Wage Dispersion, On-the-Job Search, and Stochastic Match Productivity

A Mean Field Game Approach

I. Sebastian Buhai

UC Chile–Instituto de Economia · Stockholm University–SOFI · Minho University–NIPE

Banco Central de Chile, Santiago · 4 March 2026

Why this matters for a central bank

- Wage dispersion and mobility shape **inequality**, **human capital allocation**, and **policy transmission**.
- Job ladders discipline how shocks and reforms propagate through **EE flows** and **wage growth**.
- A tractable equilibrium mapping from **policy primitives** to **(hazards, wages, dispersion)** is valuable.

One-line goal

Build a fully solved equilibrium search model where **match quality evolves stochastically within jobs**, generating an endogenous job ladder and a free-boundary separation rule.

Empirical targets (joint discipline)

1. **Large wage dispersion** even within narrow worker $\times$ job cells.
2. **Wage growth is a mobility story:** big gains often occur at EE moves.
3. **Tenure hazards are front-loaded:** high early separation/EE rates, then steep decline (long tail).

Key identification idea

A single latent state (match surplus) must jointly rationalize *wage levels*, *EE hazards*, and *tenure durations*.

Benchmark tension: where standard equilibrium search strains

- In classic BM/PVR-style models, wage dispersion is typically driven by **firm heterogeneity** and **search frictions**.
- Matching the *level* of dispersion often pushes models toward **very large** productivity heterogeneity.
- Matching *tenure patterns* simultaneously is harder: hazards and wage-tenure slopes are tightly linked.

This paper's move

Shift dispersion-generating power from permanent heterogeneity toward **within-job stochastic match productivity** plus equilibrium ladder effects.

Where this fits I: equilibrium search, wage posting, job ladders

- Canonical equilibrium on-the-job search and job ladders: Burdett–Mortensen (1998); Postel-Vinay–Robin (2002).
- Macro matching backbone and flows: Mortensen–Pissarides (1994).
- Quantitative benchmark and "frictional wage dispersion" puzzle: Hornstein–Krusell–Violante (2011).
- Wage dynamics in structural search with heterogeneity: Lentz–Mortensen (2010); Bagger et al. (2014).

This paper's angle (relative to this literature)

Endogenize the ladder using a **diffusive match state** plus **equilibrium outside options**, and keep the mapping from primitives to hazards and wage moments transparent.

Where this fits II: stochastic match quality, selection, mean field games

Stochastic match quality and selection

- Learning and turnover: Jovanovic (1979, 1984)
- Diffusion + efficient separation and tenure hazards: Buhai–Teulings (2014); Abbring (2012).

Mean field games as equilibrium technology

- Monotone MFG framework: Lasry–Lions (2007); Carmona–Delarue (2018); Cardaliaguet–Porretta (2020).
- Stopping in MFGs: Bertucci (2018); Nutz (2018).

Method in one sentence

Stationary equilibrium solves a coupled **HJB (optimality)** and **Kolmogorov (invariant law)** system; the separation threshold is a **free boundary**.

Why central banks might care

It yields a disciplined mapping from **policy wedges** (firing costs, search incentives, volatility) to **hazards, EE flows, and the wage distribution**.

Warm-up: stochastic match productivity as optimal stopping

State: match surplus $Z_t \equiv P_t - R_t$.

$$\mathrm{d}Z_t = \mu_Z \,\mathrm{d}t + \sigma_Z \,\mathrm{d}B_t.$$

Worker chooses separation time τ (stop when Z_t is low enough).

- Optimal separation is a **threshold**: separate when $Z_t \leq z^*$.
- Hitting-time structure generates **declining hazards** and a **long tail**.

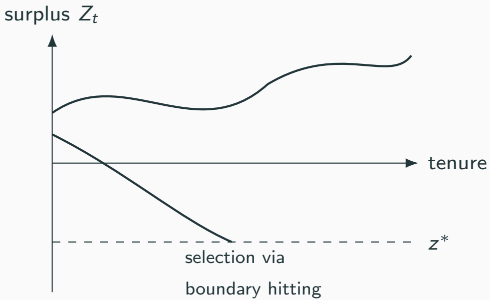

What diffusion + stopping buys you (in one slide)

- **Front-loaded hazards:** many matches start near the boundary; early sorting is intense.
- **Survivorship selection:** conditional on survival, the cross-section shifts toward higher Z.
- **Tenure profiles:** declining hazards and concave wage-tenure can be selection-driven.
- **Macro implication:** duration dependence does not equal structural state-dependent destruction.

Key identification message

Observed hazards and wage-tenure slopes are jointly informative about the latent Z process and the stopping rule.

Why general equilibrium? Outside options and offers are endogenous

- Outside value V^U depends on: unemployment payoff, job-finding prospects, and the wage distribution.
- On-the-job search implies an endogenous **offer distribution**: what you meet is drawn from the economy's cross-section.
- Firms' wage policies affect retention and poaching $\Rightarrow$ equilibrium feedback.

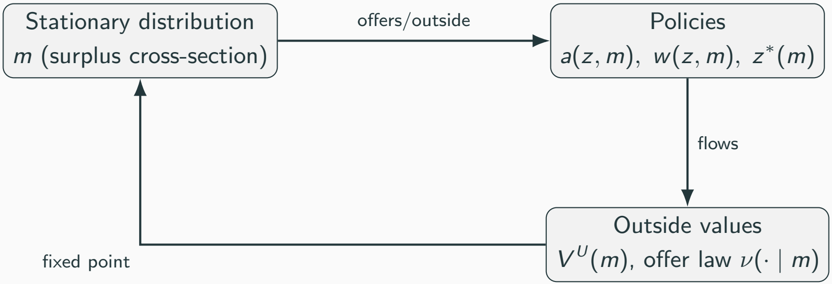

Model overview (macro-labor translation)

Within a match

- State: surplus z (1D).
- Diffusion: $\mathrm{d}z = \mu(z, a, m)\,\mathrm{d}t + \sigma(z, m)\,\mathrm{d}B$.
- Worker chooses:
  - search intensity a (costly) $\Rightarrow$ offer arrivals,
  - separation when z is too low.

Across matches (the market)

- Offers arrive at rate $\lambda(a, m)$; offer states drawn from $\nu(\cdot \mid m)$.
- Firms post Markov wage schedule $w(z, m)$ (retention vs profits).
- Stationary cross-section m must be invariant under diffusion + quits + separations + entry.

What exactly is the stationary distribution m?

- m is the **cross-sectional distribution of ongoing matches** over the state z.
- Think: at a point in (stationary) time, sample all employed matches; each has a current surplus z; m is the distribution of those z's.
- It is *not* "the distribution of realized wages" (though wages are an equilibrium function of z).
- It is also *not* "the distribution of worker types" (here z is the sufficient statistic).

Empirical translation

m is the latent-state analogue of the observed cross-section of *wages, hazards, and durations*; matching the data means matching moments implied by m after mapping $z \mapsto (w, \text{hazards})$.

What is the job ladder here? Bottom, top, and movement

- A **job** is a state z plus an associated wage/continuation value.
- **Bottom of the ladder:** z near the separation boundary z^*.
- **Top of the ladder:** high z matches (high wages, low separation risk).
- **Moving up:** receiving an offer with higher continuation value and switching (EE move).
- **Endogenous ladder:** the offer distribution is induced by m; thus the set of "rungs" is equilibrium-determined.

Concrete picture

Low-z workers search more and accept many upgrades; high-z workers rarely move (already near the top).

Deep dive

Stationary equilibrium (definition in words)

A stationary equilibrium is a **fixed point** in the cross-sectional distribution m:

1. **Worker optimality:** given (m, w), choose search $a(z, m)$ and separation to maximize lifetime utility.
2. **Firm optimality:** given m and worker behavior, choose wage policy $w(z, m)$.
3. **Outside values consistent:** unemployment/vacancy values are those generated by the same environment.
4. **Stationarity:** the induced flows keep the distribution at m.

The MFG backbone I: worker problem = HJB obstacle (continue vs separate)

Continuation vs separation (one state z, one distribution m)

Continuation region $(z > z^*(m))$:

$$rV^W(z;m) = \max_{a\in A}\Big\{ w(z,m) - c(a) + \mathcal{L}_{a,m}V^W(z;m) + \lambda(a,m)\,\mathbb{E}_{z'\sim\nu(\cdot|m)}\big[(V^W(z';m) - V^W(z;m))_+\big]\Big\}.$$

Stopping region $(z \le z^*(m))$: $\quad V^W(z;m) = V^U(m)$.

- $\mathcal{L}_{a,m}$ is the diffusion generator (drift/diffusion of surplus).
- The last term is the **option value of on-the-job search**: you only take better offers.
- The "obstacle" $V^U(m)$ encodes the **separation option** $\Rightarrow$ a threshold rule.

Deep dive

The MFG backbone II: distribution = stationary Kolmogorov (flow balance)

Forward equation (who sits where on the ladder in steady state)

Define the probability **flux**

$$J(z) \equiv \mu(z, a^*(z,m), m)\, m(z) \;-\; \tfrac{1}{2}\,\partial_z\big(\sigma^2(z,m)\, m(z)\big).$$

Stationarity = local flow balance:

$$0 \;=\; -\partial_z J(z) \;-\; \underbrace{\text{exit}}_{\substack{\text{separation at } z^*(m)\\ \text{and EE moves}}} \;+\; \underbrace{\text{entry}}_{\substack{\text{new matches / re-entry}\\ \text{and upgrades from offers}}} \quad .$$

- Think "mass conservation": drift/diffusion moves workers along z; the boundary removes mass; entry injects new mass.
- In equilibrium, the same m that agents take as given is exactly the one generated by these flows.

Interpretation: backward equation pins down *policies*; forward equation pins down the *steady-state cross-section*. *Equilibrium is their fixed point.*

Why the boundary is "free" in a free-boundary problem

- The worker problem is an **optimal stopping** problem for a diffusion.
- The *value function* solves a differential equation only on the continuation set.
- The *stopping threshold* z^* is **unknown ex ante** and must be solved *jointly* with the value function.

Two boundary conditions pin it down

Value matching: $V^W(z^*; m) = V^U(m)$. **Smooth fit (diffusion):** $\partial_z V^W(z^*; m) = 0$.

Equilibrium as a fixed point in the distribution

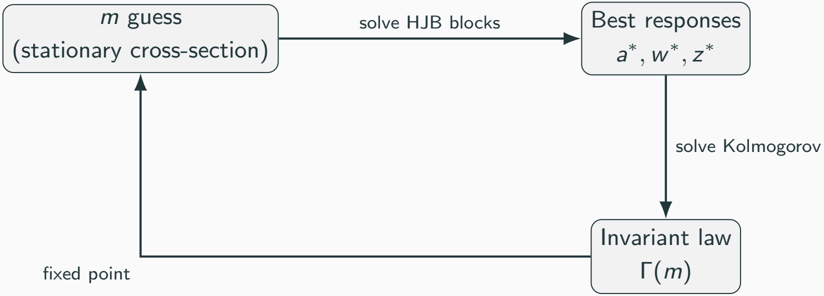

What "stationary equilibrium" means

m^* satisfies $m^* = \Gamma(m^*)$: if everyone optimizes taking m^* as given, the induced cross-section is again m^*.

Theory roadmap (what is proved, why you should believe it)

1. **Optimal policies exist** and separation is a **threshold** (free boundary) for any given environment m.
2. For any m, the induced stationary law $\Gamma(m)$ exists and is unique.
3. The equilibrium fixed point exists (continuity + compactness arguments).
4. Uniqueness follows from **Lasry–Lions monotonicity**: no self-fulfilling "ladders".

Lasry–Lions monotonicity: intuition (vs "simple monotonicity")

- **Simple monotonicity** is pointwise: if $m_1 \geq m_2$ then $f(m_1) \geq f(m_2)$.
- **Lasry–Lions monotonicity** is an *integral* one-sided condition that rules out positive feedback in the coupling:

$$\int \big(F(z, m_1) - F(z, m_2)\big)\,(m_1 - m_2)(\,\mathrm{d}z) \;\leq\; 0.$$

Economic reading

If the cross-section shifts toward "better jobs", the induced incentives (wages/outside value) should **not** amplify that shift into even stronger upward drift; i.e., no strategic complementarity strong enough to sustain multiple ladders.

Deep dive

Computation: solving the coupled HJB–Kolmogorov system

- 1D state $z \Rightarrow$ fast, reliable numerics.
- Discretize z on a grid: diffusion becomes a continuous-time Markov chain.
- Solve:
    - worker HJB obstacle via policy iteration + projection onto stopping set,
    - Kolmogorov via conservative flux form (mass preservation).
- Iterate the distribution fixed point $m^{(\ell+1)} \leftarrow \Gamma(m^{(\ell)})$ until convergence.

Convergence and stability of the algorithm (benchmark)

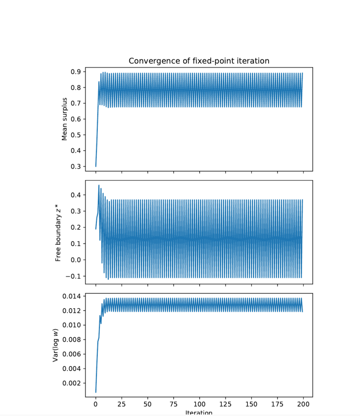

Interpretation: the fixed-point iteration quickly settles: mean surplus, the free boundary, and implied wage dispersion stabilize. The small oscillations reflect discrete-time relaxation / grid-based iteration, not instability.

Calibration: mapping parameters to moments (high level)

Key primitives

- Diffusion of match productivity: drift μ, volatility σ.
- Search technology: offer rate $\lambda(a)$ and cost $c(a)$.
- Wage-setting rule (sharing): how w loads on surplus vs outside value.
- Entry/re-entry conditions for new matches.

Disciplined by

- Separation hazards by tenure.
- EE transition rates and wage gains at moves.
- Wage dispersion by tenure (levels and slope).

Why identification is sharp here

The same latent state z must fit *simultaneously* wages and hazards; the free boundary strongly pins down how σ and V^U shape tenure dynamics.

Benchmark equilibrium objects: value, boundary, and search

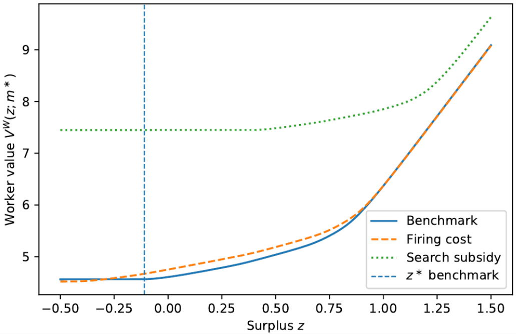

The dashed vertical line is the endogenous separation threshold z^*. Policies (firing costs, search subsidies) shift z^* and tilt the continuation value, reshaping the ladder.

Stationary surplus distribution and job ladder

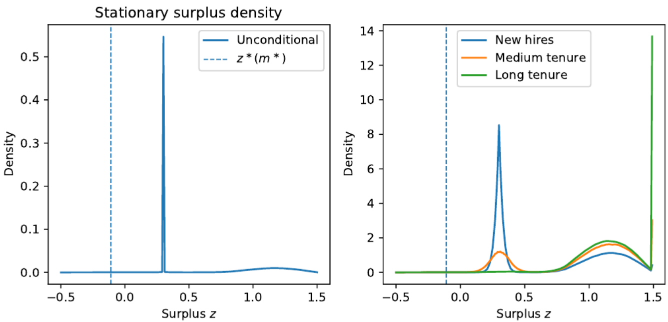

Mass concentrates above z^*: new matches enter near the boundary; surviving long-tenure matches are overrepresented in the upper tail; a diffusion-and-selection job ladder.

Wage distribution by tenure and mobility

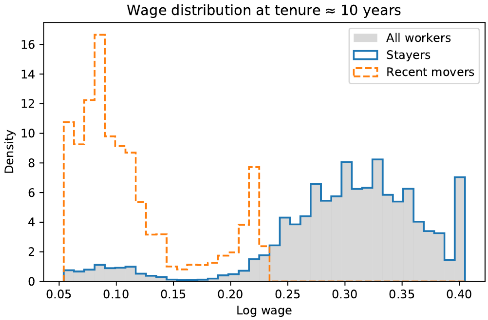

Holding tenure fixed, recent movers are shifted right and more dispersed: EE upgrades are the main way workers reach high-surplus, high-wage states.

Tenure hazards: front-loaded separations

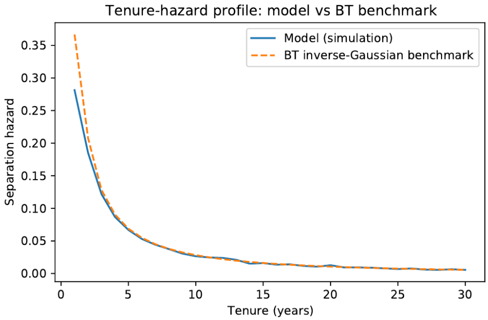

The early-tenure hazard is high because many new matches start near z^*; the hazard declines as diffusion + selection move survivors away from the boundary.

Wage-tenure profiles and selection

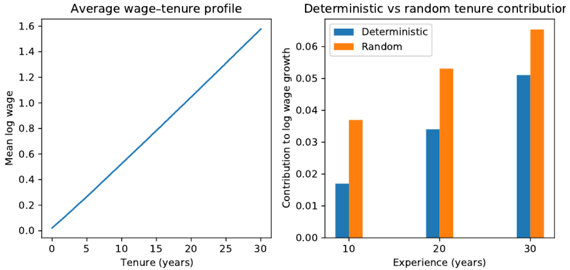

Even with parsimonious wages, part of wage growth with tenure is selection: low-z matches exit early and the surviving cohort shifts up the ladder.

Quantitative decomposition: what drives wage dispersion by tenure?

Three economies

1. **Baseline:** search + stochastic match productivity (full model).
2. **No on-the-job search:** isolate selection/stopping from mobility.
3. **No stochastic match productivity:** revert toward standard ladder dispersion mechanisms.

- Compare implied dispersion profiles and hazards across economies.
- Interpret differences as contributions of **selection**, **mobility**, and **within-job shocks**.

Decomposition results

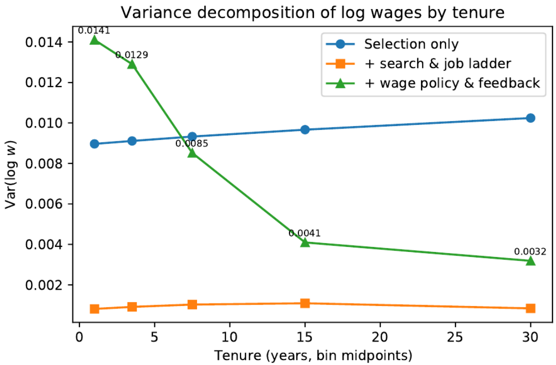

At short tenure, selection and rapid upgrading matter most; at longer tenure, within-job shocks and wage feedback increasingly shape dispersion.

Policy experiments: think "shift the boundary" + "reshape the ladder"

- Policy primitives affect the worker's continuation vs stopping tradeoff:
  - employment protection / firing costs,
  - unemployment benefits / outside value,
  - search subsidies / matching efficiency,
  - volatility of match productivity (macro uncertainty).
- In the model, they operate through:
  - the separation cutoff z^* (free boundary),
  - search effort $a(z)$ and acceptance (ladder speed),
  - the stationary cross-section m (offer distribution).

Policy 1: Employment protection / firing costs

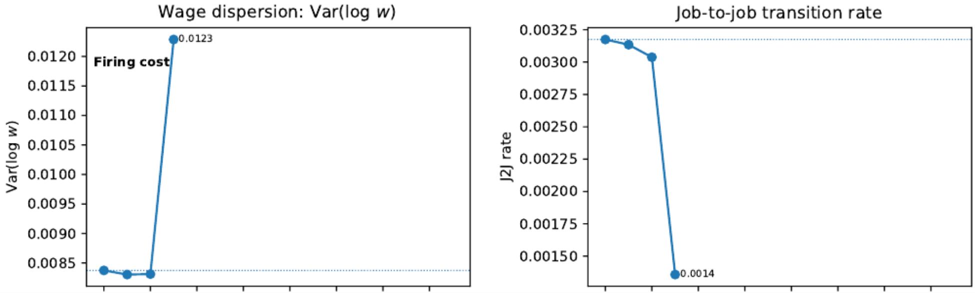

Mechanism (one sentence)

Higher firing costs effectively lower separation incentives (or make separation more costly), shifting z^* and reshaping the cross-section.

This changes EE upgrading and wage dispersion.

Policy 2: Search subsidies / matching efficiency

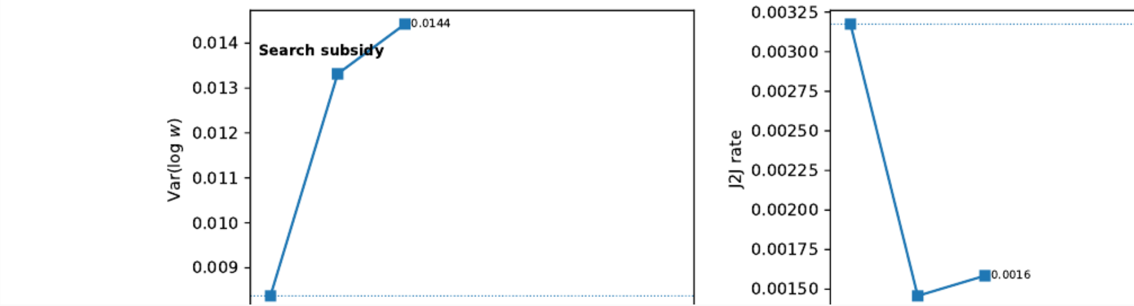

Mechanism

Lower search costs or higher matching efficiency speeds up ladder climbing (more EE moves).
This can raise wage growth at moves while changing within-tenure dispersion through selection.

Policy 3 / environment: higher volatility of match productivity

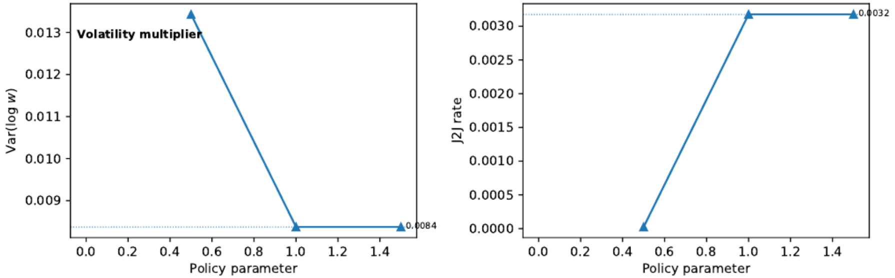

Mechanism

Higher σ raises the option value of waiting and changes the frequency of boundary hits.
Equilibrium adjusts via z^*, search, and the cross-section.

Macro takeaways (what to remember)

1. **Job stability profiles** can emerge from diffusion + endogenous stopping (selection), not only time-varying destruction shocks.
2. **Wage dispersion by tenure** decomposes into: within-job shocks + ladder mobility + selection.
3. **Policy effects are equilibrium effects:** shifting z^* changes the cross-section m, hence the ladder that everyone faces.

Operational summary for policy

When evaluating reforms, track: (i) separation boundary, (ii) ladder speed (EE), and (iii) cross-sectional reweighting (selection).

Limitations & extensions (one slide)

Limitations (what is abstracted from)	Extensions (natural next steps)
• No permanent worker/firm types (all heterogeneity is in z).	• Add types (θ, ϕ) to separate "level dispersion" from "tenure dynamics".
• Stationary environment (no aggregate shocks).	• Add aggregate shocks/common noise: time-varying m_t and policy rules.
• Wage-setting discipline uses a parsimonious sharing rule.	• Richer contracting: renegotiation, counteroffers, endogenous vacancy creation.

Deep dive

Conclusion (three headline contributions)

1. A tractable equilibrium search model with **stochastic within-match productivity** and **endogenous stopping**.
2. A **stationary mean field game** characterization with existence/uniqueness (free boundary + Lasry–Lions).
3. Quantitatively: a **structural decomposition** of wage dispersion by tenure and transparent policy counterfactuals through the boundary and ladder.

Stop point

Happy to take questions anywhere; especially on the ladder mechanism and the policy mappings.

Thank you.

Comments, critiques, and extensions are very welcome.

Deep dive: acceptance rule and ladder speed

Back to main flow

- Offer at state z' is accepted iff $V^W(z'; m) \geq V^W(z; m)$.
- This induces an endogenous **ranking of jobs** by continuation value (rungs of the ladder).
- **Ladder speed** is governed by:
  - search intensity $a(z)$ (endogenous; typically higher near the bottom),
  - the upper tail of the offer law $\nu(\cdot \mid m)$ (how many "better rungs" exist).

Deep dive: worker obstacle problem (compact form)

Back to main flow

$$0 = \max\left\{ \underbrace{V^W(z;m) - V^U(m)}_{\text{stop}},\ \underbrace{rV^W(z;m) - \max_a \left[w(z,m) - c(a) + \mathcal{L}_{a,m} V^W + G(z, V^W, m; a) \right]}_{\text{continue}} \right\}.$$

- 1D diffusion $\Rightarrow$ stopping set is an interval $(-\infty, z^*]$.
- **Smooth fit**: for diffusions without jumps, the value is C^1 across z^* (value and first derivative match).

Deep dive: stationary Kolmogorov in flux form

Back to main flow Let $J(z) = \mu(z)\, m(z) - \frac{1}{2}\partial_z\big(\sigma^2(z)m(z)\big)$ be the probability flux.

$$0 = -\partial_z J(z) - \underbrace{\delta(z)\, m(z)}_{\text{outflow}} + \underbrace{\iota(z)}_{\text{inflow}}.$$

- Intuition: stationarity is a **mass-conservation identity**; divergence of flux is offset by entry/exit.
- The boundary at z^* is "killing": mass that hits it exits and (in steady state) is replaced by inflow through $\iota(z)$.

Deep dive: Lasry–Lions monotonicity in one example

Back to main flow Suppose the coupling is $F(z,m) = f_0(z) + \beta\, z\, \bar{z}(m)$ where $\bar{z}(m) = \int z\, m(\mathrm{d}z)$. Then

$$\int \big(F(z,m_1) - F(z,m_2)\big)\,(m_1 - m_2)(\mathrm{d}z) = \beta\,\big(\bar{z}(m_1) - \bar{z}(m_2)\big)^2.$$

- If $\beta > 0$: positive feedback (can support multiple steady states / multiple “ladders”).
- Lasry–Lions monotonicity requires the opposite sign: the coupling **dampens** cross-sectional shifts.

Deep dive: adding types without breaking the mechanism

Back to main flow

- Add worker types θ and/or firm types ϕ; state becomes (z, θ, ϕ).
- Diffusion + stopping still generates **front-loaded hazards** and **selection-driven tenure profiles** within each type cell.
- Types mainly add *level* dispersion; diffusion/ladder dynamics shape *tenure* profiles.
- Empirically: helps separate "who you are" (types) from "where you are on the ladder" (state).