The Health Channel of Business Cycles*

Thorsten Drautzburg
FRB of Philadelphia

Grey Gordon
FRB of Richmond

Pablo Guerrón-Quintana
Boston College

Alexey Khazanov
Hebrew University of Jerusalem

March 26, 2026

Abstract

We document that economic contractions causally worsen health among working-age adults and that poor health predicts negative labor market outcomes. These findings reveal a health channel of business cycles. To quantify its magnitude, we build a dynamic general equilibrium model with incomplete markets where agents differ in their health, labor productivity, and wealth. The health channel captures the two-way feedback between pure health shocks—which raise the risk of downward health transitions—and other aggregate shocks, namely, demand and productivity. Our novel estimation strategy identifies the shock correlations, pinning down the health channel. We find the health channel accounts for 10% of employment variance over the cycle and 6% of the employment decline in the Global Financial Crisis.

JEL Codes: C62, E32, E24, D91, I15
Keywords: business cycles, computational methods, health, heterogeneity, labor supply, general equilibrium

*Drautzburg: Federal Reserve Bank of Philadelphia, tdrautzburg@gmail.com. Gordon: Federal Reserve Bank of Richmond, Grey.Gordon@rich.frb.org. Guerrón Quintana: Boston College, pguerron@gmail.com. Khazanov: Hebrew University of Jerusalem, alexey.khazanov@mail.huji.ac.il. We benefited from discussions with Susanto Basu and Sam Richardson and several conference and seminar audiences. Claire Jellison provided excellent research assistance. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia, the Federal Reserve Bank of Richmond, or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Introduction

During the Global Financial Crisis (GFC), self-reported health markedly declined in the U.S. among prime-aged individuals, with about 10% more individuals reporting to be in bad health. More generally, self-reported bad health has been consistently countercyclical since data begins in the 1970s, suggesting this pattern extends well beyond the GFC. What shocks drive health fluctuations over the cycle, and how much do they amplify the business cycle?

After documenting the robust countercyclicality of health, we begin answering these questions by examining the causal impact of aggregate shocks on health empirically. We find worse health outcomes in response to a variety of adverse demand shocks from the literature. For example, a contractionary 1 percentage point (pp.) monetary policy shock increases the fraction of workers with bad health (measured as in Capatina, 2015 and De Nardi et al., 2025) by 0.25 pp. We then assess the impact of pure health shocks, such as bad flu seasons or medical advances, on labor market outcomes. Using state-level data, we estimate that an unexpected, 1 pp. increase in bad health depresses employment by about 0.5 pp. for several years after the shock. Additionally, holding employment rates fixed conditional on health status, a back-of-the-envelope calculation shows that health deterioration during the GFC explains roughly one-eighth of that episode's 4 pp. employment decline. Together, these findings provide empirical support for a health channel of business cycles, where real activity affects health and health affects real activity.

We then build a model where these linkages can be shut down, allowing quantification of the health channel. The model incorporates stochastic health transitions and endogenous labor supply into an Aiyagari (1994) model. Worse health has four costs: lower labor productivity, medical expenditures, and time costs in the form of health investments and *malaise* that induces more disutility for a given amount of labor. We measure the first three costs directly from micro data and infer the fourth. The calibrated model closely matches joint health and employment transition dynamics, such as the flow from good health and employment into poor health and nonemployment. We find that productivity effects and the time costs of health sharply reduce labor supply and savings. Medical spending matters far less. At the aggregate level, deteriorating health reduces employment, capital, and GDP.

The health channel—as the two-way feedback from macro to health—is fundamentally a combination of real shocks that propagate to health and pure health shocks that spill over to real activity. We identify this combination using a new and broadly applicable methodology. Our Wold-Enabled Dynamic General Equilibrium (WEDGE) estimation combines model-based impulse response functions (IRFs) computed from MIT shocks (Boppart et al., 2018) with the Wold decomposition. By mapping IRFs and aggregate shock (co-)variances into aggregate dynamics, we can estimate the covariance between pure health shocks, demand shocks (discount factor shocks), and supply shocks (TFP shocks) while solving the model only once.

Shutting down the health channel—by eliminating the effects of health on real activity and of real activity on health—reduces employment volatility by 10% and makes health significantly less countercyclical. A variance decomposition shows this is almost entirely driven by real activity's impact on health. Productivity and time costs explain most of the health channel, as in the stationary analysis. A complementary historical counterfactual shows that health shocks account for 0.23 pp. of the data's 4.17-percentage-point decline in prime-age employment between 2007 and 2012. Here, the general equilibrium effects reduce the employment effect of pure health shocks by one third. Together, these findings demonstrate that the health channel is quantitatively meaningful: Health significantly amplifies the business cycle.

Related Literature

Our paper contributes to a growing body of work that studies the connection between health and the macroeconomy. De Nardi et al. (2025) studies the lifetime welfare costs of bad health and health shocks using a rich lifecycle model. Capatina (2015) also analyzes the effects of health risk on outcomes such as inequality and savings using a lifecycle model. Capatina and Keane (2025) use a rich lifecycle model of health to study labor supply with endogenous human capital and inequality. Abramson et al. (2024) build a model of mental health and use it to analyze the social cost of mental health problems. In contrast to these papers, we focus on aggregate health risks and the health channel of business cycle transmission, analyzing business cycle statistics rather than long-run outcomes. Health in

our model is an aggregate of mental and physical health. Finkelstein et al. (2025) analyze how countercyclical mortality would affect the welfare cost of business cycles, but their cyclicality is identified from cross-sectional variation. In contrast, we use local projections to identify the aggregate effects of changes in real activity on aggregate health outcomes over time.[1]

This paper contributes to the empirical literature linking health and economic conditions. Most notably, Ruhm (2000) established that mortality is procyclical – falling during economic downturns – using U.S. state-level data. Stevens et al. (2015) and Finkelstein et al. (2025) have confirmed this pattern, with Stevens et al. (2015) showing it partly reflects changes in nursing home staffing rather than individual employment status. However, other studies find the opposite relationship for severe economic shocks. Pierce and Schott (2020) show the China trade shock increased mortality, while Arthi et al. (2022) find similar effects from cotton trade disruptions in historical Britain. Bianchi et al. (2023) also find that unemployment shocks increase mortality in aggregate U.S. data, using timing assumptions for identification.

We contribute to resolving these conflicting findings in several ways. First, we focus on self-reported health among prime-aged adults rather than mortality. This intensive margin measure captures health changes affecting roughly 10% of the population, compared to mortality rates of only 0.2% over two years in our sample. Second, we provide causal identification using multiple established demand shocks – excess bond premia from Gilchrist and Zakrajsek (2012), defense spending from Drautzburg (2020) following Ramey (2011), and monetary policy (Romer and Romer, 2004, Bauer and Swanson, 2023) – making our results more readily interpretable than timing-based approaches. Third, we show that the health-economy relationship varies by shock type: while we confirm Ruhm (2000)'s mortality findings for monetary policy shocks, this relationship does not generalize to financial shocks. Our findings align with Currie and Tekin (2015), who show that local foreclosures increase hospital admissions, consistent with stress linking economic downturns to health deterioration.

Our work also contributes to the literature linking health and labor market out-

[1]Other work in this area includes Fonseca et al. (2023), who study cross-country differences in health care expenditures in a heterogeneous agents framework, though they focus on price mechanisms rather than business cycle dynamics.

comes. A growing body of evidence suggests health improvements causally increase labor market performance. Stephens and Toohey (2022) find that cardiac treatment increased earnings following a 1970s randomized trial, while Garthwaite (2012) shows that the unexpected removal of a popular NSAID reduced the employment-to-population ratio by 0.35 pp. Several studies document similar effects for mental health: Mintz et al. (1992) and Berndt et al. (1998) find that depression treatments improve work performance, with effects emerging after six months or longer. This literature suggests individual health shocks can have substantial labor market consequences. Some health changes even operate at aggregate scales – Krueger (2017) argues that increased opioid prescriptions account for about 40% of the decline in male labor force participation from 1999 to 2016, while Garthwaite (2012)'s NSAID removal affected millions of workers simultaneously. Building on this evidence, we estimate how cyclical changes in population health affect aggregate employment. Our approach complements the existing literature by focusing on surprise changes in the fraction of prime-aged adults in bad health, providing an estimate of the average labor market effect across different types of health deterioration. These estimates directly inform our quantitative model of the health channel in business cycles.

The WEDGE estimation procedure we propose is related to Auclert et al. (2021). They build or approximate a likelihood function using IRFs obtained from their Sequence-Space Jacobian Method. In contrast, our approach works for any solution method and uses simulated method of moments rather than calculating likelihoods. It could be incorporated, however, in quasi-likelihood methods such as the Bayesian estimation in Christiano et al. (2011).

The paper is structured as follows. Section 2 establishes the cyclicality of bad health among prime-aged adults, provides causal estimates using established demand shocks, estimates the effects of pure health shocks, and provides a benchmark for the plausible employment effects of health deterioration during the GFC. Section 3 develops the macroeconomic model and calibrates its stationary version. Section 4 describes and quantifies the model mechanisms. Section 5 estimates the business cycle version of the model and quantifies the health channel. Section 6 concludes.

2 Health and the business cycle: Empirics

We provide three sets of empirical facts. First, we document the cyclical behavior of health in the U.S., both unconditionally and following identified macro shocks. Second, we show that in the cross-section of U.S. states, surprise increases in population health today predict future labor market outcomes. Third, we show that during the GFC several labor market outcomes were correlated with bad health and that the decline in bad health extends beyond the narrow measure of bad health.

2.1 Data sources

Our main source of health information in the time series is the National Health Interview Survey (NHIS, see Blewett et al., 2019), a repeated cross-section with quarterly data. The survey itself begins in 1962. Because of the 2019 redesign, we use data through 2018. Self-reported health "in general" is available going back to 1972, albeit with a change from a 4-category scale to a 5-category scale in 1982.[2] We follow the literature, e.g. Capatina (2015) and De Nardi et al. (2025), for most of our analysis by recoding health as a binary indicator of "bad health," which corresponds to the consistently available categories of "fair" or "poor" bad health. In addition, we splice the time series in 1982 to account for a question wording change.[3] This leaves us with about 1.9 million observations for prime-aged (age 25 to 54) individuals with annual sample sizes of approximately 27,000 to 50,000. To address seasonality and sampling noise, we use 4-quarter moving averages. Also beginning in 1972, the NHIS contains information on the labor force status of respondents. More recent waves also contain information on health insurance status, doctor and hospital visits, and specific medical conditions, which allows us to construct a frailty index (Hosseini et al., 2026) that measures health continuously.

We supplement these data with data from the Current Populations Survey (CPS, Flood et al., 2020), the National Longitudinal Survey of Youth 1997 (NLSY97, U.S. Bureau of Labor Statistics, 2021), the Medical Expenditure Panel Survey (MEPS,

[2]The survey question is a minor variant of the following: "Would you say [fill: your/ALIAS's] health in general is excellent, very good, good, fair, or poor?" with the following answer choices: "Excellent," "Very good," "Good," "Fair," "Poor," "Don't know" and a "Refused" category.

[3]Since 1982, the survey asks about a person's general health. Prior to 1982, the questions asks to compare to others of the person's age. Blewett et al. (2019) cautions that this may impair comparability for some age groups. We address this potential concern by focusing on prime-aged individuals.

Blewett et al., 2023), and the Behavioral Risk Factor Surveillance System Survey Data (BRFSS, CDC, 2023). All four surveys use the same general health status question as the NHIS but have different advantages. The CPS March supplement and the NLSY survey self-reported health annually. The former has detailed labor market outcomes and permits state-level analysis; the latter provides a long panel (covering 2000 to 2019) for a single cohort. The MEPS surveys a rotating group of individuals five times over a two-year period with information on self-reported health, labor market status, and detailed usage of medical services and expenditures. Our state level analysis leverages the BRFSS, a large, state-level, monthly survey that contains employment status, demographics, and additional self-reported health information. Due to changes in the survey methodology, we stop using the data at the end of 2008.

Our macroeconomic data on inflation and interest rates come via FRED (documented in Appendix A). We use shock proxies for excess bond premia from Gilchrist and Zakrajsek (2012) (updated by Favara et al., 2016), monetary policy shocks from Romer and Romer (2004) (updated by Wieland and Yang, 2020) and Bauer and Swanson (2023), and defense spending surprises from Drautzburg (2020).

2.2 The cyclicality of health

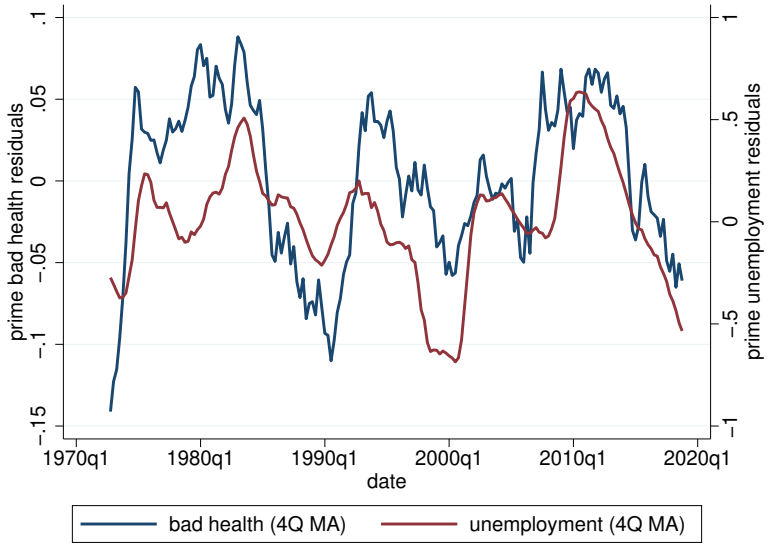
We now show bad health is strongly countercyclical. We begin by showing this for self-reported health, which has the longest time series. Appendix B shows this finding is robust to alternative health measures, conditioning on health insurance status, usage of alternative surveys, inclusion of person fixed effects, and finer disaggregations of health.

Figure 2.1 compares cubic-detrended quarterly unemployment and NHIS bad-health rates from 1972 to 2019. Bad health (the blue line) is strongly countercyclical, moving closely with unemployment (the red line) with a correlation of 0.60.[4] This tight comovement establishes a clear link between health and real activity.

What is driving the comovement of health and employment? We will now show that shocks to real activity causally affect health. Later, we will establish the con-

[4]These results are robust to alternative detrending (including not detrending) as reported in Appendix B.

Figure 2.1: Countercyclicality of bad health



Note: Log-deviations (in percent) from a cubic trend of the fraction of the prime-aged population in bad health and prime-age unemployment, 1972 to 2019; Appendix B shows alternative detrending.

verse is also true, that health shocks—unexplained variation in health—predict real activity. Together, these two linkages constitute our health channel.

2.3 The health response to demand shocks

Overall correlations reflect all shocks affecting the economy, potentially complicating the interpretation of the countercyclical bad health. Here, we show the responses of bad health and employment to three well-understood demand shocks: excess bond premiums, defense spending, and monetary policy.

For each shock, we use a local projection framework, estimating the following set of regressions:

$$y_{t+h} = \alpha_h + \beta_{y,h} d_t + \sum_{\ell=1}^{\text{1 or 2 yrs}} [y_{t-\ell}, d_{t-\ell}]\delta_{\ell,h} + \sum_{u=1}^{\tau} \theta_u t^\tau + e_{t,h}, \tag{2.1}$$

where d_t is our shock measure or proxy, and the outcome y_{t+h} is one of the following in our baseline analysis: (1) the demand or credit friction measure, (2) the employment-to-population ratio, (3) bad health. Following Montiel Olea and

Plagborg-Møller (2021), we always control for lagged values of all outcomes of interest as well as lagged shock measures, as reflected in the $[y_{t-\ell}, d_{t-\ell}]\delta_{\ell,h}$ terms. We then use White-standard errors to account for heteroskedasticity. The coefficients $\{\beta_{y,h}\}_h$ estimate the IRFs of interest. We report pointwise 68% and 90% confidence intervals. In long samples we also control for a deterministic time trend represented by $\sum_{u=1}^{\tau} \theta_u t^\tau$.

Our first shock uses the excess bond premium (EBP) from Gilchrist and Zakrajsek (2012) for 1974 to 2019. Panel (a) of Figure 2.2 shows a 1 pp. increase in the EBP induces a sharp and persistent contraction in employment: The trough of -1.5% (or 1.2 pp.) occurs around quarters 6-8, before beginning a gradual recovery.[5] Bad health rises significantly: The fraction of individuals in poor health increases by about 1% (or 0.1 pp.), peaking around quarters 8-12 before slowly reverting.

Our second shock is an exogenous increase in defense expenditure identified following Drautzburg (2020) (similar to Ramey (2011)) using data from 1974 to 2018. Panel (b) of Figure 2.2 shows defense spending remains 0.5 to 1% above trend for roughly two years. This expansionary shock increases employment gradually, generating a peak of just over 0.1% after 11 quarters. Mirroring these gradual dynamics, the fraction in bad health declines by about 0.5% after 10 quarters. Results are similar for employed individuals (shown) or all prime-aged individuals.

Finally, we examine contractionary Romer and Romer (2004) monetary policy shock using data from 1974 to 2007. Panel (c) of Figure 2.2 shows that the federal funds rate rises 1 pp. on impact and remains elevated for six quarters. Prime-age employment declines only gradually after a substantial lag.[6] This long-and-variable lag of monetary policy contrasts with bad health's response among the employed, which increases significantly beginning at the 2-quarter horizon and peaks at 2.5% (0.25 pp.) two years after the shock. A possible mechanism operates through house prices (Currie and Tekin, 2015), and we show in Appendix C.3 these begin to decline after just two quarters. That appendix also shows the results are robust to Bauer and Swanson (2023) monetary policy shocks.

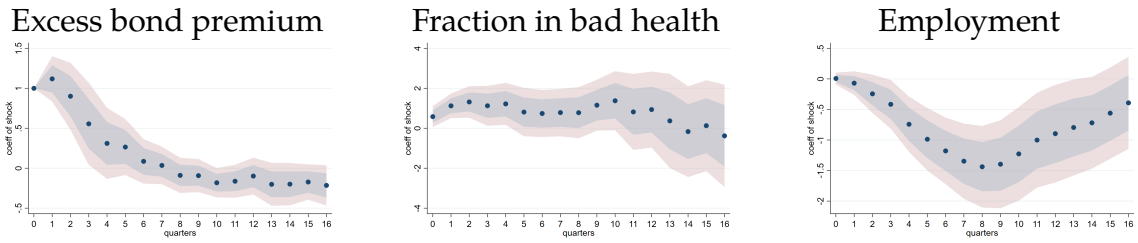
These results establish a clear causal effect of demand shocks on health. Is this effect because economic downturns also drag down health? Or does this effect arise

[5] In Appendix C.2, we show that the responses to the excess bond premium shock are qualitatively similar but quantitatively heterogeneous: The prime aged without access to health care experience larger increases in bad health and larger declines in employment.

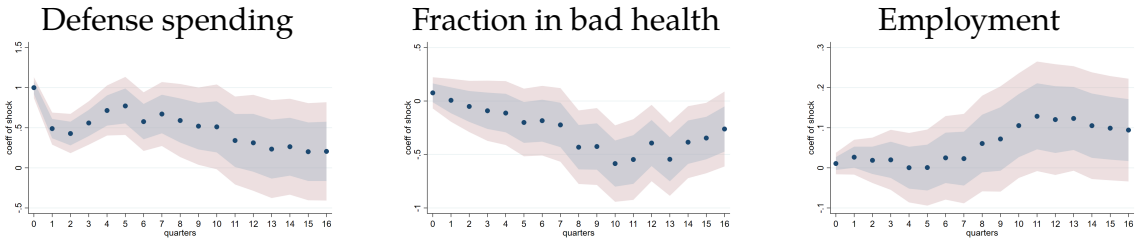
[6] Aggregate unemployment deteriorates faster. See Figure C.4 in Appendix C.3.

Figure 2.2: Health and employment impulse responses to demand shocks

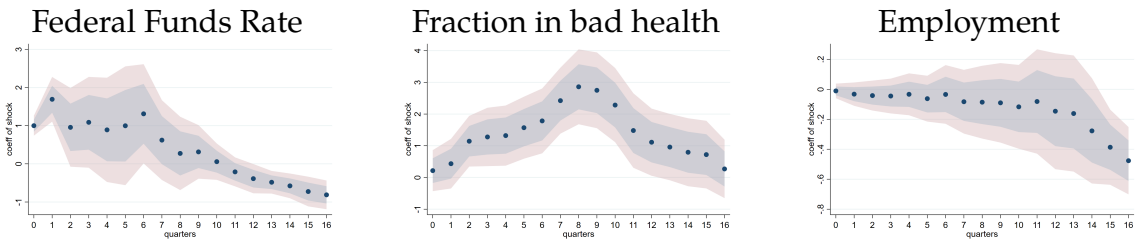
(a) Excess bond premium shocks (contractionary)



(b) Defense spending shocks (expansionary)



(c) Monetary policy shocks (Romer and Romer, contractionary)



Note: Responses are in percent (log deviations); the excess bond premium and monetary policy shocks are contractionary, while the defense spending shock is expansionary; data are for prime-aged workers in the CPS; local projections are estimated with a cubic trend and 4 lags.

because of a change in health input usage? Or does quality of care decline in recessions (Stevens et al., 2015)? As discussed in Drautzburg and Jellison (2025), stress is one possible channel, but also environmental factors such as pollution, healthy habits, or alcohol and drug abuse could play a role. Data limitations prevent us from answering these questions, so we remain agnostic about the underlying mechanisms. Instead, we quantify the health channel by correlating standard aggregate shocks with pure health shocks to recover the comovements between health and real economic activity implied by the data. We now turn to estimating the effects of these health shocks on real activity.

2.4 The labor market response to health shocks

We now seek to quantify the effect of pure health shocks (or health-specific shocks) on real activity. While specific examples exist—such as sudden changes in treatment technologies (Garthwaite, 2012), the emergence of new harmful substances (Krueger, 2017), or health-related regulatory changes—it is difficult to build a broadly applicable empirical strategy around such episodic events. We therefore take a complementary approach, exploiting unexplained health variation across U.S. states to proxy for these shocks. This cross-state variation does not admit a clean causal interpretation comparable to the demand shocks analyzed above. However, the disaggregated structure allows us to control for local labor market characteristics and to mitigate compositional biases through within-market comparisons.

To estimate the effects of health shocks, we use the following local projection specification, similar to the one used for demand shocks:

$$e_{s,t+h} = \mu_{s,h} + \nu_{t,h} + \gamma_{e,h} e_{s,t} + \beta_{e,h} h^b_{s,t} + \sum_{\ell=1}^{4} [e_{s,t-\ell}, h^b_{s,t-\ell}] \delta_{\ell,h} + e_{s,t,h}. \tag{2.2}$$

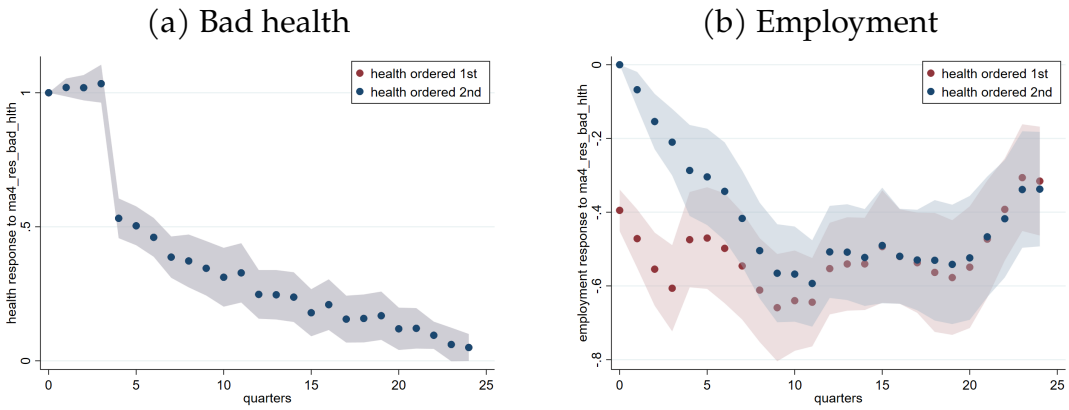
Here, $\mu_{s,h}$ denotes the state s-fixed effects and $\nu_{t,h}$ denotes optional time t-fixed effects. $e_{s,t}$ and $h^b_{s,t}$ denote 4-quarter moving averages of state-level employment and bad health, respectively. Causal orderings analogous to Cholesky-orderings in SVARs are controlled by restrictions on $\gamma_{e,h}$: Constraining $\gamma_{e,h} = 0$ corresponds to ordering bad health first—allowing health shocks to affect employment on impact, but not vice versa—while estimating $\gamma_{e,h}$ freely amounts to ordering health second.

Our sample comes from the BRFSS and spans 1993 to 2008.[7]

Figure 2.3 shows the estimated responses to pure health shocks controlling for state fixed effects. Panel (a) shows the response of prime-aged bad health, and panel (b) shows the response of prime-aged employment. The blue markers denote estimates that order health first ($\gamma_{e,h} = 0$), and red markers order it second ($\gamma_{e,h} \neq 0$). The figure also reports 90% confidence intervals using two-way clustering by state and year.

By construction, bad health rises by 1 pp. It remains close to 1 pp. for the first year because we use 4-quarter moving averages, after which it drops to about 0.5 pp. before declining gradually.[8] This adverse health shock induces a drop in employment of about 0.5 pp. whether or not we constrain the initial employment response to be zero. This effect is persistent, with the employment decline lasting 20 quarters before recovering.

Figure 2.3: Health and employment responses to pure health shocks



Note: State-level impulse responses of prime-aged employment and bad health; 4-quarter moving averages in percentage points, expressed as deviations from state fixed effects; Figure C.9 reports results with state and time fixed effects.

Controlling additionally for time fixed effects, which identifies using only idiosyncratic variation over time within states, results in less power with qualitatively similar (though quantitatively weaker) results. Figure C.9 in Appendix C.5 shows that bad health now drops to about 0.3 pp. after one year and becomes insignificant

[7] A break in the survey methodology prevents a longer sample.

[8] The response without moving averages drops on impact sharply but stays noticeable persistently. This behavior may be driven by a combination of seasonal factors (like bad flu seasons) that decay quickly with permanent changes (like drug inventions) that decay more slowly.

after about two years. The lower persistence is likely due to health shocks diffusing across states, eventually becoming aggregate in nature and absorbed by the time effects. The employment effects are correspondingly also muted, with employment declines of around 0.1 pp. that persist up to 22 quarters after the initial bad health shock.

The feedback from health to employment also holds for both mental and physical health shocks. A shock to mental health induces an employment decline of -0.35 pp., while physical health effects can be up to -0.70 pp. Effects of both are highly persistent, with troughs at around 12 quarters after the shock and noticeable weakness even 24 quarters later. These findings are shown in Figure C.10 in Appendix C.5.

2.4.1 A back-of-the-envelope health shock quantification

The above health shock analysis relied on state-level variation. An alternative perspective comes from individual-level variation: We can mechanically shift the composition of health status while holding fixed employment rates for each health status bin. This back-of-the-envelope calculation abstracts from equilibrium changes in labor supply conditional on health and from other general equilibrium effects, which the model we develop in the next section will capture.

Table 2.1 reports the population share of all five health categories for prime-aged individuals, the employment-to-population ratio, and the health composition change from 2008 to 2010. Employment rates increase monotonically in health, from 22% for those in poor health up to 86%. Holding these fixed and applying the shift in shares implies a 0.49 pp. drop in employment, about one-eighth of the overall 4 pp. drop in employment (see Figure B.6 in Appendix B.6).

This magnitude is robust on at least two dimensions. First, a similar calculation using Table 2.1 with a shift in binary health shares ("Δ2-state") implies a 0.43 pp. employment decline, only slightly smaller. Second, varying the time window results in similar magnitude as well. For instance, for 2007-2010 composition shifts, employment declines by 0.59 pp., while for 2007-2013 shifts, employment declines by 0.40 pp. Longer windows lower the employment decline as health recovers (and employment with it). These and more results are shown in Table B.2 in the Appendix.

Our benchmark calculation thus suggests that the health channel could account

Table 2.1: GFC health deterioration and back-of-the-envelope calculation

Health Status	Share	E/pop [%]	ΔShare [%]	Δ 2-state [%]
Excellent	29.6	86.3	-2.2	
Very good	35.7	85.7	0.1	-1.1
Good	24.8	77.7	1.0	
Fair	7.1	51.7	0.9	
Poor	2.6	21.8	0.1	1.1
Implied emp change [pp]			-0.49	-0.43
Implied bad hlth change [pp]			1.07	1.07

Note: Back-of-the-envelope calculation of the effects of the 2008 to 2010 change in health status on employment, prime-aged individuals, average share and employment rate by health status, change in shares, and implied employment effect from 2008 to 2010.

for around one-eighth of the employment drop in the Global Financial Crisis—holding wages, labor supply, and other responses fixed. In contrast, the quantitative model we develop in the next section allows these margins to adjust endogenously.

2.5 Discussion

Before turning to the model, we interpret our empirical findings in light of the existing literature showing mortality is procyclical and briefly summarize our main results. While our focus is on health rather than mortality, Appendix C.4 shows that most mortality responses to demand shocks in our data are statistically insignificant or mildly countercyclical. This difference reflects both populations and shock types. The existing mortality literature naturally includes retirees, while our focus is on the labor force—prime-aged adults for whom death is uncommon—so statistically insignificant mortality responses are to be expected. Mortality responses are also known to differ by shock type. For example, Pierce and Schott (2020) find mortality worsened after the China trade shock, and we find procyclical mortality in response to Bauer and Swanson (2023) monetary policy shocks.

Taken together, our empirical results establish a clear two-way health channel of business cycles. First, health is strongly and robustly countercyclical. Second, identified demand shocks causally worsen health. Third, health shocks predict persistent declines in employment. Finally, back-of-the-envelope calculations suggest that the magnitude of these health fluctuations is quantitatively meaningful. These empirical facts closely align with existing micro evidence linking health to

economic outcomes and clearly establish the existence of a health channel of business cycles. We now develop an equilibrium model of health that will allow us to quantify the health channel's magnitude.

3 A macro model of the health channel

To analyze what role the health channel plays, we incorporate health in the incomplete markets model of Aiyagari (1994) and calibrate it to the U.S. economy. As in the empirical analysis, we focus on prime-aged individuals, abstracting from lifecycle considerations.

3.1 The model

Individuals differ in health, productivity, and wealth and choose consumption, savings, and labor supply. Production is neoclassical. The government collects proportional labor income taxes and rebates them lump sum.

Individuals value expected flow utility $u(c,x) = (1-\beta)(\log(c) - G(x))$ discounted by β where c is consumption and x is exertion—a combination of labor supply $\ell \in \{0,1\}$ and time costs of health described below. G is isoelastic with $G = \lambda_x x^{1+1/\psi}/(1+1/\psi)$. Households receive iid taste shocks $\sigma_\ell \varepsilon_\ell$ tied to their labor supply choice with ε_ℓ distributed Type-I Extreme Value and σ_ℓ a scale parameter. $\varepsilon = (\varepsilon_0, \varepsilon_1)$ stacks these, which facilitate computation by making employment rates continuous but otherwise play no role.

Health h follows a discrete Markov chain with states $H = \{h_1, \ldots, h_5\} \subset [0,1]$, $h_i < h_{i+1}$, with larger values meaning health is better. Potential labor efficiency $e = \exp(f + \eta)$ has both a fixed component $f \sim N(0, \sigma_f^2)$ and an AR(1) component $\eta = \rho_\eta \eta_{-1} + \epsilon_\eta$, $\epsilon_\eta \sim N(0, \sigma_\eta^2)$.

Health has four direct effects. First, it reduces labor efficiency from its potential e to a realized level $e\psi_H(h)$. Second, it induces medical expenditures $p_H(h,\ell)$. The dependence of expenditures on h and ℓ captures both health risk and insurance. Third, it generates time investments in health, increasing exertion by $\kappa_H(h)$. Fourth, poor health generates malaise κ_M that increases exertion by $(1-h)\kappa_M$. Total exertion is $x(h,\ell) = \ell + \kappa_H(h) + (1-h)\kappa_M$. We measure the first three effects directly in the data and calibrate malaise.

A neoclassical firm operates the production technology $zK^{\alpha}N^{1-\alpha}$ where z is productivity, K is capital, and N is labor in efficiency units. It rents capital services K at rate $r_k = r + \delta$ and hires N efficiency units of labor at unit price w. During production, δK units of capital depreciate. (Competitive behavior ensures profits are zero, so we do not specify them.) The firm solves

$$\max_{K,N} \; zK^{\alpha}N^{1-\alpha} - r_k K - wN, \tag{3.1}$$

which gives factor prices r_k and w from the first order conditions. The government collects proportional labor income taxes $\tau_w wN$ and rebates them with a lump sum transfer T.

The household budget constraint is

$$\begin{gathered} c + a' = (1-\tau_w)we\psi_H(h)\ell - p_H(h,\ell) + a(1+r) + T, \\ c \geq 0, \ell \in \{0,1\}, a' \geq \underline{a} \end{gathered} \tag{3.2}$$

where $r\,(= r_k - \delta)$ is the net return to households from savings a and $\underline{a} \leq 0$ is a borrowing limit. The household's value function solves

$$V(a,h,\eta,f,\varepsilon) = \max_{c,a',\ell} \; u(c,x(h,\ell)) + \sigma_\ell \varepsilon_\ell + \beta \mathbb{E}_{\varepsilon',h',\eta'|h,\eta}\left[V(a',h',\eta',f,\varepsilon')\right] \tag{3.3}$$

subject to (3.2). Defining $\mathbf{s} = (a,h,\eta,f,\varepsilon)$, the solution to (3.3) gives policy rules $\ell(\mathbf{s})$, $c(\mathbf{s})$, and $a'(\mathbf{s})$.

Let $\mu(\mathbf{s})$ denote a joint distribution over individual states. Factor market clearing requires

$$K = \int a(\mathbf{s}) d\mu(\mathbf{s}) \text{ and} \tag{3.4}$$

$$N = \int e\psi_H(h)\ell(\mathbf{s}) d\mu(\mathbf{s}). \tag{3.5}$$

Consistency also requires the distribution μ be generated by itself when applying stochastic transitions and policy rules,

$$\mu(X) = \int \mathbf{1}[(a'(\mathbf{s}),h',\eta',f,\varepsilon') \in X] \mathbb{P}(h',\eta',\varepsilon'|h,\eta) d\mu(\mathbf{s}), \tag{3.6}$$

for measurable X. Goods market clearing holds by Walras' law.

A steady-state recursive general equilibrium is $V, \ell, c, a', K, N, r_k, r, w, T$, and μ such that all the following hold:

- r_k and w are given as by the first order conditions of (3.1) and $r = r_k - \delta$;
- V solves (3.3) with associated policy functions ℓ, c, a';
- the government budget constraint holds, $\tau_w w N = T$;
- factors markets clear, (3.4) and (3.5); and
- the distribution is consistent, (3.6).

A general equilibrium transition path is similar. Letting $\mathbf{p}$ give the parameters of the model that vary according to some law of motion,[9] all variables become conditioned on μ (except for μ itself) and $\mathbf{p}$. (3.6) is modified to give a law of motion from μ to μ' (conditional on $\mathbf{p}$ transitions). Finally, the dynamic equations (3.3) and (3.6) are modified to account for the joint law of motion for $(\mathbf{p}, \mu)$.

Computation of the model is described in Appendix D.1.

3.2 Micro moments, calibration, and validation

We calibrate the steady state of the model to match micro moments. Section 5 discusses how our novel methodology allows us to calibrate with aggregate shocks. A model period is a year.

Most preference and technology parameters are set a priori. The time discount factor β is set to 0.96, delivering a real interest rate just under 3% in the benchmark. Curvature ψ for exertion disutility is set to 1/3, which delivers an equilibrium extensive margin labor supply elasticity close to micro-based estimates from Chetty et al. (2021). The taste shock scale parameter σ_ℓ is set to 10^{-4}, a small value.[10] The capital share is $\alpha = 0.36$, depreciation is $\delta = .08$, and the labor income tax rate is $\tau_w = 0.3$. Total factor productivity z is chosen to target average GDP equal to 1. We

[9] Alternatively, $\mathbf{p}$ can be a time index.

[10] With no disutility from exertion, constant consumption, and $\sigma_\ell \approx 0$, lifetime utility is approximately $\log(c) + \sigma_\ell \varepsilon_\ell$ given our normalizations. From this, a typical value like $\varepsilon_\ell = 1$ has a consumption equivalent effect of 1 basis point.

identify the scale parameter on disutility from exertion λ_x by targeting the nonemployment rate in good health. The borrowing limit $-\underline{a}$ is identified by the fraction of households with negative net worth from Kaplan et al. (2014) (Table 3, column HtM, p. 105).

Potential labor efficiency is set to PSID-based estimates from Storesletten et al. (2004). The fixed effect standard deviation is set to $\sigma_f = 0.378$. Transitory persistence is $\rho_\eta = 0.952$, and $\sigma_\eta = 0.168$ is the average of innovation standard deviations over expansions and recessions.

Health states $H = \{h_i\}_{i=1,\dots,5}$ are determined by regressing log real hourly wages in the CPS on dummies for the five health categories, with "excellent" omitted. The estimated coefficients $(\beta_i) = (-0.409, -0.247, -0.125, -0.045, 0)$, $\beta_5 = 0$ by construction, map wage penalties to health states via $h_i = \frac{e^{\beta_i} - e^{\beta_1}}{e^{\beta_5} - e^{\beta_1}}$. This procedure delivers $H = \{0, 0.35, 0.65, 0.87, 1\}$, similar to the interval midpoints in Halliday et al. (2021). Appendix D.2 reports additional regression results. We report some statistics using a coarser classification of health as being either good, $h \in \{h_3, h_4, h_5\}$, or bad, $h \in \{h_1, h_2\}$.

Of the four direct effects of health, three are taken directly from the data. Health productivity effects $\psi_H(h_i) = 1 + (h_i - 1)(e^{\beta_5} - e^{\beta_1}) = e^{\beta_i}$ deliver the CPS-estimated log real wage penalties. Health time investments $\kappa_H(h_i)$ are set to average health investment time in the ATUS (Hofferth et al., 2020) relative to full-time work (model time units are in multiples of full-time employment $\ell = 1$). Medical expenditures $p_H(h_i, \ell)$ are set to average annual MEPS medical spending (in 2010 dollars) deflated by nominal GDP per capita in 2010 (model numeraire units are in terms of average annual income since $Y = 1$) conditioned on employment status. The remaining direct effect of health, malaise κ_M, is jointly calibrated and identified by the employment rate of those in bad health.

Table 3.1 summarizes the calibration fit, estimated parameters, and some untargeted statistics. The model closely reproduces the targets. Malaise is equivalent to 5% of a working day for those in poor health. Given our normalizations, the borrowing limit allows borrowing of up to 30% of average annual income.

The model also does very well in producing untargeted statistics. The implied real interest rate and capital-GDP ratio are close to typical values. More importantly, the EE (employed-to-employed) and NN (nonemployed-to-nonemployed) transition rates are close to the data counterparts. The model misses on wage dy-

namics as measured in the CPS, but short panels with measurement error likely mean the model is closer than the targets would suggest.

The joint employment and health transition rates in Table 3.2 are the key untargeted moments, since they will strongly influence the model's predictions for both health and non-health shocks. Overall, the model closely reproduces these transition rates. However, the model understates the persistence of the non-employed and bad health state, possibly due to our assumption of a frictionless labor market. This may lead our counterfactuals to understate the persistence of labor market effects associated with the health channel.

4 Health mechanisms in the model

Health has four direct effects in the model: productivity, medical expenditures, and time costs including health investments and malaise. We now show how each mechanism shapes labor supply, savings, and production, and find that productivity losses and time costs are the dominant forces.

Productivity reductions depress labor supply through a substitution effect, which is partially offset by an income effect from lower transfers in general equilibrium. Time costs produce a similar aggregate pattern but operate as preference shifters, making work more onerous rather than reducing wages. Medical expenditures generate only a small pure income effect.

To insure against these costs, high-health individuals precautionarily build buffer stocks and draw down on them once their health deteriorates. Consistent with these mechanisms, aggregate deteriorations in health reduce employment, savings, and GDP.

Table 4.1 examines how these affect the economy by disabling mechanisms and computing how statistics change relative to the benchmark.

4.1 Productivity costs of health

Shutting down productivity costs of health in column NP sharply increases GDP and labor by increasing labor supply of those in bad health via substitution effects. This can be seen by comparing the labor supply policy functions in the baseline and without health productivity costs. Figure 4.1 does this, tracing employment

Table 3.1: Calibration summary

Targeted statistics

Statistic	Model	Target	Parameter	Value
GDP	1.00	1.00	Productivity (z)	0.73
Non-employment rate in bad health (%)	60.07	60.11	Malaise (κ_M)	0.05
Non-employment rate in good health (%)	16.86	16.89	Exertion scale (λ_x)	2.17
Negative net worth (%)	31.57	31.20	Debt limit ($-\underline{a}$)	0.30

Untargeted statistics

Statistic	Model	Target
Real interest rate (%)	2.75	3.00
Capital-GDP ratio	3.35	3.00
EE transition rate (%)	90.33	93.53
NN transition rate (%)	64.45	75.47
Extensive-margin labor supply elasticity	0.39	0.26
Bad health rate (%)	10.44	10.89
Employment rate (%)	78.63	78.41
Log observed wage standard deviation	0.68	1.28
Log observed wage autocorrelation	0.96	0.67

Note: N stands for nonemployed; E stands for employed; ext. marg. stands for extensive margin; wages are conditional on working. The implied stationary distribution has 4.3% of agents in the employed-bad-health state and 6.2% of agents in the non-employed-bad-health state, implying an employment rate of those in bad health of 40.2%. The weighted average of health status shares and employment rates in Table D.2 is 43.7%, a similar rate.

Table 3.2: Untargeted employment-health transition dynamics

Data					Model				
	To					To			
From	EB	NB	EG	NG	From	EB	NB	EG	NG
EB	25.7	6.4	61.3	6.6	EB	37.3	11.3	50.4	1.0
NB	3.7	59.1	8.8	28.4	NB	13.4	40.2	32.4	14.1
EG	3.6	0.8	90.3	5.3	EG	2.2	3.1	87.7	7.0
NG	1.7	11.1	27.1	60.1	NG	0.4	6.9	32.8	59.9

Note: Data transition probabilities are from the CPS; E (N) stands for employed (nonemployed) and B (G) stands for bad (good) health, respectively.

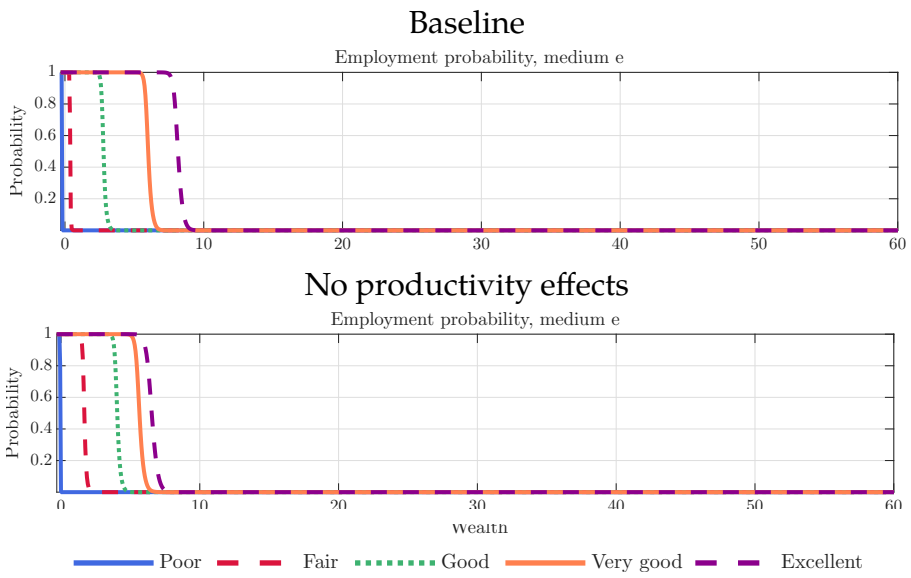
Table 4.1: Model mechanism evaluation by steady state comparison

		Change from baseline (level differences)				
Statistic	B	NP	NT	NX	NH	ZBL
GDP	1.0	0.08	0.03	-0.00	0.11	0.01
Labor (efficiency units)	0.83	0.07	0.03	-0.00	0.10	0.00
Capital-GDP ratio	3.3	-0.01	-0.04	-0.00	-0.05	0.02
N rate bad health (%)	60	-19	-31	1.5	-45	-0.08
N rate good health (%)	17	0.01	-3.0	0.52	-2.2	-0.07
Negative net worth (%)	32	7.3	10	0.06	13	-22
EE transition rate (%)	90	2.1	4.3	-0.39	5.5	-0.22
NN transition rate (%)	64	4.1	6.2	-0.11	11	-0.96
Employment rate (%)	79	2.0	5.9	-0.62	6.7	0.07
Bad health rate (%)	10	0.00	0.00	0.00	0.00	0.00

Note: The table shows the baseline B and change in levels from the benchmark to the indicated model; NP is no health productivity effects; NT is no time costs of health (both investment and malaise); NX is no health expenditures; NH is no health costs of any kind; ZBL sets the borrowing limit to zero; N stands for non-employment; E stands for employment.

probabilities $\mathbb{E}_{\varepsilon_\ell}[\ell(a, h, \eta_{\text{median}}, f_{\text{median}}, \varepsilon_\ell)]$ for wealth on the horizontal axis with one line for each health state. In the baseline, wealth effects decrease labor supply as wealth increases. Employment rates increase in health as realized efficiency $e\psi_H(h)$ increases, creating a substitution effect. When productivity effects are shut down $\psi_H = 1$, employment increases at low to medium health levels. However, the 8% increase in GDP also generates an 8% increase in transfers (since in equilibrium $T = \tau_w wN = \tau_w(1-\alpha)Y$). This generates a positive income effect that reduces labor at high health levels (where ψ_H was already close to 1). While productivity costs of health have an effect on the extensive margin, they affect efficiency units of labor even for those whose labor decision is unaffected. This is why eliminating productivity costs increases labor supply by 8% $(= .07/.83)$ measured in efficiency units, much more than the 2.0% employment rate change.

Figure 4.1: Labor supply functions with and without productivity channel



Note: The permanent type f is set to the median type. Since $e = \exp(f_{\text{median}} + \eta)$, the different levels of e correspond to transient productivity types η.

4.2 Time costs of health

Even without productivity costs, Figure 4.1 shows poor health still strongly reduces employment. This is not attributable to health expenditures, as these are pure income effects pushing labor up else equal (and, which column NX of Table 4.1

shows, has a small effect). So it is necessarily because of time costs, both malaise and health investment time. The extra exertion these cause incentivize individuals to rely on transfers and savings rather than labor income. Consistent with this, column NT of Table 4.1 shows eliminating time costs increases GDP and labor 3% and 4%, respectively. It also has a disproportionately large effect on nonemployment in bad health. Productivity reductions of poor health drive down labor, but the substitution effect is attenuated by an income effect. In contrast, time costs of poor health essentially act like a preference shock, reducing the desire to work. That is why it causes EE and NN transition rates to fall (from NT to the baseline) as time costs make health an important determinant of labor supply.

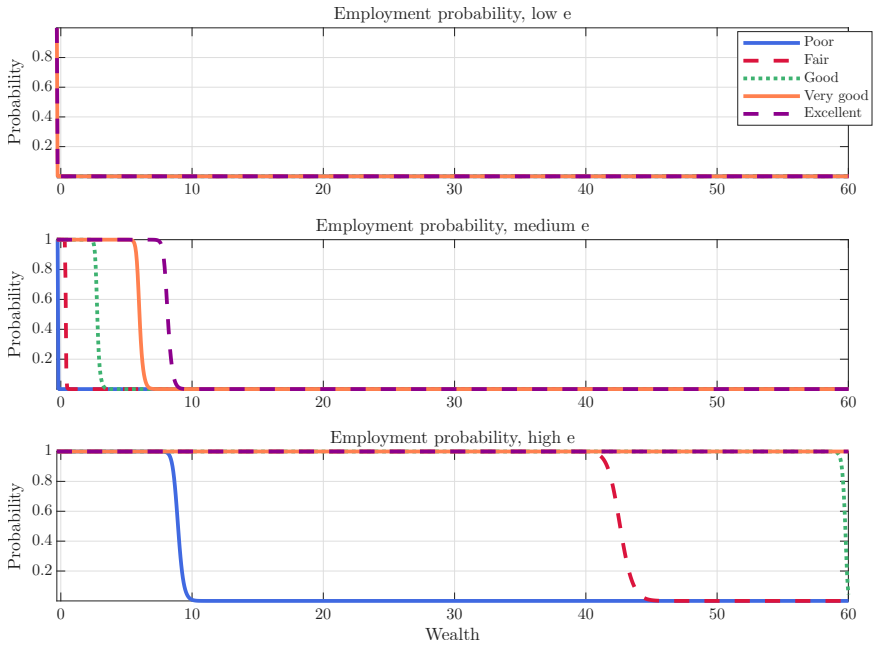
Time costs of health operate on the lower end of the income distribution, changing employment rates much more than efficiency units of labor, as Table 4.1 indicates. But this can be seen directly in employment rates by health, wealth, and labor efficiency, which Figure 4.2 displays. Individuals with high e keep working even if their health deteriorates significantly, unless they have ample assets. For instance, an individual with poor health but high e works unless wealth is more than 25 times average annual income (which is one).

4.3 Other mechanisms

Figure 4.2 shows that an individual with poor health and median e works only if they are borrowing constrained. This result indicates borrowing constraints could play a significant role in transmitting health effects to employment. To examine this, column ZBL of Table 4.1 gives steady state statistics when individuals cannot borrow ($\underline{a} = 0$). The quantitative effects of this channel are mostly small and affect those with bad and good health similarly. For instance, shutting down borrowing constraints lowers non-employment rates of bad and good health by 8 and 7 basis points, respectively. These are an order of magnitude smaller in absolute value than the effects of medical expenditures in NX, and those effects are already small compared to productivity and time costs.

Although we find a small role for medical expenditures, one caveat is that we assumed deterministic medical expenditures conditional on one's health status and labor choice. While we did this to not overstate the role of precautionary savings in our full-commitment model, stochastic expenditures would surely have different

Figure 4.2: Labor supply function: baseline calibration



labor efficiency $e = \exp(f_{\text{median}} + \eta)$

Note: The permanent type f is set to the median type. Since $e = \exp(f_{\text{median}} + \eta)$, the different levels of e correspond to transient productivity types η.

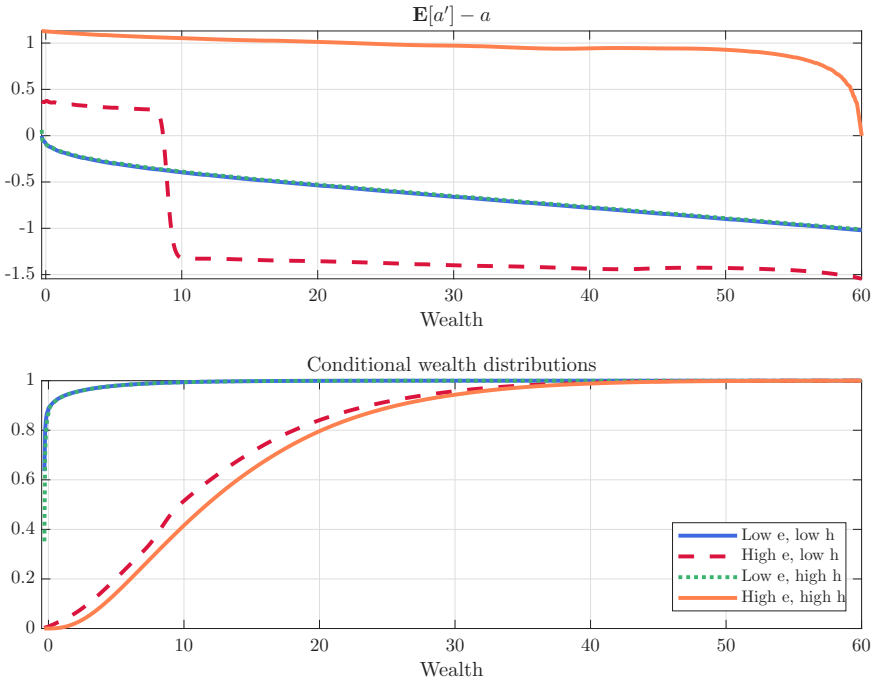
implications. So we do not make a general claim that medical expenditures are unimportant.

4.4 Combined health effects

The joint effect of health costs is substantial. Eliminating all health costs, as done in column NH of Table 4.1, increases GDP and effective labor by more than 10% and increases employment rates by 7 pp. It also increases savings, as the 1.5% ($= -.05/3.3$) decline in capital to GDP is far smaller than the 11% increase in GDP. Health induces two offsetting effects on savings. The first is precautionary savings. That is why when any health costs are eliminated, the share with negative net worth increases, despite large wealth and income increases. The second is that realized poor health causes dissavings as individuals draw down on any buffer stock built precautionarily. These effects can best be seen in Figure 4.3. Its top panel plots average net savings $\int a'(a, h, \eta, f_{median}, \varepsilon_\ell) d\mathbb{P}(\varepsilon_\ell) - a$ with wealth on the horizontal axis and differing levels of h and η reflected in separate lines. For high e (and η) states, individuals save if their health is high and—at moderate to large wealth

levels—individuals dissave if health is low.[11] At lower levels of wealth, the high e agent switches to build assets up until wealth is around 8.5 times average income. The jump from savings to dissavings occurs as the individual moves from working to not working.[12] In contrast, a low e agent almost always finds it optimal to not work, even at high health levels, so a meaningful buffer stock is not built up and the individual gravitates towards the borrowing constraint.

Figure 4.3: Savings policy functions



Note: The permanent type f is set to the median type. Since $e = \exp(f_{\text{median}} + \eta)$, the different levels of e correspond to transient productivity types η.

The figure's bottom panel, which gives wealth distributions conditional on these states, shows how these policies ultimately affect wealth. High efficiency individuals in excellent health have substantial wealth that is used to smooth through health shocks. Individuals in poor health are much more likely to be borrowing constrained irrespective of potential efficiency. However, very close to the constraint,

[11] The upper bound on wealth necessitated by the computation clearly binds in this figure. However the mass of individuals close to this constraint is endogenously very small.

[12] The usefulness of taste shocks can be made clear in this example. Absent taste shocks, this agent's savings behavior would bring them to a point of near indifference in the employment decision (supposing they remained high e low h). This means excess supply of labor can jump discretely for small changes in market clearing prices. Taste shocks effectively convexify this decision.

low e high h individuals precautionarily switch to working (self-insuring against a low h realization), which means they are borrowing constrained at a rate 20 pp less than individuals with low h.

Together, these results indicate that a deterioration of health across individuals will generate less employment and savings—except close to the borrowing constraint—and lower GDP. These predictions are verified in a "health crisis" experiment whose results are given in column HC of Table 4.1. In the health crisis, we assume health transition rates permanently follow the pattern observed in the GFC, which raises the share of population with bad health by 1.2 pp (or 12%, see Table 2.1). This pure health shock lowers GDP, labor and employment rates, and capital. Individuals close to the constraint work a bit more, reducing the share with negative net worth.

While useful for confirming the intuition built up in this section, the health crisis experiment is clearly unsatisfactory. The health deterioration in the GFC was not permanent, and the implications of a transitory shock could be quite different. Moreover, it did not occur in isolation but was accompanied with demand and supply shocks, creating a health channel of business cycles. We now turn to estimating the business cycle version of the model and quantifying the health channel.

5 Quantifying the health channel

We model the health channel by correlating standard business cycle shocks – TFP and discount factor shocks – with a pure health shock calibrated to match changes in CPS health transitions during the GFC. We compute impulse-response function via deterministic transition paths, which Boppart et al. (2018) show is equivalent to a first-order approximation. In two experiments, we fit data on the GFC and overall U.S. business cycles, allowing shocks to health transition probabilities to correlate with business cycle shocks. Eliminating the health shock yields a counterfactual economy without a health channel, isolating its contribution to cyclical fluctuations.

5.1 Impulse-responses

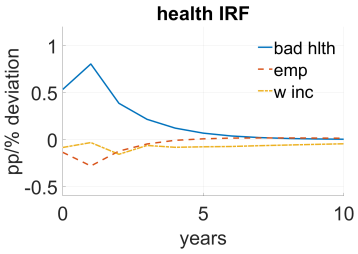
For the health shock, we calibrate a shock to the transition matrix to approximately match the change in the health distribution during the GFC from 2008 to 2010 re-

ported in Table 2.1. After two years, the transition matrix reverts. While other periods could be used, this allows for a straightforward comparison of the pure health shock and the back-of-the-envelope calculation in Section 2.4.1.

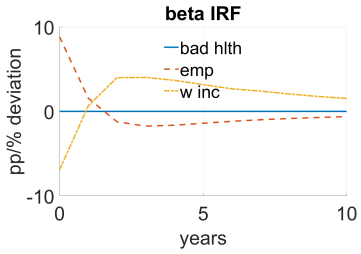
Figure 5.1(a) reports the baseline responses of bad health, employment, and earnings to a pure health shock. By construction, health rises 1.1 pp. to match the response in the data at the two-year horizon before reverting.[13] It falls to less than 0.2 pp. above its mean after 5 years. Employment falls by 0.5 pp. at the two year horizon, before reverting somewhat faster than bad health itself, hitting zero after 5 years. Earnings, in contrast, reach their peak fall of just above 0.2% after three years and remain at about -0.1% from year four to year 10. Appendix D.4 shows time costs primarily drive the employment response, while productivity costs account for most of the wage effects; expenditures play only a minor role. Earnings conditional on employment rise after two years due in part to composition effects as the marginal workers who exit the labor force tend to have lower earnings.

Figure 5.1: Deterministic transitions following health, demand, and supply shocks

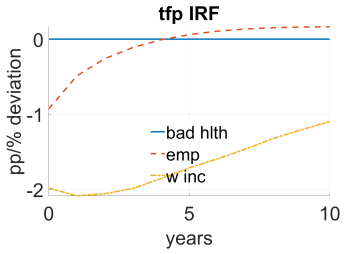
(a) Pure health shock



(b) Demand (β) shock



(c) Supply (z) shock



Note: Responses are in percentage points for employment and bad health rates (levels), and in percent (log deviations) for wage income (w inc). emp is the employment rate; bad hlth is the population bad health rate.

We consider two business cycle shocks in our model: a demand shock in the form of a discount factor β-shock, and a supply shock via TFP z. The demand-shock persistence is 0.41, the annualized value of the 0.80 quarterly persistence in Fernandez-Villaverde et al. (2015). The supply shock persistence is 0.81, annualized from the standard quarterly value of 0.95.

Panels (b) and (c) show the computed impulse responses. By construction, both shocks lower employment by 4 pp.–the fall in prime-age employment in the

[13] We scale the shock to match the response in the data. Note that since the initial distribution is not the model's steady state, the unscaled response is otherwise not exactly identical.

GFC—and have no effect on bad health. The two shocks differ in their wage income responses. A TFP shock reduces earnings because of a lower wage and fewer hours. A β shock, however, increases earnings initially because capital is fixed in the short run, causing elevated K/N that increases the wage by about 1%. Moreover, less productive workers disproportionately work less, inducing a composition shift that increases earnings per worker.

5.2 WEDGE estimation of the business-cycle model

To quantify the health channel, we need the volatility and covariance of pure health shocks, demand shocks, and supply shocks. Directly estimating these inside the full dynamic model would require repeatedly solving the economy with aggregate shocks. Our Wold-Enabled Dynamic General Equilibrium (WEDGE) estimator avoids this computational burden by using impulse responses from a single model solution.

The basic idea is simple. We view the business-cycle version of our model economy as being buffeted by health, demand, and supply shocks, with their dynamic effects traced out by the IRFs in the preceding section. Given a candidate covariance matrix Ω, we can simulate the model's aggregate fluctuations by linearly combining the precomputed IRFs. We then choose Ω to match key business-cycle moments.

Formally, the Wold decomposition allows any covariance-stationary time series to be written as a deterministic component (which we discard) plus an infinite moving average of structural shocks. For each outcome y and shock s, let $\theta_{y,s,l}$ denote the impulse response at horizon $\ell \geq 0$, and let $\Theta_{s,l} = [\theta_{e,s,l}, \theta_{h^b,s,l}, \theta_{\text{earnings},s,l}]'$ stack the responses for employment, bad health, and earnings, respectively. Given a covariance matrix Ω for the health, demand, and supply shocks, the stochastic component of y_t is given by

$$\sum_s \sum_\ell \Theta_{s,l} \text{chol}(\Omega)_{\circ,s} \epsilon_{s,t-\ell} = \sum_\ell \Theta_{\circ,l} \text{chol}(\Omega) \epsilon_{t-\ell}, \tag{5.1}$$

which we truncate in practice. Here, $\text{chol}(\Omega)$ is the lower triangular matrix from the Cholesky-decomposition of Ω, and $\epsilon_{t-\ell} = [\epsilon_{s,t-\ell}]'_s$ is iid $N(0, I)$.

We then choose $\text{chol}(\Omega)$ via simulated method of moments to match the volatility of employment, bad health, and wage income as well as the correlation of em-

ployment and bad health, placing an order of magnitude more weight on the first two targets. We impose zero correlation between demand and supply shocks. This yields the following $\text{chol}(\Omega)$ matrix:

$$\text{chol}(\Omega) = \begin{bmatrix} \underset{\sigma_\beta}{0.17} & 0 & 0 \\ 0 & \underset{\sigma_z}{0.65} & 0 \\ \underset{\rho_{h\beta}}{-0.88 \times \omega_h} & \underset{\rho_{hz}}{-0.12 \times \omega_h} & \underset{\sigma_h}{0.37} \end{bmatrix},$$

where $\omega_h \equiv \frac{\sigma_h}{\sqrt{1-\rho_{h\beta}^2-\rho_{hz}^2}}$ is the combined standard deviation of health surprises from the pure health shocks and the two business cycle shocks.

To eliminate the causal impact of demand and supply shocks on health, we set $\rho_{h\circ} = 0$. To completely eliminate the health channel, we set $\sigma_h = 0$. So going from the benchmark to $\rho_{h\circ} = 0$ eliminates real activity's impact on health, and going from $\rho_{h,\circ} = 0$ to $\sigma_h = 0$ eliminates the feedback from health to real activity.

Table 5.1 reports business-cycle statistics based on CPS data in the first column, using cubic-detrended log-series.[14] The calibration matches these moments almost perfectly, see column 2. The model also matches the other correlations (not shown) closely: The wage and employment correlation of -0.25 is close to the -0.22 in the data, and the correlation of bad health and wages is 0.05 versus 0.09 in the data. It does so by calibrating a correlation of bad health with the demand shock of -0.88 and of bad health with the TFP shock of -0.12.

5.3 Measuring the health channel

We can now quantify the health channel, which is the two-way feedback where real activity affects health and health affects real activity. To do this, we shut down health shocks in the model ($\sigma_h = 0$), which eliminates both directions of feedback. The health channel explains 10.2% ($= 100 - 89.8\%$) of the employment variance, as seen in the right-most column (top line) of Table 5.1.

We can also measure the contribution of each direction of this channel by eliminating the correlation between the health shock and other aggregate shocks. In this counterfactual, health shocks still affect real activity, but real activity no longer influences health. Thus, moving from the baseline to this counterfactual removes the

[14] See Figure D.1 for time series plots.

effect of real activity on health. Moreover, moving from this counterfactual to no health shocks eliminates the second feedback, where health influences real activity. This decomposition implies nearly all of the health channel's effect on employment operates through the real activity to health direction, with the remainder coming from health feeding back into real activity.

Table 5.1: Calibrated business-cycle statistics (standard deviations in percent of cubic-detrended log-series), the contribution of the health channel, and counterfactuals shutting down components of the channel

		Health channel directions Relative variance/correlation (%)		
	Baseline SD [data, %]	Both $\sigma_h > 0, \rho_{h,\circ} \neq 0$	Health to macro $\sigma_h > 0, \rho_{h,\circ} = 0$	Neither $\sigma_h = 0$
Employment	2.45	100.0%	90.2%	89.8%
Bad health	7.95	100.0%	20.1%	0%
Wage income	4.45	100.0%	98.8%	98.7%
Corr(Emp, bad hlth)	-0.44	100.0%	85.6%	–

The health channel is mostly attributable to productivity and time costs of health. Table D.3 in Appendix D.5 shows that disabling productivity costs of health shrinks the health channel by 73%, while disabling time costs reduces it by 63%. Medical expenditures have little effect. These magnitudes are consistent with the contributions seen in Section 4.

Finally, we estimate the sequence of shocks hitting the economy during the GFC to quantify the contribution of the health channel to that recession. Fitting the model to match the paths of bad health, employment, and wage income, Appendix D.5 finds innovations in excess of two standard deviations to each of health, demand, and supply (technology). In a counterfactual without the health channel, the peak employment decline is 0.23pp smaller. Thus, the health channel accounts for about 6% of the peak employment decline during the GFC.

This model-implied contribution of health to employment fluctuations during the GFC is therefore only about half as big as the back-of-the-envelope estimate of one-eighth from Section 2.4.1. However, comparing the GE effect to the back-of-the-envelope calculation and a partial equilibrium effect sheds light on the different mechanisms. In the back-of-the-envelope calculation employment rates by

health status are fixed. In the PE version, wages and interest rates are fixed, but employment rates are still endogenous. In this case we our model attributes an employment drop of 0.33pp to the health channel, or 8% of the observed decline. The fact that this is smaller than in the back-of-the-envelope calculation means that employment rates conditional on health status increase even for given wages. This may reflect a negative wealth effect. In general equilibrium the increase in wages incentivizes households to supply even more labor, reducing the impact of the health channel. Figure D.6 in the appendix illustrates this point by contrasting the IRFs to a demand shock in partial equilibrium and general equilibrium. The employment response is amplified in partial equilibrium, but the earnings for the employed is more muted.

These results demonstrate that the health channel is quantitatively meaningful: It explains roughly 10% of employment fluctuations over the business cycle. We conclude by discussing the implications of this channel and connecting back to the empirical evidence.

6 Conclusion

This paper shows that health is a quantitatively important channel of business cycles. Using fifty years of U.S. data, we document that health among prime-aged adults is strongly countercyclical, that established demand shocks causally worsen health, and that health shocks predict persistent declines in employment. These findings provide clear evidence of a two-way feedback between real activity and health.

We quantify this feedback using a heterogeneous-agent general equilibrium model disciplined by micro evidence on health transitions and employment dynamics. Our WEDGE estimation procedure identifies the covariance between health shocks and standard aggregate shocks without repeatedly solving the model.

The resulting estimates imply that the health channel accounts for roughly ten percent of employment fluctuations. Most of the health channel is derived from real activity's effect on health. Productivity losses and time costs of poor health emerge as the key mechanisms linking health to labor supply and earnings.

Overall, our results highlight that fluctuations in population health are not merely a consequence of the business cycle but also an amplifier of it. Incorporating health

into macroeconomic analysis can therefore improve our understanding of aggregate dynamics and inform policies that jointly stabilize economic activity and health.

References

B. Abramson, J. Boerma, and A. Tsyvinski. Macroeconomics of Mental Health. Cowles Foundation Discussion Papers 2387, Cowles Foundation for Research in Economics, Yale University, Apr 2024.

S. R. Aiyagari. Uninsured Idiosyncratic Risk and Aggregate Saving. Quarterly Journal of Economics, 109(3):659–684, 1994. doi: https://doi.org/10.2307/2118417.

V. Arthi, B. Beach, and W. W. Hanlon. Recessions, Mortality, and Migration Bias: Evidence from the Lancashire Cotton Famine. American Economic Journal: Applied Economics, 14(2):228–55, April 2022. doi: 10.1257/app.20190131. URL https://www.aeaweb.org/articles?id=10.1257/app.20190131.

A. Auclert, B. Bardóczy, M. Rognlie, and L. Straub. Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. Econometrica, 89(5):2375–2408, 2021. doi: https://doi.org/10.3982/ECTA17434. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA17434.

M. D. Bauer and E. T. Swanson. An Alternative Explanation for the "Fed Information Effect". American Economic Review, 113(3):664–700, March 2023. doi: 10.1257/aer.20201220.

E. R. Berndt, S. N. Finkelstein, P. E. Greenberg, R. H. Howland, A. Keith, A. Rush, J. Russell, and M. B. Keller. Workplace Performance Effects from Chronic Depression and Its Treatment. Journal of Health Economics, 17(5):511–535, 1998. ISSN 0167-6296. doi: https://doi.org/10.1016/S0167-6296(97)00043-X.

F. Bianchi, G. Bianchi, and D. Song. The Long-Term Impact of the COVID-19 Unemployment Shock on Life Expectancy and Mortality Rates. Journal of Economic Dynamics and Control, 146:104581, 2023. ISSN 0165–1889. doi: 10.1016/j.jedc.2022.104581.

L. A. Blewett, J. A. R. Drew, M. L. King, and K. C. Williams. IPUMS Health Surveys: National Health Interview Survey, Version 6.4 [dataset]. Minneapolis, MN: IPUMS, 2019. https://doi.org/10.18128/D070.V6.4.

L. A. Blewett, J. A. R. Drew, D. Backman, A. Chen, G. Cooper, M. Schouweiler,

S. Richards, and M. Westberry. IPUMS Health Surveys: Medical Expenditure Panel Survey, Version 2.3 [dataset]. Minneapolis, MN: IPUMS, 2023.

T. Boppart, P. Krusell, and K. Mitman. Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative. Journal of Economic Dynamics and Control, 89:68–92, 2018. doi: https://doi.org/10.1016/j.jedc.2018.01.002.

E. Capatina. Life-Cycle Effects of Health Risk. Journal of Monetary Economics, 74: 67–88, 2015. doi: https://doi.org/10.1016/j.jmoneco.2015.06.002.

E. Capatina and M. Keane. Health Shocks, Health Insurance, Human Capital, and the Dynamics of Earnings and Health. Journal of Political Economy, 2025. doi: https://doi.org/10.1086/739831. forthcoming.

CDC. Behavioral Risk Factor Surveillance System Survey Data, 2023. URL https://www.cdc.gov/brfss/annual_data/annual_data.htm. Accessed: 11/11/2023.

R. Chetty, A. Guren, D. Manoli, and A. Weber. Micro vs. Macro Labor Supply Elasticities: Resolving the Inconsistency. American Economic Review: Papers & Proceedings, 111:285–291, 2021. URL https://rajchetty.com/wp-content/uploads/2021/04/micro_macro_aerpp.pdf.

L. Christiano, M. Trabandt, and K. Walentin. Handbook of Monetary Economics. Chapter DSGE Models for Monetary Policy Analysis, pages 285–367, 01 2011. doi: 10.1016/S0169-7218(11)03007-3.

J. Currie and E. Tekin. Is There a Link between Foreclosure and Health? American Economic Journal: Economic Policy, 7(1):63–94, February 2015. doi: 10.1257/pol.20120325. URL https://www.aeaweb.org/articles?id=10.1257/pol.20120325.

M. De Nardi, S. Pashchenko, and P. Porapakkarm. The Lifetime Costs of Bad Health. The Review of Economic Studies, 92(3):1987–2026, None 2025. doi: None.

T. Drautzburg. A Narrative Approach to a Fiscal DSGE Model. Quantitative Economics, 11(2):801–837, 2020. doi: https://doi.org/10.3982/QE1083. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/QE1083.

T. Drautzburg and C. Jellison. Bad Times, Bad Health. Economic Insights, 10(1): 7–17, March 2025.

G. Favara, S. Gilchrist, K. F. Lewis, and E. Zakrajsek. Recession Risk and the Excess Bond Premium. FEDS Notes, (2016-04-08), Apr 2016. doi: 10.17016/2380-7172.1739.

J. Fernandez-Villaverde, G. Gordon, P. Guerron-Quintana, and J. F. Rubio-Ramírez. Nonlinear Adventures at the Zero Lower Bound. Journal of Economic Dynamics and Control, 57(C):182–204, None 2015. doi: 10.1016/j.jedc.2015.05.014.

A. Finkelstein, M. J. Notowidigdo, F. Schilbach, and J. Zhang. Lives Versus Livelihoods: The Impact of the Great Recession on Mortality and Welfare. The Quarterly Journal of Economics, 140(3):2269–2328, 05 2025. doi: 10.1093/qje/qjaf023.

S. Flood, M. King, R. Rodgers, S. Ruggles, and J. R. Warren. Integrated Public Use Microdata Series, Current Population Survey: Version 7.0 [dataset]. Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D030.V7.0.

R. Fonseca, F. Langot, P.-C. Michaud, and T. Sopraseuth. Understanding Cross-Country Differences in Health Status and Expenditures: Health Prices Matter. Journal of Political Economy, 131(8):1949–1993, 2023.

C. L. Garthwaite. The Economic Benefits of Pharmaceutical Innovations: The Case of Cox-2 Inhibitors. American Economic Journal: Applied Economics, 4(3): 116–37, July 2012. doi: 10.1257/app.4.3.116. URL https://www.aeaweb.org/articles?id=10.1257/app.4.3.116.

S. Gilchrist and E. Zakrajsek. Credit Spreads and Business Cycle Fluctuations. American Economic Review, 102(4):1692–1720, 2012.

G. Gordon and S. Qiu. A Divide and Conquer Algorithm for Exploiting Policy Function Monotonicity. CAEPR Working Paper 2017-006, Center for Applied Economics and Policy Research, Indiana University, April 2017.

G. Gordon and S. Qiu. A Divide and Conquer Algorithm for Exploiting Policy Function Monotonicity. Quantitative Economics, 9(2):–, 2018. doi: https://doi.org/10.3982/QE640.

T. Halliday, B. Mazumder, and A. Wong. Intergenerational Mobility in Self-Reported Health Status in the US. Journal of Public Economics, 193:104307, 2021. doi: https://doi.org/10.1016/j.jpubeco.2020.104307.

S. L. Hofferth, S. M. Flood, M. Sobek, and D. Backman. American Time Use Survey Data Extract Builder: Version 2.8 [dataset]. College Park, MD: University of Maryland and Minneapolis, MN: IPUMS, 2020.

R. Hosseini, K. Kopecky, and K. Zhao. How important is health inequality for lifetime earnings inequality? Review of Economic Studies, 93(1):556–599, 2026.

G. Kaplan, G. L. Violante, and J. Weidner. The Wealthy Hand-to-Mouth. Brookings

papers on economic activity, Brookings Institution, 2014.

A. B. Krueger. Where Have All the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate. Brookings Papers on Economic Activity, 48(2 (Fall)):1–87, 2017.

J. Mintz, L. I. Mintz, M. J. Arruda, and S. S. Hwang. Treatments of Depression and the Functional Capacity to Work. Archives of General Psychiatry, 49(10):761–768, 10 1992. ISSN 0003-990X. doi: 10.1001/archpsyc.1992.01820100005001. URL https://doi.org/10.1001/archpsyc.1992.01820100005001.

J. L. Montiel Olea and M. Plagborg-Møller. Local Projection Inference Is Simpler and More Robust Than You Think. Econometrica, 89(4):1789–1823, 2021. doi: 10.3982/ECTA18756.

J. R. Pierce and P. K. Schott. Trade Liberalization and Mortality: Evidence from US Counties. American Economic Review: Insights, 2(1):47–64, March 2020. doi: 10.1257/aeri.20180396.

V. A. Ramey. Identifying Government Spending Shocks: It's all in the Timing. The Quarterly Journal of Economics, 126(1):1–50, 2011. doi: https://doi.org/10.1093/qje/qjq008.

C. D. Romer and D. H. Romer. A New Measure of Monetary Shocks: Derivation and Implications. American Economic Review, 94(4):1055–1084, September 2004. doi: 10.1257/0002828042002651. URL https://www.aeaweb.org/articles?id=10.1257/0002828042002651.

C. J. Ruhm. Are Recessions Good for Your Health? Quarterly Journal of Economics, 115(2):617–650, 2000.

J. Stephens, Melvin and D. Toohey. The Impact of Health on Labor Market Outcomes: Evidence from a Large-Scale Health Experiment. American Economic Journal: Applied Economics, 14(3):367–99, July 2022. doi: 10.1257/app.20180686.

A. H. Stevens, D. L. Miller, M. E. Page, and M. Filipski. The Best of Times, the Worst of Times: Understanding Pro-cyclical Mortality. American Economic Journal: Economic Policy, 7(4):279–311, November 2015. doi: 10.1257/pol.20130057.

K. Storesletten, C. I. Telmer, and A. Yaron. Cyclical Dynamics in Idiosyncratic Labor Market Risk. Journal of Political Economy, 112(3):695–717, June 2004. doi: 10.1086/383105. URL https://doi.org/10.1086/383105.

U.S. Bureau of Labor Statistics. National Longitudinal Survey of Youth 1997 cohort,

1997-2017 (rounds 1-19). Produced and distributed by the Center for Human Resource Research (CHRR), The Ohio State University. Columbus, OH, 2021. URL https://www.nlsinfo.org/.

J. F. Wieland and M.-J. Yang. Financial Dampening. Journal of Money, Credit and Banking, 52(1):79–113, February 2020. doi: 10.1111/jmcb.12681.

Online Appendix for

The Health Channel of Business Cycles

Thorsten Drautzburg
FRB of Philadelphia

Grey Gordon
FRB of Richmond

Pablo Guerrón-Quintana
Boston College

Alexey Khazanov
Hebrew University of Jerusalem

March 26, 2026

A Data appendix

This appendix reports summary statistics (Section A.1) and shows how health metrics comove generally and in the GFC (Section A.2).

A.1 Summary statistics

Table A.1 reports summary statistics of the data.

Table A.1: Summary statistics

	Good health		Bad health		All	
	Mean	SD	Mean	SD	Mean	SD
Age (years)						
Men	39.18	8.53	42.88	8.24	39.48	8.57
Women	39.25	8.53	42.69	8.21	39.58	8.56
All	39.22	8.53	42.78	8.22	39.53	8.56
Doctor visits (annual, per person)						
Men	0.09	0.29	0.26	0.44	0.11	0.31
Women	0.16	0.37	0.36	0.48	0.18	0.38
All	0.13	0.33	0.32	0.47	0.14	0.35
Employed (fraction)						
Men	0.90	0.30	0.51	0.50	0.87	0.34
Women	0.76	0.43	0.42	0.49	0.73	0.45
All	0.83	0.38	0.46	0.50	0.80	0.40
Bad health (fraction)						
Men	0.00	0.00	1.00	0.00	0.08	0.27
Women	0.00	0.00	1.00	0.00	0.09	0.29
All	0.00	0.00	1.00	0.00	0.09	0.28
Frailty (average)						
Men	0.06	0.07	0.23	0.17	0.07	0.09
Women	0.07	0.08	0.27	0.18	0.09	0.11
All	0.06	0.07	0.25	0.18	0.08	0.10
Mortality (per 1,000)						
Men	1.90	0.06	14.04	0.17	2.88	0.08
Women	1.01	0.04	9.50	0.14	1.82	0.06
All	1.45	0.05	11.55	0.15	2.34	0.07

A.2 Health measure comovement

Figure A.1 shows the Z-score (mean differences relative to the standard deviation) of various health metrics over time. The different series display a significant degree of correlation among each other. Table A.2 focuses specifically on how the health categories move over the GFC and finds a similar comovement. It also reports the implied contribution from different categories, showing in the detrended series that depression plays the largest role, followed by limitations and then conditions.

Figure A.1: Relationship between health indicators over time

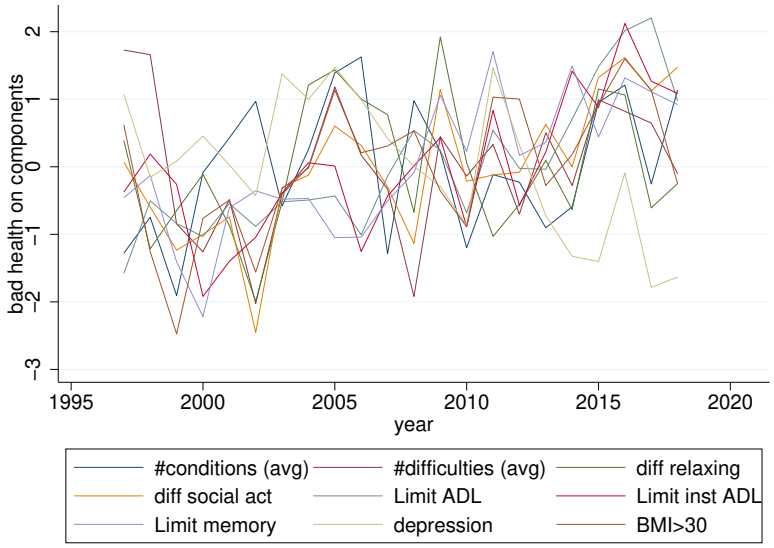


Table A.2: Breakdown of frailty changes: 2007q4 to 2010q4

	Levels		Difference		Detrended	
	2007Q4	2010Q4	change	contribution(%)	change	contribution(%)
conditions						
angipecev	0.007	0.010	0.003	1.095	0.004	1.307
emphysemev	0.007	0.009	0.002	0.730	0.002	0.654
cancerev	0.036	0.046	0.010	3.650	0.009	2.941
diabeticev	0.051	0.061	0.010	3.650	0.007	2.288
cheartdiev	0.015	0.014	-0.001	-0.365	-0.002	-0.654
heartattev	0.012	0.012	0.000	0.000	0.000	0.000
strokev	0.009	0.015	0.006	2.190	0.006	1.961
asthmaev	0.106	0.129	0.023	8.394	0.018	5.882
hypertenev	0.184	0.217	0.033	12.044	0.022	7.190
ulcerev	0.057	0.051	-0.006	-2.190	0.001	0.327
total	0.484	0.565	0.081	29.562	0.068	22.222
limitations						
lawalk	0.027	0.025	-0.001	-0.365	-0.002	-0.654
flstoop	0.161	0.179	0.018	6.569	0.025	8.170
flgrasp	0.050	0.055	0.005	1.825	0.009	2.941
flreach	0.061	0.058	-0.003	-1.095	-0.002	-0.654
flclimb	0.076	0.075	-0.001	-0.365	-0.003	-0.980
flwalk3bl	0.097	0.103	0.005	1.825	0.005	1.634
flcarry	0.063	0.077	0.014	5.109	0.018	5.882
flgoout	0.056	0.065	0.009	3.285	0.010	3.268
flpushlar	0.084	0.098	0.014	5.109	0.027	8.824
flrelax	0.031	0.036	0.005	1.825	0.007	2.288
flsocial	0.048	0.055	0.007	2.555	0.009	2.941
ladl	0.009	0.012	0.004	1.460	0.003	0.980
laiadl	0.024	0.025	0.002	0.730	0.000	0.000
laother	0.009	0.007	-0.002	-0.730	0.000	0.000
lamemry	0.022	0.024	0.002	0.730	0.002	0.654
total	0.817	0.895	0.077	28.102	0.106	34.641
depression						
total	0.469	0.573	0.105	38.321	0.142	46.405
overweight						
total	0.269	0.281	0.011	4.015	-0.010	-3.268

B Cyclical health robustness

This appendix provides a systematic set of robustness checks demonstrating that the countercyclicality of bad health is not an artifact of measurement, sample selection, or empirical specification. We proceed by showing, first, that the core pattern is visible in levels and does not depend on detrending choices (Section B.1). We then establish that it holds for objective health indicators—not just self-reports—and that it is not driven by cyclical access to employer-provided health insurance, using both insured and uninsured samples (Sections B.2 and B.3). Next, we show the pattern is consistent across major U.S. surveys and is not explained by compositional shifts, either across cohorts or immutable individual characteristics (Sections B.4 and B.5). We further demonstrate robustness to finer health categorizations and to alternative measures of real activity beyond employment (Sections B.6 and B.7). The back-of-the-envelope quantification of health effects in the GFC is robust to alternative timing windows (Section B.8). Finally, we document limited non-robustness in timing across health dimensions, highlighting heterogeneity in the dynamics of underlying conditions rather than a failure of the overall countercyclical relationship (Section B.9).

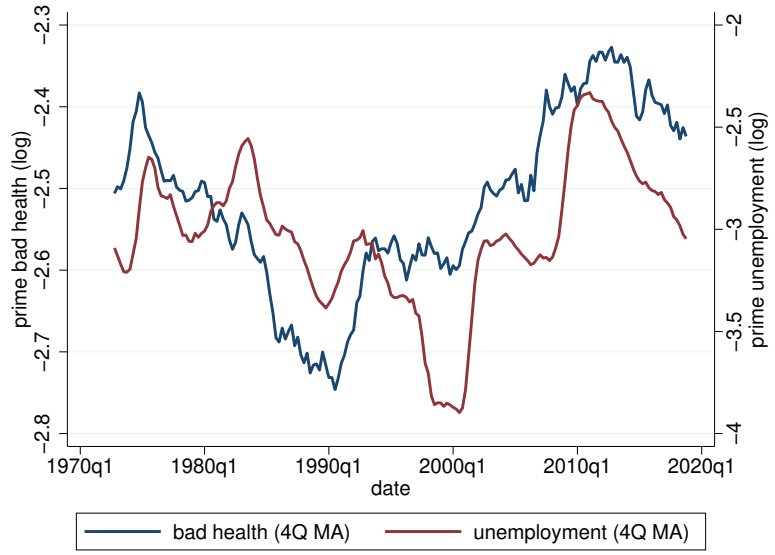
B.1 Robustness to not detrending

We first show bad health's countercyclicality is not driven by detrending. Figure B.1 shows unemployment and the fraction of population in bad health in log-levels.

B.2 Robustness to objective health metrics

Health's procyclicality is not driven by the subjectivity of self-reports. Figure B.2 shows z-scores for employment, self-reported bad health, and several objective metrics of bad health—doctor visits, inability to work, and a frailty index (Hosseini et al., 2026) that includes 27 health conditions—for the prime-aged population with health insurance coverage. (Later, Figure B.3 will show the same patterns hold for the uninsured.) Limited by objective health data coverage, the plot runs from 1998 to 2019. All the objective measures increase in the two recessions and tend to decline with employment gains. For example, bad health rises by about 1.5 standard deviations from 2006 to 2010, while frailty and doctor visits rise by 1.5 to

Figure B.1: Bad health and unemployment move together: U.S. prime-aged population (levels), 1972 to 2019



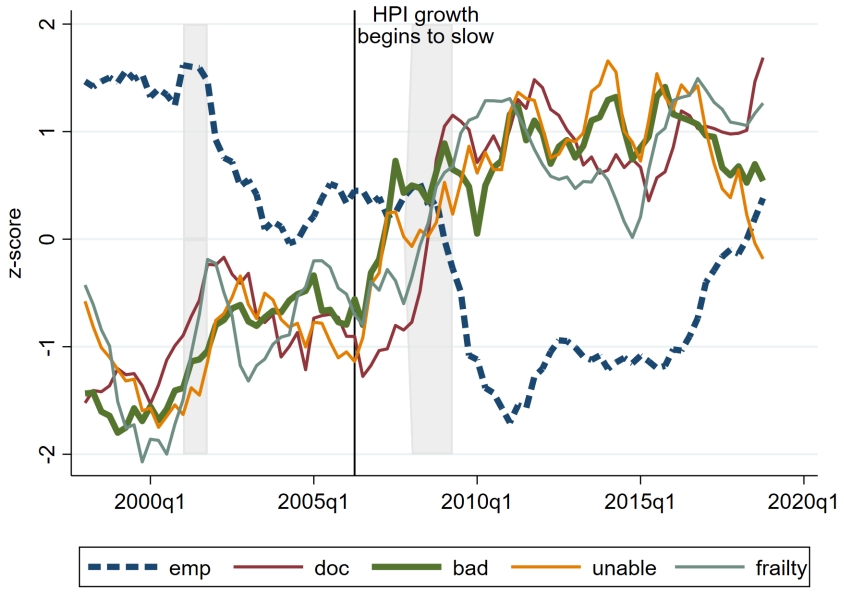
2.5 standard deviations from 2008 to 2010. That frailty—perhaps the most objective measure being determined by the prevalence of specific medical conditions—rises strongly indicates that health's cyclicality is not merely subjective.

B.3 Robustness to employer-provided health insurance

Health's procyclicality is also unlikely to be explained by cyclical changes in access to employer-provided health insurance. One piece of evidence for this is from Figure B.2 where doctor visits increase in the GFC, which would be surprising if bad health were caused by loss of access to health care. Another is that the series in Figure B.2 were all conditioned on those with health insurance. These are strengthened by Figure B.3, which shows patterns similar to those in Figure B.2 hold both for the insured and the uninsured.

A limitation of the preceding results is that, since health insurance status is unknown prior to the late 1990s, we cannot directly include earlier recessions. However, although insurance status is not available going back to the 1970s, restricting the sample to employed individuals should eliminate most cyclical variation in insurance status, provided access to health insurance *for employed workers* is acyclical. Conditioning on employed individuals, Figure 2.1 is virtually unchanged as seen

Figure B.2: Bad health increases in recessions according to different indicators in the NHIS: Comparison of standardized employment and different indicators of adverse health for prime age individuals with health insurance.



Note: We stop the NHIS time series in 2019 because of a methodological change in 2020. The CPS data come from the annual March supplement. The NHIS is a quarterly survey, and we show the 4-quarter moving average.

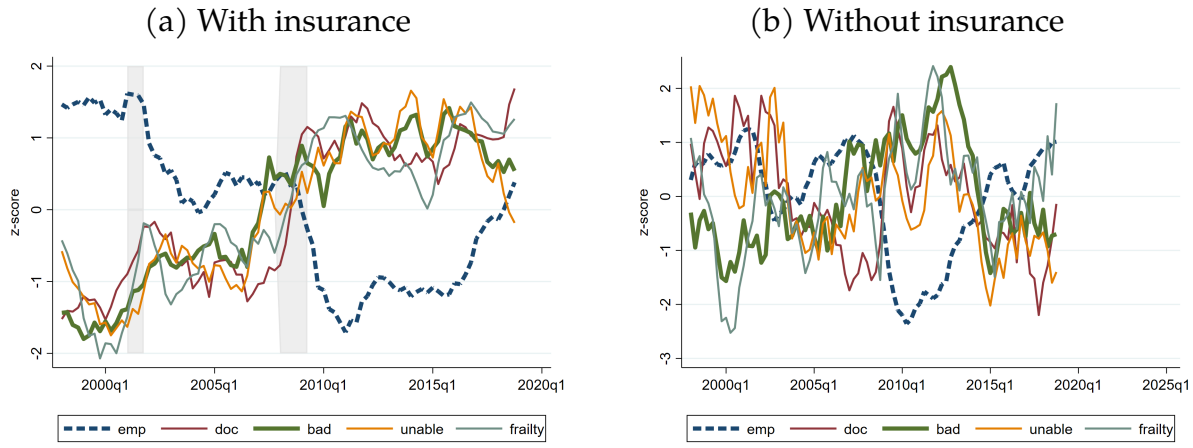
in Figure B.4. We conclude that the correlations in the NHIS have counterparts in objective measures and are unlikely to be explained by loss of access to health insurance.[15]

B.4 Robustness across surveys

The cyclical pattern in health is also robust across surveys. To show this, we put data on bad health and employment from the late 1990s to 2019 side by side in Figure B.5. Panel (a) shows the familiar NHIS data, now for all prime-aged survey respondents. Employment is shown as the dashed line and bad health as the solid line. Panel (b) shows the corresponding plot based on CPS data. Both the levels and the cyclical variation are very similar: Both surveys show a peak-to-trough drop in the employment-to-population ratio of around 5 pp. and an increase in

[15]Health insurance may, however, contribute to the aggregate patterns.

Figure B.3: Robustness of the countercyclicality of bad health in the NHIS: Comparison of different indicators of adverse health for prime age individuals in the 2000s by insurance status.



Note: We stop the NHIS time series in 2019 because of a methodological change in 2020. The CPS data come from the annual March supplement. The NHIS is a quarterly survey, and we show the 4-quarter moving average.

bad health of around 1 pp.[16] Last, the pattern is broadly similar in the NLSY97 (Figure B.5(c)), which similarly to the NHIS and CPS, has employment dropping by about 6 pp. during the GFC and bad health rising by about 1.5 pp. since 2006.

B.5 Robustness to composition effects

Health's cyclicality is also not driven by composition effects. This is evidenced by the NLSY97 series of Figure B.5 resembling the others. There, a single cohort is tracked, so composition across cohorts is not driving the result. And our inclusion of person fixed effects indicates shifts in the composition of immutable individual characteristics are not driving the cyclicality.

B.6 Robustness to finer health status bins

Bad health's countercyclicality is also robust to finer breakdowns of health. Figure B.6 shows the evolution of the share of the prime-aged across more detailed health measures from 1997 to 2023. It also shows the employment share (dashed

[16] Bad health in the CPS rises with a delay compared with the NHIS. Figure B.6 shows that the top two health categories also started to decline in 2006 in the CPS.

Figure B.4: 50 years of bad health and unemployment: U.S. prime-aged population with employment health

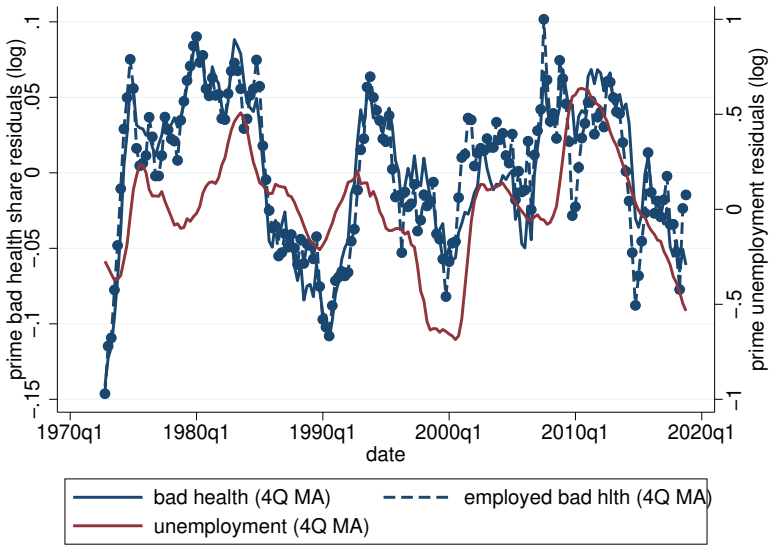
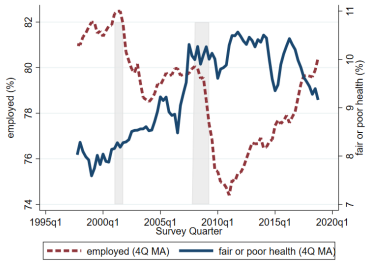


Figure B.5: Bad health increases in recessions in three separate surveys: Bad health and employment of prime-aged individuals in the NHIS, CPS, and NLSY97

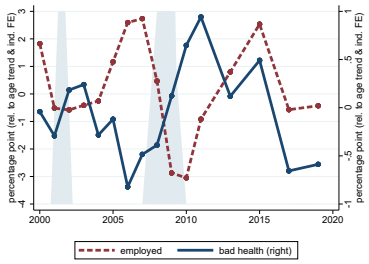
(a) NHIS



(b) CPS

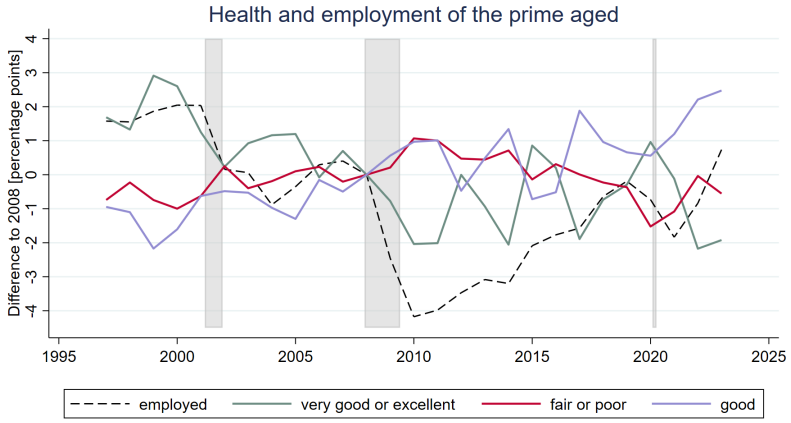


(c) NLSY97



line). All shares are expressed relative to their 2008 levels. The figure shows that the roughly 1.0 pp. increase in bad health from 2008 to 2010 is only part of the picture: The middle health category rises by a similar amount, while the share in the top two health categories falls by about 2.0 pp. The increase in bad health largely persists until about 2015.[17]

Figure B.6: While employment fell during the GFC, health worsened not only for those in bad health: Share of prime-aged individuals across health categories and in employment (dashed line): CPS, 1997–2023. The light-green line denotes "fair or poor" health; yellow denotes "good" health; red denotes the combined "very good or excellent" categories. All series are expressed as deviations (in percentage points) from their 2008 levels.



B.7 Robustness to alternative real activity metrics

Thus far we have focused on employment as a clean measure of real activity, as this is cleanly measured in our micro data and comparable across surveys. However, bad health's countercyclicality is also robust to other measures of real activity. Figure B.5(b) indicates that bad health among the prime-aged returned to its pre-recession level only in 2015, and employment remained depressed for longer still, possibly due to a downward trend. Motivated by this apparent persistence, we now turn to state level data to quantify the connection between the initial increase in bad health and subsequent labor market outcomes.

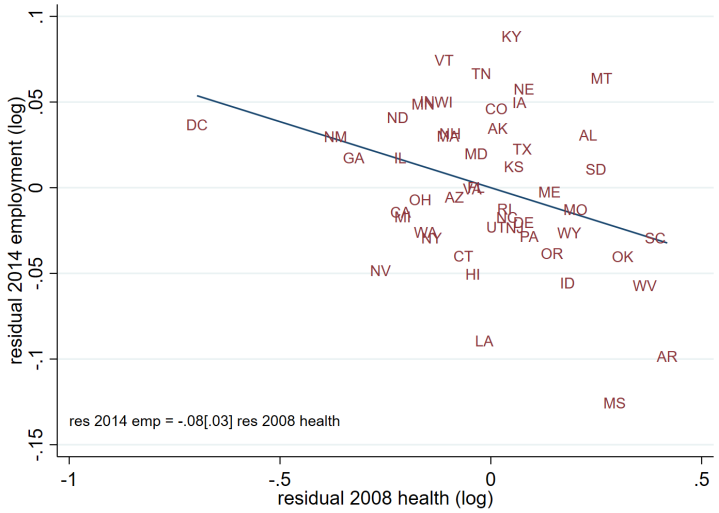
[17] The other two health categories in the chart are too volatile to tell.

The model we use to summarize the data is the following:

$$\ln y_{s,2008+h} = \alpha_h + \beta_y \ln h^b_{s,2008} + \gamma \ln y_{s,2008} + e_{s,t}. \tag{B.1}$$

$y_{s,t}$ is a labor market outcome of interest in state s at time t. $h^b_{s,2008}$ is initial bad health in the state – all conditional on individuals being of prime age.

Figure B.7: States with higher bad health in 2008 tend to have worse employment in 2014, even after controlling for 2008 employment levels. CPS data.



Bad health has predictive power for current and future labor market outcomes. Even after controlling for the severity of the recession via the 2008 state employment-to-population ratio, bad health in 2008 significantly predicts the 2014 employment-to-population ratio in a state, as Figure B.7 illustrates. The figure plots 2014 employment against bad health in 2008, both as residuals from projecting out 2008 employment. The β_e implied by this regression means that states with 10% higher bad health in 2008 had 0.8% lower employment (to population) in 2014.

The results are similar for other labor market outcomes: For example, Table B.1 (a) provides analogous regressions for total hours worked for multiple years. Health in 2008 explains about 5% of the residual variation in 2014 total hours worked, after accounting for 2008 levels of hours worked. States with a one standard deviation higher bad health in 2008 on average had total hours that were 0.24 standard devi-

ation ($= 0.28 \times (-0.05)/0.07$) lower in 2014. Similar results hold for 2012, 2013, and 2015 and also for extensive margin measures and average hourly wages.

Extensive margin measures are more sensitive to bad health than total hours. For employment and participation, Table B.1 (b) and (c), the explanatory power is slightly higher. The partial R-squared is 0.07 for both variables and the effects of a one standard deviation shock in bad health cause movements of 0.32 and 0.39 standard deviations, respectively.

Last, the effects on hourly wages in Table B.1 (d) are similar to the effects on total hours – both in terms of the magnitude and the sign. This is noteworthy because a simple drop in labor supply due to worse health might intuitively be associated with an increase in wages. Instead, these results suggest that either overall labor productivity drops or that there is a selection effect that lowers both wages and labor supply.

Table B.1: States with higher bad health in 2008 had worse future labor market. CPS data.

	Total hours				E/Pop	LFPR	wage/h
	2012	2013	2014	2015	2014	2014	2014
LHS in 2008	0.74***	0.71***	0.64***	0.67***	0.53***	0.46**	0.40***
	(0.12)	(0.13)	(0.13)	(0.09)	(0.17)	(0.19)	(0.14)
Bad health (log) in 2008	-0.06**	-0.05*	-0.06**	-0.05*	-0.08***	-0.07**	-0.09*
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)
Mean 2008 LHS	3.49	3.49	3.49	3.49	-0.22	-0.18	2.76
SD 2008 LHS	0.06	0.06	0.06	0.06	0.05	0.04	0.11
Mean 2008 health	-2.30	-2.30	-2.30	-2.30	-2.30	-2.30	-2.30
SD 2008 health	0.28	0.28	0.28	0.28	0.28	0.28	0.28
Mean LHS	3.42	3.43	3.43	3.46	-0.26	-0.21	2.83
SD LHS	0.07	0.07	0.07	0.06	0.06	0.05	0.11
Overall R-squared	0.68	0.62	0.56	0.61	0.53	0.53	0.26
Partial R-squared health	0.04	0.03	0.05	0.03	0.07	0.07	0.05
No. obs.	51	51	51	51	51	51	51

B.8 Robustness of back-of-the-envelope calculation

Table B.2 shows that the back-of-the-envelope employment effects of health deterioration during the Global Financial Crisis are robust to the choice of time window. While the magnitude of the implied employment decline varies with the horizon considered, the estimates remain of similar order across shorter and longer windows. Longer windows tend to yield smaller employment effects as health partially recovers over time, but the qualitative conclusion is unchanged: Shifts in the health

distribution plausibly account for a nontrivial share of the observed employment decline. This robustness across windows supports the interpretation that health deterioration contributed meaningfully to labor market outcomes during and after the crisis.

Table B.2: Plausible employment effects of 2008–2010 and longer term health changes: back-of-the-envelope calculations

			Change by period [pp]			
Health Status	Share	E/pop [%]	2008 to '10	2007 to '10	2007 to '13	2006-08 to '12-14
Excellent	29.6	86.3	-2.2	-1.3	-0.5	-1.2
Very good	35.7	85.7	0.1	-1.5	-1.1	-0.0
Good	24.8	77.7	1.0	1.5	1.0	0.7
Fair	7.1	51.7	0.9	1.2	0.3	0.3
Poor	2.6	21.8	0.1	0.1	0.3	0.2
Implied emp change [pp]			-0.49	-0.59	-0.40	-0.31

B.9 Non-robustness: Timing differences

While the overall cyclical pattern is robust across all these dimensions, there are some timing differences in Figures B.2 and B.5 that reflect the multi-faceted nature of health. Table A.2 in the appendix establishes that the objective health conditions underlying the frailty index respond to the GFC with differing timings. Depression and medical conditions, such as hypertension or asthma, rise first, while functional limitations rise with a delay of approximately one year.[18] Health's deterioration in Figure B.2's begins significantly before employment declines but after house price growth slows. This likely reflects Currie and Tekin (2015) that foreclosures are highly correlated with increases in urgent hospital and emergency room visits. Additionally, self-reported bad health tends to lead the other series. This likely reflects Halliday et al.'s (2021) finding that self-reported health outperforms objective metrics in predicting mortality: People are aware of health deteriorations before they materialize.

[18] Depression, which might be viewed as less objective, contributes substantially to the increased frailty in the three year period from 2007q4 to 2010q4, explaining about 40% of the increase. In contrast, limitations and medical conditions together contribute almost 60%. See Table A.2 for more detail.

C Additional Impulse Response Functions

C.1 Alternative specifications for excess bond premium shock responses

Figure C.1 displays impulse responses to the excess bond premium shock under different specifications. The responses are qualitatively and quantitatively robust to the different assumptions.

C.2 Heterogeneity of responses

Figure C.2 shows how the uninsured health and employment evolve relative to the insured health and employment, respectively, following an excess bond premium shock. The uninsured fare worse, which could be driven by a composition effect of uninsured workers tending to have lower income and being marginally attached to the labor force.

C.3 Additional monetary policy shock responses

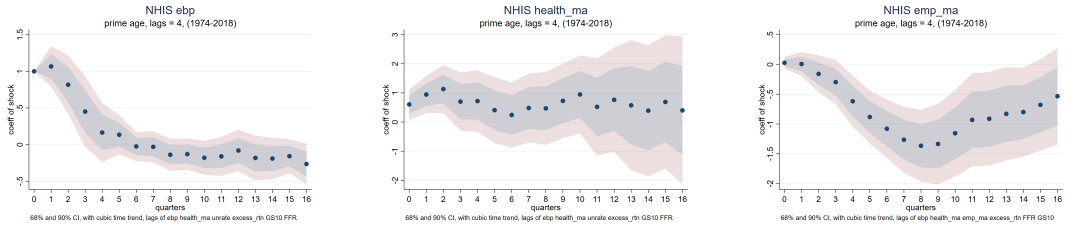
C.3.1 Response of health for the employed

Figure C.3 reports impulse responses of health *for employed workers* alongside labor market outcomes broadly after contractionary monetary policy shocks for alternative identification schemes and specifications. Panel (a) uses Romer and Romer monetary policy shocks with a cubic time trend and a longer lag structure, while Panel (b) uses Bauer and Swanson shocks with no trend and fewer lags. Across specifications, contractionary monetary policy raises the fraction of employed individuals in bad health and weakens labor market outcomes, with health responses emerging earlier and more sharply than employment. The similarity of the qualitative patterns across panels indicates that the deterioration in health following monetary tightening is not sensitive to the choice of shock measure or time-series specification.

Figure C.1: Excess bond premium shock responses under alternative specifications

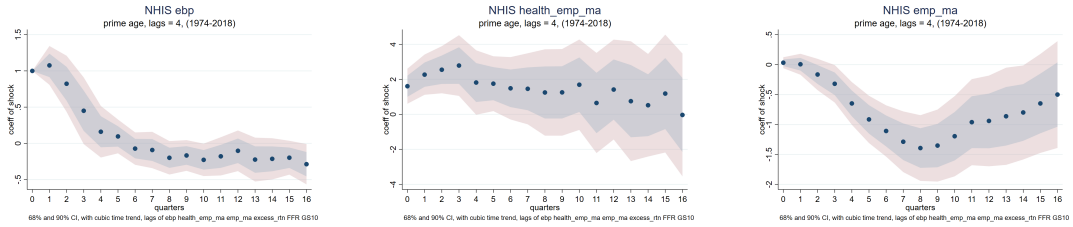
(a) Baseline: Cubic trend, 4 lags, extra controls

Excess bond premium | Fraction in bad health | Employment



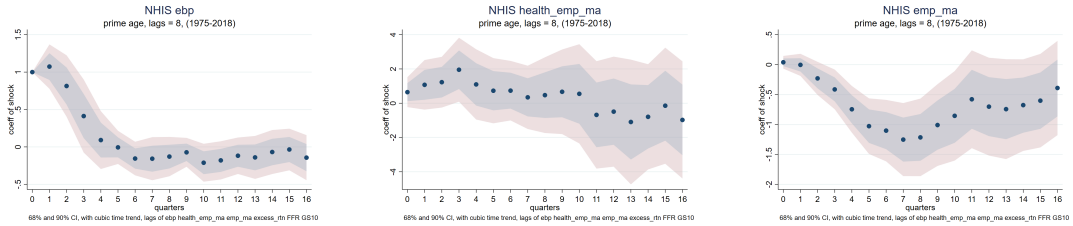
(b) Health of the employed: Cubic trend, 4 lags, extra controls

Excess bond premium | Fraction in bad health | Employment



(c) Health of the employed: Cubic trend, 8 lags, extra controls

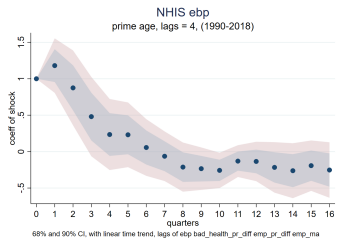
Excess bond premium | Fraction in bad health | Employment



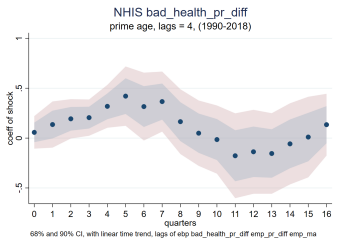
Note: Impulse responses to a 1 pp. increase in the excess bond premium, estimated using local projections. Outcomes are the excess bond premium, the fraction of prime-aged individuals in bad health, and employment. Panel (a) uses the full sample of prime-aged individuals with a cubic time trend and four lags. Panel (b) restricts the health measure to employed individuals with four lags. Panel (c) uses the employed-only sample with eight lags. All specifications include extra controls and 4-quarter moving averages. Shaded bands denote 68 and 90 percent confidence intervals.

Figure C.2: Excess bond premium IRFs: Uninsured vs insured

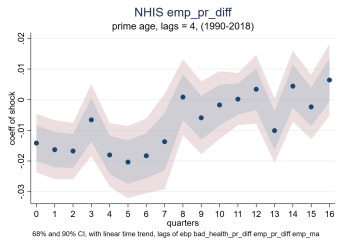
Excess bond premium



Fraction in bad health: Uninsured-insured



Emp: Uninsured-insured



Note: Impulse responses to a 1 pp. increase in the excess bond premium, estimated using local projections. Detrending is linear and 4 lags are used on the shorter sample. Outcomes are the excess bond premium, the bad health rate of the uninsured minus the bad health rate of the insured, and employment rate of the uninsured minus the employment rate of the insured. Shaded bands denote 68 and 90 percent confidence intervals.

C.3.2 Response of house prices and aggregate unemployment

Figure C.4 provides additional response variables to the Romer and Romer shocks. Unlike prime-aged employment (our main measure of real activity) that displays no response for many years, aggregate unemployment begins rising after two years. Another possibly important mechanism is through house prices, which begin declining after just a few quarters.

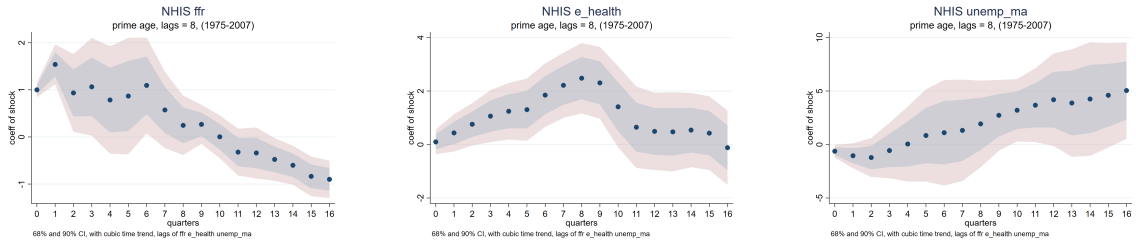
C.4 Response of mortality

The figures below examine the response of mortality to contractionary macroeconomic shocks under a range of specifications. Figures C.5 and C.6 report impulse responses to contractionary monetary policy shocks identified using Bauer and Swanson (2023), while Figures C.7 and C.8 report responses to excess bond premium shocks. Across specifications, mortality generally tends to increase following contractionary shocks, consistent with a deterioration in population health during downturns. However, the estimated mortality responses are often statistically insignificant and display sensitivity to the choice of lag length, trend specification, and the health sample used. Taken together, the figures suggest that while contractionary shocks are associated with higher mortality on average, this relationship is weaker and less robust than the corresponding responses of self-reported health.

Figure C.3: Employed response to monetary policy shocks

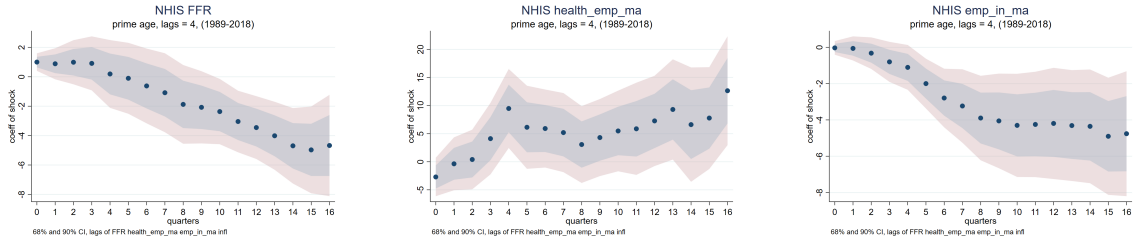
(a) Health of the employed: Cubic trend, 8 lags, Romer & Romer

Federal Funds Rate | Fraction in bad health | Unemployment



(b) Health of the employed: No trend, 4 lags, Bauer & Swanson

Federal Funds Rate | Fraction in bad health | Employment



Note: Impulse responses to a contractionary monetary policy shock estimated using local projections. Outcomes are the federal funds rate, the fraction of employed, prime-aged individuals in bad health, and labor market outcomes. Panel (a) uses Romer and Romer monetary policy shocks with a cubic time trend and eight lags. Panel (b) uses Bauer and Swanson shocks with no time trend and four lags. All specifications use 4-quarter moving averages. Shaded bands denote 68 and 90 percent confidence intervals.

Figure C.4: Impulse-response functions for the U.S. prime-aged population. Local projections estimated with a cubic trend and 4 lags

Agg unemployment | House prices

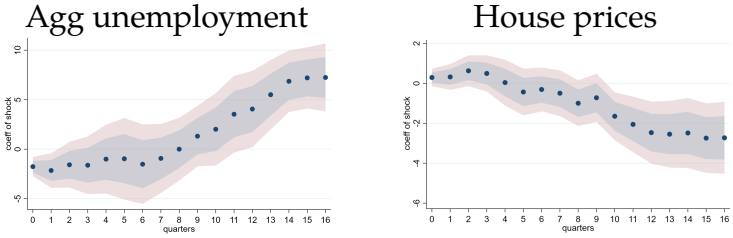
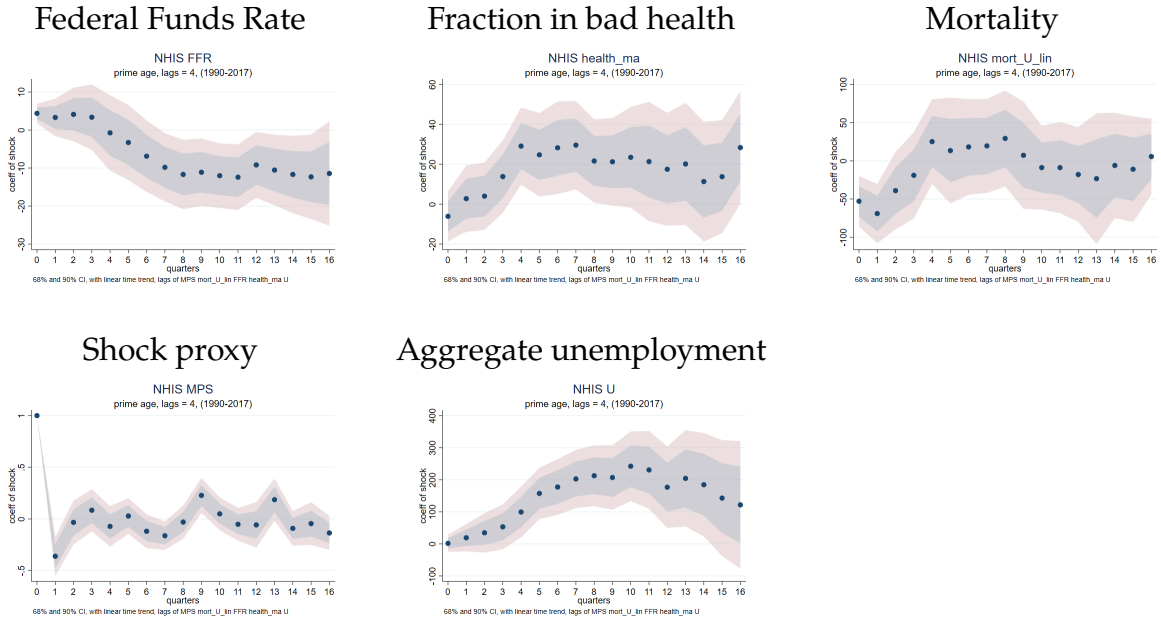


Figure C.5: Mortality responses to a contractionary monetary policy shock (Bauer & Swanson).

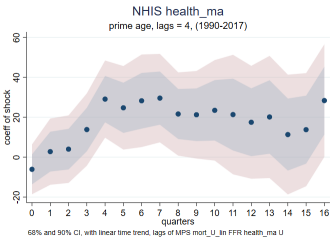


Note. Local-projection impulse responses to a contractionary monetary policy shock identified using Bauer and Swanson (2023). Specification uses a linear time trend and 4 lags. Bad health is measured for prime-aged insured individuals; mortality is measured for all prime-aged individuals. Shaded bands denote pointwise 68 and 90 percent confidence intervals.

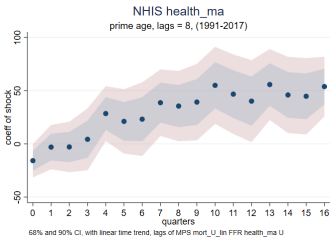
Figure C.6: Robustness of mortality responses to monetary policy shocks

(a) Fraction in bad health

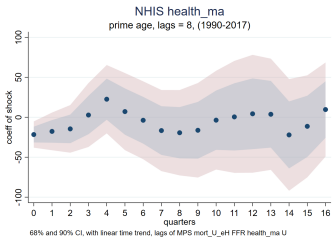
Prime-aged, insured
4 lags



All prime-aged
8 lags

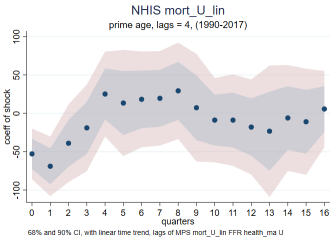


Employed only
8 lags

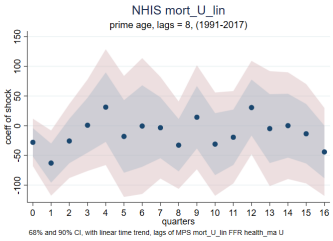


(b) Mortality

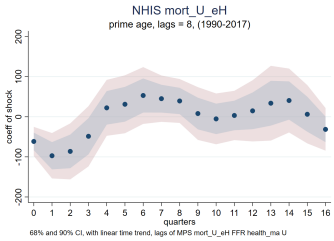
Prime-aged, insured
4 lags



All prime-aged
8 lags

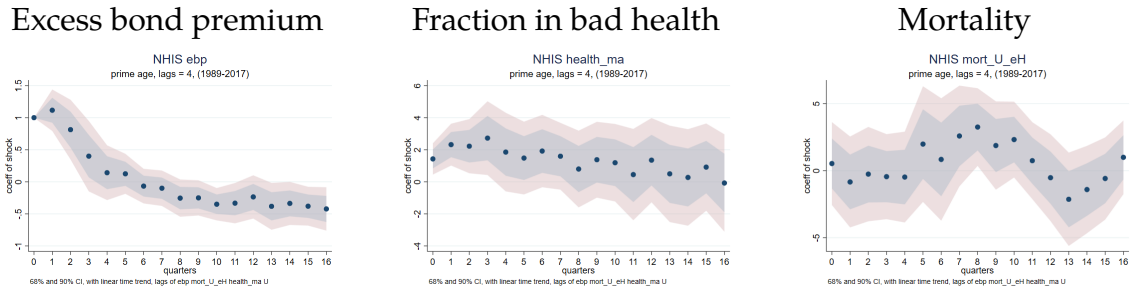


Employed only
8 lags



Note. Local-projection impulse responses to a contractionary monetary policy shock identified using Bauer and Swanson (2023). Columns vary the lag length and the health sample (all prime-aged, employed, or insured-only for the health series). Unless otherwise indicated, specifications use a linear time trend. Shaded bands denote pointwise 68 and 90 percent confidence intervals.

Figure C.7: Mortality responses to a contractionary excess bond premium shock.

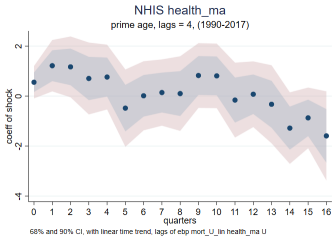


Note. Local-projection impulse responses to an excess bond premium shock. Specification uses a linear time trend and 4 lags, and controls for prime-age employment. Health is measured for employed individuals; mortality is measured for all prime-aged individuals. Shaded bands denote pointwise 68 and 90 percent confidence intervals.

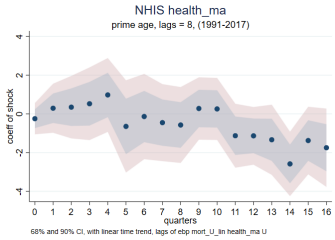
Figure C.8: Robustness of mortality responses to excess bond premium shocks

(a) Fraction in bad health

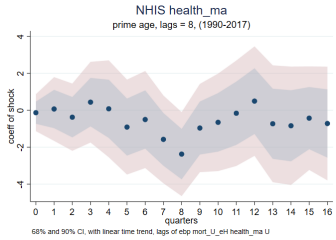
Prime-aged, insured
4 lags



All prime-aged
8 lags

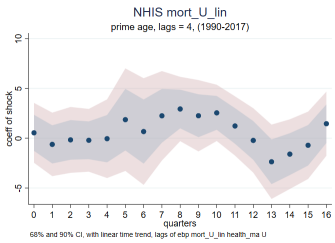


Employed only
8 lags

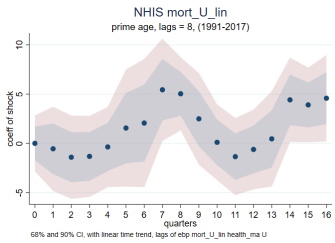


(b) Mortality

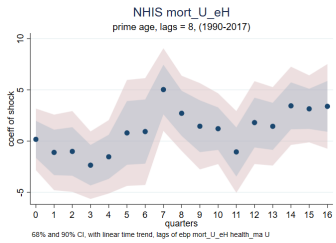
Prime-aged, insured
4 lags



All prime-aged
8 lags



Employed only
8 lags

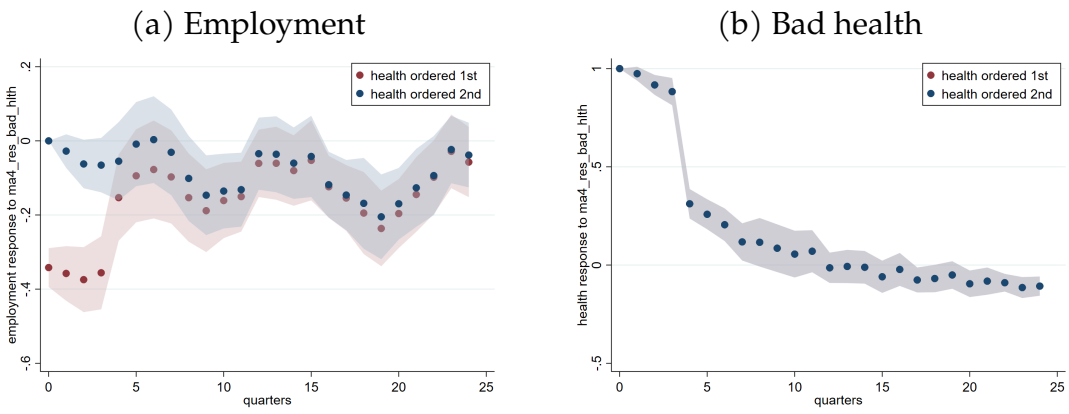


Note. Local-projection impulse responses to a contractionary excess bond premium shock. Columns vary the lag length and the health sample (all prime-aged, employed, or insured-only for the health series). Unless otherwise indicated, specifications use a linear time trend. Shaded bands denote pointwise 68 and 90 percent confidence intervals.

C.5 Health breakdown in the BRFSS

Figures C.9 and C.10 examine the labor-market effects of pure health shocks using state-level variation in the BRFSS. Figure C.9 shows that surprise increases in bad health continue to predict persistent declines in employment even when time fixed effects are included, though the estimated responses are attenuated relative to specifications without time fixed effects. This attenuation is consistent with time fixed effects absorbing common national health fluctuations, leaving only idiosyncratic within-state variation for identification. Figure C.10 further decomposes these health shocks into mental and physical components. Both types of shocks generate qualitatively similar and persistent employment declines, indicating that the adverse labor-market effects are not driven by a specific health dimension. Taken together, the figures demonstrate that the employment response to health shocks is robust to the inclusion of time fixed effects and to alternative definitions of health, albeit with reduced magnitude when identification relies solely on within-state variation over time.

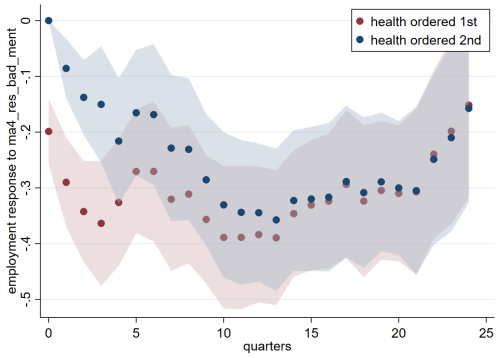
Figure C.9: Impulse responses to pure health shocks



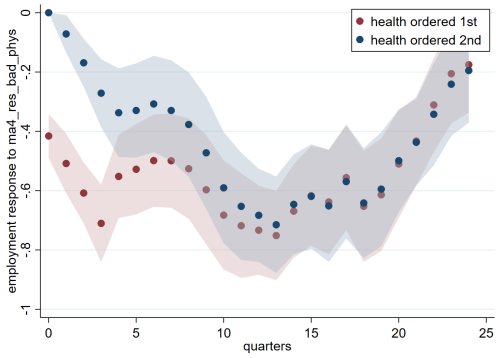
Note: State level impulse-responses of prime-aged employment and bad health; 4-quarter moving averages in percent deviations from state fixed effects; includes time fixed effects.

Figure C.10: State level dynamics of prime-aged employment and following bad mental health shocks (Panel (a)) and bad physical health shocks (Panel (b)): State fixed effects

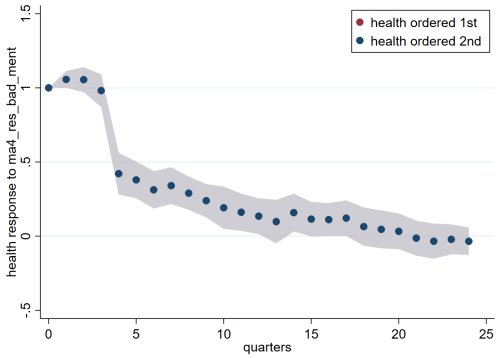
(a) Bad mental health surprise
Employment response



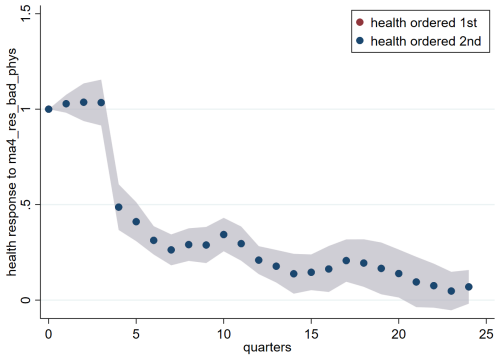
(b) Bad physical health surprise
Employment response



(c) Bad mental health surprise
Bad mental health response



(d) Bad physical health surprise
Bad physical health response



D Model appendix

D.1 Computation

We solve the model using value function iteration, applying the divide-and-conquer algorithm in Gordon and Qiu (2018), which exploits policy function monotonicity. We allow the a' choice to lie in between grid points, recovering the exact solution subject to linear interpolation as described in the earlier working paper (Gordon and Qiu, 2017). Market clearing prices are computed by iteratively updating guesses on K/N and T. We compute impulse-response function via deterministic transition paths.

D.2 Calibration

Table D.1 reports the regression results for two different sub-samples in columns (a) and (b), the normalized coefficients based on the full sample from column (a) in column (c) and a benchmark from Halliday et al. (2021) in column (d). Comparing columns (c) and (d) shows that our approach based on a wage regression yields results very similar to the mid-points of a mapping of our 5-category general health indicator to the $[0, 100]$ interval in Halliday et al. (2021).

Figure D.1 plots the detrended time series used for computing business cycle statistics.

D.3 Back of the envelope with the calibrated model

Table D.2 reports the back of the envelope calculation when using the calibrated model's shifted shares. There is a slightly smaller predicted employment effect.

D.4 Dynamic model mechanisms

Figure D.2 computes the impulse response of employment, wages, and capital to a pure health shock in the presence of various model mechanisms. The baseline B includes all mechanisms. The lines labeled PE correspond to a partial equilibrium transition where prices and transfers are held fixed. For PE, the plotted series for wages and capital are the implied values—those that correspond to K and N taking account of these decisions (the non-implied values are constant). The FP

Table D.1: Calibration, calibration target, and validation for health status

	CPS Prime-Aged		**CPS**	**Halliday et al. (2021)**
	(a) 1996–23	(b) 2007–10	Normalized h_i	Interval Midpoint
Poor / fair health	-0.409***	-0.443**	0	$0.155 = \frac{0.10+0.30}{2}$
	(-19.55)	(-4.91)		
Fair health	-0.247***	-0.243***	0.35	$0.50 = \frac{0.30+0.70}{2}$
	(-38.75)	(-18.53)		
Good health	-0.125***	-0.121***	0.65	$0.775 = \frac{0.70+0.85}{2}$
	(-37.93)	(-18.32)		
Very good health	-0.045***	-0.047***	0.84	$0.90 = \frac{0.85+0.95}{2}$
	(-16.99)	(-5.93)		
Excellent (omitted)	0	0	1	$0.975 = \frac{0.95+1}{2}$
R^2	0.44	0.44		
Within R^2	0.03	0.03		
Observations	905,889	140,849		
Years	28	4		

Normalized h_i values are adjusted using the implied regression 1996-2023 sample's estimated wage penalty scaled to $[0, 1]$: The coefficient β_i is mapped to $\frac{e^{\beta_i}-e^{\beta_1}}{e^{\beta_5}-e^{\beta_1}}$. Excellent health is the omitted category in the CPS regressions.
t-statistics in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure D.1: Detrended time series used for computing business cycle statistics

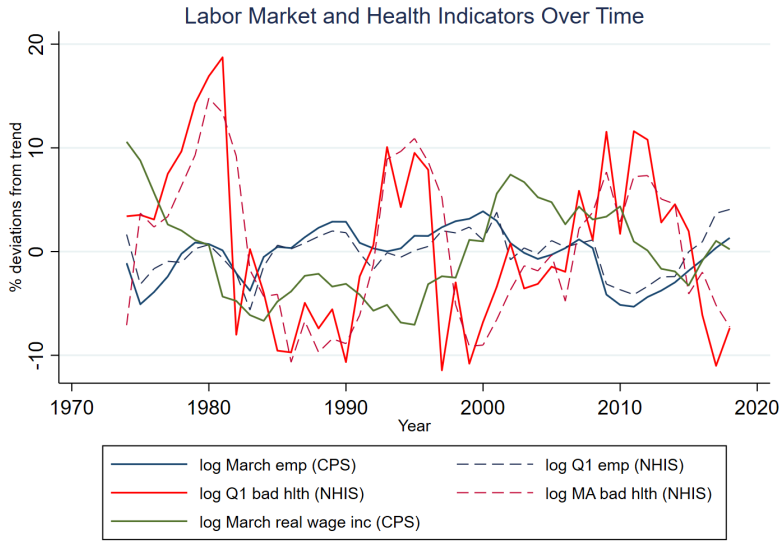
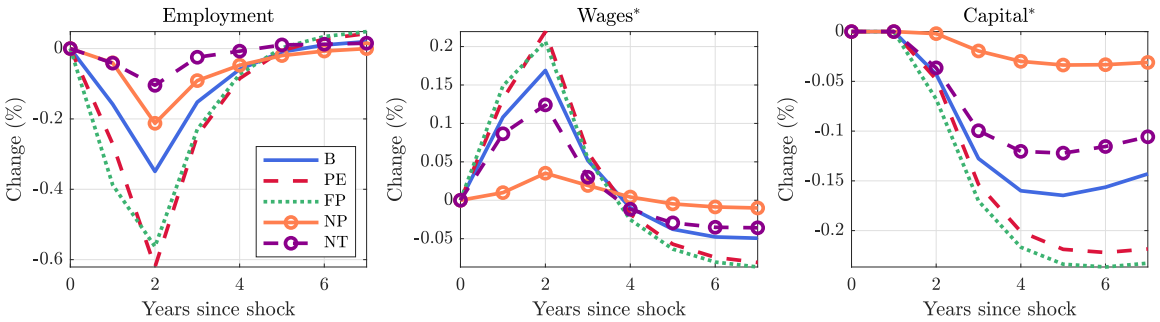


Table D.2: Back-of-the-envelope benchmark with calibrated GFC counterfactual

Health Status	Share	E/pop [%]	ΔShare [%]	Δcalibrated [%]
Excellent	29.6	86.3	-2.2	-2.2
Very good	35.7	85.7	0.1	0.4
Good	24.8	77.7	1.0	1.0
Fair	7.1	51.7	0.9	0.9
Poor	2.6	21.8	0.1	-0.1
Implied emp change [pp]			-0.49	-0.34
Implied bad hlth change [pp]			1.07	0.81
Scaled emp change [pp]			-0.49	-0.44

series is similar, but fixed policy functions at the initial steady state. This makes FP essentially identical to the back-of-the-envelope calculation from Section 2.4.1. The other lines turn off health's direct effects: NP disables health productivity effects while NT eliminates time costs (both investment and malaise). The no-health-expenditures case, as well as the zero-borrowing-limit specification, are omitted since they are very close to the baseline.

Figure D.2: Health shock response with different model mechanisms. The figure shows the baseline B and the change in levels (percentage points) from the benchmark to the indicated model: PE = partial equilibrium; FP = partial equilibrium with policy functions fixed at the initial steady state; NP = no health productivity effects; NT = no time costs of health (both investment and malaise). The no health expenditures and zero borrowing limit mechanisms have been omitted for readability as they are very close to the baseline.



Without general equilibrium effects, employment's response is nearly twice as large. Lower health in partial equilibrium causes a labor supply shortage that else equal drives up wages. This is attenuated by a reduction in savings, as individuals draw down on their wealth to smooth consumption against this persistent but temporary shock. This shows both the advantages and limitations of the back of the

envelope calculation: The response is roughly correct, but is here badly overstated. The fixed policy experiment has a larger employment drawdown on impact with less two years after the shock: Individuals respond to the shock by working more precautionarily until the health risks subside. This is also reflected in precautionary savings, with FP's implied capital stock lower than PE's.

The contribution of health costs to health shock effects is broadly in line with the steady state findings (see Table 4.1 and its discussion). Productivity costs play the largest role, time costs a large secondary role, and health expenditures (not shown) contribute only minimally. The paths are robust to a zero borrowing limit (not shown).

Table D.3 performs the same decomposition as in Table 5.1 but with different model mechanisms in place. Productivity and time costs matter most, as in the other decompositions.

Table D.3: Mechanisms and directional contributions to volatility (percent change in variance relative to the calibrated model)

	Employment variance			Bad-health variance		
Mechanism	$\Delta^{M\to H}$	$\Delta^{H\to M}$	Δ^{Total}	$\Delta^{M\to H}$	$\Delta^{H\to M}$	Δ^{Total}
Baseline (all)	9.8%	0.4%	10.2%	79.7%	20.3%	100%
No productivity cost	2.8%	0%	2.8%	97.1%	2.9%	100%
No time cost	4.0%	0%	4.0%	97.3%	2.7%	100%
No expenditure cost	9.0%	0.2%	9.2%	84.3%	15.7%	100%
Zero borrowing limit	8.1%	0.3%	8.3%	82.9%	17.1%	100%
Partial equilibrium	14.3%	3.1%	17.4%	45.8%	54.2%	100%

Notes: Baseline (Both) is $\sigma_h > 0, \rho_{h,o} \neq 0$. $\Delta^{M\to H}$ is the percent change in SD when setting $\rho_{h,o} = 0$ (holding $\sigma_h > 0$), $\Delta^{H\to M}$ is the additional percent change when setting $\sigma_h = 0$, and Δ^{Total} is the percent change in SD under $\sigma_h = 0$. By construction, $\Delta^{M\to H} + \Delta^{H\to M} = \Delta^{Total}$.

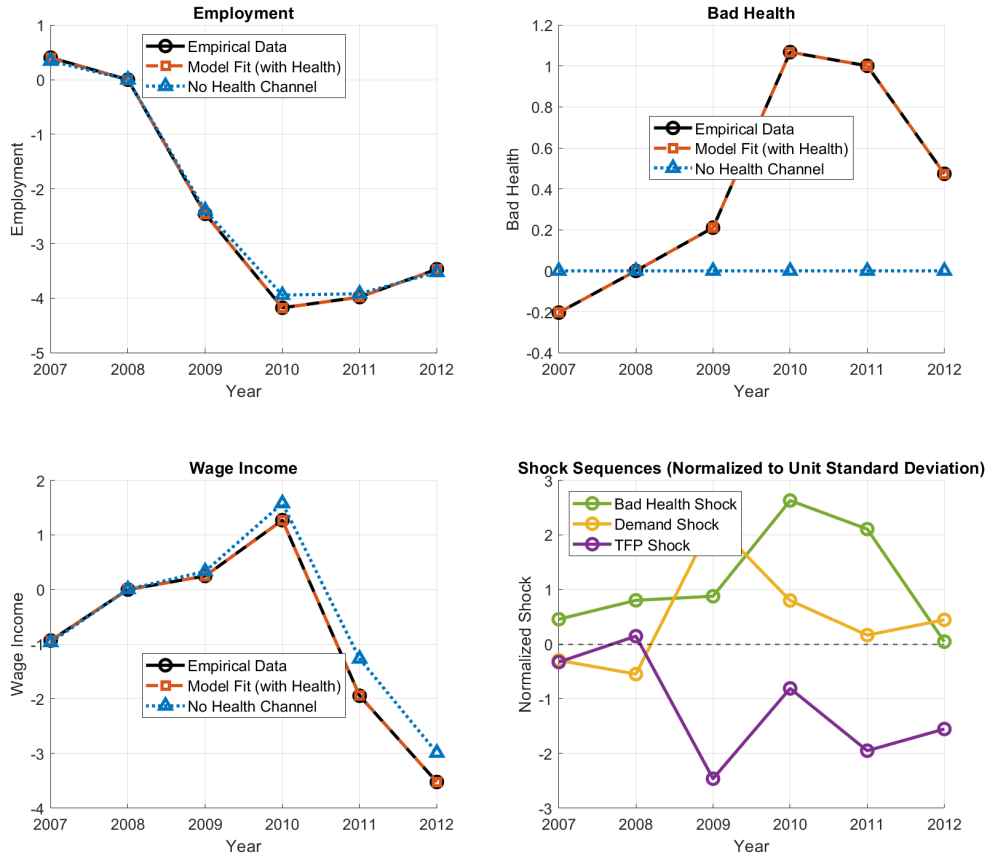
D.5 GFC analysis

Figure D.3 shows a GFC counterfactual constructed from a sequence of demand, productivity, and health shocks. By construction, the model with health is targeted to match the empirical paths of employment, bad health, and wage income: First, the health shock is chosen to match the health dynamics perfectly. Then the demand and TFP shock sequence is chosen to match the dynamics of employment

and wage income. Because we match the health dynamics perfectly and we fix the persistence of the demand and TFP shock beforehand, we are unable to match the data perfectly and we minimize the Euclidean norm between the model-implied employment and wage income and the data.

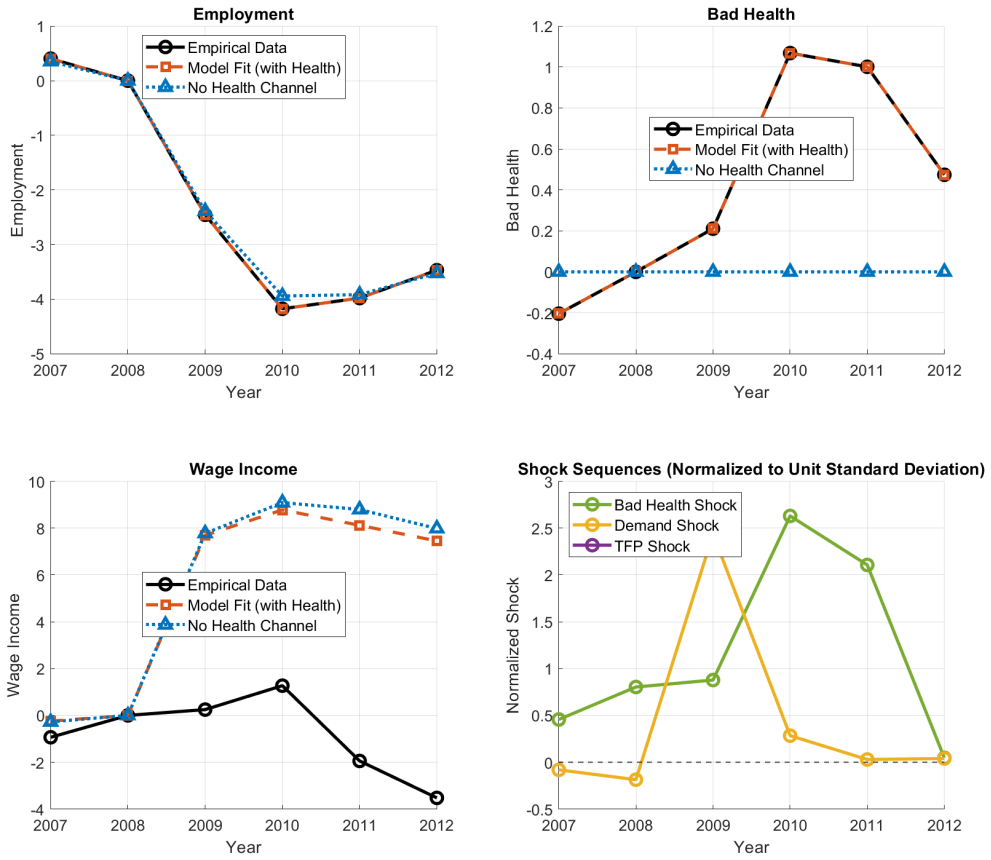
The key comparison is with the counterfactual that shuts down the health channel. Absent the health channel, the same shock sequence for demand and supply generates a substantially smaller employment decline, implying that a nontrivial portion of the observed employment contraction can be attributed to health. In this scenario, health shocks explain about 6% of the peak employment drop during the GFC. This is smaller than the 12% effect in the shift-share analysis – because general equilibrium factors dampen the responses.

Figure D.3: GFC scenario



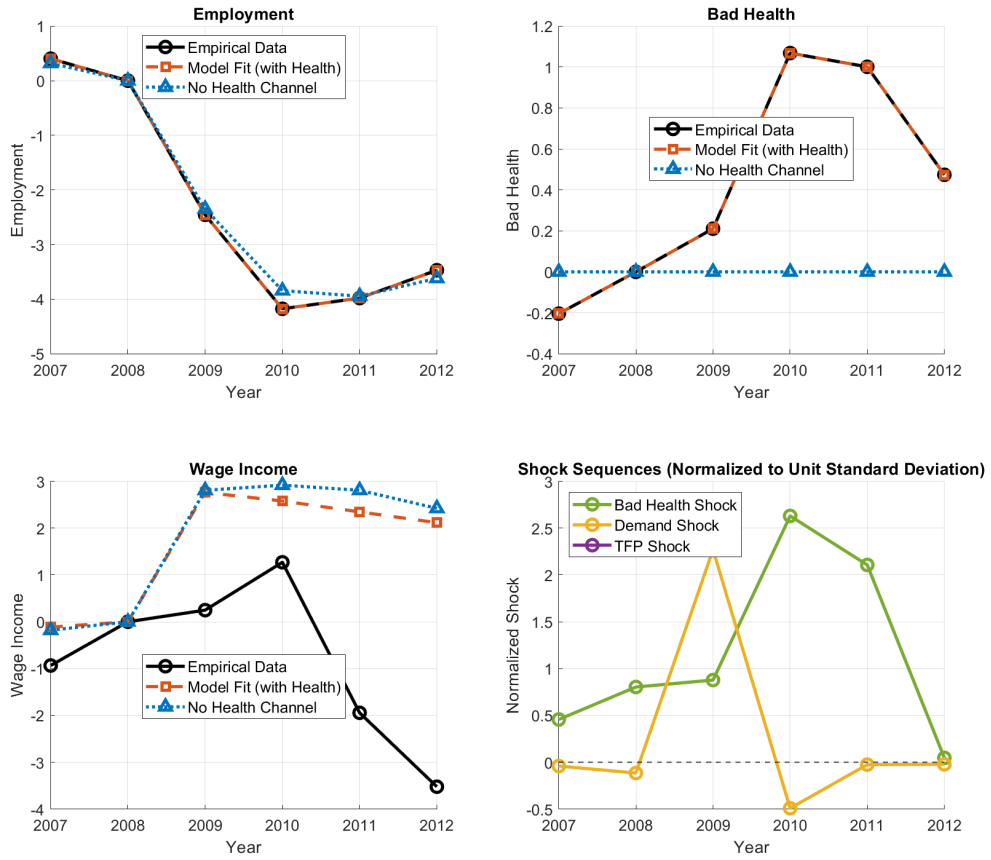
Note: The figure shows a GFC counterfactual generated by a sequence of demand, productivity, and health shocks. The model with health is targeted to match the empirical paths of employment, bad health, and wage income. We weight the targets uniformly, sacrificing some fit on employment to better match wages. The dashed lines report the implied counterfactual paths when the health channel is shut down while holding the same shock sequence fixed. In the absence of the health channel, the model generates a smaller employment contraction and no increase in bad health, implying that health shocks account for a nontrivial share of the observed employment decline during the GFC.

Figure D.4: GFC scenario in general equilibrium: 2 shocks



Note: This figure is constructed in analogy to Figure D.3 except that it only uses the bad health and the demand shock.

Figure D.5: GFC scenario in partial equilibrium: 2 shocks



Note: This figure is constructed in analogy to Figure D.4 except that here the policy functions are computed in partial equilibrium.

Figure D.6: Demand-shock IRFs in general equilibrium and partial equilibrium

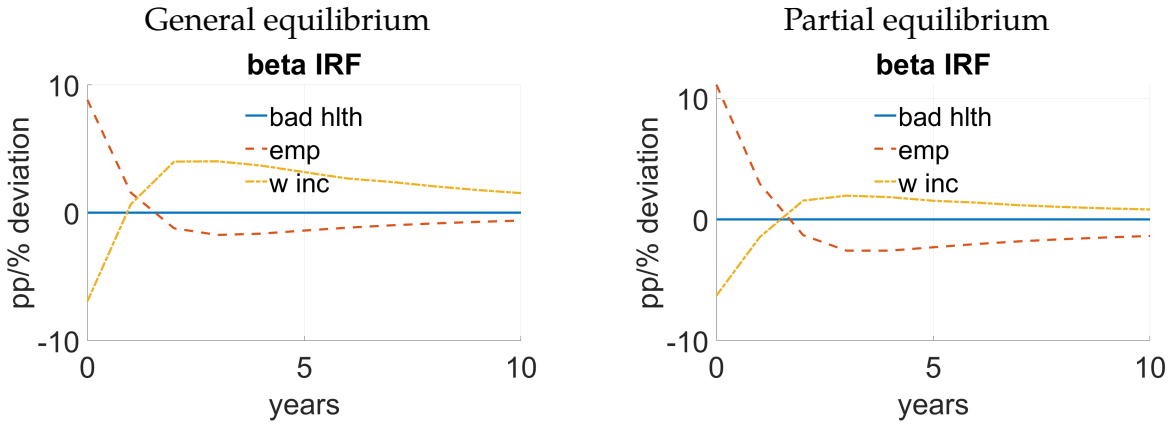


Figure D.7: Health-shock IRFs in general equilibrium and partial equilibrium

