Friends in Low Places: Network Capital and the Scarring of Labor Markets

Joseph Kopecky*

August 2026

Abstract

Why do deep recessions scar labor markets long after their shocks fade? Networks of employed contacts, built by employment, eroded by unemployment, are a depreciable capital stock that no one owns. I build this stock into a heterogeneous-agent model of the U.S. labor market, disciplined by three decades of CPS data on cyclical job-search behavior. In a calibrated China-shock experiment, the network accounts for a quarter to a third of the employment scar, concentrated on low-wealth workers. Its destruction is an unpriced externality; a state-contingent hiring credit corrects it cheaply in this economy and dominates equal-cost unemployment insurance.

JEL codes: E24, E32, J64, J68, D62, D85. *Keywords:* social networks, referral hiring, network capital, scarring, hysteresis, Pigouvian policy.

*Trinity College Dublin. Email: jkopecky@tcd.ie.

1. Introduction

Why do deep recessions scar labor markets long after their shocks have faded? The question is old, and the facts behind it are well documented. Employment in regions and industries hit by large, concentrated shocks (import competition, plant closures, the collapse of a dominant local industry) remains depressed for a decade or more, long after wages, prices, and the shocks themselves have adjusted. The leading explanations look inside the worker: skills atrophy, health deteriorates, attachment fades. This paper looks between workers instead, asking three questions. First, how much of the scarring caused by a large sectoral decline runs through the labor market's social fabric—the networks of employed contacts through which workers find jobs? Second, who bears that part of the scar? Third, if networks are the vehicle for scarring, what is the right policy instrument to address this channel and what does it cost to implement?

The mechanism at the center of the paper is a two-way feedback between employment and what I call network capital. A large share of workers find jobs through people they know, and the value of that channel depends on the employment of the people they know. An employed contact hears about vacancies, vouches for applicants, and hires. Employment therefore builds a stock of useful connections, and unemployment erodes it. This happens slowly, the way capital depreciates. In this account, a deep sectoral slump turns the feedback vicious: job losers lean harder on their networks precisely when those networks are thinnest, and every job lost thins them further. The stock this loop runs on is owned by no one.[1] No worker, firm, or government counts it on a balance sheet, and no market prices its decay.

In this paper I build that stock into a quantitative equilibrium model and measure what it does. Households face uninsurable unemployment risk and accumulate assets and human capital; sectors are connected by an empirical mixing network estimated from job-to-job flows; job-finding runs through a formal channel priced by free entry and an informal channel whose productivity depends on each sector's network capital. The laboratory experiment is a China-shock-like decline concentrated in goods production. With the network channel alive, the declining sector loses a third more employment-years than in an identical economy in which the network is frozen, and it takes fourteen years rather than ten to recover half the loss. The gap between those two economies (the same shock, with and without the mechanism) is the object this paper measures, prices, and treats.

The mechanism is disciplined by four facts that I document from three decades of Current Population Survey microdata, using the survey's own question about search through "friends or relatives." Reliance on the channel is strongly countercyclical. The

[1] Throughout, "network" means employed contacts of any kind (relatives, former colleagues, acquaintances), not only close friends: the weak-ties finding of Granovetter (1973) is one reason to read it broadly.

share of unemployed searchers using it rises from 13 percent at the 2000 peak to 29 percent in the aftermath of the Great Recession, roughly one percentage point per point of unemployment in the aggregate series and about 0.4 when identified from cross-state variation. The channel's relative effectiveness, measured as its users' job-finding relative to other searchers', does not fall when reliance rises (the opposite of the congestion logic in Beaman 2012), so the network's share of realized job-finding rises in slumps. Reliance builds for months after unemployment rises, beyond its impact response. And sustained slumps move it while spikes do not: the COVID recession, for all its violence, left the share almost untouched. These last two are the signatures of a slowly adjusting stock, and the calibrated model reproduces the sustained-versus-spike ordering as an untargeted moment, though not its spacing, and the build-up's timescale qualitatively.

In the calibrated model the network channel accounts for a quarter to a third of the cumulative employment scar from the sectoral collapse, and deepens its human-capital scar by nearly as much: depressed networks lengthen unemployment spells, and longer spells erode skills. The two scarring stocks compound each other.

Who bears the cost? In the model, the scar concentrates where the capacity to bear it is smallest. Within the declining sector, human-capital losses are deepest in the bottom quartile of the wealth distribution, whose members stay unemployed longest; consumption losses are deepest in the bottom half, among workers with too little buffer to smooth a long spell. Households in healthy sectors are largely spared on both margins. The network channel is not a uniform amplifier: it is a machine for converting a sectoral shock into a personal, persistent one for the workers least equipped to absorb it.

The remedy follows from the ownership problem. Because the stock is unowned, each hire's contribution to it is an externality, and the paper's analytical core prices that externality: a wedge between the private and social value of a job, governed by a single statistic (the loop gain) that measures the strength of the employment–network feedback and, with it, how sharply the economy amplifies a shock. The implied instrument is not unemployment insurance but a state-contingent hiring credit, switched on when the wedge is elevated. Implemented in the model at half its Pigouvian price, the credit removes roughly half of the sectoral scar at a present-value fiscal cost of about 2.5 percent of one quarter's GDP, and it dominates an equal-cost expansion of unemployment insurance, whose marginal insurance value, at the 40 percent baseline replacement rate, is largely offset by its hiring cost. The gain is not evenly shared: the credit's near-term cost falls on the poorest quartile of the shocked sector, the workers the scar hits hardest, and offsetting it takes a targeted transfer separate from the credit itself. A companion instrument (a subsidy to search across sectors, toward better-connected partners), priced by the same theory, pays about a third more than that price in general equilibrium, and only for the declining side: bridges are worth paying for because a re-employed worker rebuilds capital their whole sector searches on.

The paper adds to the macroeconomic literature by developing a new source of scarring and a way to price it: a social, between-worker channel the standard model leaves out, and the Pigouvian wedge that corrects it. Because it reduces to a single state variable, the channel is a tractable extension of the standard heterogeneous-agent New Keynesian model with search frictions, one the field can understand and build on within a familiar setting.

The claim has four boundaries. The network is an amplifier of sustained slumps, not a creator of persistence from nothing: a transitory spike in separations heals within quarters whether the channel is alive or frozen, which is consistent with COVID barely moving reliance in the data. Nor is it a doom loop: the loop gain stays well below the threshold at which slumps would become self-sustaining, in the model as calibrated and at every shock size I can compute. The interaction with monetary policy is modest: the channel amplifies through duration, not through the demand spiral, so passive monetary policy deepens slumps without disproportionately deepening the network's share of them. And the quantitative claims are scoped to separation-driven collapses: at disciplined parameters the free-entry margin cannot generate a 2008-scale demand-only slump, a limit Section 5.3 documents.

This paper sits at the intersection of four literatures. A microeconomic literature beginning with Granovetter (1973) documents that job information flows through employed contacts, and its modern branch identifies the causal effect of a contact's employment on a searcher's prospects (Cingano and Rosolia 2012; Glitz 2017) and the congestion that can attend shared networks (Calvó-Armengol and Jackson 2004; Beaman 2012). I take from it the elasticities that discipline the model's informal channel, and I contribute to it by providing an aggregate object: the sector-level stock those individual ties compose, and its cyclical dynamics. A smaller literature brings referral hiring into equilibrium search models of the business cycle (Galenianos 2014; Arbex, Caetano, and O'Dea 2016; Arbex, O'Dea, and Wiczer 2019); relative to it, the addition here is the stock itself, a state variable that slumps deplete, and its Pigouvian price.

A second literature documents scarring itself: the long shadows of displacement, graduating into a recession, and regional trade shocks (Davis and von Wachter 2011; Autor, Dorn, and Hanson 2013). Its explanations are largely individual (human capital, mobility costs, demand persistence), and where it finds scarring deepening in recessions, it reads the deepening as skills eroding faster when work is scarce. The network channel does not displace these mechanisms. In the model it compounds the human-capital one, and it offers an endogenous, between-worker source for the very recession-deepening the individual accounts build in by hand, through a shared stock that many who lose their jobs at once draw down together. It supplies what those accounts lack: a reason scarring is social, reaching workers never directly shocked through the connections of the displaced.

Third, the paper builds on heterogeneous-agent models of unemployment risk and aggregate demand (Ravn and Sterk 2017; den Haan, Rendahl, and Riegler 2018; Kaplan, Moll, and Violante 2018). I adopt their machinery (incomplete markets, search frictions, nominal rigidity) and add a state variable those models treat as constant: the efficiency of matching itself, made endogenous to the employment history of each sector's network.

Finally, a normative literature prices labor-market externalities, from Hosios (1990) to the sufficient-statistic tradition. The network externality priced here stacks on the classic congestion wedges rather than replacing them, and its price shares a statistic with the economy's fragility, which is what makes the policy analysis portable beyond this one model.

The paper proceeds as follows. Section 2 documents the four facts that motivate my approach. Section 3 builds the analytical framework and derives the externality, its price, and the loop-gain statistic. Section 4 embeds the mechanism in a quantitative model, and Section 5 takes the model to the data and validates it on untargetted moments. Section 6 conducts a China-shock experiment that answers the paper's first two questions; Section 7 studies the policy instruments that answers the third. Section 8 concludes.

2. Four Facts About Network Search over the Cycle

How much do workers lean on their networks when times are bad? The Current Population Survey has asked unemployed workers the same question since 1994: which methods have you used to look for work in the past four weeks? One of the response categories is "contacted friends or relatives." I assemble the answers from the raw monthly survey files (389 months from January 1994 through June 2026, 1.1 million unemployed workers actively searching) into a consistent panel of how search runs through the informal channel over the business cycle. Two features of this object should be kept in mind throughout. First, it is evidence of a search-method, not a hiring channel: the survey records that a worker used the channel, not that it produced a job. Second, the category is a proxy for broader measures. The model prices contact with employed connections of any kind, while the survey records use of the friends-or-relatives method, not the contact's employment status or a full range of weak ties. Appendix C documents the survey files, the construction, and every estimate in this section. Mukoyama, Patterson, and Şahin (2018) document how job-search effort and method use move over the business cycle in the same CPS data. The additions here are the stock signatures (the build-up lag and the sustained-versus-spike asymmetry) and the use of the share's slope as a disciplining moment.

Four facts emerge. Together they suggest that the informal channel behaves like a slowly adjusting stock rather than a static search technology. These discipline the model throughout the paper.

Fact 1: reliance on the network channel is strongly countercyclical. Figure 1 plots the share of unemployed searchers who contacted friends or relatives. The series falls to 13 percent at the peak of the late-1990s boom, climbs to 29 percent in the aftermath of the Great Recession, and sits there through 2013. The swing is wide enough to see in the raw series. Two questionnaire-vintage level breaks account for up to five points of the raw swing; every slope used for discipline below is break-robust. Regressing this share on the unemployment rate shows that nationally, a one percentage point rise in the unemployment rate raises the share by 0.88 percentage points, increasing to 1.10 when adding a linear time trend. Identification from cross-state variation gives a smaller slope of 0.39 to 0.47 once measurement error in monthly state unemployment rates is corrected, providing a more conservative bound.[2] The calibration targets the aggregate margin, on the maintained assumption that formal hiring channels tighten with the aggregate cycle rather than with state-level deviations from it; Appendix G reports the model under the conservative cross-state target instead. The slope is not a composition artifact of the duration distribution: Appendix C shows it little changed when reweighted to a fixed duration mix, and rising with aggregate unemployment within every duration cell.

Fact 2: the channel's relative effectiveness does not fall when reliance rises. A natural objection, going back to Beaman (2012), is that rising reliance could be undone by falling yield: more workers leaning on networks that work worse precisely when contacts are jobless. The CPS rotation structure allows for a direct test of this. I link each searcher to their labor-force status one month later and ask whether users of the friends-and-relatives method convert to employment at different rates over the cycle. They do, but in the opposite direction from the Beaman concern: the interaction between channel use and unemployment is positive (0.28 percentage points per point of unemployment), so network searchers' job-finding falls by less than everyone else's in a slump. Three boundaries define what this test identifies. First, it is relative. Every channel yields less in a slump, and the model's own eroding stock implies the informal channel is no exception. What the positive

[2] Monthly state unemployment rates computed from the CPS carry substantial sampling noise; a split-sample first stage puts their reliability at 0.38. The corrected local slope uses rotation-group instruments and state-year aggregation, which bracket 0.4. Appendix C reports every specification, including questionnaire-vintage break checks.

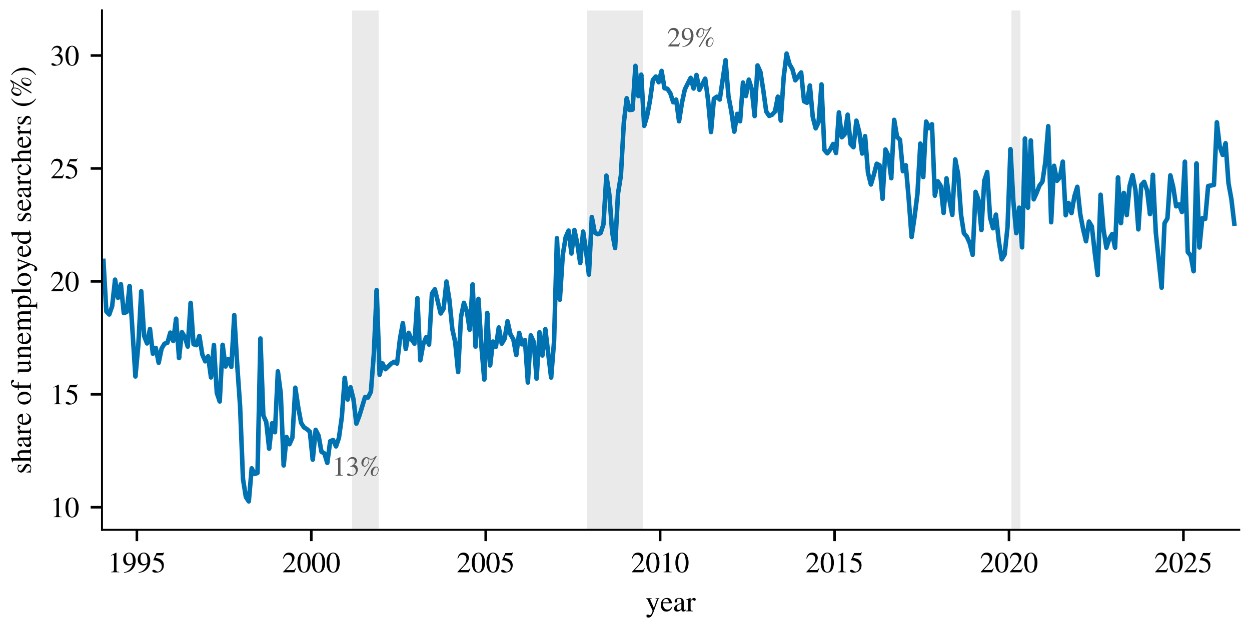

Figure 1: *Reliance on the informal search channel, 1994–2026.*
Notes: weighted share of unemployed active searchers whose reported search methods include "contacted friends or relatives," computed from the raw monthly CPS files. NBER recessions shaded.

interaction rules out is deterioration faster than formal search, the scenario in which rising reliance would be self-defeating. Second, the design cannot separate congestion within the channel from erosion of the network behind it, only bound their sum, which is the object the model prices. Finally, it is a reduced-form object. It is robust to observables and to search intensity, but cyclical selection on unobservables into who uses the channel cannot be excluded (Appendix C). The relevant result for my purposes is that the channel's share of realized transitions to employment rises from 21 percent in tight markets to 27 percent in slack ones. Rising reliance is reinforced by relative effectiveness, not offset by it.

Fact 3: reliance builds slowly. The share does not jump with the unemployment rate; it accumulates. In distributed-lag specifications (Figure 2a), unemployment enters at the contemporaneous, six-month, and twelve-month lags with coefficients of 0.80, 0.55, and an imprecisely estimated 0.47. Reliance keeps responding well after impact. A static congestion story, in which searchers substitute toward the informal channel the moment formal channels tighten, has no reason to generate this pattern. A slowly depreciating stock does. A third candidate is composition: sustained slumps build a pool of long-duration searchers with a different search mix. Appendix C checks it directly; at fixed composition the lag profile survives, present within duration cells and strongest among long spells.

Fact 4: sustained slumps move reliance; spikes do not. Estimating the slope separately

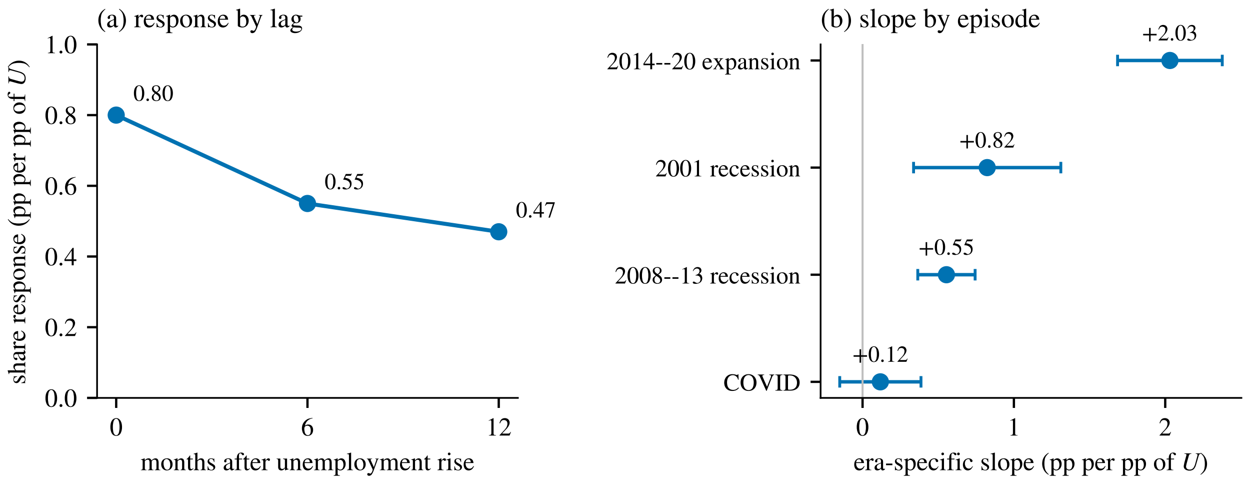

Figure 2: *The stock signatures: the reliance lag and the era ranking.*
Notes: panel (a): distributed-lag response of the share to the unemployment rate, trend-controlled (point estimates; standard errors in Appendix C). Panel (b): era-specific slopes from the era-interaction specification with a linear trend, 95% confidence intervals; the 1994–2000 expansion is excluded (1998 questionnaire redesign).

by episode (Figure 2b) delivers a clean ranking. In long low-unemployment expansions, where variation comes from the hiring margin, the slope is steepest (2.0 in the 2014–2020 expansion). In the recessions of 2001 and 2008–2013 it is close to one. In the COVID episode, a brief spike of temporary layoffs, it is indistinguishable from zero: unemployment quadrupled and reliance barely moved. An alternative reading the design cannot rule out is composition: COVID's temporary layoffs left searchers expecting recall rather than working their contacts. A placebo using the "sent out résumés" response cuts against that reading. This pure intensity margin rises even in the COVID spike (+0.44) while the friends channel does not (+0.03), and its lag profile dies out within a year while the friends channel's persists (Appendix C). Whatever the informal channel responds to, it is not the level of unemployment this month; it is how long the labor market has been slack.

These four facts represent the blueprint that the paper's mechanism must match. A stock of network capital that is built by employment, eroded by sustained unemployment, and too slow to react to a spike produces all four. Facts 3 and 4 are the signatures of its slow adjustment, and Section 5.4 will show that the calibrated model reproduces the sharpest of them out of sample—the era ranking of Fact 4, and the spike-versus-sustained asymmetry it implies. What cannot be established here is a causal effect of networks on job-finding at the individual level. These facts are statements about aggregates, and the model is disciplined by their slopes, not their levels.

3. A Theory of Network Capital

This section builds the theory of the stock behind Section 2's facts. I establish laws of motion for network capital, define the externality this object creates, and define a method for pricing it. The organizing question is what a job is worth to everyone other than the worker who holds it. The gap between the social and the private answer is the externality at the paper's core, and it has more structure than externalities usually do: a single statistic, the network loop gain, establishes the price of this inefficiency, determines the economy's amplification, and measures the distance to a self-reinforcing slump.

To begin, I work in the smallest environment that contains the mechanism: one sector, risk-neutral workers, one decision margin at a time. These choices are deliberate: risk neutrality separates the new wedge from the insurance motives the quantitative model restores, and a single margin makes it unambiguous what the market gets wrong. Formal proofs are collected in Appendix A.

3.1. Laws of Motion

Time is discrete and a period is a quarter. A unit continuum of workers discounts the future at β; a worker produces y when employed and b at home when unemployed. Two aggregate states describe the labor market: the employment rate e, and network capital $K \in [0,1]$, the density of active ties between the sector's workers and their contacts (whose current employment enters separately, below), normalized so that a healthy steady state sits at (E°, K°). An unemployed worker searching with effort a finds a job with probability $a \cdot f(e,K)$, where

$$f(e,K) \;=\; (1-\chi)\,\bar{f} \;+\; \chi\,\bar{f}\,\frac{K}{K^\circ}\Big(\frac{e}{E^\circ}\Big)^{\eta}. \tag{1}$$

A share $1-\chi$ of job-finding runs through formal channels (postings, applications, agencies), held acyclical in this benchmark. The remaining share χ runs through the network, and its productivity has two arguments: the stock of ties K, and the employment of the people at the other end of them, with elasticity η, because a useful contact must both exist and be in a position to help. The parameter to watch is χ, the share of job-finding that touches the network at all.

The aggregate states update recursively (separations first, then matching, then the stock):

$$e' \;=\; (1-s)\,e + a\,(1-e)\,f(e,K), \tag{2}$$

$$K' \;=\; K + \varphi\,e\,(1-K) \;-\; \delta\,(1-e)\,K\,\big[\,1+\kappa\,(1-e/E^\circ)\,\big]. \tag{3}$$

Equation (3) is the model's main mechanism, a coevolution of ties and employment status as in Bramoullé and Saint-Paul (2010). Employed workers build and refresh ties at rate φ (ties form most readily among the employed), with less room to build the denser the network already is. The ties of the unemployed decay at rate δ, and they decay faster (by a factor rising in κ) when the surrounding market is slack, because a tie to a contact who cannot help is harder to keep alive. Employment builds the stock, unemployment erodes it, and slack accelerates the erosion.

To a macroeconomist, the informal term in (1) is a matching-efficiency wedge. Business-cycle accounting (Chari, Kehoe, and McGrattan 2007) recovers such wedges as residuals; here the wedge is structural, moving with the network stock K, which evolves according to (3). Hold the informal term at its steady-state level (the stock K at K° and its contacts' employment at E°) and the economy is a standard fixed-efficiency search model; let the term move, and matching efficiency becomes a slow-moving function of the sector's own employment history. Counterfactuals in the later quantitative estimates make this comparison: the same economy and the same shock, with the network channel live and with it frozen.

Workers are atomistic: each takes the paths of (e, K) as given, and equations (1)–(3) embed three feedbacks that no individual internalizes. A worker's employment raises the job-finding of those who search through their ties this period: the information spillover. A graph-level version of this channel is modeled by Calvó-Armengol and Jackson (2004). The same employment builds K, which persists: the capital spillover. It slows the decay of everyone else's ties, through the κ bracket: the protection spillover. All three run in the same direction.

One further benchmark choice: in the worker-margin economy the wage equals full match output, $w = y$, so the worker appropriates the entire bilateral surplus. Nothing below is about how the surplus of a match is split.

3.2. The Externality

I give the unemployed worker a single decision, search effort a, at convex cost $c(a)$. A worker weighs the private value of employment, $\Delta \equiv W - U$, and searches until $c'(a) = \beta f \Delta$. The planner chooses the effort path under the same laws of motion and the same meeting structure (no power to rewire who knows whom, no transfers of K) and weighs instead the social value of an employed worker, λ, which travels with a companion costate ν, the shadow value of network capital. The planner's condition is $c'(a) = \beta f \lambda$. Everything turns on λ against Δ, and subtracting their recursions gives the gap in closed form:

$$\Omega \equiv \lambda - \Delta = \Delta \cdot \frac{L}{1-L}, \tag{4}$$

$$L = \frac{\beta\, a(1-e)\, f_e + \beta^2 J_{eK} J_{Ke} / (1 - \beta J_{KK})}{1 - \beta(1 - s - af)}. \tag{5}$$

The loop gain L collects everything the worker leaves out, in two pieces. The first term in the numerator is the information spillover: employment raises others' finding within the period, at rate f_e, the derivative of (1) with respect to contacts' employment. The second is the capital loop, a product of three links: employment builds the stock (J_{Ke}, with the κ-protection term inside it), the stock persists (J_{KK}), and the stock raises finding tomorrow (J_{eK}). The denominator capitalizes the flow into the value of a job. Closed forms for every entry are in Appendix A.

Theorem 1 (constrained inefficiency)**.** *Under regularity conditions stated in Appendix A (non-negative elasticities; quarterly turnover below one; discounted spectral radius of the transition below one): (i) at any interior laissez-faire steady state, a marginal uniform increase in search effort strictly raises welfare if and only if $\chi > 0$; if $\chi = 0$ the equilibrium is constrained efficient. (ii) The sign is unconditional: each of the three spillover channels enters L non-negatively, so $\lambda > \Delta$ strictly whenever $\chi > 0$. (iii) Under a convexity condition on c ensuring unique fixed points, the planner's steady state has strictly higher effort, employment, and network capital than the market's.*

Two features distinguish this wedge from the congestion externalities the search literature already prices. First, it cannot change sign within the model class. The classic search externality nets thick-market gains against congestion losses, and its direction turns on bargaining weights relative to matching elasticities (Hosios 1990); inside the network channel there is no offsetting congestion term, and the market under-provides. Second, the wedge is not appropriability in disguise. The benchmark hands the worker the entire match surplus and the wedge survives, because it falls on parties who are not at the table: the searchers who would have used the worker's tie, and the sector whose stock their employment would have built.

The stock component of the externality requires decay. With $\delta = 0$ (ties, once formed, never decay) employment can build nothing scarce: the stock saturates and the capital loop vanishes, though the contemporaneous information spillover through f_e survives whenever χ and η are positive. The same fragility that makes slumps scar is what makes the depreciable part of the market failure large.

Neither the decision margin nor the number of sectors is essential to the result. On the firm side, replace worker effort with vacancy posting:

Proposition 1 (the wedge stacks on Hosios)**.** *In the vacancy-posting variant (a matching function for formal meetings, free entry at a flow cost, referrals as in (1)), let w^* be the appropriability-neutral wage at which entry in the $\chi = 0$ economy is constrained efficient. With $\chi > 0$, equilibrium entry at w^* is strictly too low, and the correcting instrument is a hiring subsidy $\tau_h > 0$ per filled job, vanishing as $\chi \to 0$.*

Firms do not internalize that the workers they hire build the referral infrastructure that fills everyone's future vacancies, and no wage, the Hosios wage included, can price a spillover whose beneficiaries sit outside the match. The wedge is margin-portable: whichever margin the economy possesses, worker effort or firm entry, the same Ω prices its distortion. The instrument does not share this property. A worker-effort economy prices the wedge with a search subsidy, a firm-entry economy with the hiring subsidy τ_h, and each subsidy falls on the margin it corrects.

As for sectors: with S of them connected by a mixing matrix m (a searcher in sector j meets employment $\tilde{e}_j = \sum_k m_{jk}\,\bar{e}_k$), the same argument runs through the $2S$-dimensional system, and every connected sector carries a strictly positive wedge whenever $\chi > 0$ (Appendix A). χ decides whether the externality exists; the mixing matrix decides where it lands. Incidence is Section 6's subject.

3.3. One Statistic, Three Roles

The loop gain serves as more than a bookkeeping device. The social value of a job can be written as its flow value times $(1 - \beta J_{KK})/\det(I - \beta J)$, and the determinant factors (Appendix A) as

$$\det(I - \beta J) \;=\; \big(1 - \beta(1 - s - af)\big)\,\big(1 - \beta J_{KK}\big)\,\big(1 - L\big), \tag{6}$$

so the wedge diverges exactly where the costate system loses invertibility. The dynamics (2)–(3), meanwhile, lose steady-state uniqueness at a fold bifurcation, $\det(I - J) = 0$. A second, depressed steady state arises. As $\beta \to 1$ the two conditions become the same equation.

Proposition 2 (one statistic, three roles)**.** *The loop gain L simultaneously (i) prices the externality, $\Omega/\Delta = L/(1 - L)$; (ii) sets the amplification with which the economy propagates employment shocks, $1/(1 - L)$; and (iii) measures distance to the fold: the wedge's pole and the loss of steady-state uniqueness coincide in the patient limit, at $L = 1$.*

Figure 3 draws this relationship. The wedge $L/(1 - L)$ and the amplification factor $1/(1 - L)$ climb the same axis and diverge at the same point, the fold at $L = 1$ (the

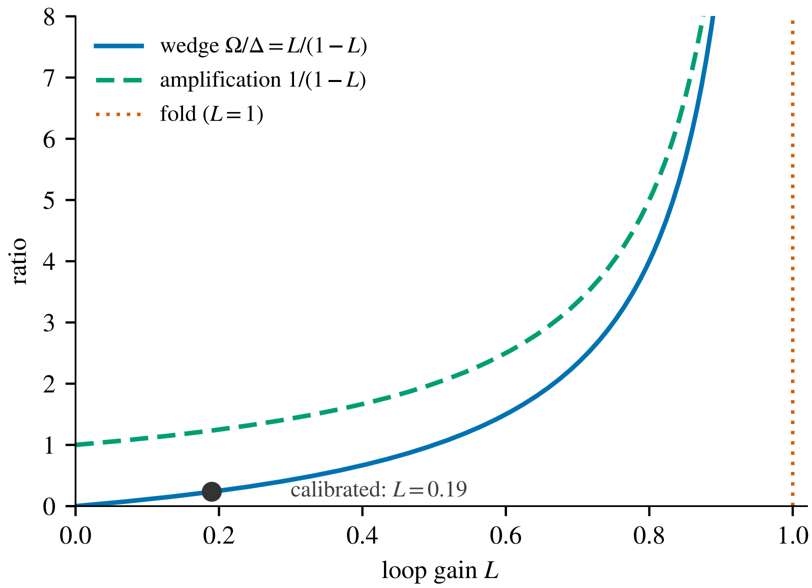

Figure 3: *The loop gain L: one statistic, three roles.*
Notes: the wedge $\Omega/\Delta = L/(1-L)$ and the amplification factor $1/(1-L)$ diverge together at the fold $L = 1$, where the steady state loses uniqueness (Proposition 2). The marked point anticipates Section 5's calibration.

coincidence exact in the patient limit). An economy can sit far from the fold and still carry a substantial wedge (inefficiency does not require fragility), but inefficiency and fragility are two readings of the same statistic, the loop gain L, which is what makes the policy analysis that follows portable beyond the present application.

The statistic also composes across channels. When a second feedback runs through the same two states, its loop gain adds to the network's, and L is their sum. The benchmark held the formal channel acyclical. Allowing for procyclicality (formal hiring that dries up with aggregate employment) makes that feedback enter f_e alongside the network terms. A channel that carries states of its own, such as the precautionary-demand loop of Ravn and Sterk (2017) in which unemployment risk depresses spending and hence hiring, composes only to first order, under an explicit accounting convention. Either way the network moves an economy toward whatever fragility it already has, by its own contribution to the loop gain, and Section 6.4 tracks that contribution along a calibrated employment collapse.

3.4. The Bridge Margin

With more than one sector, searchers face a margin the one-sector model lacks: where to search. Let there be two sectors, one declining (a permanently higher separation rate), and let sector j's unemployed direct a share ω_j of their search toward the other sector, at convex flow cost. Sectoral bridging does two things. The searcher samples the other sector's referral productivity, $g_{j'} = (K_{j'}/K^{\circ})(e_{j'}/E^{\circ})^{\eta}$, in place of their own; and their network exposure, $\tilde{e}_j = (1-\omega_j)\, e_j + \omega_j\, e_{j'}$, replaces e_j in the decay bracket of (3), so exposure to a healthier market slows the decay of ties. In equilibrium the declining sector bridges until marginal cost equals the private gain, and the healthy sector has no incentive to bridge at all.

Theorem 2 (under-bridging)**.** *At the asymmetric steady state with an interior equilibrium bridge choice, the marginal social value of bridging out of the declining sector (sector 1) is strictly positive:*

$$\frac{dW}{d\omega_1}\Big|_{CE} \propto \underbrace{\beta\chi\bar{f}\,(g_2-g_1)\,(\lambda_1-\Delta_1)}_{\textit{finding chain}} + \underbrace{\beta\,\nu_1\,\delta K_1\kappa\,(e_2-e_1)/E^{\circ}}_{\textit{capital protection}} > 0. \tag{7}$$

The planner leaves the healthy sector at its corner, $\omega_2 = 0$, exactly as the market does; and the expression vanishes identically at the symmetric steady state, growing with sectoral divergence.

The two terms are the gains that accrue to others. The first is a finding chain: bridging raises the bridger's re-employment probability, and by Theorem 1 each re-employment is worth $\lambda_1 - \Delta_1$ more to the sector than to the worker. The second is capital protection, as cross-sector exposure slows the decay of the entire home sector's ties through κ. Note what the expression omits. The private diversification gain (reaching toward the healthier market because it is healthier) is priced in full by the worker and earns no subsidy; the correction pays only for the parts that fall on others. It is one-sided, paid to the depressed side alone, and state-contingent, zero in symmetric times and growing as sectors diverge.[3]

Bridges also change where hiring subsidies should land. Once the depressed sector's searchers reach across, the healthy sector's jobs feed both sectors' job-finding, and the employment wedge migrates toward the hub: bridge policy and hiring policy are complements, to be set jointly. Section 7.4 examines the decentralization in general equilibrium.

3.5. From Wedges to Instruments

I now map the wedge directly into the policy instruments that close it.

Proposition 3 (the Pigouvian credit)**.** *A subsidy priced at the wedge (a per-worker flow with present value Ω, or its hiring-credit equivalent converted at the effective discount rate) implements*

[3] The bridge margin is directed (a choice of where to search), so no reciprocity externality arises by construction. If ties are instead two-sided, link formation adds a Jackson and Wolinsky (1996)-type externality on top, which bargaining between the two ends of a tie could in principle internalize. Both components of (7) fall on parties to no bilateral bargain (the bridger's entire home sector), so under-bridging is robust to the tie-formation protocol.

the planner allocation exactly, on whichever margin the economy possesses: paid on the worker side it corrects search effort; paid to firms on top of the appropriability-neutral wage it corrects entry.

The general-equilibrium model of Section 4 uses the entry-margin version, the credit τ_h per hire, because its labor market is built on free entry. The theory also fixes the instrument's shape. The wedge is a function of the state, not a constant, so the natural implementation is not a permanent subsidy but a standing rule: a credit that switches on when the wedge is elevated and off in normal times. Section 6 shows the wedge peaks early in a collapse, before the stock is depleted, and Section 7 implements exactly this rule.

Proposition 4 (UI and the targeting principle)**.** *A marginal increase in unemployment benefits, lump-sum financed, has a first-order welfare cost equal to the wedge times the induced reduction in search; the private moral-hazard distortion is second-order by the envelope theorem, and the transfer itself washes out. With the credit of Proposition 3 in place, the first-order cost is zero.*

The entire first-order efficiency cost of marginal UI is the network externality (each claimant's reduced search deprives their neighbors of future referral capital), and that cost vanishes once the externality is priced. UI and the credit are therefore not substitutes on a single margin. The credit prices the cause and leaves UI to insure the symptom, a role with real value once households are risk-averse and imperfectly insured (Baily 1978; Chetty 2006). Section 7.2 makes the comparison at equal fiscal cost.

The costate ν prices a broader family of interventions than the paper models. Any technology acting on the capital law (job clubs, alumni networks, outplacement that maintains a displaced worker's ties) buys K at some cost, and is worth buying whenever that cost is below ν. The market prices none of it: the stock is unowned and non-excludable, so no private actor collects the return that ν measures.

One instrument is deliberately absent: a subsidy to search effort itself. Its sign turns on a question the aggregate laws (1)–(3) do not answer: whether referral information is rival at the level of a tie, one referral crowding another out of a queue, or broadcast, helping whoever hears it. The stock instruments above are robust to that distinction; the effort instrument is not. Appendix B derives the condition its sign requires, and I leave it out of the policy analysis on that ground.

What this theory delivers is an externality signed by one parameter (Theorem 1), priced by one statistic (Proposition 2), and corrected by instruments with derived shapes (a state-contingent hiring credit and a one-sided bridge subsidy), with a clean division of labor left for unemployment insurance (Proposition 4). What it cannot deliver is size. How large is L in the United States, and how much of a real collapse does the loop explain? Those questions require a full quantitative model, which I develop in the following section.

4. A Quantitative Model of Network Scarring

The theory delivered structure, but to understand the potential scale of network effects requires a quantitative model. This section embeds the network block of Section 3 in a heterogeneous-agent New Keynesian model of the U.S. labor market with search frictions and human-capital scarring. It is a standard HANK economy with matching frictions, in the tradition of Ravn and Sterk (2017) and Kaplan, Moll, and Violante (2018), extended by one addition, a vector of sector-level network stocks K_j that enter job-finding.

4.1. Households

The economy has S sectors. A household belongs to a sector, holds liquid assets k subject to a borrowing limit, receives idiosyncratic earnings-ability shocks, and is either employed or unemployed. Preferences are CRRA over consumption; markets are incomplete, so unemployment risk is imperfectly insured.

Each worker also carries human capital h: it rises stochastically while employed, depreciates during unemployment, and lifts job-finding at elasticity ζ. This is the standard human-capital channel of the scarring literature, the route by which past unemployment lowers present re-employment; Davis and von Wachter (2011) document the persistent losses it captures. It is in the model on purpose: it is the incumbent account of the persistence the network is meant to explain, so carrying it lets Section 6 weigh the network against a familiar mechanism rather than against nothing. Its form is the literature's plainest, a constant depreciation rate, which leaves any deepening of the scar in a slump to come from the network's stock rather than from the worker's own skills eroding faster. Because the finding rate depends on h, employment status and human capital cannot be tracked as independent dimensions; they are fused into a single exogenous state whose transition matrix the labor market block rebuilds each period. Sector membership joins that fused state (the network makes sector, like h, a determinant of the finding rate), so a household's exogenous state is the triple (employment status, h, sector) alongside its earnings ability, and its distribution is defined over that state and asset holdings. The Bellman problem and the transition-matrix construction are standard given the finding rates and are collected in Appendix A; what the network changes is only where those finding rates come from.

4.2. The Labor Market

Hiring runs through two channels, composed additively as in the theory. The formal channel is a standard matching function: firms post vacancies at a flow cost, aggregate matches form over effective searchers and vacancies, and free entry pins market tightness θ, giving every searcher a direct finding rate $f_{dir}(\theta)$. The informal channel delivers f_{ref} per unit of the network multiplier defined below. The sector finding rate is

$$f_j \;=\; f_{dir}(\theta) \;+\; f_{ref} \cdot \underbrace{\frac{K_j(-1)}{K^\circ}\left(\frac{\tilde{e}_j}{E^\circ}\right)^{\eta}}_{\text{the network multiplier}}, \tag{8}$$

where (-1) denotes a variable's previous-period value, and an individual searcher in sector j with human capital h finds work at rate $f_j \cdot (h/h_{ref})^{\zeta}$, the standard finding gradient in human capital. The scale f_{ref} is set so that the informal channel's share of hires is χ at the healthy steady state. Equation (8) is Section 3's finding function with its two halves made endogenous: the acyclical formal term is now priced by free entry, and the network term now moves with each sector's own stock and exposure. Setting $\chi = 0$ (equivalently $f_{ref} = 0$) deletes the second term, and with it every trace of K from the model:

$$f_j = f_{dir}(\theta) \qquad \text{for all } j. \tag{9}$$

What remains is exactly the standard economy.

Wages follow a real-wage rule rather than period-by-period Nash bargaining: $w = w_{ss}\,(mc/mc_{ss})^{\gamma_w}$, where mc is real marginal cost. Section 5.2 shows this is not a taste choice (the paper's disciplining moment rejects the flexible-wage model on sign and identifies γ_w), and the rule is a standard specification for it. Separations are exogenous at rate s_j, the margin the experiments shock.

4.3. The Network Block

Each sector has a stock of network capital K_j evolving exactly as equation (3):

$$K_j' \;=\; K_j + \varphi\, e_j\,(1-K_j) \;-\; \delta\,(1-e_j)\,K_j\left[\,1+\kappa\,(1-\tilde{e}_j/E^\circ)\,\right], \tag{10}$$

with one object promoted from theory to measurement: the exposure $\tilde{e}_j = \sum_k m_{jk}\,\bar{e}_k(-1)$ now runs over the empirical mixing kernel m, the matrix (estimated from job-to-job flows in Section 5.1) of whose employment a sector's searchers actually meet. The kernel is row-stochastic, so the healthy steady state is symmetric and costless; its off-diagonal mass is what transmits one sector's collapse into its neighbors' finding rates. Finding rates depend only on last period's stocks and employment, so the network block is purely backward-looking in the sequence-space solution. The stock is sector-level and mean-field: each searcher faces sector averages, not specific contacts. Appendix B micro-founds the laws at the tie level and shows the aggregation is nearly exact; its one systematic bias actually understates how concentrated the scarring is. Section 6 attaches fixed-graph correction bands to the headline numbers.

4.4. Nominal Closure and Fiscal Policy

The supply side is a standard New Keynesian block. Wholesale firms produce with labor and pay the wage bill, retailers set prices subject to nominal rigidity generating a Phillips curve with slope κ_p, and monetary policy follows a Taylor rule with coefficient ϕ_π; the real rate follows from the Fisher equation. Profits are rebated to households lump-sum in the baseline; Appendix F quantifies the convention's incidence consequences against a deviation-payout alternative. The welfare gains of Section 7 are reported as ranges across the two. The government runs unemployment insurance at a fixed replacement rate inside a budget with a debt-feedback tax rule; the UI expansions and hiring credits of Section 7 enter through this budget and, for the credit, through the free-entry condition directly.

4.5. Equilibrium and Computation

An equilibrium is a path of prices, tightness, sector employment rates, and network stocks such that households optimize, free entry holds, each sector's employment flow is consistent with its finding and separation rates, goods and bond markets clear, and the fiscal rule is satisfied. The model is solved in the sequence space (Auclert, Bardóczy, Rognlie, and Straub 2021), with tightness, the sector employment paths, and inflation as unknowns, the sector-consistency and free-entry conditions as targets, and the bond market serving as the untargeted audit of Walras's law. Appendix E shows that for every experiment the residual stays under a few parts in a hundred thousand of the debt stock. Deep slumps drive market tightness toward zero, the corner where posting a vacancy stops paying. While this corner is an economic boundary, it creates numerical challenges that the solver addresses in two ways: it works in log tightness, so iterations cannot cross zero, and it approaches deep shocks gradually, each solved shock starting the next. Neither device changes the answer, only the route to it. Appendix E adds grids, accuracy checks, and the computational cost of the sector-fused household state.

The machine, then, is the standard economy plus equations (8) and (10) and a matrix measured from data, with (9) guaranteeing the ablation is exact. Section 5 describes the calibration of model parameters.

5. Taking the Model to the Data

Parameters the literature already pins down (separation rates, steady-state unemployment, the return on liquid assets, the slope of the Phillips curve) take their standard values or hit their standard targets. The network block's levels (χ and η) come from the microeconomic referral literature, points bracketed by bands; the tie-formation rate φ is backed out to hit the normalized steady-state stock; the decay acceleration κ has no moment to discipline it, so results are checked across a range; and the mixing kernel is estimated from job-to-job flows. That leaves one quantity: how much search shifts toward the network as unemployment rises. This is the slope Section 2 measured, and it disciplines exactly one parameter, the wage-rigidity coefficient γ_w. Section 5.4 then reports the model's predictions for untargeted moments.

5.1. Calibration

Table 1 lists every parameter, its value, whether it is set internally, and its target or source. Three groups deserve further comment.

The first group is the network block's levels. The share of hires that arrive through the informal channel, χ, is set to 0.45, the center of a large literature spanning countries and data types: roughly 35 to 60 percent of jobs are found through contacts. Inside the model, the scale f_{ref} is set so that the informal share of hires at the healthy steady state equals χ.[4] The mapping is an upper bound for the modeled channel, since the hires-share literature counts jobs found through contacts of any kind, including some whose value does not ride on a depletable stock. The productivity of the channel comes from Glitz (2017), whose natural experiment moves the employment rate of a worker's former coworkers and finds that 10 percentage points of contact employment buy 7.5 percentage points of re-employment probability, and no wage effect, which is what separates a network channel from a human-capital one. That elasticity implies the curvature $\eta = 2.86$ at the calibrated point; its transfer from German coworker networks to U.S. cross-sector referrals is an identifying assumption, and the band below brackets it from the linear side. Because the effectiveness evidence brackets rather than pins the curvature, every result in Sections 6 and 7 is reported across the band $\eta \in [1, 2.86]$, from linear to the Glitz-implied curvature. The band matters little

[4] Granovetter (1974); Ioannides and Loury (2004); Pellizzari (2010); Burks, Cowgill, Hoffman, and Housman (2015). The cross-country dispersion is real, and 0.45 sits at the U.S.-consistent center. This hires-share level is a different, higher object than the CPS search-method share (roughly 23 percent pooled): the first counts jobs found through contacts, the second the fraction of searchers who use the method. The two are assigned different jobs. The 45 percent disciplines the model's level; the CPS enters only through its cyclical slopes. The χ sweep of Appendix F runs downward, through 0.35, because the model has no solution for values much above 0.5.

Parameter	Value	Internal?	Target or source
Households			
Risk aversion σ	2	no	standard
Discount factor β	internal	yes	bond clearing at $r = 2\%/\mathrm{yr}$
UI replacement rate	0.40	no	standard
HC depreciation (unemp.) δ_h	0.02	no	scarring lit.
HC accumulation (emp.) φ_h	0.01	no	scarring lit.
Finding–HC elasticity ζ	0.5	no	scarring lit.
Labor market			
Separation rate s	0.035	no	JOLTS/CPS flows
Matching efficiency	internal	yes	$U_{ss} = 6\%$
Entry cost κ_v	internal	yes	$\theta_{ss} = 1$ (normalization)
Wage rigidity γ_w	0.2	yes	CPS share slope (Section 5.2)
Nominal block			
Phillips slope κ_p	0.04	no	micro-disciplined, standard
Taylor coefficient ϕ_π	1.5	no	standard
Markup μ	1.2	no	standard
Network block			
Informal share of hires χ	0.45	no	hires-share literature (fn. 2)
Referral curvature η	2.86	no	Glitz (2017); band $[1, 2.86]$
Tie decay δ	half-life 5.4q	no	tie-formation literature (range 1–3y)
Tie build φ	backed out	yes	(E°, K°) normalization
Decay acceleration κ	1.0	no	assigned baseline; swept (App. G)
Mixing kernel m	Fig. 4	no	Census J2J flows
Sectors S	9	no	J2J supersectors

Table 1: *Parameters.*

Notes: Internal? = set inside the model to hit the stated target (yes) versus taken externally (no). σ is the CRRA coefficient (EIS 0.5). Human-capital h lies on a grid $[0.1, 2.0]$; grids, the h reference level, realized values of the internal parameters, and computation are in Appendix E. The three human-capital parameters (δ_h, φ_h, ζ) take standard, conservative values from the scarring literature (Davis and von Wachter 2011). The mixing kernel m is the 9×9 matrix of Figure 4.

for the sectoral headline, which barely moves across it. The aggregate multiple is more sensitive (1.28 to 1.42), and the bridge-partner premium results in Section 6.3 depend on η outright, roughly doubling across the band.

The decay side of the stock is where published evidence is weakest. Tie decay during unemployment, δ, is set to a half-life of five and a half quarters, inside the one-to-three-year range the tie-formation literature supports qualitatively. The build rate φ is not a free parameter, but is backed out so that the healthy steady state reproduces the reference stock. The decay-acceleration parameter κ has no measured counterpart and is assigned a baseline of one. As Section 5.2 shows, the disciplining moment is flat in κ over the checked range, so I treat it as unidentified by this moment; Appendix G reports results with κ varied from zero to three.

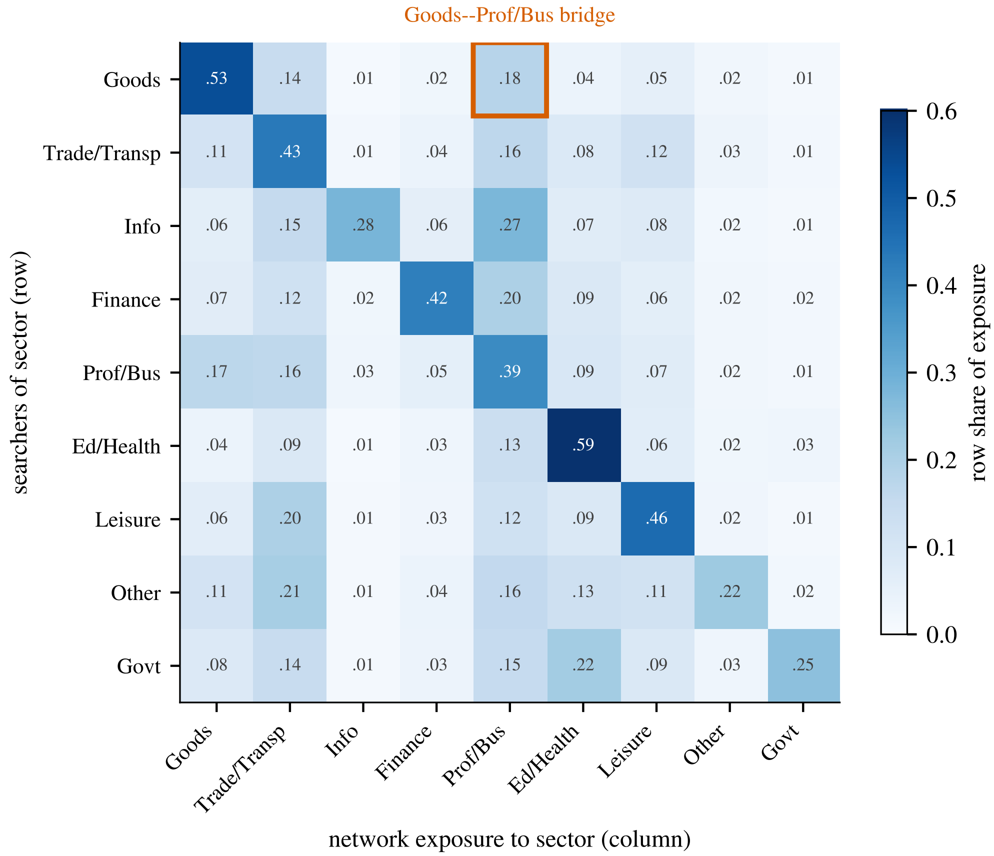

Figure 4: *The empirical mixing kernel.*
Notes: row-normalized network exposure across the nine J2J supersectors; row j, column k is the weight of sector k's employment in sector j's searchers' network exposure $\tilde{e}_j$. The outlined cell marks the Goods–Professional/Business bridge that carries Section 6.3's analysis. Construction and BEA cross-check in Appendix D.

The second group is the mixing geography. Sectors meet each other's networks through the kernel m, estimated from Census Job-to-Job flows across nine broad sectors and cross-checked against BEA input–output linkages (Appendix D documents the construction and the flows-to-ties assumption). Figure 4 plots this mapping. The diagonal dominates (most network exposure is own-sector), but the off-diagonal structure is the economics: goods production's strongest bridge is professional and business services, the economy's hub sector, a link that will carry the correlated-collapse analysis of Section 6.3. The kernel is used in mean-field form; fixed-graph (quenched) benchmarks put the mean-field approximation's error at five to twenty percent in the conservative direction.

The third group is the nominal side. The Phillips-curve slope $\kappa_p = 0.04$ and the Taylor coefficient are standard; the wage rule is not standard boilerplate here, and it gets its own subsection, because the paper's one disciplining moment turns out to identify it.

5.2. The Disciplining Moment

The moment is Section 2's aggregate slope: a one-point rise in unemployment raises the informal-search share by 0.88 to 1.10 percentage points. Its model counterpart is the referral share of job-finding, aggregated across sectors and regressed on aggregate unemployment along a demand transition.

With flexible Nash wages, the model gets the sign wrong: in a slump the share falls. The mechanics are the Shimer (2005) logic, now with a second channel attached. Under free entry with flexible wages, a slump barely moves the formal channel (wages absorb the shock and the fundamental surplus protects entry) while the network channel falls hard: contacts' employment enters the referral term raised to η. Referral finding falls faster than formal finding, and the share moves the wrong way. What the data demand is a formal channel that collapses in slumps, which is to say: wage rigidity. I adopt a standard real-wage rule, $w = w_{ss}(mc/mc_{ss})^{\gamma_w}$, and the share moment then identifies γ_w, a parameter the unemployment-volatility literature has struggled to pin down, here pinned by a cyclical search-behavior moment that has, to my knowledge, not been used this way before. The map from γ_w to the model slope is hump-shaped, and only γ_w near 0.2 reaches the data band, with 0.2 a standard new-hire rigidity value (a dense sweep puts the crossing on the single interval $[0.17, 0.33]$; the hump-shaped relationship is mapped out in Appendix F). Targeting instead the corrected cross-state slope of Section 2 selects $\gamma_w = 0.87$ and moves the scar multiples of Section 6 by at most 0.02 (Appendix G).

The anchor experiment matches the identification to the variation that produced the data. The CPS band is identified off sustained recessions, and the model's counterpart is a persistent separation slump at the scale of 2008–09, when JOLTS layoff and discharge rates rose by roughly sixty percent. On that experiment the model delivers a slope of $+0.96$, inside the band, with no other parameter moved. The curvature $\eta = 2.86$ was fixed by the Glitz (2017) microdata, κ was fixed at its prior value and turns out to be irrelevant to the moment, and the hump map means γ_w could not have bought the fit anywhere else. Solved nonlinearly at increasing depth, the moment rises above the band, to $+1.22$ at a 2.4-point unemployment rise and $+1.37$ at 3.7 points. I report the linear in-band value, a conservative reading of the model's slope. Targetted moments are reported in Table 2.

Section 3 closed by asking how large L is in the United States. The disciplined steady state provides a value of $L = 0.19$, which is far from the fold at one, and remains squarely in the wedge-bearing region of Figure 3. Sections 6 and 7 will show that this value (and the implied wedge) has important implications for amplification and for policy.

Moment	Data	Model
Steady-state unemployment (%)	6.0	6.0
Informal share of hires	0.45	0.447
Real rate (%/yr)	2.0	2.0
CPS share slope (pp per pp of U)	$+0.88$ to $+1.10$	$+0.96$

Table 2: *Targeted moments.*
Notes: the share slope is the separation-anchor experiment at the 2008–09 JOLTS scale, solved linearly at a 4pp unemployment rise; nonlinear solves at increasing depth give $+1.22$ to $+1.37$ at 2.4–3.7pp, above the band.

5.3. Why a Separation Shock Anchors the Moment

A natural alternative might be to ground this moment on a demand recession, but the model itself rules it out. At the disciplined parameters, demand levers cannot produce a deep slump: pushing the monetary shock to the edge of what the free-entry margin supports (market tightness driven nearly to zero, the economy almost out of vacancy-posting incentive) moves unemployment by only about two percentage points, and the patience lever caps in the same place. A 2008-scale demand-only slump does not exist in this model; the limit belongs to the small-surplus free-entry block, not to the network block. The data point in the same direction: the recessions in the identifying sample opened with a large layoff wave (the 2008–09 JOLTS spike), and while much of the subsequent unemployment variation ran through job-finding (Elsby, Michaels, and Solon 2009; Shimer 2012), that is the margin the model cannot deepen; the anchor experiment matches the data's variation on the margin the model can actually reach. On the demand margin at small scales the model's share response is several times too steep per unit of unemployment, a hypersensitivity recorded as a model property; its one testable implication is the era ranking of share slopes that Section 5.4 checks.

5.4. Untargeted Validation

The share moment disciplined the last free parameter. This subsection now reports model predictions that remain untargetted. Table 3 compares two of them with the data; a third shows that the timing of the mechanism is consistent with the reliance lag of Fact 3.

The first is an overidentification reading of the targeted moment itself: a joint test, not an out-of-sample prediction. The test has content because the map from γ_w to the model slope is hump-shaped, so nothing guaranteed that any value of the one free parameter would reach the data band. That the model lands inside it at the Glitz curvature provides some evidence for the network block's levels and the wage rule together.

	Data	Model
(a) Overidentification		
Aggregate share slope	$+0.88$ to $+1.10$	$+0.96$ (η, κ untouched)
(b) Era ranking (slope, 95% CI | experiment type)		
2014–20 expansion	$+2.03$ (± 0.35)	hiring-margin: several times the data slope
2001 recession	$+0.82$ (± 0.49)	sustained separation: $+0.96$
2008–13 recession	$+0.55$ (± 0.19)	
COVID	$+0.12$ (± 0.27)	8q spike: $+0.72$ (does not reach the data's zero)

Table 3: *Validation and overidentification.*

Notes: panel (b) data slopes from the era-interaction specification with a linear trend (Appendix C); parentheses are 95% confidence-interval half-widths. The model column reports the experiment whose variation matches the era: the anchor experiment's linear slope for the two sustained recessions (one experiment, hence one entry), the small-scale demand experiment for the expansion (its slope is scale-dependent and several times the data's; Section 5.3), and the eight-quarter spike of Figure 5 for COVID. The ordering is the model's prediction; the spacing is not matched (see text). The 1994–2000 expansion is excluded: the 1998 questionnaire redesign contaminates its slope (Appendix C).

The second is the era ranking. In the model, what moves the informal share is not unemployment itself but the margin the variation arrives on. The hiring-margin variation moves the share steeply, sustained separation slumps move it roughly one-for-one, and a fast spike barely engages the stock before healing. The CPS delivers exactly this ordering, estimated era by era. The clean 2014–2020 expansion is steepest at $+2.03$, while the two sustained recessions sit near one, and COVID (the largest, fastest spike in the sample) sits far below the sustained-slump slopes. The 2005–07 window lacks within-era variation and is uninformative; Appendix C documents its exclusion. The ordering is the model's prediction, and the data deliver it, though the spacing is not matched. The model's hiring-margin response overshoots the expansion slope, and its spike response does not fall all the way to COVID's zero. The table shows both.

The reliance lag of Fact 3 points the same way. In the data, reliance builds for six to twelve months after unemployment rises rather than jumping on impact. The model's stock operates on exactly that timescale: the tie-decay half-life, calibrated independently from the tie-formation literature, is about five and a half quarters, so the network capital a slump destroys is rebuilt over years, not weeks. What the model does not deliver is a clean quantitative match to the data's distributed-lag coefficients, because the informal-search *share* is a ratio that also moves with the formal channel, which collapses fast under wage rigidity; the slow-stock signature the data's lag reflects shows up in the model instead as the spike-versus-sustained asymmetry, to which I now turn.

Figure 5 closes the section with the sharpest version of the stock signature: the spike-versus-sustained asymmetry, run as a model experiment. An eight-quarter spike in separa-

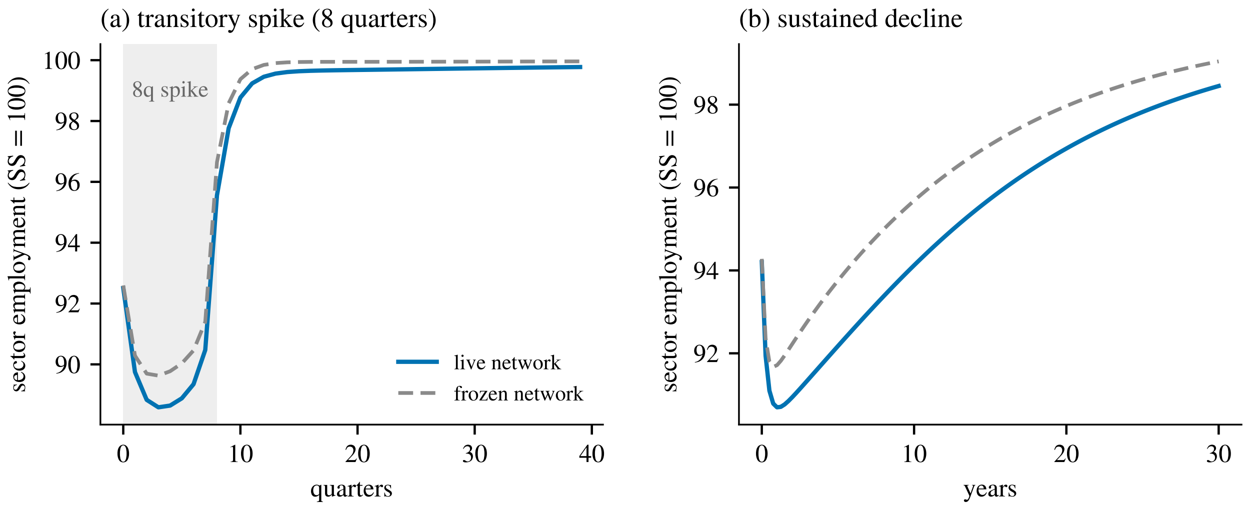

Figure 5: *Spikes heal; sustained slumps scar.*
Notes: panel (a): an eight-quarter spike in separations; the half-life from shock end is two quarters live against one frozen. Panel (b): the sustained separation decline of Section 6, first thirty years. Solid: live network; dashed: frozen.

tions (elevated at full size, then straight back to baseline) heals within quarters, and it heals within quarters whether the network is live or frozen: the half-life from shock end is two quarters against one. Fast flows dominate fast shocks, and the model does not manufacture persistence from transitory shocks; a mechanism that did should not be believed. It is sustained declines that engage the stock, erode it, and open a gap between the live and frozen economies measured in years rather than quarters, the experiment Section 6 runs at full scale. This asymmetry is the model's own explanation for the era ranking it just matched: COVID was a spike; the Great Recession was not.

6. The Anatomy of a Scar

With the model disciplined by data, I now present its main results. The central experiment is to hit the model economy with a sustained, concentrated collapse in goods production. This is the kind of shock that much of the scarring literature measures, and represents something like a "China shock" in the context of this model. The first two of the paper's three questions get their answers here—how much of the scar runs through the network, and who bears it—and the collapse itself generates the price signal that Section 7 will act on.

6.1. The Experiment, and How Much of the Scar Is the Network

Separation rates in the goods sector rise by a factor of two and a half on impact and decay with a half-life of about eight and a half years, a structural-decline ramp calibrated to the shape of the 2000s import-competition shock that Autor, Dorn, and Hanson (2013) document. Characterized by sustained pressure on one sector for over a decade, the shock spreads to the rest of the economy through the mixing kernel of Figure 4. The transition is solved nonlinearly and runs seventy-five years, ample time for the shock's decade-scale decay and the still slower recovery that follows to play out in full. Against each path I plot an identical economy, identical shock, with network capital held at its steady-state level. The vertical gap between the two lines is therefore the mechanism in question.

Figure 6 is the paper's central exhibit. Goods employment falls 9.3 percent, reaching its trough within a year (separations bite fast), and aggregate unemployment climbs from 6.0 to 7.4 percent. As the spike experiment of Figure 5 anticipated, the network margin has little effect on the depth of the downturn. The two economies reach nearly the same bottom; what the live network changes is the shape of the recovery. The frozen-network economy recovers half its loss in ten years; the live economy needs fourteen, is still 10 percent short of home after thirty, and does not put 90 percent of the loss behind it until nearly four decades out. The network channel does not deepen the wound so much as slow the healing—which, for a stock of connections that must be rebuilt by the very employment it gates, is exactly the impact one would expect.

Table 4 quantifies the gap. In cumulative discounted employment-years, the live sectoral scar is 1.33 times the frozen one at the calibrated curvature, and the aggregate scar is 1.42 times. The network channel accounts for a quarter of the sector's own scar and nearly a third of the economy's. Across the effectiveness band $\eta \in [1, 2.86]$ the sectoral multiple runs 1.25 to 1.33, and the quenched-graph corrections (the mean-field approximation replaced by fixed-graph benchmarks under both referral protocols) put it between 1.20 and 1.31; a fifth to a third of the scar, with the calibrated point near the top of the range. Two features of the decomposition have clear economic intuition behind them. The aggregate multiple exceeds the sectoral one because the spillover losses in never-shocked sectors are almost entirely network-borne. Separations there never rise; what falls is job-finding, and mostly because searchers' networks contain goods workers who are no longer employed, not because the common labor market has slackened. And the human-capital scar deepens by nearly the same factor, with goods-sector mean human capital falling 5.2 percent against 4.0 in the frozen-network economy. This is the two scarring stocks compounding: thinner networks lengthen spells, and longer spells erode skills. Households in healthy sectors, meanwhile, scar an order of magnitude less in consumption and one to two orders less in human

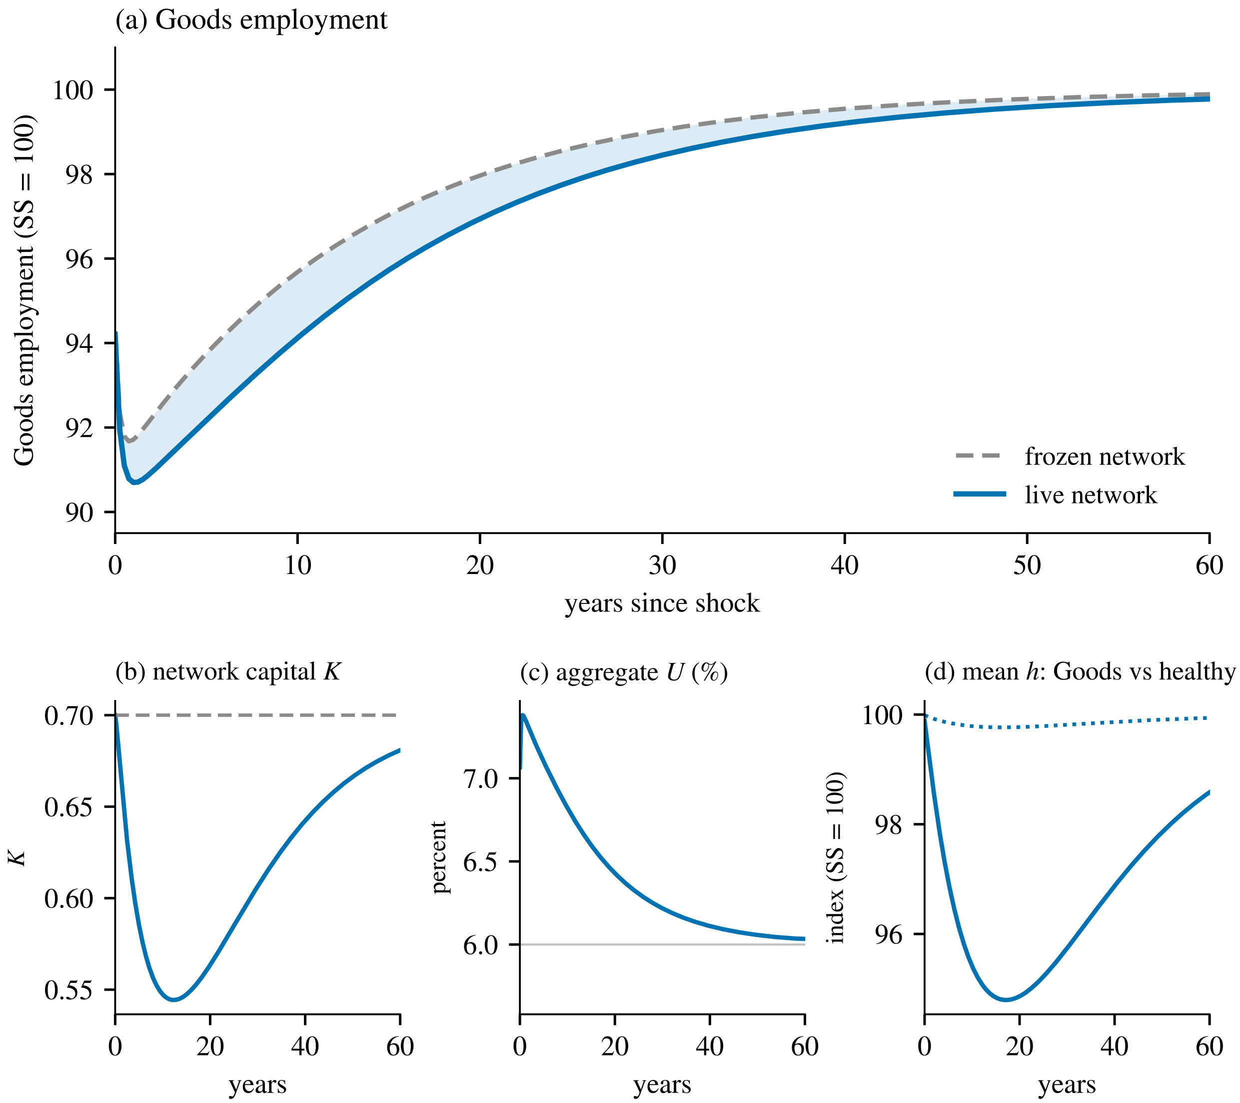

Figure 6: *The China-shock experiment: live network versus frozen.*

Notes: separations in Goods rise $\times 2.5$ *on impact and decay with a half-life of* $\sim$*34 quarters; solved nonlinearly on the empirical kernel. Panel (a): Goods employment under the live network (solid) and with network capital frozen at its steady state (dashed); the shaded gap is the network channel's contribution to the scar, and the live economy takes fourteen years to recover half its loss against ten for the frozen. Panels (b)–(d): network capital, aggregate unemployment, and mean human capital; in (b) the dashed line is the frozen reference, and in (d) the solid line is Goods and the dotted line healthy sectors. The frozen economy's human-capital scar* (-4.0% *against* -5.2% *live) is reported as a ratio in Table 4.*

	$\eta = 2.86$	η band	Quenched, broadcast	Quenched, rivalry
Sectoral employment scar	1.33	[1.25, 1.33]	[1.20, 1.31]	[1.21, 1.29]
Aggregate employment scar	1.42	[1.28, 1.42]	—	—
Goods human-capital scar	1.31	—	—	—
Implied network share (η band)	sectoral 20–25%; aggregate 22–30%			

Table 4: *Scar decomposition: live-to-frozen ratios of cumulative discounted losses.*
Notes: quenched columns replace the mean-field approximation with fixed-graph benchmarks under each referral protocol (broadcast = non-rival information, rivalry = token-passing; Appendix B); Appendix F carries the full band grid, including the aggregate and human-capital rows.

capital; the shock stays concentrated even as its network shadow spreads.

The frozen-network twin isolates the network by holding it fixed, but it leaves the rest of the scar undivided, because matching frictions and the skill channel are both alive in it. Table 5 separates them by switching the skill channel off as well. Of the full sectoral scar, matching frictions account for 72 percent, the within-worker skill channel for 4, and the network for the remaining quarter, of which 22 points are its own scarring and 3 its compounding with skills. The skill channel, in its constant-rate form, moves the scar by 4 percent; the recession-deepening that the scarring literature reads as accelerated skill loss is, in this economy, almost entirely the between-worker network. In the aggregate the split tilts further toward it: 67 percent frictions, 4 skill, 29 network.

The model's persistence can be put against the episode's own record. Autor, Dorn, and Hanson (2021) find the employment effects of import exposure in affected commuting zones present through 2019, two decades after the shock's onset and a decade past its 2010 plateau, with no visible recovery. The frozen economy is too fast for that record: it recovers half its loss within ten years. The live economy's profile, half-recovery at fourteen years and a tenth of the loss still open at thirty, sits closer to what the regional evidence shows. The comparison is between a sectoral experiment and commuting-zone outcomes that include migration and demand spillovers, so it is a consistency check rather than a test; but the persistence the literature documents is the live economy's signature, not the frozen one's.

6.2. Who Bears The Cost of Network Scarring?

Averages hide an important distributional point. Table 6 cuts the trough losses by sector and wealth quartile, separately for human capital and consumption.

Human-capital scarring is deepest at the bottom of the wealth distribution: the poorest quartile of goods-sector workers loses 1.65 percent of mean human capital at the trough, two and a half times the loss in the top quartile. The mechanism runs through the fused

Contribution to the sectoral scar	Employment-years	Share
Matching frictions (HANK+SAM)	3.42	72%
+ within-worker skill scarring	+0.17	4%
+ network scarring	+1.03	22%
+ compounding	+0.16	3%
Full model	4.78	100%

Table 5: *The scar by channel: additive decomposition of the cumulative discounted goods-sector employment scar.*
Notes: the four corners of (skill scarring $\times$ network), each on or off, under the China shock; "skill off" sets the finding gradient $\zeta = 0$, "network off" freezes K at its steady state. Entries are cumulative discounted employment-years; rows are additive contributions over the matching-frictions baseline, the network row being its effect with skills held fixed and the compounding row the residual interaction. The full-to-frozen ratio, $4.78/3.59 = 1.33$, recovers the headline multiple of Table 4. Aggregate shares: 67 percent frictions, 4 skill, 29 network. Shares are rounded.

stocks, not the budget: job-finding depends on a worker's human capital and their sector's network, not on their assets, and low human capital and low wealth coincide at the bottom of the distribution. The workers there re-employ slowest through a depressed network, and every extra quarter of unemployment is a quarter of further skill erosion. This is the double jeopardy the fused stocks were built to capture, a thin network and low human capital each lengthening the spell the other feeds on. Consumption scarring instead concentrates in the bottom half, as quarterly consumption falls 10 to 11 percent for the two bottom quartiles against 6.7 at the top, with the very bottom slightly cushioned by the UI floor. These are workers with too little buffer to smooth a spell that the network has lengthened. Households in healthy sectors lose under one percent of consumption and a rounding error of human capital. The network shadow does reach them, but only marginally.[5]

The answer to the second question, then, is not "everyone in the shocked sector." The scar sorts: it affects the workers with the least wealth to wait out a depressed network and the least slack to absorb the cost of waiting.

6.3. Correlated Collapse: Your Partners Matter

The kernel implies a sharper cross-sectional prediction: it should matter not just that a sector falls, but who else is falling. I run the same goods-sector shock twice at equal size: once with its strongest bridge partner, professional and business services, collapsing

[5] The consumption rows inherit one accounting convention: dividends are rebated lump-sum, which overstates the dividend income of poor households. Appendix F re-computes the incidence matrices under a deviation-payout convention. Qualitative patterns are unchanged, and the numbers the convention moves, chiefly the welfare results of Section 7, are reported as ranges across the two conventions.

Wealth quartile	Human capital (%) Goods	Healthy	Consumption (%) Goods	Healthy
Q1 (poorest)	−1.65	−0.12	−10.27	−0.69
Q2	−1.04	−0.07	−10.79	−0.88
Q3	−0.51	−0.01	−8.39	−0.66
Q4 (richest)	−0.66	−0.01	−6.65	−0.69

Table 6: *Who bears the scar: trough losses by sector and wealth.*
Notes: percent deviations from steady state at the employment trough, by steady-state asset quartile (cutoffs in Appendix E); "Healthy" averages the never-shocked sectors. Point estimates; the band grids of Appendix F apply. Consumption rows inherit the dividend-rebate convention (see footnote in text).

alongside it, and once with its weakest, information. I then compare the focal sector's own scar across the two runs. The gap between them is the bridge-partner premium. With its hub partner down rather than its weakest link, the identical own shock scars goods 2.2 percent more in cumulative terms, and 4.8 percent more at the post-shock snapshot. The premium roughly doubles moving across the η band, from 2.6 to 4.8 percent, which ties it to the referral curvature; the model contains no input–output channel that could produce it. One boundary limits the claim: joint scars are sub-additive in general equilibrium. The two-sector collapse costs about 3 percent less than the sum of its parts, with stabilizers outweighing the wedge's convexity at feasible shock sizes. The premium is a statement about where spillovers land, not a super-additive amplification of the collapse itself.

6.4. Fragility and the Loop Gain

Figure 7 tracks two series along the collapse. The wedge prices the externality, the gap between a job's social and private value; the loop gain measures the distance to the fold, where a slump would become self-sustaining.

At the steady state, the social value of a goods-sector job exceeds its private value by 26.6 percent. Shutting off cross-sector mixing lowers the wedge to 23.9 percent; the difference is the value a goods-sector job creates for searchers in other sectors. Along the collapse the shocked sector's wedge nearly doubles, peaking at 48.7 percent in the second quarter, before the employment trough arrives in the fourth, while healthy sectors' wedges jump to 27 percent on impact and then drift below their steady-state level. The externality is largest early, while the stock is still intact and each job saved still protects it; by the trough, much of what the subsidy would have preserved is already gone. That timing is the most policy-relevant fact in the figure, and will be further built on in Section 7.

The loop gain rises from 0.19 to a peak of 0.24 along the transition, and never exceeds 0.33 in the experiments computed for this paper (deeper slumps, joint collapses, passive

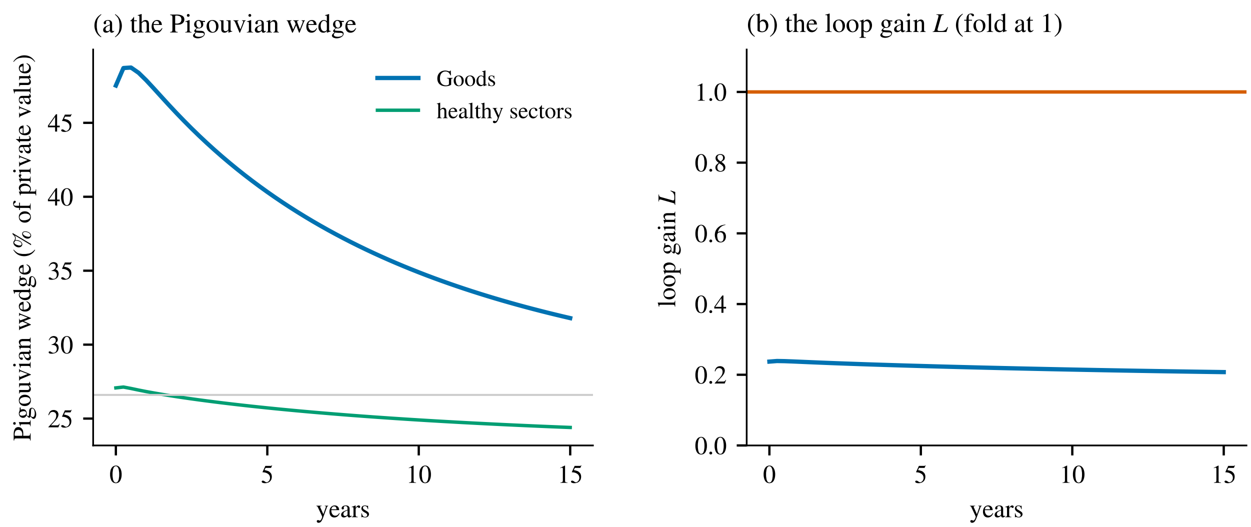

Figure 7: *The wedge and the loop gain along the collapse.*
Notes: panel (a): the Pigouvian wedge Ω/Δ for the shocked (Goods) and healthy sectors along the China-shock transition; the thin horizontal line marks the steady-state wedge (26.6%). The shocked sector's wedge roughly doubles, peaking near 49% in the second quarter, before the employment trough in the fourth, while healthy sectors' wedges jump to about 27% on impact and then drift below their steady-state level. Panel (b): the loop gain L along the same path (orange line: the fold at $L = 1$); L rises from 0.19 at the steady state to a peak of 0.24, and 0.33 is the largest value attained in any experiment in the paper.

monetary policy). The fold at $L = 1$ is far away. While a self-sustaining slump is theoretically possible, it remains far from the calibrated values explored here.

The same statistic also disposes of a natural conjecture: that the network channel should deepen demand-driven slumps by feeding the precautionary spiral of the heterogeneous-agent literature. At this calibration it does not. Under passive monetary policy the slump itself deepens substantially, but the network's share of the amplification falls. The live-to-frozen multiple drops from roughly 1.27–1.30 under active policy to 1.10–1.13 under passive, for separation and demand shocks alike. The reason is the same stock logic that runs through the whole paper: the network multiplier feeds on sustained slack, which erodes capital slowly, while passive-money slumps are deeper but front-loaded price-spiral events whose extra depth the frozen economy shares almost fully. The network is a hysteresis amplifier, not a spiral amplifier, and its contribution to fragility is the loop-gain contribution Figure 7 plots, not a monetary interaction.

7. The Policy Instrument and Its Cost

The theory priced the correction. This section implements it in the calibrated economy, asking three practical questions. Does the priced credit undo the scar, and at what fiscal cost? Does it beat the instrument governments typically reach for in downturns? And finally, is the formula anywhere near the optimum?

7.1. The Entry-Margin Credit

The instrument is the entry-margin credit of Proposition 3: a payment per hire inside the free-entry condition, priced at the entry-margin wedge (10.5 percent of the value of a filled job at the steady state; the 26.6 percent figure of Section 6.4 is the worker-margin wedge on a different base, converted to the entry margin in Appendix F.7) and scaled along the transition by the wedge's elevation, so that it pays nothing in normal times, switches on as the collapse begins, and peaks with the wedge at 7.8 percent of the quarterly wage. It is financed through the model's fiscal rule, so the outlay accrues to debt and is serviced by the rule's feedback. One implementation detail matters for reading these results: the nonlinear solver supports the credit at up to half its priced scale (which the diagnostics of Appendix E indicate is a reach limit of the current solver rather than nonexistence), so the paths below run the half-priced rule, whose peak is 3.9 percent of the quarterly wage. Section 7.3 will show the disclosure costs little, as the half-priced rule already captures most of what the priced rule can.

Figure 8 shows what this credit buys. Under the half-priced credit the goods sector's cumulative scar shrinks by 48.3 percent (roughly half the damage of Figure 6 undone), while the network stock itself is visibly protected, peaking about eight percent above its no-policy path. The hiring response is a boom, not an explosion: job-finding stays well inside feasible range. And the bill is small. The present-value fiscal cost is about 2.5 percent of one quarter's GDP, because a state-contingent credit pays per hire only while the wedge is elevated; the same timing that makes the externality large early is what makes pricing it cheap overall. One caution from the implementation belongs in the record: the worker-side subsidy that decentralizes the same allocation on the effort margin does not transplant to the entry margin (free entry compounds a persistent per-hire payment many times over), and Appendix F works the conversion out as a cautionary note on margin-specific prices.

7.2. The Horse Race

Unemployment insurance (UI) makes for a natural comparison benchmark as it is already a widely used tool. I give UI the same money (a replacement-rate increase of one percentage point over the same window, matched to the credit's present-value cost) and let the two run on the same shock. On employment, UI does approximately nothing. The sectoral scar is unchanged to a tenth of a percent; the wage rule's outside-option channel (the benefit

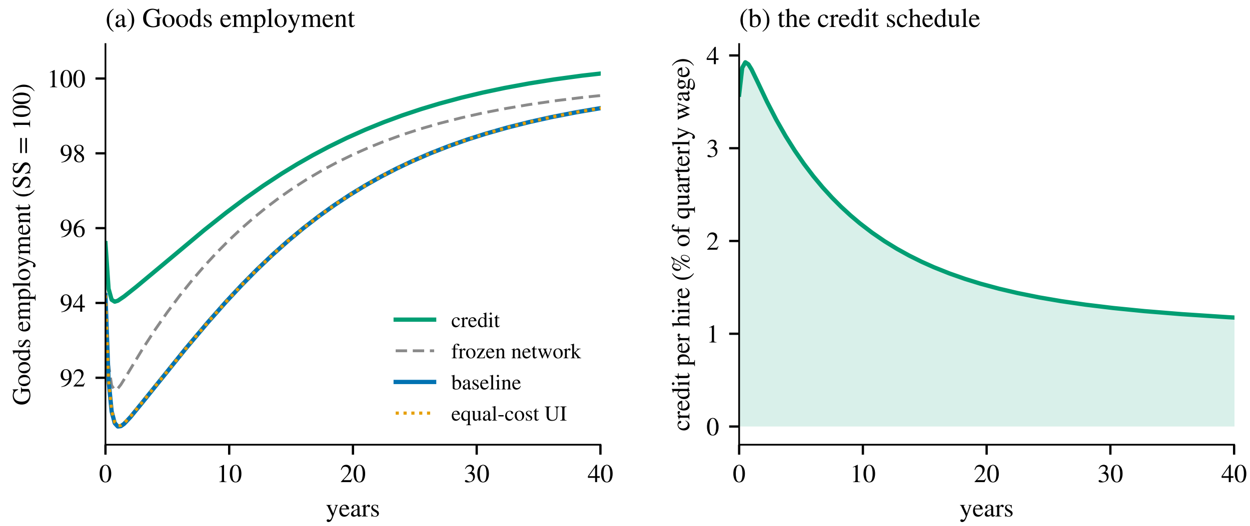

Figure 8: *Paths under the credit and under equal-cost UI.*
Notes: panel (a): Goods employment under the no-policy baseline, the half-priced credit, and the equal-cost UI expansion (indistinguishable from the baseline); the frozen-network path is retained as a reference. The credit economy can overtake the frozen reference because the credit stimulates entry rather than merely neutralizing the wedge. Panel (b): the credit per hire as a share of the quarterly wage; it pays only while the wedge is elevated.

enters the wage level) is visible but inert at this size. Table 7 reports the welfare verdict in consumption-equivalent units, and it is lopsided: the credit is worth between +1.40 and +1.66 percent of lifetime consumption depending on the dividend-payout convention, while equal-cost UI is worth between +0.00 and +0.05 percent. The credit dominates equal-cost UI outright. The horse race prices UI's entry-crowding cost; the quantitative model has no search-effort margin, so the classic moral-hazard cost of UI is absent from the comparison, and including it would cut against UI. The UI result should be read carefully, as the insurance is not worthless. At the 40 percent baseline replacement rate the insurance gain of another point of replacement almost exactly cancels its entry-crowding cost; the optimal-insurance margin of Baily (1978) and Chetty (2006) is close to saturated. Proposition 4's division of labor shows up quantitatively on the transition as well. Evaluated along the collapse, the network term in UI's marginal cost rises roughly fivefold exactly as the credit's value doubles. In a slump, the case for pricing the cause strengthens at the same moment the case for another unit of symptom insurance weakens.

The bottom wealth quartile loses under both policies, between 2.0 and 3.4 percent under the credit depending on the payout convention. The mechanism is fiscal incidence, not the labor market, as the hiring boom is paid for out of dividends and debt whose service falls partly on households that live on current income, while the employment gains accrue later. Roughly 40 percent of the loss is the lump-sum dividend-rebate convention alone: rebate

	Hiring credit		UI expansion	
	lump-sum	dev.-payout	lump-sum	dev.-payout
Total (all households)	+1.40	+1.66	+0.00	+0.05
Goods, Q1 (poorest)	−3.38	−1.99	−0.46	−0.80
Goods, Q2	+4.90	+4.09	+0.43	+0.40
Goods, Q3	+10.28	+9.85	+0.36	+0.70
Goods, Q4 (richest)	+2.82	+3.89	−0.07	+0.84

Table 7: *Equal-cost welfare: credit versus UI (CE % of lifetime consumption).*
Notes: both policies matched to the same present-value fiscal cost. Columns give the two dividend-payout conventions (lump-sum rebate; deviation payout re-computed on the solved paths, Appendix F); the range across conventions is the quoted object. Healthy-sector rows in Appendix F.

dividends where equity is actually held, at the top, and the loss falls to 2.0 percent (Table 14). Financing the credit progressively rather than by a flat tax removes little more. I show in Appendix F that re-computing household welfare with the credit's fiscal cost routed onto wealth moves the figure by about two-tenths of a point, the credit's tax bill being modest and the share of it the poorest would pay smaller still. Under every convention I can compute the bottom quartile comes out behind, and the two amendments are far too small, even combined, to reverse the sign. The clear winners are middle-wealth workers in the shocked sector, at around +10 percent. This group captures the re-employment gains without being dependent on the dividend stream that funds them.

A loss that survives every financing convention is not a flaw in the credit's design, but rather a distributional side effect. Rectifying these distributional costs relies on its own instrument. The credit's can deliver improvements in efficiency, the closing of the network wedge; the incidence is a matter of timing, borne by households who cannot smooth a temporary dividend cut. The remedy is separate and cheap. A direct transfer to the liquidity-constrained during the credit's window, funded out of an aggregate surplus that is many times its size, could in principle offset the loss without touching the credit. I do not design it here as it is orthogonal to the externality and efficiency instruments.

7.3. Is the Formula the Optimum?

So far the exercise has evaluated the theory's rule, not the best rule. To ask how much is left on the table, I treat the credit's scale and time profile as free parameters and solve the model at each setting: a numerical Ramsey exercise in the credit's two-parameter family (grid and solver details are deferred to Appendix F). Figure 9 plots welfare against the credit's scale. It is concave through the computable range, rising from +0.85 percent at a quarter of the priced scale to +1.57 at five-eighths, with marginal gains falling by roughly forty percent per step. The exact peak lies beyond the current solver's reach, but the curvature places it near the priced scale, and the half-priced rule already delivers just under 90 percent of the best-found gain. Being roughly right on scale is nearly free. Timing is less forgiving. At equal scale, front-loading the credit harder than the wedge's own path loses badly, and the flattened profile does no better than the wedge's own shape per fiscal dollar; nothing in the computed grid improves on the wedge's timing. Per fiscal dollar the ranking inverts in an instructive way: small, early credits have the highest return of anything in the table, because the earliest hires are the most valuable. The optimization thus recovers on its own the timing logic of the wedge path in Figure 7. How much of the credit's value is the externality, and how much is entry stimulus that rigid wages would reward in any case? The direct test is to run the identical credit path in the frozen-network economy, where the stock cannot respond. It still helps: the same path removes 43 percent of the frozen economy's scar and is worth $+0.94$ percent of lifetime consumption, against 48 percent and $+1.40$ in the live economy. The split is stable across metrics: about a third of the credit's value is the network externality, and two thirds is the correction of the entry inefficiency that rigid wages create in any slump. Appendix F reports the four arms. The wedge earns its role twice over even so: it sizes the third that is the externality's, and its time profile remains the best computed timing for the whole instrument, stimulus component included.

The placebo also disciplines the headline return. Across the computed grid the credit pays 27 to 40 times its present-value fiscal cost, a ratio that invites scrutiny. Two thirds of it is the entry-margin correction of a rigid-wage slump, a return any model with this wage rule would deliver; the network externality's own component returns roughly ten times cost. Both components scale roughly linearly with the wedge, and hence with the imported χ and η; the one-sided χ sweep of Appendix F bounds the mapping risk from above.

7.4. The Bridge Subsidy Decentralizes

The last instrument is that of Theorem 2. The quantitative model does not endogenize the direction of search (that margin belongs to the theory), so the priceable experiment is the policy's direct effect: a state-contingent tilt of the shocked sector's search exposure toward its hub partner, professional and business services, switched on over the same wedge window. Three results close the loop with the theorem. The tilt works. The shocked sector's employment improves and its network stock is protected, in proportion to the modest size of the tilt. The hub is not hurt, as its own employment rises slightly (the one-sidedness of Theorem 2 in action); bridging feeds the depressed sector without taxing the healthy one. The realized general-equilibrium welfare gain comes in at 1.35 times the theorem's costate formula. The first-order price is correct, and general equilibrium adds to

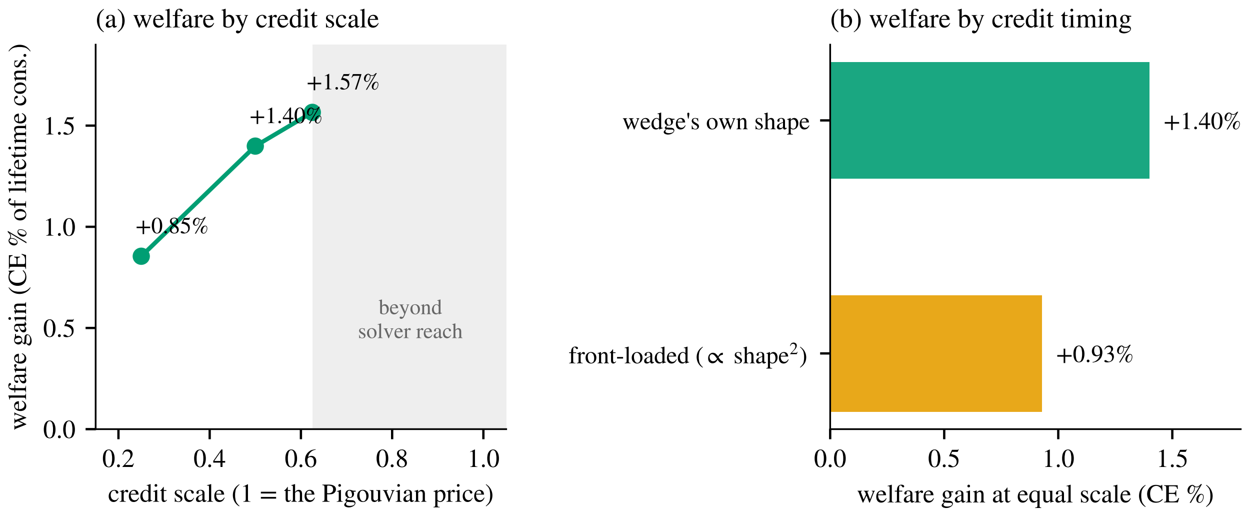

Figure 9: *The Ramsey check: welfare in the credit scale and shape.*
Notes: panel (a): welfare gain at each credit scale along the wedge's own time profile; the shaded region lies beyond the current solver's reach. Panel (b): equal-scale comparison of the wedge's profile against a front-loaded alternative. Flattening is compared per fiscal dollar in the text; the equal-scale flattening comparison awaits the queued grid extension (Appendix F).

it through channels the formula does not have, chiefly the human-capital gains of shorter spells. On the empirical kernel the capital-protection share of the bridge's value is about a third, larger than in the theory's illustration. The instrument's clock differs from the credit's. The bridge subsidy peaks in mid-collapse, after the employment trough, when cross-sector divergence is widest. The two instruments form a schedule (credits early, while the wedge is spiking; bridges through the middle, while the sectors are far apart), and the schedule is not a design choice but the wedge geometry's own timing.

8. Conclusion

Why do deep recessions scar labor markets long after their shocks fade? This has a large existing literature of answers located inside the worker: skills, health, attachment. This paper has argued that a quantifiable part of the answer lives between workers. On the evidence assembled here, the networks through which a large share of jobs are found behave like a capital stock: built by employment, eroded slowly by unemployment, and owned by no one. Three decades of CPS data point to a reliance that builds over months and responds to sustained slumps but not spikes, and a disciplined equilibrium model that takes those stylized facts seriously attributes a quarter to a third of the employment scar from a China-shock-like collapse to this one channel. Damage is concentrated on the low-wealth workers of the declining sector, and adds years to the recovery. The mechanism is an amplifier of sustained slumps, not a creator of persistence from nothing, and one statistic, the network loop gain, prices its damage, its amplification, and its distance from the self-reinforcing region throughout.

For policy, the ownership failure is the operative fact. Because no one owns the stock, a collapse destroys capital no market participant is paid to protect, and the correction is Pigouvian, not redistributive. In the calibrated model, a hiring credit priced by the wedge, switched on early in a collapse while the capital still exists, undoes roughly half the scar at a fiscal cost of a few percent of one quarter's GDP and dominates an equal-cost expansion of unemployment insurance. The case for state-contingent hiring support in sectoral collapses is part Pigouvian and part repair of an entry margin that rigid wages distort in any slump; the model prices the split at a third and two thirds. Its timing follows the wedge throughout: credits early, bridge subsidies to better-connected partner sectors through the middle of the divergence. All of it is priced by a statistic a policymaker could, in principle, monitor.

The boundaries of these claims deserve consideration. The CPS object is a search-method, not a record of how jobs were found, and the paper uses it only for cyclical slopes, taking levels from the direct-hiring literature. No individual-level causal effect of networks is estimated here, with all elasticities imported from the quasi-experimental literature, and the facts are aggregate patterns consistent with the mechanism rather than proof of it. On the model side, the free-entry block cannot generate a 2008-scale demand-only slump at disciplined parameters, so the paper's conclusions are strongest for the separation-driven collapses its calibration anchors on. The binding constraint is the entry margin's small surplus, exhausted after roughly a two-point rise in unemployment; endogenizing the separation margin, or enriching the entry technology, could give demand shocks realistic depth and put the network mechanism to the test in demand-driven slumps. The network block is mean-field, with fixed-graph corrections carried as bands, and the distributional results inherit a dividend-rebate convention whose variants Appendix F reports. The search-effort instrument is left unpriced, needing evidence on whether referral information is rival. Two projects follow directly. An administrative event study around displacement could test the calibration's tie-decay assumption directly: ties die with a half-life near five and a half quarters, faster while the sector stays slack, and Appendix B maps this to an observable object, the survival gap between the contacts of unemployed and employed workers. And the mobility margin the model holds fixed, whether to bring networks to workers or workers to networks, is the natural question for a successor.

The scar of a recession is not only the internal margin of *what* workers forget; it is also the external margin of *whom* they lose, and that margin compounds when the connections who remain are struggling themselves. Recessions are long and painful when you've got friends in low places.

References

Arbex, Marcelo, Sidney Caetano, and Dennis O'Dea. 2016. "The Implications of Labor Market Network for Business Cycles." *Economics Letters* 144:37–40.

Arbex, Marcelo, Dennis O'Dea, and David Wiczer. 2019. "Network Search: Climbing the Job Ladder Faster." *International Economic Review* 60 (2):693–720.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. 2021. "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models." *Econometrica* 89 (5):2375–2408.

Autor, David, David Dorn, and Gordon H. Hanson. 2021. "On the Persistence of the China Shock." *Brookings Papers on Economic Activity* Fall:381–447.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6):2121–2168.

Baily, Martin Neil. 1978. "Some Aspects of Optimal Unemployment Insurance." *Journal of Public Economics* 10 (3):379–402.

Beaman, Lori A. 2012. "Social Networks and the Dynamics of Labour Market Outcomes: Evidence from Refugees Resettled in the U.S." *Review of Economic Studies* 79 (1):128–161.

Bramoullé, Yann and Gilles Saint-Paul. 2010. "Social Networks and Labor Market Transitions." *Labour Economics* 17 (1):188–195.

Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman. 2015. "The Value of Hiring through Employee Referrals." *Quarterly Journal of Economics* 130 (2):805–839.

Calvó-Armengol, Antoni and Matthew O. Jackson. 2004. "The Effects of Social Networks on Employment and Inequality." *American Economic Review* 94 (3):426–454.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. 2007. "Business Cycle Accounting." *Econometrica* 75 (3):781–836.

Chetty, Raj. 2006. "A General Formula for the Optimal Level of Social Insurance." *Journal of Public Economics* 90 (10-11):1879–1901.

Cingano, Federico and Alfonso Rosolia. 2012. "People I Know: Job Search and Social Networks." *Journal of Labor Economics* 30 (2):291–332.

Davis, Steven J. and Till von Wachter. 2011. "Recessions and the Cost of Job Loss." *Brookings Papers on Economic Activity* 2011 (2):1–72.

den Haan, Wouter J., Pontus Rendahl, and Markus Riegler. 2018. "Unemployment (Fears) and Deflationary Spirals." *Journal of the European Economic Association* 16 (5):1281–1349.

Elsby, Michael W. L., Ryan Michaels, and Gary Solon. 2009. "The Ins and Outs of Cyclical Unemployment." *American Economic Journal: Macroeconomics* 1 (1):84–110.

Galenianos, Manolis. 2014. "Hiring through Referrals." *Journal of Economic Theory* 152:304–323.

Glitz, Albrecht. 2017. "Coworker Networks in the Labour Market." *Labour Economics* 44:218–230.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78 (6):1360–1380.

———. 1974. *Getting a Job: A Study of Contacts and Careers*. Cambridge, MA: Harvard University Press.

Hosios, Arthur J. 1990. "On the Efficiency of Matching and Related Models of Search and Unemployment." *Review of Economic Studies* 57 (2):279–298.

Ioannides, Yannis M. and Linda Datcher Loury. 2004. "Job Information Networks, Neighborhood Effects, and Inequality." *Journal of Economic Literature* 42 (4):1056–1093.

Jackson, Matthew O. and Asher Wolinsky. 1996. "A Strategic Model of Social and Economic Networks." *Journal of Economic Theory* 71 (1):44–74.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108 (3):697–743.

Kopecky, Joseph. 2026. "Rival Referrals: Business-Stealing, Countercyclical Congestion, and the Limits of Search-Effort Policy in Job-Search Networks." Trinity Economic Papers, Trinity College Dublin.

Mukoyama, Toshihiko, Christina Patterson, and Ayşegül Şahin. 2018. "Job Search Behavior over the Business Cycle." *American Economic Journal: Macroeconomics* 10 (1):190–215.

Pellizzari, Michele. 2010. "Do Friends and Relatives Really Help in Getting a Good Job?" *Industrial and Labor Relations Review* 63 (3):494–510.

Ravn, Morten O. and Vincent Sterk. 2017. “Job Uncertainty and Deep Recessions.” *Journal of Monetary Economics* 90:125–141.

Shimer, Robert. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies.” *American Economic Review* 95 (1):25–49.

———. 2012. “Reassessing the Ins and Outs of Unemployment.” *Review of Economic Dynamics* 15 (2):127–148.

A. Proofs and Derivations

This appendix derives the costate system behind Section 3, proves Theorems 1 and 2 and Propositions 1–4, and states the household problem and equilibrium definition for the quantitative model of Section 4. Every closed form and every claim proved here is also verified numerically in the replication package, by independent routes: closed forms against finite-difference Jacobians of the actual transition map, and costate formulas against brute-force simulation of discounted welfare.

A.1. The Costate System and the Wedge in Closed Form

The environment is that of Section 3.1: laws of motion (1)–(3), search effort a at convex cost $c(a)$ paid by the unemployed, wage $w = y$. Worker values satisfy $W = w + \beta[(1-s)W + sU]$ and $U = b - c(a) + \beta[afW + (1-af)U]$, so the private value of employment $\Delta \equiv W - U$ obeys

$$\Delta \;=\; \big(w - b + c(a)\big) + \beta\,(1 - s - af)\,\Delta, \tag{A.1}$$

where the $c(a)$ term appears because employment saves the search cost. The worker's first-order condition is $c'(a) = \beta f \Delta$.

The planner chooses the effort path subject to the same laws and the same mean-field meeting structure. Writing $V(e,K)$ for the planner value, $\lambda \equiv V_e$ and $\nu \equiv V_K$, the envelope conditions at a stationary allocation are

$$\lambda \;=\; \big(y - b + c(a)\big) + \beta\big[J_{ee}\,\lambda + J_{Ke}\,\nu\big], \tag{A.2}$$

$$\nu \;=\; \beta\big[J_{eK}\,\lambda + J_{KK}\,\nu\big], \tag{A.3}$$

and the planner's effort condition is $c'(a) = \beta f \lambda$. Here J is the Jacobian of the state transition $(e,K) \mapsto (e',K')$ implied by (2)–(3):

$$\begin{aligned}
&J_{ee} = (1-s) - af + a(1-e)\,f_e, \qquad J_{eK} = a(1-e)\,f_K,\\
&J_{Ke} = \varphi(1-K) + \delta K\big[1 + \kappa(1 - e/E^\circ)\big] + \delta(1-e)K\,\kappa/E^\circ,\\
&J_{KK} = 1 - \varphi e - \delta(1-e)\big[1 + \kappa(1 - e/E^\circ)\big],
\end{aligned} \tag{A.4}$$

with $f_e = \chi\eta\,\bar{f}\,(K/K^\circ)\,e^{\eta-1}/(E^\circ)^\eta$ and $f_K = \chi\,\bar{f}\,(e/E^\circ)^\eta/K^\circ$ from (1). At the reference steady state (E°, K°) these reduce to $f_e = \chi\eta\bar{f}/E^\circ$ and $f_K = \chi\bar{f}/K^\circ$, and the decay bracket equals one.

Lemma A.1 (the wedge in closed form)**.** *At an interior stationary allocation with $\beta\rho(J) < 1$*

($\rho(\cdot)$ the spectral radius),

$$\nu = \frac{\beta J_{eK}}{1 - \beta J_{KK}} \lambda, \qquad \lambda = \frac{\Delta}{1 - L}, \qquad \Omega \equiv \lambda - \Delta = \Delta \frac{L}{1 - L}, \tag{A.5}$$

with L as in (5). Equivalently $\lambda = (y - b + c(a))\,(1 - \beta J_{KK}) / \det(I - \beta J)$.

Proof. Solve (A.3) for ν and substitute into (A.2): $\lambda\big[1 - \beta J_{ee} - \beta^2 J_{Ke} J_{eK} / (1 - \beta J_{KK})\big] = y - b + c(a)$. From (A.4), $1 - \beta J_{ee} = 1 - \beta(1 - s - af) - \beta a(1 - e) f_e$, so the bracket equals $(1 - \beta(1 - s - af))(1 - L)$ with L defined in (5). Under $w = y$ the private flow in (A.1) is the same $(y - b + c(a))$, so $\Delta = (y - b + c(a))/(1 - \beta(1 - s - af))$ and $\lambda = \Delta/(1 - L)$ follow, as does the resolvent form. □

A.2. Proof of Theorem 1

Part (i). Perturb the competitive equilibrium by a uniform increase da in search effort in every period, holding the equilibrium objects at their laissez-faire values. The perturbation costs $(1 - e)\, c'(a)\, da$ per period and creates $(1 - e) f\, da$ additional matches per period, each valued at $\beta\lambda$ by the planner's costate. Using the worker's first-order condition $c'(a) = \beta f \Delta$,

$$\frac{dW}{da} = \sum_t \beta^t \,(1 - e) \big[\beta f \lambda - c'(a)\big] = \frac{(1 - e)\, \beta f\, \Omega}{1 - \beta}. \tag{A.6}$$

By Lemma A.1, $\Omega > 0$ if and only if $L > 0$. Every term of L carries a factor of χ: the static term through $f_e \propto \chi\eta$ and the capital loop through $J_{eK} \propto \chi$. Hence $dW/da > 0$ if and only if $\chi > 0$. At $\chi = 0$, $\lambda = \Delta$, the planner's and the worker's first-order conditions coincide at every state, and the laissez-faire allocation solves the planner problem: the equilibrium is constrained efficient.

Part (ii). Under the maintained conditions — $\eta, \kappa \geq 0$, $\varphi, \delta > 0$, quarterly turnover below one ($s + af \leq 1$ and $\varphi e + \delta(1 - e)[1 + \kappa(1 - e/E^\circ)] \leq 1$) — every entry of J in (A.4) is non-negative. With $\beta\rho(J) < 1$ the Neumann series $(I - \beta J)^{-1} = \sum_{k \geq 0} (\beta J)^k$ converges and is entrywise non-negative, so each of the three spillover channels enters L non-negatively, and $L > 0$ strictly whenever $\chi > 0$. There is no negative entry anywhere in the system: the sign cannot reverse within the model class.

Part (iii).

Assumption A.1. c is increasing and sufficiently convex that (a) the worker's and planner's effort problems have unique interior solutions, and (b) the fixed-point maps from an effort schedule to the stationary state (e, K) implied by (2)–(3) have unique, stable solutions.

By parts (i)–(ii), $\lambda > \Delta$ at every interior state when $\chi > 0$, so the planner's effort schedule lies strictly above the worker's at every state: $c'(a^P) = \beta f \lambda > \beta f \Delta = c'(a^{CE})$ and c convex imply $a^P > a^{CE}$ pointwise. The stationary maps are increasing in the effort schedule: e' rises with a directly, and K' rises with e through $J_{Ke} > 0$. Under Assumption A.1 the unique stable fixed points are therefore ranked: $a^P > a^{CE}$, $e^P > e^{CE}$, $K^P > K^{CE}$. (The ranking is also verified computationally at every calibration used in the paper.) □

A.3. Proof of Proposition 1

In the vacancy-posting variant, direct meetings form according to $M(u,v) = \mu_m u^\alpha v^{1-\alpha}$, so with tightness $\theta \equiv v/u$ a vacancy fills with probability $q(\theta) = \mu_m \theta^{-\alpha}$ and a searcher meets a vacancy with probability $f_{dir}(\theta) = \mu_m \theta^{1-\alpha}$; referral finding is as in (1) and creates hires without posted vacancies. The wage is fixed at w; a filled job is worth $J^f = (y-w)/(1-\beta(1-s))$ to the firm, and free entry at flow cost κ_v requires $\kappa_v = \beta q(\theta) J^f$.

Define w^* as the wage at which laissez-faire entry is constrained efficient in the $\chi = 0$ economy; this is the fixed-wage analogue of the Hosios condition, equating the firm's entry margin to the planner's at $\chi = 0$, and it exists and is unique because both margins are continuous and monotone in w. Now let $\chi > 0$ and hold $w = w^*$. The planner's value of an additional hire exceeds the private joint value of the match by the wedge of Lemma A.1: the hire builds the sector's stock (J_{Ke}), raises contemporaneous referral finding (f_e), and slows others' tie decay (κ) — none of which accrues to the matched pair at any wage, because these components fall on agents outside the match. The planner's entry condition therefore holds at a strictly higher θ than the free-entry condition: $\theta^P > \theta^{CE}$. A hiring subsidy

$$\tau_h \;=\; \frac{\kappa_v}{\beta q(\theta^P)} - J^f \;>\; 0 \tag{A.7}$$

restores free entry at θ^P, and $\tau_h \to 0$ as $\chi \to 0$ because the wedge vanishes. □

A.4. Proof of Proposition 2

Direct computation of the 2×2 determinant gives

$$\begin{aligned}\det(I - \beta J) &= (1-\beta J_{ee})(1-\beta J_{KK}) - \beta^2 J_{eK} J_{Ke} \\ &= \big(1 - \beta(1-s-af)\big)\big(1-\beta J_{KK}\big)\big(1-L\big),\end{aligned} \tag{A.8}$$

where the second line substitutes $1 - \beta J_{ee} = \big(1-\beta(1-s-af)\big) - \beta a(1-e) f_e$ and collects terms using the definition (5) of L. This is (6). The three roles follow. *(i) Price:* Lemma A.1.

(ii) Amplification: the discounted employment created by a marginal job is the (e, e) entry of the resolvent,

$$\left[(I - \beta J)^{-1}\right]_{ee} = \frac{1 - \beta J_{KK}}{\det(I - \beta J)} = \frac{1}{1 - \beta(1 - s - af)} \cdot \frac{1}{1 - L}, \tag{A.9}$$

the private duration multiplier scaled by $1/(1-L)$. *(iii) Distance to the fold:* a fold bifurcation of the steady state of (2)–(3) occurs where $\det(I - J) = 0$. Setting $\beta = 1$ in (A.8), the first two factors equal $s + af > 0$ and $\varphi e + \delta(1-e)[\,\cdot\,] > 0$ under the turnover bounds, so $\det(I - J) = 0$ if and only if $L|_{\beta=1} = 1$. The wedge's pole and the fold are the same condition in the patient limit. $\square$

A.5. The S-Sector System

With S sectors connected by a row-stochastic mixing matrix m, the state collects the employment rates $(e_1, \ldots, e_S)$ and the stocks $(K_1, \ldots, K_S)$, exposure is $\tilde{e}_j = \sum_k m_{jk}\, e_k$, and the laws of motion apply sector by sector with $\tilde{e}_j$ replacing e in the informal term and in the decay bracket. The $2S \times 2S$ transition Jacobian has the same structure as (A.4), with every cross-sector entry proportional to $m_{jk} \geq 0$: the system remains entrywise non-negative under the turnover bounds. With β times the spectral radius below one, the Neumann argument of Theorem 1(ii) applies to the stacked system, and $\lambda_j - \Delta_j > 0$ for every sector j in the same communicating class as any sector with $\chi > 0$. Since χ is common across sectors and the empirical kernel of Section 5.1 has all entries strictly positive, every sector carries a strictly positive wedge: χ decides existence, and m decides only relative magnitudes — incidence.

A.6. Proof of Theorem 2

At the asymmetric steady state, sector 1 (higher s_1) has lower employment and, through (3), a lower stock, so $g_1 < g_2$ and $e_1 < e_2$. The interior equilibrium bridge choice satisfies $c_B'(\omega_1) = \beta\chi\bar{f}\,(g_2 - g_1)\,\Delta_1$: the worker prices the private diversification gain in full.

Perturb ω_1 uniformly by $d\omega$ across sector 1's unemployed, holding the equilibrium fixed. Three first-order effects arise each period. (a) Finding: re-employment probability rises by $\beta\chi\bar{f}\,(g_2 - g_1)\,d\omega$ per searcher, valued at λ_1 by the planner. (b) Cost: search cost rises by $c_B'(\omega_1)\,d\omega = \beta\chi\bar{f}\,(g_2 - g_1)\,\Delta_1\,d\omega$ per searcher, by the equilibrium condition. (c) Protection: exposure $\tilde{e}_1$ rises by $(e_2 - e_1)\,d\omega$, which slows the decay of sector 1's stock by $\delta K_1 \kappa\,(e_2 - e_1)/E^\circ\,d\omega$ per unit of unemployment, valued at $\beta\nu_1$. Summing over the unemployed mass $(1-e_1)$ and over time,

$$\left.\frac{dW}{d\omega_1}\right|_{CE} = \frac{1-e_1}{1-\beta}\Big[\,\beta\chi\bar{f}\,(g_2-g_1)\,(\lambda_1-\Delta_1)\;+\;\beta\nu_1\,\delta K_1\kappa\,(e_2-e_1)/E^{\circ}\,\Big], \tag{A.10}$$

which is (7). Both bracketed terms are strictly positive: $g_2 > g_1$ and $e_2 > e_1$ at the asymmetric steady state, $\lambda_1 > \Delta_1$ by the two-sector version of Theorem 1, and $\nu_1 > 0$ by the Neumann argument. For the healthy sector every sign in (A.10) reverses ($g_1 - g_2 < 0$, $e_1 - e_2 < 0$), so the planner remains at the corner $\omega_2 = 0$, exactly as the market does. At the symmetric steady state $g_1 = g_2$ and $e_1 = e_2$, and the expression vanishes identically. □

A.7. Proofs of Propositions 3 and 4

Proposition 3. A flow subsidy τ_w paid while employed shifts the private value to $\Delta_{\tau_w} = \Delta + \tau_w/(1-\beta(1-s-af))$. Setting

$$\tau_w^* \;=\; \Omega\cdot\big(1-\beta(1-s-a^P f^P)\big) \tag{A.11}$$

evaluated at the planner allocation gives $\Delta_{\tau_w^*} = \lambda$, so the worker's first-order condition coincides with the planner's at the planner allocation, which is therefore an equilibrium of the subsidized economy. On the firm margin, the hiring credit τ_h of (A.7) implements θ^P directly in the free-entry condition. The two implementations correct whichever margin the economy possesses; their incidence differs because the margins do.

Proposition 4. Let τ_{UI} be a marginal, lump-sum financed increase in the benefit paid while unemployed. The transfer itself nets to zero in welfare. The induced change in search effort, $da/d\tau_{UI} < 0$, has no first-order private cost by the envelope theorem applied to the worker's problem; its only first-order welfare effect runs through the externality:

$$\frac{dW}{d\tau_{UI}} = \frac{da}{d\tau_{UI}}\cdot\frac{(1-e)\,\beta f\,(\lambda-\Delta)}{1-\beta}\;<\;0, \tag{A.12}$$

by (A.6), where Δ is the prevailing private surplus. The credit of Proposition 3 raises that surplus to $\Delta_{\tau_w^*} = \lambda$, so the wedge $\lambda - \Delta_{\tau_w^*}$ and the right-hand side vanish: the first-order cost of marginal UI is exactly the unpriced network externality, and it disappears once the externality is priced. □

A.8. The Household Problem and Equilibrium (Section 4)

A household in the quantitative model carries liquid assets k, earnings ability ε (a discrete Markov chain), and the fused exogenous state $x = (z, h, j)$: employment status $z \in \{e, u\}$, human capital h on a discrete grid, and sector j. The Bellman equation is

$$V(k,\varepsilon,x) = \max_{c,\,k'\geq \underline{k}} \; u(c) + \beta\,\mathbb{E}\big[V(k',\varepsilon',x')\,\big|\,\varepsilon,x\big], \qquad c + k' = (1+r)\,k + y(\varepsilon,x) - T, \tag{A.13}$$

where y is labor income $w\,\varepsilon\,h$ when employed, replacement income at rate b_{ui} when unemployed, plus lump-sum transfers and rebated dividends, and T collects taxes under the fiscal rule. The transition of x is the fused Markov matrix built each period from the sector finding rates (8) and separation rates s_j: employed households climb the h grid stochastically and separate at rate s_j; unemployed households descend the grid and find at rate $f_j \cdot (h/h_{ref})^{\zeta}$. Because f_j depends on $(K_j, \tilde{e}_j)$ and on h, status, human capital, and sector cannot be tracked as product dimensions and are fused into a single exogenous chain, as described in Section 4.1.

An equilibrium is a path of prices (w, r, π, mc), tightness θ, sector employment rates $\{\bar{e}_j\}$, and network stocks $\{K_j\}$ such that: (i) households solve (A.13); (ii) free entry prices vacancies at the firm discount factor β_f, $\kappa_v = \beta_f\, q(\theta)\, J^f$; (iii) each sector's employment rate is consistent with aggregation of the household distribution under the fused transition; (iv) each K_j obeys (10); (v) inflation satisfies the Phillips curve, monetary policy follows the Taylor rule, and the real rate obeys Fisher; (vi) the government budget holds under the debt-feedback rule, with UI and any credits inside it; and (vii) the goods market clears. Bond-market clearing is then implied by Walras's law and is used as the untargeted numerical audit reported with every experiment (Appendix E).

A.9. Verification

Every closed form in this appendix is checked against finite-difference Jacobians of the actual transition map over random parameter draws spanning the disciplined ranges; every costate result is computed twice, by the closed form or backward recursion and by brute-force simulation of discounted welfare; the decentralization of Proposition 3 is verified to machine precision; and the S-sector positivity premise is verified on the stacked system. The replication package reports the full check ledger.

B. Micro-Foundation of the Network Block

This appendix derives the laws of motion (1)–(3) from an explicit tie-and-referral process. Its purpose is to verify the reduced-form laws, not to add results: the laws the paper works with are the mean-field aggregate of the micro process below, exact in the large-population limit, and both the decay-acceleration bracket κ and the curvature η come out of it directly.

The micro model itself is developed in full in a companion paper (Kopecky 2026), the source of the decay-bracket aggregation of Section B.2 (its §2.3 and Appendix A), of the sole-candidate share and the convexity results of Section B.4, and of the fixed-network validation quoted in Section B.5. All results here are machine-verified in the replication package (checks M1–M7).

B.1. Primitives

Workers i hold a tie portfolio $n_i \in [0,1]$ (tie capacity normalized; $c \cdot n_i$ is the expected number of active contacts at capacity c). Employment status follows the aggregate flows: separations at rate s, finding equal to a direct component plus the referral hazard below. Matching is *annealed* — each period, the partner at the end of each unit of tie stock is a tie-weighted random draw from the population — which is the mean-field assumption itself. The operative aggregate is the tie-weighted employment rate $\tilde{e}_w \equiv \sum_i n_i z_i / \sum_i n_i$, the micro counterpart of the model's $\tilde{e}$.

Ties form through the workplace: employed workers add $\tilde{\varphi}\,(1-n_i)$ per period; the unemployed build nothing (the asymmetry of Bramoullé and Saint-Paul 2010). Ties decay at rates keyed to the joint status of their endpoints: employed–employed ties self-maintain ($d_{ee}=0$), employed–unemployed ties decay at d_{eu}, and unemployed–unemployed ties decay fastest, at $d_{uu} = d_{eu}(1+\tilde{\kappa})$.

Referral information takes one of two polar protocols: *broadcast* (non-rival: an employed worker's job information reaches each unemployed contact independently) and *rival token-passing* (in the spirit of Calvó-Armengol and Jackson 2004: one usable lead, passed to one uniformly drawn unemployed contact). Reality is a blend; the blend share is an estimable structural object, discussed below.

B.2. The Decay Bracket Is Exact Aggregation

Proposition B.1 (the κ bracket). *Under annealed matching, an unemployed worker's per-tie decay rate is* $d_{eu}\,\tilde{e}_w + d_{uu}\,(1-\tilde{e}_w) = d_{eu}\big[1+\tilde{\kappa}(1-\tilde{e}_w)\big]$*, which equals the model's* $\delta\big[1+\kappa(1-\tilde{e}/E^\circ)\big]$ *under the affine reparametrization* $\delta(1+\kappa) = d_{eu}(1+\tilde{\kappa})$*,* $\delta\kappa/E^\circ = d_{eu}\tilde{\kappa}$*. The mapping is exact.*

Two consequences. First, κ — with no measured counterpart in the aggregate data (Table 1) — corresponds to an observable micro object: the gap between u–u and e–u tie survival, estimable from any panel recording tie persistence and both endpoints' employment. At the calibrated (δ,κ) the implied rates are $d_{eu} = 0.112$ and $d_{uu} = 0.240$ per quarter: ties between two unemployed workers rot roughly twice as fast. Second, the sentence "your contacts' unemployment makes your own ties decay faster" is exact aggregation of the assumed tie process.

B.3. Accuracy of the Two-State Reduction

Proposition B.2 (closure)**.** *A simulation of the micro process with $N = 50{,}000$ workers, mapped to the model's parameters, reproduces the aggregate laws: steady-state employment 0.938 against the model's 0.940; mean tie stock 0.713 against 0.700; and along a severe shock transition, maximum deviations of 0.006 (employment) and 0.014 (stock).*

The closure works because of timescale separation: status churn is roughly $25\times$ faster than tie dynamics, so a worker's tie stock averages over many employment spells and the status–tie covariance stays small. The same slowness that makes the stock the persistence engine is what validates the aggregation. Two emergent properties deserve note. The friendship paradox appears without being assumed: employment histories alone generate $\mathrm{Cov}(n_i, z_i) > 0$, so contacts are more employed than the average worker. And the one systematic closure bias is economics rather than error: the unemployed pool selects toward low-tie workers, so sector-level aggregation slightly understates how concentrated the scarring is — the same selection that drives the distributional results of Section 6.2. It does not misstate aggregates.

B.4. Curvature Is the Protocol

Proposition B.3 (finding function)**.** *Under the broadcast protocol, the referral hazard of an unemployed worker aggregates exactly to the model's informal term, multiplicative in the stock and in network employment, with $\eta = 1$. Under rival token-passing, the measured cross-sectional elasticities are $\eta_{\text{eff}} \in [2.7, 7]$ and a stock elasticity below one, varying with contacts-at-capacity c (Kopecky 2026):*

contacts at capacity c	ε_K (SS)	η_{eff} (SS)	η_{eff} (slump trough)
5	0.90	2.7	2.2
10	0.80	4.3	3.1
20	0.64	7.1	4.2

Convexity table reproduced from Kopecky (2026), §4.

The mechanics: under rivalry, higher network employment supplies more referral tokens *and* thins the queue of competing searchers, and the competition-relief term generates curvature well above one. The paper's $\eta > 1$ is therefore not a free amplifier; it is what the standard description of referrals (one vacancy, one referral) produces, while $\eta = 1$ is the broadcast benchmark. The relevant structural question becomes the blend of the two protocols — the non-rival share of referral information, μ_{NR} — an estimable object with direct hooks to referral microdata.

B.5. Instruments Under Rivalry, and the Effort-Sign Condition

Holding network composition fixed, adding one percent of searcher tie-mass moves incumbent searchers' referral hazards by -0.99 percent under rivalry and exactly zero under broadcast: rivalry makes search partially business-stealing among searchers. On a fixed graph the congestion elasticity is attenuated roughly fivefold by locality (to about -0.2 at the steady state; Kopecky 2026), so the annealed figure is an upper bound. Three consequences for the paper's policy analysis:

1. *The wedge survives both protocols.* The Pigouvian wedge is positive under broadcast and larger under rivalry; the constrained-inefficiency verdict of Theorem 1 is protocol-free.

2. *Stock instruments are protocol-robust.* Hiring credits and the bridge subsidy act on stocks and on the cross-sector allocation of stocks; their externality components finance no rival transfer. Under rivalry the bridge subsidy gains a third positive component: a bridged searcher exits the home-sector referral queue, raising every remaining home searcher's hazard. Theorem 2 is robust a fortiori.

3. *The effort instrument is not.* Under rivalry a search-effort subsidy partly finances token-stealing. A searcher's effort creates a match only when they are the token holder's sole usable candidate; the surviving fraction of the effort margin is that sole-candidate share, and the sign of the corrective effort instrument turns on the blend: a net effort subsidy is warranted only if the non-rival share μ_{NR} exceeds roughly 0.18 at the steady state — and roughly 0.71 in a slump, because congestion is countercyclical (thresholds computed in Kopecky 2026 at that paper's illustrative 20 percent match-value wedge). This is the condition referenced in Section 3.5, and the reason the effort instrument is excluded from the policy analysis.

The cyclical effectiveness evidence of Section 2 (Fact 2) bears directly on the protocol: rival congestion predicts that friends-users' conversion deteriorates *faster* than non-users' in slack, and the data show the opposite. The positive interaction favors a high non-rival share — the protocol under which the reduced-form laws hold exactly and congestion is zero — while the quantitative results in Sections 6 and 7 are reported under both protocols via the quenched bands of Appendix F.

C. Data Appendix

This appendix documents the CPS pipeline behind Section 2: the construction of the search-method panel, every specification behind Fact 1 and Figure 2, the era-interaction regressions, the questionnaire-vintage break checks, and the month-linked effectiveness panel behind Fact 2.

C.1. Survey Files and Construction

The panel is assembled from the raw Census monthly basic CPS files: 389 months from January 1994 through June 2026 (October 2025 does not exist on the Census server — the shutdown gap; immaterial). Twenty-five data dictionary vintages are parsed; the search-method battery, the person weight, and the labor-force recode sit at byte-identical positions in all of them, and the method codes are unchanged over the whole sample: code 4 is "contacted friends or relatives," asked of unemployed active searchers. The pooled file holds 1.30 million unemployed person-records, of which 1.11 million are active lookers with a valid method battery (battery validity is 100 percent in every year). A searcher "uses the channel" if any of the six method slots records code 4; monthly shares weight by the CPS final person weight.

Unemployment-rate regressors are computed from the CPS microdata itself (the standard download services were unavailable to this build; the CPS-computed rates are internally consistent by construction). The national series matches published annual averages to the first decimal in every benchmark year checked (1994: 6.12 vs. 6.1; 2000: 4.05 vs. 4.0; 2009: 9.32 vs. 9.3; 2020: 8.08 vs. 8.1; 2024: 4.05 vs. 4.0). State rates are monthly, not seasonally adjusted, and unsmoothed by design; seasonality is absorbed by month or time fixed effects in every specification, and the sampling noise in monthly state rates is treated explicitly below.

C.2. Reliance Regressions: Every Specification

Table 8 reports the full specification set behind Fact 1. The individual-level regressions put the friends/relatives indicator on the unemployment rate with composition controls (age, sex, education, race, hispanicity, unemployment duration, occupation, industry) and the stated fixed effects; standard errors are clustered by state.

The gap between the aggregate and local margins is real, and part of it is measurement. Monthly state unemployment rates computed from CPS microdata carry substantial sampling noise: regressing one rotation-half state rate on the other (independent households,

Specification	Identification	Slope (pp/pp)	s.e.
Micro, national U, month FE + composition	aggregate cycle	$+0.88$	0.07
Micro, national U + linear trend	aggregate cycle	$+1.10$	0.09
National time series (NW-12)	aggregate cycle	$+1.30$	0.33
Micro, state U, state + time FE + composition	local slack	$+0.16$	0.12
Split-half rotation IV (errors-in-variables corrected)	local slack	$+0.39$	0.24
State $\times$ year panel	local slack	$+0.47$	0.22
State $\times$ month aggregate share	local slack	$+0.26$	0.13
Broad outcome (all method slots)	local slack	$+0.15$	0.12
Pre-2020 only	local slack	$+0.10$	0.13
Raw (no composition controls)	local slack	$+0.26$	0.12

Table 8: *Reliance on the friends/relatives channel: specification set.*
Notes: dependent variable: indicator for using the friends/relatives search method (individual-level rows) or the monthly share (time-series rows). Slopes in percentage points of the share per percentage point of unemployment. Standard errors clustered by state (micro rows); Newey–West with 12 lags (time-series rows). The pre-2020 local slope is approximately $+0.27$ after the errors-in-variables (EIV) correction.

same state-month, within state and time fixed effects) gives a reliability of 0.376, so the naive state slope recovers roughly 38 percent of the true coefficient. Two independent corrections — a split-half rotation-group IV and state-year aggregation — bracket 0.4 ($+0.39$ and $+0.47$). The remaining aggregate-local gap (roughly 0.9–1.1 against 0.4–0.5) is economics rather than measurement: formal hiring channels dry up with the aggregate cycle, not with state-level deviations from it, and cross-state job search dilutes the local experiment. The aggregate margin is the one the paper's experiments discipline against; the corrected local slope is the conservative bound.

C.3. The Distributed Lag

The lag profile of Figure 2(a) is the time-series regression of the monthly share on the national unemployment rate at lags 0, 6, and 12 months, with a linear trend, Newey–West standard errors with 12 lags, 377 months:

	lag 0	lag 6 months	lag 12 months
Coefficient (pp/pp)	$+0.80$	$+0.55$	$+0.47$
HAC s.e.	(0.36)	(0.24)	(0.30)

The contemporaneous and six-month coefficients are individually significant at the 5 percent level; the twelve-month coefficient is not ($t = 1.5$). The informative object is the profile: conditional on current unemployment, lagged unemployment continues to predict reliance, with roughly 56 percent of the cumulative response arriving after impact.

Era	Micro, era FE slope	s.e.	Micro, + trend slope	s.e.	Time series (NW-12, + trend) slope	s.e.
1994–2000 expansion	+2.16	0.15	+3.00	0.23	+3.17	0.36
2001 recession	+0.91	0.26	+0.82	0.25	+1.19	0.42
2005–07 expansion	−1.16	0.72	−0.49	0.70	−0.95	0.93
2008–13 recession	+0.54	0.14	+0.55	0.10	+1.03	0.21
2014–20 expansion	+1.38	0.14	+2.03	0.18	+2.29	0.28
COVID	+0.06	0.13	+0.12	0.14	+0.04	0.07
Post-2022	+0.93	0.49	+0.25	0.50	+0.30	0.55

Table 9: *Era-specific share slopes (pp per pp of U).*

Notes: the micro columns are individual-level regressions with composition controls; slopes are the U× era interactions with era fixed effects absorbed. The main text quotes the trend-controlled micro column. The 1994–2000 slope is contaminated by the 1998 questionnaire redesign (Section C.6) and is excluded from the body's comparison. The 2005–07 window contains almost no within-era unemployment variation, hence the unstable sign and large standard error. The 2001 and 2008–13 slopes are statistically indistinguishable (difference −0.27, s.e. 0.27).

C.4. Era-Interaction Regressions

Fact 4 and Figure 2(b) come from interacting the national unemployment rate with era indicators, eras absorbed, so each slope is identified from within-era time variation. Era windows: 1994m1–2001m2 ("1994–2000 expansion"), 2001m3–2004m12 ("2001 recession"), 2005m1–2007m11, 2007m12–2013m12 ("2008–13 recession"), 2014m1–2020m2 ("2014–20 expansion"), 2020m3–2021m12 ("COVID"), 2022m1– ("post-2022").

The ranking the model predicts — hiring-margin expansions steepest, sustained recessions near one, the COVID spike near zero — is the pattern in every column of Table 9 save the variation-starved 2005–07 window, set aside above.

C.5. Duration Composition

Sustained slumps shift the unemployed pool toward long-duration searchers, whose search-method mix differs, so the aggregate facts could in principle be composition rather than behavior. Table 10 removes it. Shares are computed within five duration bins (0–4, 5–14, 15–26, 27–52, and 53+ weeks) and reaggregated at the full-sample duration mix, so the fixed-composition series moves only when behavior within cells moves. The aggregate slope survives, at +1.10 plain and +0.77 with trend and breaks against +1.30 and +0.95 unadjusted; composition accounts for roughly 15 to 20 percent of the cyclicality. The decisive cut is within cells: the share rises with aggregate unemployment inside every duration bin, at +1.0 to +1.2 in the plain specification. The distributed lag survives reweighting, with a cumulative response of +1.54 against +1.81, and the build-up is present within cells and

	Original	Fixed composition
(a) Aggregate slope (pp/pp)		
plain	+1.30 (0.33)	+1.10 (0.30)
+ trend	+1.49 (0.31)	+1.28 (0.28)
+ trend + breaks	+0.95 (0.29)	+0.77 (0.25)
(b) Within-cell slopes, plain (pp/pp)		
0–4 weeks	+1.03 (0.28)	
5–14 weeks	+1.10 (0.29)	
15–26 weeks	+1.24 (0.27)	
27–52 weeks	+1.13 (0.34)	
53+ weeks	+1.08 (0.42)	
(c) Distributed lag (+ trend)		
impact	+0.80 (0.36)	+0.73 (0.35)
6 months	+0.55 (0.24)	+0.41 (0.23)
12 months	+0.46 (0.30)	+0.40 (0.29)

Table 10: *Duration-composition adjustment.*
Notes: fixed-composition series reweights within-cell shares by the full-sample duration distribution (weights 27.9, 28.0, 16.0, 15.3, 12.8 percent); Newey–West (12) standard errors in parentheses; era ordering on the fixed-composition series is unchanged, with COVID at +0.07 (0.06). Construction in `cps_duration_adjusted.py`.

strongest for long spells. The era ordering is unchanged, with COVID at +0.07 (s.e. 0.06). Controlling instead for the aggregate long-term-unemployment share does absorb the lag coefficients, but mechanically: that share is itself a distributed lag of unemployment, so the control over-fits the object under test; the fixed-composition series is the clean design. One number is reported plainly because it moves the wrong way: the composition-adjusted break-robust slope, +0.77, sits just below the +0.88 floor of the band that disciplines γ_w. The anchor moment sits inside the unadjusted band, which remains the target, and the conservative recalibration of Appendix G shows the headline is insensitive over exactly this range.

C.6. Questionnaire-Vintage Breaks

The national share jumps at two questionnaire-vintage boundaries (1997→98: −3.6pp; 2006→07: +4.6pp). Byte positions of the method battery are verified identical across vintages, so the jumps are in the data, not the parsing; instrument changes at those dates cannot be fully excluded. Three mitigations: the state + time fixed-effects design is immune to any national-level break by construction and still delivers a corrected local slope of at least 0.4; allowing level shifts at both boundaries, the trend-controlled time-series slope is +0.95 (s.e. 0.29), the center of the micro band, so the unadjusted trend specification (+1.50) is partly break-inflated while the disciplining moment itself is break-robust; and the era design of Section C.4 absorbs era means entirely, so within-era slopes are unaffected by level shifts at era boundaries.

C.7. Placebo Channels

If the friends-channel patterns reflected generic search intensification (desperate searchers using more methods of every kind) rather than a network-specific stock, other search methods should show the same signatures. They do not. Computing the same monthly weighted shares for two methods with no plausible network content — contacting employers directly and sending résumés or filling applications — and running the identical distributed-lag and era-interaction specifications:

	Distributed lag (pp/pp)			Era slopes (pp/pp)			
Channel	0m	6m	12m	exp. 14–20	GR 08–13	2001	COVID
Friends/relatives	+0.80	+0.55	+0.47	+2.28	+1.03	+1.22	+0.03
Direct employer contact	−0.37	−0.35	−0.28	−1.53	−0.33	+0.84	−0.38
Sent résumés/applications	+0.32	+0.37	−0.04	+1.54	+0.35	+1.41	+0.44

Direct employer contact shows no stock signature at all (it falls in slack). Sending résumés — the natural intensity margin — rises in every slump *including* the COVID spike (+0.44, standard error 0.19), and its lag profile dies out within a year. The friends channel is the only method whose response persists at twelve months and vanishes in the spike: the pair of signatures a slowly depreciating stock produces and generic intensification does not. Time-series specifications with trend, Newey–West standard errors (12 lags); the friends-channel era slopes here are the time-series analogues of Table 9's micro estimates.

C.8. The Effectiveness Panel (Fact 2)

Fact 2 requires linking each searcher to their labor-force status one month ahead. For each consecutive month pair, active lookers at t in rotation groups that can be followed (MIS 1–3 and 5–7) are matched to their $t+1$ records among all persons — keeping every labor-force status, so exits to nonparticipation stay in the denominator. Match keys follow the standard household-identifier design: household identifiers plus person line number, with the second household identifier from 2004m5 onward (recovered for 2004m5–2005m7, where it sits at a byte position the base parser misses) and state as a fallback key before; a match requires the rotation index to advance by exactly one and is validated on sex, race, age (within one year), and state.

The linked file holds 808,223 looker-records from 1994 onward, 724,725 with a valid battery and working age, and 499,296 in the primary sample (2004m5 onward, where keys are clean: validation-failure rates 0.1–0.6 percent, against roughly 0.3–1.5 percent under the pre-2004 fallback keys). Match rates run 90–94 percent in almost every year, with documented soft spots at the 1995 sample redesign, the 2004 identifier phase-in, and 2020–2026 response decline. Two external checks: linked-sample monthly job-finding is textbook procyclical and sits on the published series (14–16 percent at the 2009–13 trough, 29–30 percent at the 2000 peak); and linked-versus-full looker composition is near-identical on age, sex, channel use, and search intensity, so the link selects almost nothing beyond movers.

Three objects, all on the primary sample, weighted, clustered by state:

Level. Friends-users convert to employment 2.1pp worse raw and -0.95pp with composition and state and time fixed effects — but $+0.01$pp (s.e. 0.24) once the *number* of search methods is controlled. Friends-users search more intensively (each additional method predicts -0.70pp conversion, a difficulty marker); the friends method itself carries no conversion premium or penalty at the level.

Cyclicality (the Fact-2 interaction). The friends $\times$ unemployment interaction on conversion is positive in every specification: $+0.28$ (s.e. 0.07) in the aggregate-margin baseline; $+0.27$ with state and time fixed effects; $+0.19$ on the local margin and $+0.21$ after the split-half EIV correction of both the rate and its interaction; $+0.34$ controlling for search intensity and its cyclicality; $+0.25$ pre-COVID; $+0.30$ on the full 1994 sample; $+0.31$ on the broad method battery (t-statistics 3.3–4.6 throughout). The congestion prediction is a negative interaction; the data reject it in sign, not just magnitude.

Net. The channel's share of realized U$\rightarrow$E transitions, ψ, rises $+0.80$pp per pp of unemployment (s.e. 0.09): 21.1 percent of new hires in the lowest unemployment tercile had used the channel, 27.0 percent in the highest. Absolute flows through every channel fall in slumps (total job-finding falls 1.13pp per pp of unemployment); the network's *relative* role is what rises.

Caveats: the CPS records search methods, not hiring channels, so these are footprint objects; ψ conditions on the employment outcome and is an accounting share of realized hires, not a causal contrast; and the positive interaction, though robust to all observables and to search intensity, could in principle reflect cyclical selection on unobservables into channel use. The claim the paper rests on is the reduced-form one: no evidence of relative degradation, positive evidence of relative resilience.

D. The Empirical Mixing Kernel

D.1. Construction

The kernel m is built from the Census LEHD Job-to-Job Flows origin–destination files (public-use release R2026Q2, national coverage 2000Q2–2025Q2): job-to-job moves plus hires from persistent nonemployment, all workers, all firm ages and sizes, pooled over all available quarters so that the kernel is a long-run average mixing structure. The public files stack many nested geography-by-industry aggregation layers in one file; summing rows without pinning a single layer overcounts every flow (by a factor of roughly three in the raw totals). The build therefore pins exactly one layer — national origin by state destination, sector by sector — and sums it across the 51 state files, which reconstructs the national origin–destination sector matrix without double counting, because each job's destination lies in exactly one state.

Flows are aggregated from 20 NAICS sectors to the 9 BLS supersectors used in the model and row-normalized with the diagonal kept, so m_{jk} is the share of sector j's realized job transitions that go to sector k, and the diagonal is the within-sector (homophily) share. The model uses this row-stochastic matrix directly as the exposure operator in (10); because m is row-stochastic, the symmetric healthy steady state is exact, and the kernel matters only off the steady state.

D.2. Properties

Within-sector shares average 40 percent at $S = 9$ (33 percent at the native $S = 20$), so roughly 60 percent of job transitions cross sector lines: homophily with substantial bridging, the regime the mechanism needs. The diagonal by sector, largest to smallest: education/health 59 percent, goods 53, leisure/hospitality 46, trade/transport/utilities 43, financial 42, professional/business 39, information 28, government 25, other services 22. Every sector's single largest destination is itself; professional and business services, and trade/transport, are the receiving hubs.

Cross-sector centrality is dispersed and skewed — coefficient of variation 0.65, Gini 0.37, top-decile share 0.22 at $S = 9$ — and the dispersion survives coarsening: aggregating 20 to 9 sectors moves the statistics only modestly (CV 0.76 to 0.65, Gini 0.42 to 0.37) rather than collapsing them toward the uniform-mixing null. At these dimensions a power-law tail is not statistically identifiable, and the paper claims only measurable, coarsening-robust dispersion. One further caveat: the sector-level "friendship paradox" (a random sector's partners are more central than itself) is present unweighted but nets to approximately zero when weighted by employment, because the largest sectors are themselves central; no result in the paper leans on the weighted version.

D.3. The BEA Cross-Check

As a validation object, an input–output kernel is built from the BEA Make–Use framework (producers' prices, after redefinitions, 2013 vintage): direct-requirements and Leontief total-requirements matrices at the Sector (15) and Summary (71) levels, with technical coefficients computed against industry gross output. The computed Leontief inverse matches BEA's published total-requirements tables with correlation 0.9997–0.9998, which validates the construction. The labor and goods kernels correlate positively but weakly (off-diagonal Pearson 0.10, Spearman 0.27; hub rankings Spearman 0.35): related enough to cross-validate, distinct enough to be different objects. The goods kernel is much flatter (CV 0.24 at $S = 15$), reinforcing that the labor kernel, not input–output linkages, carries the mixing structure the mechanism runs on. The BEA kernel is used only for this cross-check, never as the social baseline.

D.4. Flows to Ties

The maintained assumption is that the cross-sector distribution of realized job transitions proxies the cross-sector distribution of employed contacts: where a sector's workers move is where their ties are. This is the natural revealed-preference reading, and the pooled long-run average is the robust object for an exposure operator, but it is an assumption, not a measurement of ties. Two mitigations: the kernel enters only as the weights in $\tilde{e}_j$, so results depend on its broad geography (homophily, the hub structure) rather than fine cells; and public-use suppression of small cells biases only the finest cells toward zero, not the sector aggregates. A time-varying kernel and direct tie measurement are left to future work.

E. Computation

E.1. Implementation

The model is solved in the sequence space (Auclert et al. 2021). The household block is a standard incomplete-markets HANK block: the fused exogenous state (employment status $\times$ human capital $\times$ sector) enters as a single Markov chain whose transition matrix is rebuilt each period from the sector finding rates (8) and separation rates, alongside an independent earnings-ability chain. Grids: 10 human capital points on $[0.1, 2.0]$ (reference level 1.8), 5 ability states, 50 asset points; at $S = 9$ the fused exogenous dimension has

$2 \times 10 \times 9 = 180$ states. All sector-count-dependent blocks (the exposure operator, the K_j laws, sector consistency conditions) are generated programmatically per S. Timing is recursive throughout: last period's stocks and employment enter this period's finding rates.

The dynamic unknowns are market tightness, the sector employment paths, the wage-bill auxiliary, and inflation; the targets are free entry, sector consistency, and goods-market clearing. Bond-market clearing is left untargeted and serves as a dynamic audit of Walras's law: it holds to roughly 10^{-8} on linear solutions and a few times 10^{-5} to 10^{-4} (absolute; a few times 10^{-5} relative to the debt stock) on nonlinear transitions, the residual reflecting horizon truncation at $T = 300$.

E.2. The Log-Tightness Reparametrization

Realistic-size slumps drive tightness far from the steady state, and the standard full-step quasi-Newton iteration can step through $\theta \leq 0$ and fail on shocks of empirically relevant size. The solver therefore uses $\log \theta$ as the dynamic unknown. The reparametrization is exact — it reproduces pre-existing solutions to machine precision — and converts the positivity constraint into an unconstrained problem.

E.3. Damped Homotopy for Deep Slumps

The library's nonlinear solver is a full-step chord iteration on the steady-state-factored Jacobian. On deep slumps the true Jacobian drifts far from its steady-state factorization and the cold-start iteration diverges even where an equilibrium exists. The production solver adds three elements: step damping with backtracking to the best iterate, homotopy continuation in shock size (each solution seeds the next larger shock), and a final full-output harvest pass. The fixed point is identical — on cases the unmodified library solver handles, the damped solver reproduces its solution to 5×10^{-9} — only the iteration path differs. Every number in the paper comes from a converged solve at its stated size.

E.4. The Free-Entry Corner

The economically meaningful boundary of the solvable region is the entry margin, not the network. Pushing demand shocks to increasing depth drives market tightness monotonically toward zero (minimum attained $\theta \approx 0.04$, vacancy-filling value nearly exhausted) while the loop gain never exceeds 0.33 — far from the fold at $L = 1$. Deep demand-driven slumps thus fail to exist in this model for the classic small-surplus reason: free entry shuts down before the network channel could tip anything. This is the scope limit reported in Section 5.3 — demand levers move unemployment by only about two percentage points at the disciplined parameters — and it belongs to the free-entry block, not the network block. Separation-driven slumps of the size the paper studies solve comfortably; the deepest aggregate separation experiments converge with tightness well inside the feasible region and the wedge machinery far from any singularity.

E.5. Accuracy and Cost

The steady state solves once (about 18 minutes at $S = 9$) and is cached; the household Jacobian computes once per horizon (roughly 70–120 seconds) and serves every dynamics-only variant, since the network, wage-rule, and entry-credit dials live outside the household block; policies that change the household budget directly (UI, the tax-financing convention, the dividend payout) instead re-run the household block on the altered paths, the replays of Section F.5. Linear transitions solve in seconds; nonlinear transitions in 25–260 seconds each; the Ramsey exercise of Section 7.3 is a laddered sequence of nonlinear solves. The steady-state asset-quartile cutoffs used in Tables 6 and 7 are 0.8, 2.3, and 7.1 in model units — roughly one, three, and nine quarters of wage income. Realized values of the internally set parameters: $\beta = 0.978$ (firm discount $\beta_f = 0.99$), matching efficiency 0.305, steady-state wage $w_{ss} \approx 0.78$, with the entry cost backed out to normalize $\theta_{ss} = 1$. Replication code, cached objects, and the full check ledger accompany the paper.

F. Additional Results

F.1. The Wage-Rigidity Hump Map

Section 5.2 states that the share moment identifies γ_w through a hump-shaped map. The map, on the anchor experiment at its linear 4pp scale with all other parameters fixed:

γ_w	0.05	0.2	0.5	1.0
Share slope (pp/pp)	$+0.61$	$+0.96$	$+0.21$	$+0.48$

Only γ_w near 0.2 reaches the CPS band ($+0.88$ to $+1.10$). A dense sweep of thirty-three values over $[0.02, 1.5]$ places the crossing on a single contiguous interval, $\gamma_w \in [0.17, 0.33]$, which brackets the standard new-hire rigidity value of 0.2 the calibration adopts; the moment reaches the band nowhere else. The roughness away from the interval is an artifact of the peak-ratio moment (the argmax of the unemployment path shifts between local peaks as γ_w changes), immaterial to a selection that turns on where the interval sits rather than on the values outside it.

	$\eta = 1$	$\eta = 2.86$	Quenched, broadcast	Quenched, rivalry
Sectoral scar ratio	1.25	1.33	[1.20, 1.31]	[1.21, 1.29]
Aggregate scar ratio	1.28	1.42	—	—
Goods h-scar ratio	—	1.31	—	—
Partner premium, snapshot (%)	$+2.6$	$+4.8$	$[+2.0, +4.5]$	$[+2.2, +4.3]$
Worker-margin τ_w^* peak, % of wage	—	23.9	[18.9, 22.4]	[19.8, 21.3]

Table 11: *Band grid for the main results.*
Notes: scar ratios are live-to-frozen cumulative discounted losses (Table 4); the premium is Section 6.3's strongest-versus-weakest partner contrast; the τ_w^* row is the theory-priced ($1.0\times$) peak of the wedge-scaled worker-margin *subsidy of Proposition 3 — the entry-margin credit's priced path peaks at 7.8 percent of the quarterly wage (Section 7.1), and the two are different margin-specific prices (Section F.7). Quenched corrections from the fixed-graph benchmarks of Appendix B, applied to the amplification component; dashes denote cells not separately computed.*

F.2. Effectiveness and Fixed-Graph Bands

Where both are computed, the main magnitudes are reported with two disciplined bands: the effectiveness band $\eta \in [1, 2.86]$ (from linear to the Glitz-implied curvature) and the quenched-graph corrections, which replace the mean-field (annealed) approximation with fixed-graph benchmarks under each referral protocol, scaling the network's amplification contribution by factors of $[0.79, 0.94]$ (broadcast) and $[0.83, 0.89]$ (rivalry).

The scar multiple is robust to the tie decay rate δ. The microdata place δ's half-life only loosely, between one and three years, and across that whole range the sectoral multiple barely moves: 1.27 at an eleven-quarter half-life, 1.33 at the calibrated five and a half, 1.38 at two and a half, edging up as faster decay lets a sustained slump bite the stock harder.

The multiple is robust to the referral share χ too: from 0.35 to the calibrated 0.45 the sectoral multiple rises smoothly from 1.24 to 1.33. The share cannot be pushed much past 0.45 against the same shock, for a reason in the economics rather than the numerics. Job-finding through the formal channel relies on firms posting vacancies, and in a deep enough slump a job is worth so little that firms stop posting (the free-entry limit of Section 5.3). The more of hiring that runs through the network, the thinner the formal channel is, so a smaller slump exhausts it; by $\chi = 0.5$ the full China decline is enough to collapse formal hiring, and the model has no solution left in which to compare the live and frozen paths.

F.3. The Ramsey Grid

Table 12 reports every point of the numerical Ramsey exercise of Section 7.3: the credit path $\tau(\text{scale}, p) = \text{scale} \times \text{peak} \times \text{shape}^p$, where shape is the wedge's own normalized time profile, $p = 1$ is the priced rule, and each cell is a full nonlinear solve.

Scale	Shape power p	CE gain (%)	PV cost	Scar removed (%)	Gain/cost
0.25	1.0	$+0.85$	0.021	27.0	40.3
0.50	1.0	$+1.40$	0.044	48.3	31.7
0.625	1.0	$+1.57$	0.056	57.4	27.9
0.25	0.5	$+1.09$	0.031	38.4	34.8
0.25	2.0	$+0.54$	0.011	14.3	47.0
0.50	2.0	$+0.93$	0.024	26.4	39.3

Table 12: *The Ramsey grid.*
Notes: CE gain in percent of lifetime consumption; PV cost in model units (0.044 ≈ 2.5% of one quarter's GDP); scar removed is the share of the baseline sectoral scar. Scales above 0.625 on the wedge's shape lie beyond the chord solver's reach (Appendix E); the curve is concave through the computed range, with the marginal gain falling roughly 40 percent per eighth-step, so the half-priced rule attains 89 percent of the best-found gain. The flattened shape at scale 0.50 is not in the computed grid; at equal cost the flattened profile does no better than the wedge's own shape. Front-loaded small credits have the highest return per fiscal dollar: the earliest hires are the most valuable.

	Sectoral cum. loss	Scar removed (%)	CE gain (%)	PV cost
Live baseline	4.776	—	—	—
Live credit	2.469	48.3	$+1.40$	0.044
Frozen baseline	3.591	—	—	—
Frozen credit	2.047	43.0	$+0.94$	0.044

Table 13: *The frozen-credit placebo: four arms under the identical credit path.*
Notes: losses are cumulative discounted sectoral employment-years; scar removed and CE gains are relative to each economy's own baseline; the frozen arms use the frozen-network dials of Section 6.1 with the live economy's implemented τ_h path. Dynamic Walras residual 1.4×10^{-4} in every arm. Computed by `paper_a_placebo_frozen_credit.py`.

F.4. The Frozen-Credit Placebo

The identical implemented credit path (the half-priced rule of Section 7.1) is run in the frozen-network economy of Section 6.1. Table 13 reports the four arms. The live baseline and live credit reproduce the published numbers exactly, the frozen baseline matches the China-shock counterfactual to machine precision, and the frozen credit converges at the full live path. The network share of the credit's value, one minus the frozen-to-live ratio, is 32.4 percent by welfare, 33.1 by scar removal, and 32.9 per fiscal dollar.

F.5. Payout and Financing Conventions

The baseline rebates profits lump-sum, which overstates poor households' dividend income. The variant re-runs the household block on the solved general-equilibrium paths under a deviation-payout convention (dividend deviations accrue to asset returns rather than

	Lump-sum rebate			Deviation payout		
	Credit, Goods	Credit, healthy	UI, Goods	Credit, Goods	Credit, healthy	UI, Goods
Q1 (poorest)	−3.38	−2.60	−0.46	−1.99	−1.51	−0.80
Q2	+4.90	+2.06	+0.43	+4.09	+2.87	+0.40
Q3	+10.28	+8.44	+0.36	+9.85	+5.98	+0.70
Q4 (richest)	+2.82	+0.87	−0.07	+3.89	+1.83	+0.84
Total (all)	credit +1.40; UI +0.00			credit +1.66; UI +0.05		

Table 14: *Equal-cost welfare by payout convention (CE %).*
Notes: the payout convention explains roughly 40 percent of the bottom quartile's loss under the credit ($-3.38 \to -1.99$); the residual runs through the debt-financed credit's lump-sum tax path. Progressive financing shrinks the cell by a further fifth of a point ($-3.38 \to -3.16$, computed by the same replay); neither amendment alone, nor the two combined, reverses the sign. Total welfare gains rise under the realistic convention: moving payout risk from hand-to-mouth households to wealth-holders raises the value of the same policy.

to current income; total payout conserved). With the lump-sum convention the replay reproduces the full GE welfare totals to 3×10^{-5}, and its scope is the pure re-incidence of dividend risk holding macro paths fixed — the general-equilibrium feedback of the redistribution is not in these numbers. An exact-GE version of the deviation-payout economy is borderline non-contractive for the steady-state-anchored solver and did not converge at full payout intensity; it is left to future work, and the tables report the replay. Table 14 reports the full matrices.

The same replay answers the financing side. The credit's cost is met by a tax the baseline levies lump-sum; the progressive variant routes the tax deviation onto wealth instead, sparing hand-to-mouth households, with the aggregate tax conserved. The wiring check holds (at a flat tax the replay reproduces the credit's welfare matrix to 3×10^{-5}), and the effect is small: progressive financing moves the bottom quartile's Goods-sector credit CE from -3.38 to -3.16 percent, about a fifth of a point, because the credit's tax bill is modest and the poorest would bear little of it under either convention. The financing convention is the minor lever and the dividend convention the major one; neither alone, nor the two combined, reverses the sign.

F.6. Spike Versus Persistent Shock Shapes

The share moment itself distinguishes shock shapes, in the direction the era evidence requires. On the eight-quarter spike of Figure 5, the model's share slope is +0.72 at a 1.4pp unemployment rise, against +0.96 on the persistent anchor shape at 4pp (linear); solved nonlinearly at increasing depth, the persistent moment rises above the CPS band, from +1.22 (2.4pp) to +1.37 (3.7pp). Sectoral employment after the spike heals with a half-life of two quarters live against one frozen, from shock end.

F.7. Margin Conversion: A Cautionary Note

The theory prices the worker-margin correction as a flow employment subsidy τ_w^*; its one-off hiring-bonus equivalent is $\tau = \tau_w^*/(1-\beta(1-s))$; capitalizing the steady-state τ_w^* (11.3 percent of the quarterly wage) gives roughly three quarters' pay per hire, because a bonus paid at every future re-employment is capitalized into the value of unemployment as well. That worker-margin bonus does *not* transplant to the firm's entry margin: free entry capitalizes a persistent per-hire payment again, compounding it by roughly $1/(1-\beta_f(1-s)\rho)$ — a further factor of seven to eight — and the implied tightness explodes (finding rates exceed one; the solver correctly finds no equilibrium). The entry margin's own Pigouvian price is Proposition 1's τ_h, about 10.5 percent of the value of a filled job, an order of magnitude smaller, with the wedge-scaled path peaking at 7.8 percent of the quarterly wage. In general, margin-specific prices do not transplant across margins; results are reported in flow form.

G. Additional Robustness

G.1. The Decay-Acceleration Parameter

κ has no measured counterpart and, as Section 5.2 notes, the disciplining moment is flat in it over the checked range: the anchor slope is $+0.96$ at both $\kappa = 1.0$ and $\kappa = 1.5$, to two decimals, at the calibrated curvature. Its influence on the theory side is second-order at the steady state — the analytical wedge moves from 11.9 to 13.8 percent of the private value as κ runs from 0 to 3 in the disciplined one-sector benchmark — though it matters for the state-contingent protection component of the bridge subsidy. The gap between that benchmark and the quantitative model's 26.6 percent is curvature and geometry, not machinery: at the Glitz curvature the same one-sector benchmark prices the wedge at 21.2 percent, the full mean-field model at 23.9, and the empirical kernel at 26.6. Its micro counterpart (the u–u versus e–u tie-survival gap) is observable in principle (Appendix B), which is where discipline should eventually come from.

G.2. Alternative Anchors

The calibration's conclusions do not depend on the specific anchor cell. Across alternatives: the linear anchor at 4pp gives $+0.96$; nonlinear solves at increasing depth give $+1.22$ to $+1.37$ (2.4 to 3.7pp), bracketing the CPS band from above; the low-curvature band endpoint

($\eta = 1$) gives $+1.47$, so the effectiveness band brackets the anchor from above as well; and the spike shape gives $+0.72$ (Appendix F). Demand-margin moments are several times the band at sub-2pp unemployment responses and turn over at depth as the share saturates; as Section 5.3 records, this margin cannot reach a 2008-scale slump and is not used as an anchor. Under every alternative in this list, $\gamma_w \approx 0.2$ remains the only wage-rigidity value reaching the band, and the headline scar decomposition moves within the bands of Table 11.

Targeting the corrected cross-state band ($+0.39$ to $+0.47$) instead of the aggregate band selects $\gamma_w = 0.87$, with the anchor moment at $+0.43$. The steady state, and with it the loop gain and the wedge, is invariant to γ_w by construction; on the China shock the sectoral multiple is 1.32 against the baseline 1.33, the aggregate 1.40 against 1.42, and healing times move by one quarter. The disciplining moment is a share-response object that wage rigidity prices; the scar itself runs through the stock, and recalibrating the one barely touches the other.

G.3. Pre-2020 Subsamples

Dropping COVID-era data changes no sign. The corrected local reliance slope is approximately $+0.27$ pre-2020 (against $+0.39$ to $+0.47$ on the full sample), and the uncorrected slope loses statistical significance without the COVID window ($+0.10$, s.e. 0.13); the effectiveness interaction is $+0.25$ pre-COVID against $+0.28$ on the primary sample; and the aggregate margin's main identifying episode (2008–13) predates COVID entirely. The COVID months matter for Fact 4 — they supply the spike — but for none of the disciplining moments.

G.4. Cross-Sector Spillovers

The experiment's damage is heavily concentrated in the shocked sector: healthy sectors lose an order of magnitude less consumption and one to two orders less human capital than the shocked sector (Table 6), and their wedges peak at roughly half the shocked sector's (Figure 7). The partner-composition contrast of Section 6.3 — the same own shock scarring the focal sector more alongside its strongest partner than its weakest, with the predicted sign surviving in relative terms even where joint scars are sub-additive in levels — is the sharpest spillover test the kernel admits. Full per-sector spillover matrices for every experiment are in the replication package.