The Life Cycle, Ageing and Longevity

Lecture 2 (of 2): Income Uncertainty, Health Shocks, and Mortality Risk

Yannick Schindler

EIT Oxford

29 October 2025

Recap & roadmap

- **Recap from Day 1:**
  - We built a **discrete-time** life-cycle model.
  - We added endogenous labour supply and an absorbing retirement state.
- **Today's roadmap:**
  1. **From discrete to continuous Time (CT):** the "how" and "why".
  2. **Motivating uncertainty:** income, health, and mortality risk.
  3. **Building a fuller lifecycle model:** constructing a lifecycle model with mortality, health, and income uncertainty.
  4. **Key results:** how lifecycle uncertainty affects the "value of prevention".

Practical information

- **A "cookbook" style lecture**
  - The purpose is to demonstrate the dynamic optimization problem, briefly discuss how to solve it, and then move quickly to the results.
- **What we will *not* cover today**
  - Detailed mathematical derivations.
  - In-depth exposition on numerical solution methods.
  - Distributions and the Kolmogorov Forward Equation (KFE).
- **The main goal:**
  - To get you **thinking about the models** and their economic implications.
- **Want to learn more?**
  - For those who want to get into the details, these are excellent resources to start with:
  - Ben Moll's Lecture Notes
  - Pontus Rendahl's Lecture Notes

Why bother with continuous time?

1. **Computational speed:**
    - Solving the HJB system (with finite differences) involves solving a system of linear equations
    - This is often *much faster* than discrete-time counterpart (see. e.g. Ben Moll's note for analysis)
2. **The KFE is "free":**
    - Once HJB is solved, solving the Kolmogorov Forward equation (KFE) is straightforward
    - Solution to KFE gives us distribution (e.g. across income and assets)

The canonical consumption-saving problem (discrete vs. continuous time)

Discrete Time (Lecture 1)

$$\max_{\{c_t\}_{t=0}^{T}} \quad \sum_{t=0}^{T} \beta^t \, u(c_t)$$

$$\text{s.t.} \quad a_{t+1} = (1+r)a_t + y_t - c_t,$$

$$a_{t+1} \geq \underline{a}, \quad a_0 \text{ given.}$$

- Time moves in discrete steps $t, t+1, \ldots$
- β is the discount factor.
- Constraint is a *difference equation*.

Continuous Time Analogue

$$\max_{\{c_t\}_{t=0}^{T}} \quad \int_0^T e^{-\rho t} \, u(c_t) \, \mathrm{d}t$$

$$\text{s.t.} \quad \frac{\partial a_t}{\partial t} \equiv \dot{a}_t = r a_t + y_t - c_t,$$

$$a_t \geq \underline{a}, \quad a_0 \text{ given.}$$

- Time t flows continuously.
- ρ is the discount rate (where $\beta \approx e^{-\rho \Delta t}$).
- Constraint is a *differential equation*.

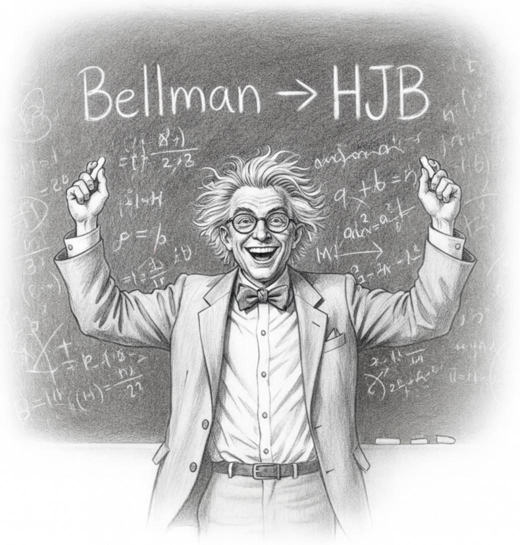
Bellman → HJB

From discrete time (Bellman) to continuous time (HJB)

- **Recall the DT Bellman (from Day 1):**

$$V_t(a_t) = \max_{c_t}\{u(c_t) + \beta V_{t+1}(a_{t+1})\}$$

- **Let's formalise the time step:**
  - Let the period length be Δt. The problem is:

$$V(a,t) = \max_{c}\{u(c)\Delta t + e^{-\rho\Delta t} V(a + \dot{a}\Delta t, t + \Delta t)\}$$

  - Here, $\beta = e^{-\rho\Delta t}$ and $a_{t+1} = a_t + \dot{a}\Delta t$.
- **The "Trick":** Taylor expand $V(a + \dot{a}\Delta t, t + \Delta t)$ around (a,t) and Taylor expand $e^{-\rho\Delta t}$ around 0. Re-arrange and then take $\Delta t \to 0$
- **End result: The Hamilton-Jacobi-Bellman (HJB) Equation**

$$\rho V(a,t) = \max_{c}\{u(c) + V_t(a,t) + V_a(a,t)\cdot\dot{a}\}$$

Recursive formulation (Bellman vs. HJB)

Discrete Time: Bellman Equation

$$V_t(a_t) = \max_{c_t} \left\{ u(c_t) + \beta\, V_{t+1}(a_{t+1}) \right\}$$

$$\text{s.t.} \quad a_{t+1} = (1+r)a_t + y_t - c_t$$

- This is a **recursive functional equation**.
- It relates the value at time t to the value at time $t+1$.

Continuous Time: HJB Equation

$$\rho V(a,t) = \max_{c_t} \left\{ u(c_t) + V_t(a,t) + V_a(a,t) \cdot \dot{a}_t \right\}$$

$$\text{s.t.} \quad \dot{a}_t = r a_t + y_t - c_t$$

- This is a **partial differential equation (PDE)**.
- It relates the "required return" on value (ρV) to its "flow benefits".

Motivating uncertainty in lifecycle models

- Life is risky. Our simple model from Day 1 was deterministic.
- To understand behaviour across the lifecycle, we should model the "Big 3" uncertainties:

1. **Income risk:** Productivity shocks, promotions, job loss.
   - *Drives precautionary saving, does it crowd out investment in better health?*
2. **Mortality risk:** Uncertainty about the length of life.
   - *Affects expected length of life and hence the entire dynamic lifecycle optimisation.*
3. **Health (morbidity) risk:** Getting sick, chronic illness, disability.
   - *A key driver of labour exit and consumption changes.*

Question: Why might health matter for lifecycle choices?

Estimating effect of chronic disease on labour supply (Schindler & Scott (2025))

- **Data:** UK Household Longitudinal Study (UKHLS), 2009–2019, individuals observed annually with detailed health and labour outcomes.
- **Design:** Local Projection Difference-in-Differences (LP-DiD)

$$y_{i,t+h} - y_{i,t-1} = \beta_h D_{i,t} + X'_{i,t-1}\gamma + \lambda_{t-1} + \varepsilon_{i,t+h},$$

  where $D_{i,t}$ marks first diagnosis and h indexes years since diagnosis.
- **Sample construction:**
  - Treated: individuals first diagnosed with a specific chronic condition (CVD, cancer, etc.).
  - Controls: individuals who remain event-free for all available waves.
- **Outcomes:** (i) Exit to inactivity; (ii) Transition from full-time to part-time work.
- **Check of pre-trends:** Can run regression for $h < -1$ to see if effect of diagnosis shows up before diagnosis.

Why might health matter for the lifecycle?

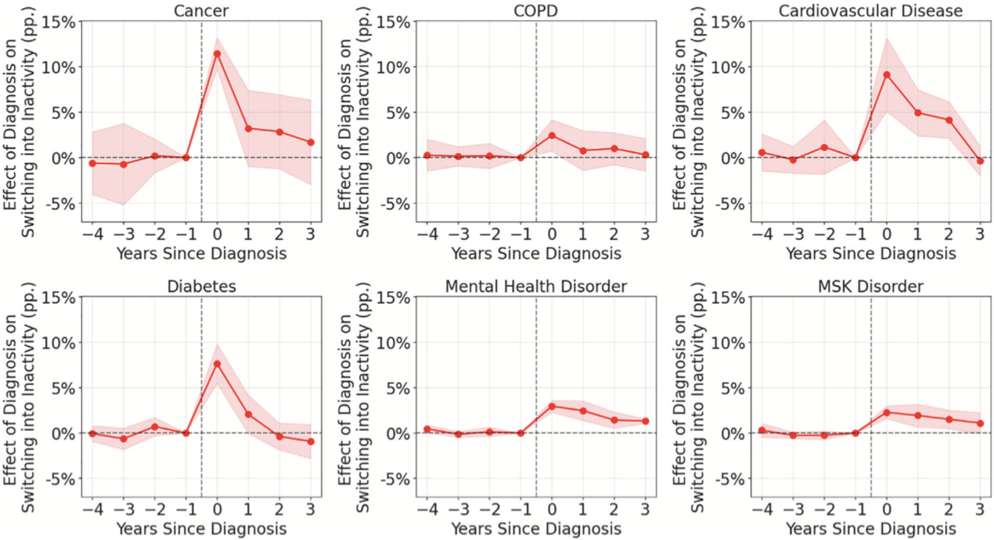

Why considering health in our models is important

- **Demographic shifts:** As populations age, the prevalence of chronic, work-limiting conditions becomes a first-order economic question (see e.g. Schindler & Scott (2025)).
- **Health improvements** have two effects:
  1. **Better averages:** People live longer and are healthy for longer.
  2. **Reduced uncertainty:** Medical advances (e.g., lower cancer incidence) don't just improve the average across the population; they reduce uncertainty at the individual level.
- **Today's Goal:** Build a model to quantify how health, mortality, and income uncertainty affect the "value of prevention"

Model ingredients (1/4): preferences

Preferences

Standard utility over consumption c and labour l:

$$u(c, l) = \frac{c^{1-\gamma}}{1-\gamma} - \frac{l^{1+1/\eta}}{1+1/\eta}$$

γ = Risk Aversion , η = Frisch elasticity

Productivity

An agent's wage $e(h, z_t, t)$ has 3 parts:

$$e(h, z_t, t) = \underbrace{p(t)}_{\text{Hump-shape}} \cdot \underbrace{z_t}_{\text{Stochastic}} \cdot \underbrace{\mathbb{I}_h}_{\text{Health}}$$

- $p(t)$: Deterministic age-profile (hump-shaped).
- z_t: A persistent stochastic shock.
- $\mathbb{I}_h$: Health status multiplier.

Model Ingredients (3/4): Income shocks

Income shocks: Ornstein-Uhlenbeck (OU) process

The log-productivity process z_t is mean-reverting :

$$d\log(z_t) = -\theta \log(z_t)dt + \sigma dW_t$$

θ is the speed of mean-reversion, σ is the volatility, and W_t is a standard Wiener process (Brownian Motion).

- **What this looks like:** It's a "persistent" shock. If you are lucky ($\log(z_t) > 0$) or unlucky ($\log(z_t) < 0$), you tend to stay that way, but you are always pulled back towards the mean ($\log(z_t) = 0$).

Question: Why not model process for z_t instead of $\ln(z_t)$?

Example Ornstein-Uhlenbeck processes with lifecycle profile

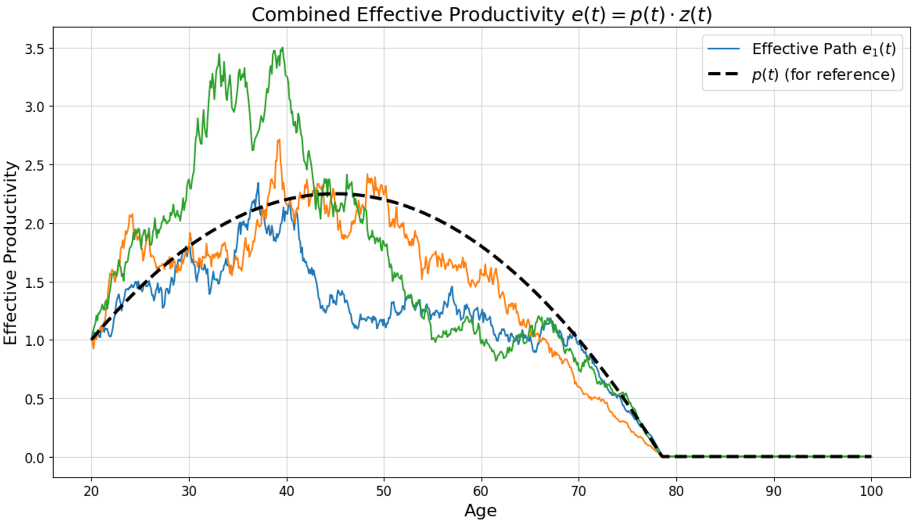

Can study income distributions *across* and *within* cohorts

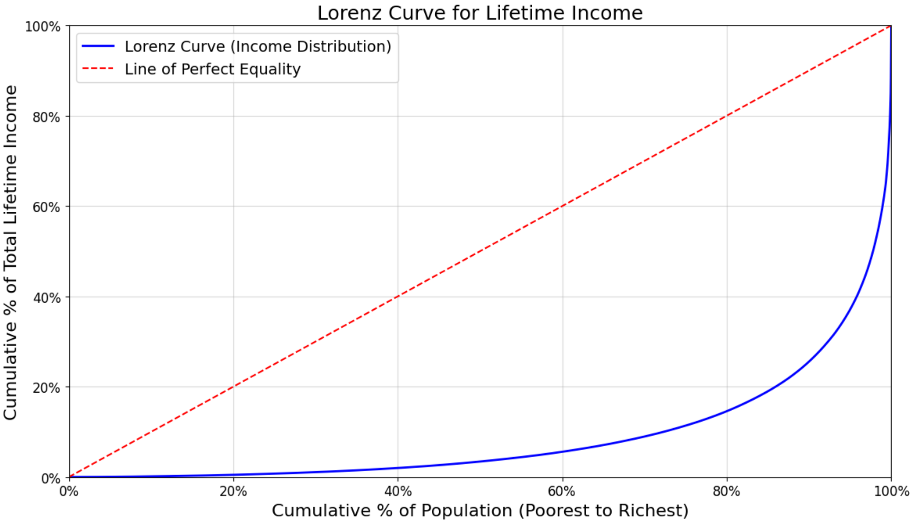

Mortality Risk Agents face an *age-dependent* instantaneous probability of death $\lambda_m(t)$.

Health/morbidity risk Healthy agents face an *age-dependent* sickness hazard, $\lambda_s(t)$.

If the shock hits, they transition to a **permanent, absorbing** sick state.

The "Sick" State Productivity is permanently reduced by a factor $\psi \in (0, 1)$.

$$\mathbb{I}_h = \begin{cases} 1 & \text{if } h = \textit{Healthy} \\ \psi & \text{if } h = \textit{Sick} \end{cases}$$

Discussion: in what other ways could sickness enter a model like this?

The coupled HJB system

- Because sickness is absorbing, we have a *system* of two coupled value functions: V^H (Healthy) and V^S (Sick) .

HJB for Healthy Agent

$$(\rho + \lambda_m(t) + \lambda_s(t))V^H = \max_{c,l}\{u(c,l) + V_t^H + V_a^H \dot{a}^H + \mathcal{A}V^H + \lambda_s(t)V^S\}$$

- **Key terms:** Total discount rate includes $\lambda_s(t)$ (risk of getting sick). Flow benefits include the value of transitioning to the sick state, $\lambda_s(t)V^S$.

HJB for Sick Agent

$$(\rho + \lambda_m(t))V^S = \max_{c,l}\{u(c,l) + V_t^S + V_a^S \dot{a}^S + \mathcal{A}V^S\}$$

- **Key difference:** Sickness hazard $\lambda_s(t)$ is gone. Productivity is permanently lower.

Impact of health shocks on lifecycle behaviour

- **What happens when an agent gets sick?**

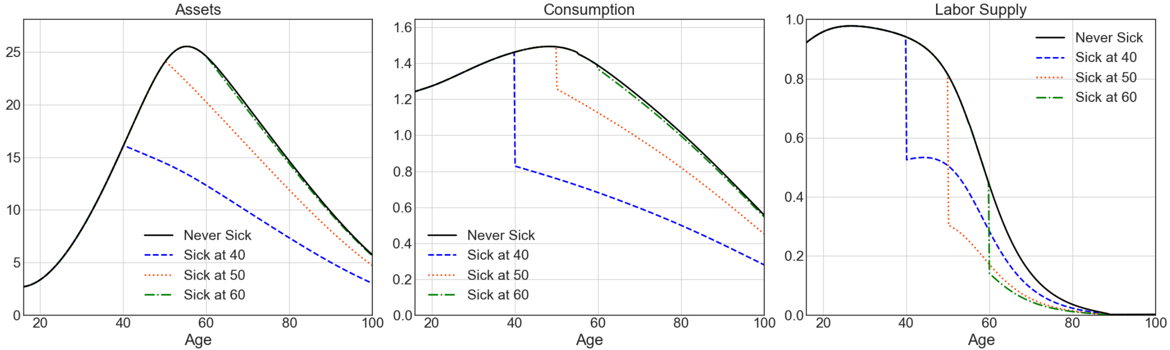

Source: Schindler (2025), Figure 1.

- **Assets (Left) & Consumption (Middle):** Both drop immediately and permanently .
- **“Morbidity millstone”** Person who become sick early is working *the most* later in life!

What is "Willingness to Pay (WTP)"?

- **Definition:** WTP is how much of your income (or assets) you would be willing to give up *today* in exchange for a permanent 1 percentage point reduction in the sickness hazard $\lambda_s(t)$.

WTP: The Indifference Condition

The WTP is the amount of assets an agent would pay to be indifferent between the baseline world (high risk) and the prevention world (low risk):

$$\underbrace{V^H(a - \text{WTP}, z, t; \lambda_s^{\text{new}})}_{\text{Value with Prevention, but less assets}} = \underbrace{V^H(a, z, t; \lambda_s^{\text{baseline}})}_{\text{Value in Baseline World}}$$

where $\lambda_s^{\text{new}} < \lambda_s^{\text{baseline}}$.

- It's the "price" (paid from assets a) that makes you indifferent.
- The model allows us to compute this WTP for *any* agent, based on their age, assets, productivity, and health .

The *young* and *asset-poor* value prevention most

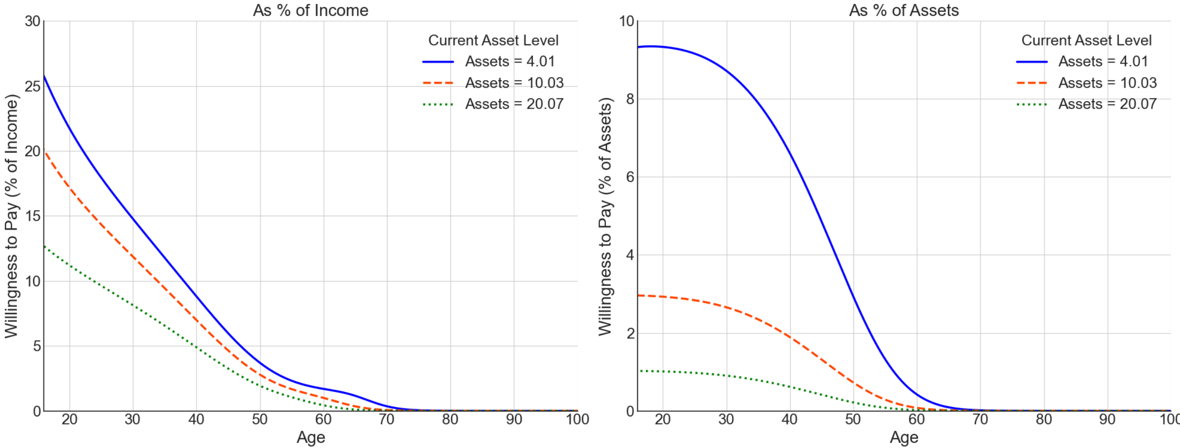

Source: Ashwin & Schindler (2025)

Question: why might asset-rich value prevention less?

Core tension: insurance vs. prevention

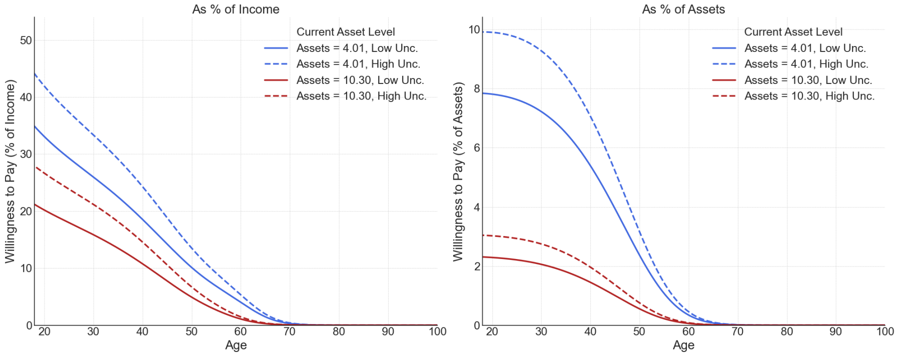

Source: Ashwin & Schindler (2025)

- **Result:** Higher income uncertainty (dashed lines) *raises* the WTP for prevention.
- **Why?** One dollar spent on prevention improves chance of high-income-healthy state and reduces chance of low-income-unhealthy state > using one dollar to self-insure.
- **Question:** when might self-insurance be the more dominant force?

Effect of health shocks on income inequality

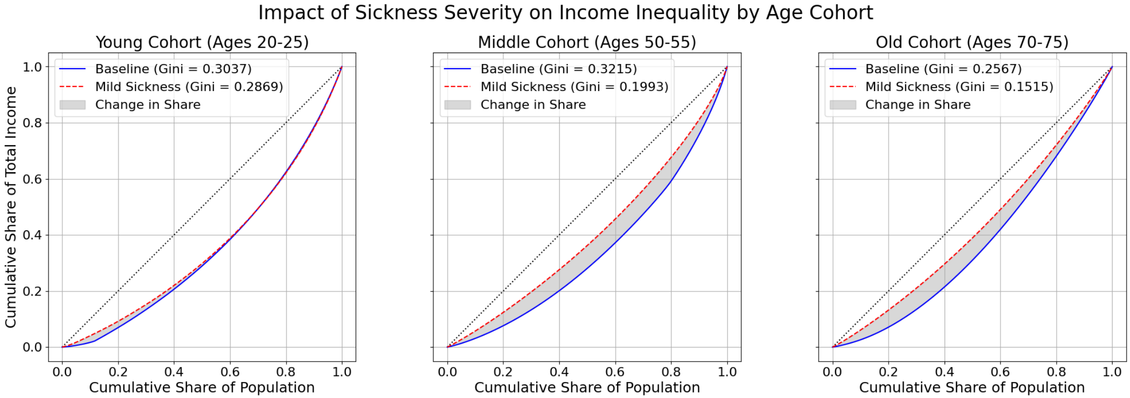

Bonus: the Kolmogorov-Forward equation (KFE)

Main idea: The change in the density at a given point $(a, z, t, \mathbb{I}_h)$ is the sum of all inflows and outflows. If stationary (i.e. no change as time changes), then inflows equal outflows.

KFE for Healthy (g^H)

$$\frac{\partial g^H}{\partial t} = -\underbrace{\frac{\partial}{\partial a}(\dot{a}^H \cdot g^H)}_{\text{Net Flow from Asset Drift}} + \underbrace{\mathcal{A}^* g^H}_{\text{Net Flow from Productivity Shocks}} - \underbrace{\lambda_s(t) g^H}_{\text{Outflow (Become Sick)}} - \underbrace{\lambda_m(t) g^H}_{\text{Outflow (Mortality)}}$$

- This equation tracks the density $g^H(a, z, t, \mathbb{I}_h)$
- Agents leave this group by either becoming sick (λ_s) or through mortality (λ_m)
- Some get "flung" out (in) to (from) other parts of the distribution: $\mathcal{A}^* g^H$
- There is net flow, due to asset drift, coming from points on either side of a: $\frac{\partial}{\partial a}(\dot{a}^H \cdot g^H)$

KFE for sick agents (g^S)

Two KFEs: One for healthy and one for sick. KFE for sick is only slightly different.

KFE for Sick (g^S)

$$\frac{\partial g^S}{\partial t} = -\underbrace{\frac{\partial}{\partial a}(\dot{a}^S \cdot g^S)}_{\text{Net Flow from Asset Drift}} + \underbrace{\mathcal{A}^* g^S}_{\text{Net Flow from Productivity Shocks}} + \underbrace{\lambda_s(t) g^H}_{\text{Inflow (from Healthy)}} - \underbrace{\lambda_m(t) g^S}_{\text{Outflow (Mortality)}}$$

Total Wealth Distribution $g(a)$

To find the distribution across assets a only, we integrate out the productivity state z and sum across health types:

$$g(a, t, \mathbb{I}_h) = \underbrace{\int g^H(a, z, t, \mathbb{I}_h)\, dz}_{\text{Wealth dist. of Healthy}} + \underbrace{\int g^S(a, z, t, \mathbb{I}_h)\, dz}_{\text{Wealth dist. of Sick}}$$