inequality." *The Quarterly Journal of Economics*, 134(4): 1793–1844.

Anderson, M.L. and Matsa, D.A. 2011. "Are Restaurants Really Supersizing America?" *American Economic Journal: Applied Economics*, 3(1): 152–188.

Archer, E., Hand, G.A. and Blair, S.N. 2013. "Validity of US Nutritional Surveillance: National Health and Nutrition Examination Survey Caloric Energy Intake Data, 1971–2010." *PloS one*, 8(10): e76632.

Autor, D.H., Katz, L.F. and Kearney, M.S. 2006. "The Polarization of the US Labor Market." *American economic review*, 96(2): 189–194.

Baum, C.L. and Ruhm, C.J. 2009. "Age, Socioeconomic Status and Obesity Growth." *Journal of Health Economics*, 28(3): 635–648.

Bramoullé, Y. and Ghiglino, C. 2022. "Loss Aversion and Conspicuous Consumption in Networks." CEPR Discussion Paper DP17181.

Christakis, N. and Fowler, J.H. 2007. "The Spread of Obesity in a Large Social Network over 32 Years." *The New England journal of medicine*, 357(4): 370–9.

Cohen-Cole, E. and Fletcher, J.M. 2008. "Is Obesity Contagious? Social Networks vs. Environmental Factors in the Obesity Epidemic." *Journal of Health Economics*, 27(5): 1382–1387.

Cutler, D.M., Glaeser, E.L. and Shapiro, J.M. 2003. "Why Have Americans Become More Obese?" *Journal of Economic Perspectives*, 17(3): 93–118.

Datar, A. and Nicosia, N. 2018. "Assessing Social Contagion in Body Mass Index, Overweight, and Obesity Using a Natural Experiment." *JAMA Pediatrics*, 172(3): 239.

de Gonzalez, A.B., Hartge, P., Cerhan, J.R., Flint, A.J. et al. 2010. "Body-Mass Index and Mortality among 1.46 Million White Adults." *New England Journal of Medicine*, 363(23): 2211–2219.

DeSalvo, K.B., Bloser, N., Reynolds, K., He, J. et al. 2006. "Mortality Prediction with a Single General Self-Rated Health Question." *Journal of General Internal Medicine*, 21(3): 267–275.

Duffey, K.J. and Popkin, B.M. 2011. “Energy Density, Portion Size, and Eating Occasions: Contributions to Increased Energy Intake in the United States, 1977–2006.” *PLOS Medicine*, 8(6): e1001050.

Glover, A., Heathcote, J., Krueger, D. and Ríos-Rull, J.V. 2023. “Health versus Wealth: On the Distributional Effects of Controlling a Pandemic.” *Journal of Monetary Economics*, 140: 34–59.

Griffith, R., Lluberas, R. and Lührmann, M. 2016. “Gluttony and Sloth? Calories, Labor Market Activity and the Rise of Obesity.” *Journal of the European Economic Association*, 14(6): 1253–1286.

Guner, N., Kaygusuz, R. and Ventura, G. 2014. “Income Taxation of US Households: Facts and Parametric Estimates.” *Review of Economic Dynamics*, 17(4): 559–581.

Hall, R.E. and Jones, C.I. 2007. “The Value of Life and the Rise in Health Spending.” *The Quarterly Journal of Economics*, 122(1): 39–72.

Lakdawalla, B.D., Philipson, T. and Bhattacharya, J.A.Y. 2005. “Welfare-Enhancing Technological Change and the Growth of Obesity.” *The American Economic Review*, 95(2): 253–257.

Mancino, L., Todd, J. and Lin, B.H. 2009. “Separating What We Eat from Where: Measuring the Effect of Food Away from Home on Diet Quality.” *Food Policy*, 34(6): 557–562.

Margaris, P. and Wallenius, J. 2023. “Can Wealth Buy Health? A Model of Pecuniary and Non-Pecuniary Investments in Health.” *Journal of the European Economic Association*, jvad044.

Ozkan, S. n.d.. “Preventive vs. Curative Medicine: A Macroeconomic Analysis of Health Care over the Life Cycle.” 64.

Tudor-Locke, C., Ainsworth, B.E., Washington, T.L. and Troiano, R. 2011. “Assigning Metabolic Equivalent Values to the 2002 Census Occupational Classification System.” *Journal of Physical Activity and Health*, 8(4): 581–586.

Tudor-Locke, C., Washington, T.L., Ainsworth, B.E. and Troiano, R.P. 2009. “Linking the American Time Use Survey (ATUS) and the Compendium of Physical Activities: Methods and Rationale.” *Journal of Physical Activity and Health*, 6(3): 347–353.

Zhao, K. 2017. “Social Insurance, Private Health Insurance and Individual Welfare.” *Journal of Economic Dynamics and Control*, 78: 102–117.

Appendices

A Calibration

A.1 Spatial Analysis

The first step is to show that there is spatial correlation in obesity. We do this by estimating a simple linear model and calculating the Moran's I statistic for global spatial correlation using the residuals.

We regress the prevalence of obesity on county level on the prevalence of inactivity, no physical activity during leisure time, current smoking, a dummy for the classification of the county as "high amenity", recreational facilities per 1000, percent of workers employed in manufacturing, percent of black individuals, median household income, percent of 65 years and older, level of taxation for food and soda, and state fixed effects.

We test the hypothesis that the residuals are not spatially correlated ($H_0 : E\left(uu'\right) = \sigma^2 \mathrm{I}$) using the Moran's I statistic. The Moran's I statistic is defined as:

$$I = \frac{\hat{u} W \hat{u}}{\hat{\sigma}\left[tr\left\{W' + W\right\} W\right]^{1/2}}$$

with $I^2 \sim \chi^2\left(1\right)$. W is the weight matrix that defines how close counties are to each other. We use the inverse of the distance between population centers in each county.

The test rejects the null hypothesis of no spatial correlation in the residuals. Next, we use spatial regression to fit the same model but with the addition of: (i) a spatial lag of obesity (obesity prevalence of other counties), (ii) the spatial lag of a subset of covariates and (iii) spatially correlated errors:

$$\mathbf{y}_{ob} = \beta_0 + \alpha_1 \mathbf{x} + \beta_2 \mathbf{W} \mathbf{y}_{ob} + \alpha_2 \mathbf{W} \mathbf{x}' + \left(\mathbf{I} - \rho W\right) \varepsilon$$

where $\mathbf{W}\mathbf{y}_{ob}$ is the weighted obesity prevalence of the other counties, $\alpha_2 \mathbf{W}\mathbf{x}'$ are the weighted covariates of the other counties and ρW is the weighted correlation of the errors to be estimated.

The coefficient of interest is β_2 which shows whether obesity in other counties affects obesity within the county, even after controlling for the spatial lag of other covariates (for example, amenities in a nearby county might lower obesity in the county of interest because of the ease of visiting) and the potential unobservable common factors or spillovers from nearby counties by controlling for State fixed effects and spatially correlated errors.

β_2 is positive and significant, suggesting that obesity in other counties increases obesity in the county of interest. The coefficient is also economically significant. One percentage point increase in the spatial lag of obesity increases obesity by 0.357 percentage points.

The next step is to identify hot spots of obesity. We use the local Moran statistic to identify counties with high or low obesity prevalence. The statistic is defined as:

Table A1: Spatial Lag OLS

	(1) Obesity Rate
Obesity Rate	
Physical Inactivity	0.213***
No Leisure-Time Activity	0.154***
Education	-0.0802***
Unemployment Rate	0.111***
Current Smoking	0.267***
High Amenities	-0.207**
Metro Area	0.124
Median HH Income	-0.405
Employment in Manufacturing	0.0351***
Recreation Facilities	-3.355***
% Age 65+	-0.0879***
% Non-Hispanic Black	0.0431***
Low Food Access	0.000152
Food Tax	1.116***
Soda Tax (Stores)	4.514***
W	
% Non-Hispanic Black	-0.120**
High Amenities	-7.102***
Metro Area	-0.645
Employment in Manufacturing	-0.474***
Recreation Facilities	-39.42**
Obesity Rate	0.400***
e.Obesity Rate	5.964***
Observations	2976

Notes: Spatial lag regression of county-level obesity rates. The top panel shows direct effects of county characteristics on local obesity rates. The W panel shows spatial lag coefficients, capturing spillover effects from neighboring counties' characteristics. The coefficient on spatial lag of obesity rate (0.400) indicates that a one percentage point increase in neighboring counties' obesity increases local obesity by 0.40 percentage points, controlling for spatial correlation in covariates and errors. State fixed effects included.
Source: CDC BRFSS, USDA Food Access Research Atlas.

$$I_i = \frac{(x_i - \bar{x})}{\sum_{k=1}^{n} (x_k - \bar{x})^2 / (n-1)} \sum_{j=1}^{n} w_{ij} (x_j - \bar{x})$$

With the local Moran statistic, we obtain one statistic per county and the z-score determines the cluster (high or low) that the county belongs to.

Due to data limitations we can only use data on the state level. We classify a state as a high or low obesity hot spot, if 75% of its counties belong to the corresponding hot spot.

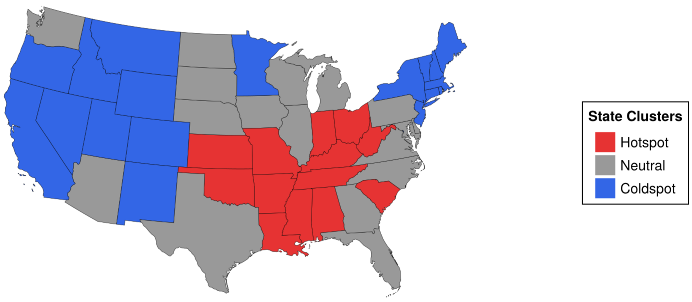

FIGURE 15: HOT SPOTS OF HIGH AND LOW OBESITY PREVALENCE - STATE LEVEL

Notes: The map shows the classification of states according to obesity prevalence. %An index of 1 are hot spots of high obesity prevalence, an index of 2 are cold spots of low obesity prevalence, and an index of 0 are states that are classified as neither.

The states classified as hot spots are: Alabama, Arkansas, Indiana, Kentucky, Louisiana, Mississippi, Missouri, Oklahoma, South Carolina, Tennessee, and West Virginia.

The states that are classified as cold spots are: California, Colorado, Connecticut, District of Columbia, Idaho, Massachusetts, Minnesota, Montana, Nevada, New Hampshire, New Jersey, New Mexico, New York, Oregon, Rhode Island, Utah, Vermont, and Wyoming.

The rest of the states are classified as neither hot spots nor cold spots.

A.2 Labor Income

We estimate the labor income process using pooled CPS–ATUS data merged with the geographic cluster classification described in the Spatial Analysis Appendix. We restrict the sample to working-age individuals (ages 25–65) with positive wage income and non-missing BMI and occupational MET values.

Occupational strenuousness is classified into three categories based on occupational MET: sedentary (MET ≤ 1.5), moderate ($1.5 <$ MET ≤ 2.5), and strenuous (MET > 2.5). Weight classes are defined relative to the BMI distribution within each cluster-education cell. For each cell, we compute the 25th, 50th, and 75th percentiles of BMI and assign individuals to four quartile-based categories. This captures relative body mass within each local and educational reference group rather than relying on absolute BMI thresholds.

For each education group $e \in \{\text{High School}, \text{College}\}$, we estimate the following survey-weighted regression:

$$\ln \tilde{w}_i = \beta_0 + \beta_1 \text{Stren}_i + \beta_2 \text{WeightClass}_i + \beta_3 \text{Cluster}_i + \beta_4 \text{Age}_i + \beta_5 \text{Age}_i^2 + Z_i' \gamma + \varepsilon_i, \quad (13)$$

where $\tilde{w}_i$ denotes normalized real labor income, Stren_i is a set of occupational strenuousness dummies (with sedentary as the reference), WeightClass_i is a set of BMI quartile dummies (with the lowest quartile as the reference), Cluster_i is a set of geographic cluster dummies (with the hot spot as the reference), and Z_i includes controls for marital status, race, and sex.

The full regression results are reported in Table A2. Occupational strenuousness is associated with significantly lower income for both education groups, with the penalty increasing monotonically with strenuousness. The cluster coefficients indicate that individuals in the neutral and cold spot regions earn more than observationally similar individuals in the hot spot. Weight class effects are small and mostly insignificant, with the exception of a modest penalty for the heaviest college graduates.

To construct the income profiles used in the model, we implement a two-step residualization procedure that isolates the structural components of labor income. In the first step, we re-estimate the baseline specification excluding weight class indicators. We then partition the estimated coefficients into two groups: *structural* coefficients—those on strenuousness and cluster, which correspond to mechanisms present in the model—and the remaining coefficients on age, demographic controls, and the intercept. We set the structural coefficients to zero and compute predicted log income using only the non-structural coefficients. This yields a "clean" income measure that is purged of the effects of occupational strenuousness and geographic region.

In the second step, we regress this clean income measure on age and age squared. The resulting age profile coefficients capture the life cycle shape of income at the appropriate level, net of the structural components. We then recombine the age profile coefficients from the second step with the structural coefficients from the first step. This procedure yields income profiles that: (i) have the correct age shape and level for each education group, and (ii) preserve the estimated effects of strenuousness and cluster, which enter the model as separate shifters of the income process.

The model income for an agent of education e, cluster k, strenuousness s, and age j is then constructed as:

$$\ln w_e(j,k,s) = \hat{\beta}_0^e + \hat{\beta}_4^e j + \hat{\beta}_5^e j^2 + \hat{\beta}_{1,s}^e + \hat{\beta}_{3,k}^e, \quad (14)$$

where the age profile coefficients $(\hat{\beta}_0^e, \hat{\beta}_4^e, \hat{\beta}_5^e)$ come from the second-step regression and the structural shifters $(\hat{\beta}_{1,s}^e, \hat{\beta}_{3,k}^e)$ come from the first-step regression.

A.3 Social Security Benefits.

Since it is computationally infeasible to track the lifetime earnings of each agent, we compute average log income by education, cluster, and strenuousness cell and calculate pension benefits using a stylized approximation of the US Social Security system. The replacement rate schedule follows the progressive structure described in Table 8. For each cell, we compute the ratio of average group income to overall average income and apply the corresponding marginal replacement rates. The resulting pension benefits vary by education, cluster, and strenuousness, reflecting the systematic differences in lifetime earnings across these dimensions.

Table A2: Survey-Weighted Regression Results

	Dependent variable:	
	Log Real Income	
	High School or Less	College
	(1)	(2)
stren2	−0.321***	−0.322***
	(0.024)	(0.021)
stren3	−0.497***	−0.661***
	(0.027)	(0.048)
weight_class2	0.008	0.024
	(0.031)	(0.028)
weight_class3	−0.0001	0.007
	(0.030)	(0.029)
weight_class4	−0.014	−0.054*
	(0.029)	(0.031)
StateCluster2	0.080***	0.059**
	(0.025)	(0.028)
StateCluster3	0.143***	0.155***
	(0.026)	(0.028)
age_group5	0.216***	0.222***
	(0.024)	(0.023)
I(age_group5ˆ2)	−0.021***	−0.021***
	(0.003)	(0.003)
marst2	−0.371***	−0.162**
	(0.135)	(0.078)
marst3	−0.194***	−0.183**
	(0.048)	(0.088)
marst4	−0.055**	−0.065**
	(0.025)	(0.032)
marst5	−0.142**	0.023
	(0.068)	(0.078)
marst6	−0.252***	−0.177***
	(0.029)	(0.027)
racen2	−0.059**	−0.043
	(0.027)	(0.044)
racen3	−0.301***	−0.238
	(0.077)	(0.176)
racen4	−0.267***	−0.034
	(0.077)	(0.041)
racen5	0.030	−0.245**
	(0.087)	(0.099)
sex	−0.510***	−0.461***
	(0.021)	(0.021)
Constant	1.777***	2.168***
	(0.070)	(0.069)
Observations	12,399	11,311
Log Likelihood	−17,168.700	−15,421.480
Akaike Inf. Crit.	34,377.400	30,882.960
Note:		*p<0.1; **p<0.05; ***p<0.01

Notes: Survey-weighted regression of log real labor income on individual characteristics. stren2 and stren3 represent occupation strenuousness categories (moderate and strenuous relative to sedentary). weight_class2-4 represent BMI categories. StateCluster2 and StateCluster3 represent neutral and cold spot regions relative to the hot spot baseline. Controls include age (linear and quadratic), marital status, race, and sex. Standard errors in parentheses.
Source: CPS (2010-2024).

A.4 Medical Spending

We estimate the distribution of medical spending shocks using data from the Medical Expenditure Panel Survey (MEPS, 2010–2024). Following Margaris and Wallenius (2023), we partition the distribution of medical expenditures into three bins: below the 50th percentile, between the 50th and 95th percentiles, and above the 95th percentile. Within each bin, we regress medical spending shocks on age, age squared, and BMI, separately by education group. Tables A3 and A4 report the results for high school and college graduates, respectively.

TABLE A3: Medical Spending Shock: High School

	Dependent variable:		
	Medical Spending Shock		
	Bottom 50%	50-95	Top 5%
Age	0.004*	0.097***	0.121***
	(0.002)	(0.008)	(0.021)
Age Squared	0.002***	−0.002***	−0.005***
	(0.0002)	(0.001)	(0.001)
BMI	−0.004	0.122***	1.302***
	(0.004)	(0.021)	(0.068)
Observations	3,497	3,122	363
Note:			*p<0.1; **p<0.05; ***p<0.01

Notes: Regression of medical spending shocks on age, age squared, and BMI for high school graduates, estimated separately by position in the medical spending distribution. Standard errors in parentheses.
Source: MEPS (2010-2024).

A.5 MET

Table A5 reports the distribution of occupational strenuousness by education and geographic cluster, estimated from merged CPS–ATUS data. Occupations are classified into sedentary, moderate, and strenuous categories based on MET values from Tudor-Locke et al. (2011). College graduates are overwhelmingly concentrated in sedentary occupations, while high school graduates are more evenly distributed across strenuousness categories. Across clusters, hot spot regions exhibit a slightly higher share of strenuous occupations for both education groups.

Table A4: Medical Spending Shock: College

	Dependent variable:		
	Medical Spending Shock		
	Bottom 50%	50-95	Top 5%
Age	0.048***	0.050	−0.001
	(0.017)	(0.045)	(0.001)
Age Squared	0.002*	−0.001	0.001***
	(0.001)	(0.003)	(0.0001)
BMI	0.357***	1.594***	0.004
	(0.045)	(0.119)	(0.003)
Observations	760	102	3,693
Note:			*p<0.1; **p<0.05; ***p<0.01

Notes: Regression of medical spending shocks on age, age squared, and BMI for college graduates, estimated separately by position in the medical spending distribution. Standard errors in parentheses.
Source: MEPS (2010-2024).

Table A5: Occupational Strenuousness by Cluster and Education

Education	Cluster	Sedentary (%)	Moderate (%)	Strenuous (%)
HS	Hot Spot	32.3	43.4	24.3
	Neutral	34.0	43.3	22.7
	Cold Spot	34.1	43.1	22.8
College	Hot Spot	55.1	37.9	7.0
	Neutral	57.7	36.6	5.7
	Cold Spot	58.5	36.1	5.4

Notes: Distribution of occupational strenuousness categories by education and geographic cluster. Occupations are classified as sedentary (low MET), moderate, or strenuous (high MET) based on their physical intensity.
Source: CPS-ATUS (2003-2018).

A.6 Exercise

We estimate the average time spent on sports and exercise by cluster and education using the American Time Use Survey (ATUS, 2003–2018). College graduates spend substantially more time exercising than high school graduates across all clusters. Exercise time also increases when moving from hot spot to cold spot regions for both education groups, consistent with the lower obesity prevalence observed in cold spots.

Table A6: Average Time Spent on Sports by Cluster and Education

Cluster	High School	College
Hot Spot	16.0	20.7
Neutral	15.7	24.2
Cold Spot	19.1	26.1

Notes: Average minutes per week spent on sports and exercise by cluster and education among individuals who report any time spent exercising.
Source: Authors' calculations using CPS–ATUS data.

A.7 Health Process

We estimate the health transition probabilities using panel data from the Panel Study of Income Dynamics (PSID). We restrict the sample to male household heads aged 25 and older. Health status is defined as a binary variable based on self-reported health: individuals reporting fair or poor health are classified as unhealthy, and individuals reporting good, very good, or excellent health are classified as healthy.

We estimate two separate random-effects probit models, one for each prior health state. The first model estimates the probability of transitioning from good to bad health, conditioning on the subsample of individuals who were healthy in the previous wave. The second model estimates the probability of remaining in bad health, conditioning on those who were previously unhealthy. In both models, the dependent variable is an indicator for being unhealthy in the current period, and the regressors are age group, age group squared, a college education dummy, current BMI, and cluster dummies (with the hot spot as the base category):

$$\Pr\left(h_t = b \mid h_{t-1} = \varphi\right) = \Phi\left(\beta_0^{\varphi} + \beta_1^{\varphi}\text{Age}_t + \beta_2^{\varphi}\text{Age}_t^2 + \beta_3^{\varphi}\mathbf{1}[\text{College}] + \beta_4^{\varphi}\text{BMI}_t + \gamma^{\varphi}\mathbf{d}_k + u_i\right) \tag{15}$$

where $\varphi \in \{g, b\}$ denotes the prior health state, $\Phi(\cdot)$ is the standard normal CDF, $\mathbf{d}_k$ is a vector of cluster dummies, and $u_i \sim N(0, \sigma_u^2)$ captures unobserved individual heterogeneity.

Table A7 reports the estimation results. Higher BMI significantly increases the probability of being in bad health regardless of prior health status, though the effect is stronger for individuals who were previously healthy. College education is associated with substantially lower probability of bad health in both models. The age profile is U-shaped for the good-to-bad transition, with the quadratic term significant at the 5% level. The cluster dummies are not statistically significant in either model, indicating that after controlling for BMI and education, geographic location does not independently predict health transitions. In the model, we use the estimated coefficients excluding the cluster dummies, so that health transitions depend on age, education, and BMI.

Table A7: Health Transition Probabilities

	(1) Good to Bad b/se	(2) Bad to Bad b/se
Unhealthy		
Age Group	-0.0446	0.1248*
	(0.0375)	(0.0645)
Age Group^2	0.0082**	-0.0021
	(0.0034)	(0.0056)
College	-0.4758***	-0.2510***
	(0.0560)	(0.0871)
BMI	0.0469***	0.0266***
	(0.0043)	(0.0057)
Neutral	-0.0890	-0.0203
	(0.0599)	(0.0971)
Cold Spot	-0.0679	-0.0053
	(0.0698)	(0.1182)
Constant	-3.0109***	-1.2827***
	(0.2021)	(0.2569)
Observations	13144	1856
Groups	6728	1333
Log-Likelihood	-3406.45	-1222.07
χ^2	161.79	52.61

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Random-effects probit estimated on PSID data.
Base cluster category: Hot Spot.

A.8 Cohort Correction

Since our model is a steady-state model, we need to remove cohort effects from the BMI data. We regress BMI on age, age squared, birth cohort, and cohort squared using NHIS data. The estimated cohort effects are used to adjust the data to a common reference cohort, ensuring that the cross-sectional age profiles of BMI used as calibration targets are not contaminated by secular trends across birth cohorts.

Table A8: OLS - Cohort Correction

	(1)
Age	0.319
	(0.002)
Age Sq.	−0.002
	(0.000)
Cohort	0.650
	(0.046)
Cohort Sq.	−0.000
	(0.000)
Num.Obs.	835469
R2	0.069
R2 Adj.	0.069
Log.Lik.	−2352401.688
RMSE	4.04

B Decompositions

The following tables present the full decomposition results underlying the compact tables in the main text. In the education decomposition, we report the change in BMI separately for high school and college graduates, as well as the population-weighted average, within each cluster. In the cluster decomposition, we analogously report results by education group within each cluster. These tables allow the reader to assess how the aggregate effects reported in the main text are distributed across education groups and clusters.

Table A9: Education Decomposition

Counterfactual	Hot Spot			Cold Spot		
	HS	College	Avg	HS	College	Avg
Disutility Exercise	-0.39	0.39	-0.13	-0.66	0.74	-0.04
MET	0.51	-0.45	0.2	0.43	-0.41	0.05
Food Prices	-0.08	0.07	-0.03	-0.09	0.06	-0.03
Relative Weight	-0.39	0.47	-0.1	-0.33	0.41	0.0
Labor Income	-0.04	0.01	-0.02	-0.06	0.03	-0.02
GHI Offer	0.01	-0.01	0.01	-0.01	0.0	-0.01
Medical Spending	-0.04	0.01	-0.02	-0.01	0.0	-0.01
Initial BMI	-0.14	0.11	-0.06	-0.15	0.11	-0.04

Notes: Change in BMI by education group and cluster when assigning the college type's parameter values to the high school type (negative values in HS column) and vice versa (positive values in College column). Avg shows the population-weighted average effect within each cluster. Each row represents a counterfactual where a single parameter is varied while holding all other parameters at their benchmark values. Displayed separately for hot spot and cold spot regions.
Source: Simulation results.

Table A10: Cluster Decomposition

Counterfactual	Hot Spot			Cold Spot		
	HS	College	Avg	HS	College	Avg
Reference BMI	-0.56	-0.45	-0.52	0.63	0.34	0.49
Disutility Exercise	-0.16	-0.6	-0.3	0.28	0.45	0.35
MET	0.03	0.05	0.04	-0.04	-0.04	-0.04
Food Prices	0.01	-0.01	0.01	0.0	0.01	0.0
Education CDF	-0.01	-0.01	-0.14	-0.01	0.01	0.18
Labor Income	0.0	0.01	0.01	-0.02	-0.02	-0.02
Initial BMI	0.0	-0.04	-0.02	-0.02	0.04	0.01

Notes: Change in BMI by education group when assigning the cold spot's parameter values to the hot spot (negative values in Hot Spot columns) and vice versa (positive values in Cold Spot columns). Avg shows the population-weighted average effect within each cluster accounting for the education composition. Each row represents a counterfactual where a single parameter is varied while holding all other parameters at their benchmark values. Education CDF changes the distribution of education types within each cluster.
Source: Simulation results.