Can Workers Hedge AI by Owning It? Labor Adjustment and the Limits of Financial Insurance*

Murray Z. Frank Vojislav Maksimovic

August 27, 2026

Abstract

Can workers hedge AI displacement by owning the firms that benefit from AI? We study a quantitative model in which wages are determined by a labor-market institution, firms choose worker retention after a technology shock, and households choose portfolios. The labor market determines how much of the shock stays common across workers, and so contractible in aggregate securities, and how much becomes a lottery over which worker is displaced. Across a historically informed range of wage flexibility, the welfare cost of that lottery falls by two thirds, but the share of it that complete aggregate markets can remove stays near 2%. The natural hedge does not close the gap: AI-sector equity explains 99.6% of the variation in household consumption but only 1.3% of the variation in the welfare-relevant displacement loss, and less still when AI capability is allowed to raise the value of the firms that adopt it. Wage flexibility shrinks the displacement lottery; it does not make the lottery insurable. The insurable share of technological displacement is set in the labor market, not in the security market.

*Frank is from University of Minnesota. E-mail: murra280@umn.edu. Maksimovic is from University of Maryland. E-mail: max@rhsmith.umd.edu. Generative AI assistants were used during the preparation of this manuscript to search for errors, to improve the writing, and to help with the code. We reviewed and edited every suggestion, and the substantive scientific decisions are ours alone. We take full responsibility for the final manuscript, including any remaining mistakes or misinterpretations.

1 Introduction

Technological change creates a natural hedging argument. The same innovation that lowers the value of some workers' human capital can raise the value of the firms that develop or adopt it. Artificial intelligence makes the argument especially salient: if workers lose labor income in the same aggregate states in which AI-related capital appreciates, owning that capital would let workers share in the gains from the technology that threatens their jobs. The premise has empirical support—firms whose workforces are exposed to generative AI gained value after ChatGPT (Eisfeldt et al., 2026)—and it underlies both a household-finance question and a policy intuition about broadening ownership of AI-related capital. But it presumes that the labor market delivers the technology shock to workers in a form that an aggregate security can hedge.

In this paper, we study an economy in which AI capability raises the automated share of tasks, firms choose how many workers to retain conditional on a wage-setting regime, and households choose financial-market participation and portfolios. In that economy we show two results. First, the labor market determines the *amount and form* of worker risk that arrives at financial markets: greater wage responsiveness routes more of a technology shock through pay, whereas adjustment through separations creates different outcomes for otherwise identical workers in the same aggregate state. Second, once worker-specific displacement risk exists, aggregate financial markets reach only a small part of its welfare-relevant cost, and the failure does not depend on how equity moves with AI. Whether AI-sector equity rises or falls with capability, a sign the announcement-return evidence does not settle, equity explains roughly one percent of the welfare-relevant displacement loss, and less still when capability is allowed to raise firm value. Wage flexibility shrinks the displacement lottery; it does not make the lottery insurable.

The mechanism is simple but changes the insurance question. The aggregate state of the economy has two components, the productivity of the AI sector and the capability of the technology, and both are common to every worker. Conditional on that state, an exposed worker is either retained or displaced, and the probability of displacement is chosen by the firm. A common technology shock therefore generates a worker-specific outcome after the aggregate state is known. The transfer that would equalize consumption between a retained and a displaced worker takes from one and gives to the other within the same state, so it averages to zero across workers in that state. Any security whose payoff depends only on the aggregate state pays the two workers the same amount and cannot reproduce that transfer. An aggregate technology shock need not generate an

aggregate insurable risk. That observation is a definition restated. The economic content is that firms' labor decisions determine how large the within-state lottery becomes, so the insurable share of displacement is an equilibrium object rather than a primitive.

We first ask how the labor-adjustment regime changes the risk that must be insured. Holding technology and preference primitives fixed, we vary the responsiveness of the retained wage to the technology shock across a historically informed range, from a regime in which wages do not respond at all to the most flexible regime the evidence supports, and re-solve equilibrium capital at each point. The within-state identity component of worker risk falls to about one third of its benchmark value, and the welfare cost of displacement under the baseline asset menu falls by roughly two thirds, from 84.33 to 29.76 annual basis points of consumption. Labor-market adjustment has a first-order effect on the size of the displacement problem that financial markets confront. Section 4 reports the displacement rates and the full path.

We next ask whether existing equity provides the hedge the opening intuition suggests. For total consumption the answer appears to be yes: the AI-sector equity payoff explains 99.6% of the variation in lottery-averaged consumption for participants. But that span is driven by common exposure to AI-sector productivity, and it is not insurance against displacement. When the target is the marginal-utility-weighted displacement loss, equity explains only 1.3% of its variance for participants. Allowing an arbitrary payoff on the full modeled aggregate state raises that figure only to a few percent. The nonparticipant figures, reported in Section 5, tell the same story. The main financial gap is not between existing equity and a better aggregate security; it is between aggregate-state risk and the worker-specific displacement lottery.

We then give aggregate finance its strongest opportunity in equilibrium. Starting from the baseline menu, we let households trade a complete set of zero-net-supply Arrow claims on all eighteen aggregate states, holding the installed capital stock fixed so that the exercise isolates risk sharing. The displacement burden falls by 2.16 annual basis points. Relative to the baseline burden, this is a 2.56% reduction. For comparisons across wage-setting regimes, we also report a companion self-normalized measure, equal to 2.21% at the benchmark, which remains close to 2% over the experiment. The two normalizations answer slightly different accounting questions and are reported separately. The share and the span statistic are different objects, one a measurability statistic at the benchmark allocation and the other an equilibrium certainty-equivalent comparison after portfolios and prices adjust, but they carry the same implication: completing trade in the modeled aggregate state leaves most of the displacement burden in place. What makes this a result

about institutions rather than about one calibration is that the boundary does not expand when the wage regime changes. Across the full range of wage responsiveness the completion share stays near 2% while the burden itself falls by two thirds. Greater wage responsiveness does not turn a large worker-specific lottery into an insurable one; it makes the lottery smaller.

The quantitative environment is disciplined by evidence rather than fitted to the welfare result. Because mature AI displacement has not yet been observed, we use wage and employment responses to an earlier large labor-demand shock to construct a historically motivated scenario range for the wage-versus-employment adjustment margin. Announcement returns around AI capability events delimit the equity return loading, which is a nuisance parameter here rather than the mechanism. Eight moments discipline six freely chosen parameters, and an untargeted displaced-worker earnings loss of 32.9% lies inside the empirical range of roughly 30–40%. Preference curvature, the scale of capability risk, and related welfare inputs are treated as external calibrations. The magnitudes should be read accordingly. The absolute welfare cost of displacement is sensitive to those inputs and is only weakly disciplined by the estimated moments, and across deliberately wide external-calibration grids the completion share can change sign when small and can rise to about 10%. These are criterion-set and calibration ranges, not sampling intervals. The paper's claim is about an economic ordering, that the reach of aggregate finance is modest in every environment studied and does not grow with wage flexibility, not a precisely estimated universal number.

These findings separate two stages that are often combined in discussions of technological risk. Labor-market institutions and firm choices determine how the technology shock is transmitted to workers, through wages, through employment, or through both. Financial markets then determine how well the aggregate component of the resulting household risk can be shared. Greater wage responsiveness does make capability-linked aggregate claims more useful for *total* consumption, because more capability risk passes through wages. It does not make the retained-versus-displaced lottery more aggregate-state measurable. The policy boundary follows. Broader equity ownership can let households participate in aggregate gains from AI, but a security written on an aggregate outcome cannot transfer resources between two workers in the same aggregate state who realize different employment outcomes. Reaching that margin requires employment-contingent contracts or transfers.

The paper contributes first to the literature on human capital and household portfolio choice. A large literature studies portfolio decisions when labor income is nontradable or imperfectly insured (Mankiw and Zeldes, 1991; Constantinides and Duffie, 1996; Cocco et al., 2005; Benzoni et al., 2007;

Catherine, 2022), and related work measures the covariance of occupation-level income shocks with traded returns (Davis and Willen, 2013). Most directly, Huang et al. (2026) construct equity projection portfolios for skill-specific human-capital shocks, including an AI portfolio, and evaluate their hedging and welfare properties. We ask a logically prior question. Rather than taking the worker's labor-income process as given and asking which portfolio hedges it, we let an aggregate technology shock pass through firms' wage and retention decisions and ask how much of the resulting worker risk lies in the aggregate financial payoff space at all. We then compare existing equity with the complete set of aggregate claims rather than asking only whether a particular projection portfolio can hedge the observed process.

The paper also connects the automation and displacement literature with work on who owns the capital that benefits. Task-based models show how automation changes labor demand (Acemoglu and Restrepo, 2018, 2020). In finance, Gârleanu et al. (2012) show that when innovation displaces the human capital of existing cohorts and that risk cannot be fully shared, displacement risk is priced in equilibrium asset returns; Kogan et al. (2020) trace the losses from creative destruction to the workers left behind, and Eisfeldt et al. (2026) measure the revaluation of exposed firms after ChatGPT. We ask the complementary question. Taking as given that displacement risk exists and may be priced, we ask whether the securities that carry it can be used by the exposed workers to insure the displacement itself. In Gârleanu et al. (2012) the missing risk sharing is across cohorts; here the component no aggregate security can reach lies within the aggregate state, between retained and displaced workers who are identical before the shock resolves, and its size is set by the labor market. Recent work measures heterogeneity in worker- and firm-level AI exposure (Hampole et al., 2025; Kogan et al., 2023), and Rockall et al. (2025) document that AI exposure is concentrated among higher-income workers, who also hold more wealth and larger risky-asset shares. That evidence is what makes the natural-hedge argument plausible: the workers most exposed to AI may also be unusually well placed to benefit from its capital gains. Our contribution is to distinguish ownership from insurance. Even where exposed households own equity, those holdings insure displacement only to the extent that equity pays on the component of worker welfare that displacement generates, and the model lets that component be measured directly.

Third, the paper connects work on insurance inside the employment relationship with work on incomplete markets and security design. Guiso et al. (2005) show that firms absorb part of firm-level shocks rather than passing them fully to workers, and Meeuwis et al. (2026) model how wage contracts, limited commitment, and financial conditions jointly determine earnings and separation

risk. We take the next step in the insurance chain: conditional on the wage and separation functions the labor market delivers, what can outside financial markets insure? On the financial side, additional securities can improve risk sharing when relevant aggregate states are not spanned by existing assets (Duffie and Rahi, 1995; Shiller, 1993; Athanasoulis and Shiller, 2001). Our complete Arrow-menu experiment deliberately gives that logic its strongest opportunity within the modeled aggregate state. The distinction is between a *menu gap* and a *measurability gap*: a missing traded security can be repaired by financial innovation, whereas a payoff that is identical for retained and displaced workers within the same aggregate state cannot directly insure their employment-outcome difference. We do not claim that idiosyncratic labor-income risk is new, that firms have not previously been viewed as insurers, or that additional aggregate securities cannot improve risk sharing. The new object is the *endogenous insurable share* of technological displacement, and the new result is that the labor market sets its size while leaving its insurable fraction essentially unchanged.

The remainder of the paper proceeds as follows. Section 2 develops the labor-adjustment mechanism and the measurability result. Section 3 describes the quantitative discipline, identification, and validation. Section 4 studies how labor-market adjustment changes the risk to be insured, and Section 5 measures the limits of financial insurance. Section 6 discusses sensitivity, scope, and what could change the conclusions. Section 7 concludes.

2 Labor Adjustment and Worker-Specific Displacement Risk

This section asks how a common technology shock becomes a risk borne by an individual worker. The sequence matters for the financial question studied in the paper. Firms observe the technology state before households can insure the labor-market consequences of that state. They respond through wages and retention. Wage adjustment leaves workers exposed to a common change in pay; employment adjustment instead creates uncertainty about which otherwise identical worker is displaced. The latter risk remains heterogeneous even conditional on the aggregate state and therefore cannot be reproduced by a payoff that depends only on that state.

The theory establishes this boundary, not its quantitative importance. In particular, it does not imply that changing the wage-setting regime must raise the *share* of displacement risk that aggregate financial markets can reach. Section 3 disciplines the wage and financial margins of the model. Section 4 then asks how changing the wage-setting regime changes the amount and form of worker risk, and Section 5 measures how much of the remaining displacement risk aggregate

financial markets can span.

2.1 Wage setting and endogenous retention

At date 0, the firm enters with installed technology capital K and a workforce that can be retained or displaced at date 1. At date 1, the aggregate state $s = (\varepsilon, \theta)$ is realized. The shock ε is an AI-sector productivity shock. The shock θ indexes the capability of the new technology and therefore the extent to which installed technology capital can automate tasks. Keeping the shocks distinct will matter later: equity can be highly exposed to AI-sector productivity even when worker displacement is driven primarily by technology capability.

Technology capital and capability determine the share of tasks that are automated,

$$z(K, \theta) = \bar{z}\big(1 - e^{-aK\theta}\big), \tag{1}$$

where $0 < z < \bar{z}$. A higher θ raises z and shifts the production technology away from labor and toward installed technology capital.

The firm does not freely reset its retained wage after observing the technology shock. Instead, the wage is determined by the wage-setting regime,

$$\bar{w} = \bar{w}_0 \, \varepsilon^{\varpi} (1 - z)^{\nu_\theta}, \tag{2}$$

where $\nu_\theta \geq 0$ measures how strongly the retained wage responds to automation. At $\nu_\theta = 0$, the retained wage does not respond directly to the capability shock. As ν_θ rises, a larger part of the adjustment to higher automation occurs through the wage. We treat ν_θ as a feature of the wage-setting environment—for example, bargaining arrangements, contract length, legal restrictions, or norms surrounding wage cuts—and the firm takes the wage schedule as given.

Conditional on that wage schedule, the firm chooses retention. Let N denote retained employment and

$$\lambda = 1 - N \tag{3}$$

denote the displacement probability for an incumbent worker. At an interior solution, retained employment satisfies

$$\mathrm{MPL}(N, K; s) = \bar{w}\big(1 - \psi\lambda\big), \qquad \lambda = 1 - N, \tag{4}$$

where MPL is the value marginal product of labor. The firm also pays separation costs $\frac{1}{2}\psi\bar{w}(1-N)^2$,

so (4) equates the value of retaining another worker to her wage net of the marginal separation cost that retention avoids. The problem is strictly concave in N, and Appendix A treats the full-retention corner explicitly.

A displaced worker receives the outside option

$$b(s, K; \lambda) = \max\Big\{b_0,\ \beta_T A_T(\bar{L}_T + \lambda)^{\beta_T - 1} - \phi_R\Big\}. \tag{5}$$

The second term is the marginal product of labor in a traditional sector that absorbs displaced workers, net of a reallocation cost ϕ_R. Because $\beta_T < 1$, additional displacement congests that sector and lowers the outside option. Define the retained–displaced consumption gap in this two-date representation as

$$\Delta c(s) = \bar{w}(s) - b(s, K; \lambda(s)). \tag{6}$$

Both parts of this gap are therefore endogenous conditional on the wage-setting regime: automation affects the retained wage, while the firm's retention choice affects the displaced worker's outside option.

The distinction between the institution and the firm's choice is important. The paper does not explain why an economy has a particular ν_θ. It asks what follows for worker risk once that wage-setting environment is given. Conditional on ν_θ, the firm chooses N, and the equilibrium determines λ, b, and Δc. The first object we need is therefore not an asset return. It is the worker risk produced by this labor-market adjustment.

2.2 From aggregate shocks to worker-specific risk

Let Y^L denote date-1 labor income. Conditional on the aggregate state s, an otherwise identical worker is retained with probability $1 - \lambda(s)$ and displaced with probability $\lambda(s)$. The law of total variance separates two kinds of labor-income risk:

$$\mathrm{Var}(Y^L) = \underbrace{\mathrm{Var}\big(\mathbb{E}[Y^L \mid s]\big)}_{\text{variation across aggregate states}} + \underbrace{\mathbb{E}\big[\mathrm{Var}(Y^L \mid s)\big]}_{\text{variation across workers within a state}}. \tag{7}$$

The first term is conventional aggregate risk. The second remains after s is known: two workers facing the same technology state can receive different labor outcomes because one is retained and the other is displaced.

With only the retained and displaced outcomes, the within-state component is

$$\mathcal{I}(\nu_\theta) \equiv \mathbb{E}\big[\mathrm{Var}(Y^L \mid s)\big] = \mathbb{E}\big[\lambda(s)\big(1-\lambda(s)\big)\Delta c(s)^2\big]\,. \tag{8}$$

We call $\mathcal{I}$ the *conditional identity component*. It is large when displacement is likely and when the retained–displaced gap is large. Crucially, neither input is primitive: both are generated by the wage-setting environment and the firm's equilibrium retention decision.

Equation (8) is a diagnostic for the mechanism, not a welfare measure. To connect the same lottery to insurance, let $c_R(s)$ and $c_D(s)$ denote consumption in the retained and displaced outcomes, with $\Delta c(s) = c_R(s) - c_D(s)$. Full within-state insurance would give both workers the resource-average consumption

$$c^F(s) = \big(1-\lambda(s)\big)c_R(s) + \lambda(s)c_D(s). \tag{9}$$

The resource-balanced transfer $d = c^F - c$ that implements this allocation is

$$\begin{aligned} d_R(s) &= -\lambda(s)\Delta c(s), \\ d_D(s) &= \big(1-\lambda(s)\big)\Delta c(s), \end{aligned} \tag{10}$$

and hence

$$\mathbb{E}[d \mid s] = 0. \tag{11}$$

The transfer needed to insure the employment lottery therefore moves resources between retained and displaced workers *within* an aggregate state. This observation leads directly to the financial boundary.

2.3 What aggregate securities can insure

Lemma 1 (Aggregate-state measurability)**.** *Let $X = X(s)$ be the payoff of any security whose payoff depends only on the aggregate state. Conditional on s, any portfolio of such securities pays the same amount to a retained and a displaced worker. It therefore cannot reproduce the employment-contingent transfer in* (10) *when $0 < \lambda(s) < 1$ and $\Delta c(s) \neq 0$. Equivalently,*

$$\mathrm{Var}\big(Y^L + X(s) \mid s\big) = \mathrm{Var}(Y^L \mid s). \tag{12}$$

The proof is in Appendix A. The argument uses only conditioning: once s is fixed, an aggregate-state payoff is a constant across the retained and displaced outcomes.

Lemma 1 is deliberately modest. It is the familiar incomplete-markets observation that a security cannot condition on information outside its payoff state. The paper's economic contribution is not this measurability fact. It is to place the fact *after* the firm's labor-market decision: a common technology shock first passes through wages and retention, and that equilibrium response determines the size of the worker-specific lottery confronting financial markets.

The lemma also does not imply that aggregate securities have zero welfare value for displaced workers. They can transfer resources across aggregate states and thereby change consumption, marginal utility, prices, and equilibrium portfolios. Indeed, although the raw insurance transfer is conditionally balanced, its marginal-utility-weighted counterpart generally is not:

$$\mathbb{E}\big[u'(c)d \mid s\big] = \lambda(1-\lambda)\Delta c\left[u'(c_D) - u'(c_R)\right] > 0, \tag{13}$$

when $c_D < c_R$ and utility is concave. Aggregate trading can therefore affect the *welfare valuation* of the employment lottery without directly insuring worker identity. This is why the quantitative analysis will examine both the aggregate-state measurable component of $u'(c)d$ and the equilibrium change in the displacement burden when the aggregate security menu is completed.

This distinction also separates two possible failures of financial insurance. Existing traded securities may load on the wrong aggregate shock, in which case a better aggregate security could help. Or the relevant loss may vary mainly across workers within the same aggregate state, in which case changing the aggregate security menu cannot reproduce the required transfer. Section 5 measures the importance of these two margins rather than assuming either one is quantitatively dominant.

2.4 How wage adjustment changes displacement

The remaining theoretical question is how the wage-setting regime changes the employment response that creates the lottery. Wages and retention are substitute margins of adjustment.

Proposition 1 (Wage adjustment and the displacement response)**.** *At an interior retention solution, the sensitivity of displacement to the capability shock is*

$$\frac{\partial \lambda}{\partial \theta} = \frac{z_\theta}{D}\left[\frac{1-\nu_\theta}{1-z} - \ell_N\right], \qquad z_\theta = aK(\bar{z}-z) > 0, \tag{14}$$

where

$$D = \frac{z}{N} + \frac{\psi}{1-\psi\lambda} > 0, \qquad \ell_N = \ln\left(\frac{KN_0}{K_0 N}\right). \tag{15}$$

For a given state and local retention point, the term in brackets is strictly decreasing in ν_θ, with derivative $-1/(1-z) < 0$. Thus a more responsive retained wage weakens the direct displacement response to a capability shock. The bracket is zero at the state-dependent threshold

$$\nu^*(s) = 1 - \ell_N(1-z). \tag{16}$$

Appendix A derives Proposition 1, treats the full-retention corner, and verifies the expression against the solved retention rule. The qualification "for a given state and local retention point" is essential. The objects N, D, ℓ_N, the outside option, and equilibrium capital can themselves change with ν_θ. The proposition therefore identifies a direct substitution between the wage and headcount margins; it does not sign every general-equilibrium response.

The intuition is immediate. A larger ν_θ makes the retained wage fall more after a capability-induced increase in automation. Holding the other local equilibrium objects fixed, retention becomes cheaper and the firm responds less through displacement. But three stronger conclusions do *not* follow from the proposition. It does not imply that the identity component $\mathcal{I}(\nu_\theta)$ is globally monotone, because λ, the outside option, and the retained–displaced gap all move in equilibrium. It does not imply that lower displacement raises worker welfare, because wage adjustment can reduce the consumption of retained workers. And it does not imply that the *fraction* of the remaining displacement loss measurable with respect to the aggregate state must rise with ν_θ.

Those distinctions determine the order of the quantitative analysis. Section 3 first uses wage, employment, and financial-payoff evidence to discipline the model. Section 4 then holds technology and preference primitives fixed and changes the wage-setting regime to ask how much worker-specific displacement risk the labor market produces. Section 5 asks the logically separate question: conditional on the risk that remains, how much can existing equity, the full aggregate state, and complete aggregate financial markets actually reach? The theory tells us why these are different questions; the quantitative model determines their magnitudes.

The mechanism is not specific to artificial intelligence. It applies whenever a common technology shock can be absorbed through both wages and employment. AI is the application in which the paper disciplines the separate productivity, capability, equity-payoff, and displacement margins. Likewise, the reduced-form wage schedule in (2) is not essential to the measurability result. A richer

contracting model can alter wage, separation, and outside-option schedules and thereby change the size of the worker-specific lottery. Once two workers nevertheless realize different employment outcomes in the same aggregate state, an aggregate-state payoff still cannot condition on which worker received which outcome. Section 6 returns to this scope and to the missing worker quit margin.

3 Quantitative Discipline, Identification, and Validation

Section 2 established two facts that are qualitative: firms can turn a common technology shock into a worker-specific displacement lottery, and aggregate-state securities cannot directly reproduce the transfer that insures that lottery. Neither fact tells us how large the lottery is, how much of its welfare-relevant component remains aggregate, or how much households gain from trading a richer security menu. Answering those questions requires a quantitative economy that jointly determines labor adjustment, firm payoffs, household portfolios, and welfare.

This section describes only the parts of that economy needed to interpret the results. We first define the quantitative objects that Sections 4 and 5 will use. We then separate parameters fitted to targeted moments from objects disciplined by external evidence or maintained calibration. Finally, we ask whether weakly disciplined parameter directions can materially change the paper's financial conclusion and whether the calibrated displacement process reproduces a labor-market fact that was not targeted in estimation. The complete environment, parameter inventory, and numerical algorithm are reported in the Internet Appendix.

3.1 Quantitative implementation and the objects of interest

The aggregate state remains $s = (\varepsilon, \theta)$. The productivity shock ε has two realizations and the capability shock θ is represented by nine nodes, giving eighteen aggregate states. Firms enter the period with installed technology capital and choose retention state by state as in Section 2.1. The same equilibrium therefore determines firm cash flows, retained wages, displacement probabilities, and the outside option of displaced workers.

Households differ in initial wealth and in a cost of participating in the equity market. A participant may hold the risk-free asset and existing AI-sector equity; a nonparticipant holds only the risk-free asset. Equity positions are long-only and unlevered in the benchmark. These portfolio restrictions matter for incidence, but they are not imposed as the explanation for limited insurance:

Section 5 will compare existing equity with the entire payoff space generated by the modeled aggregate state.

We use three asset menus for the central comparisons. Let A_0 denote the benchmark allocation with the observed participation friction and existing equity. Let A_θ augment the financial opportunity set with claims that span the capability state θ. Let A_s provide zero-net-supply Arrow claims on the full aggregate state $s = (\varepsilon, \theta)$. The last menu is the paper's aggregate-financial-market-completion experiment. In these financial counterfactuals, installed capital is held at its benchmark equilibrium value while security prices and household portfolios adjust. Thus the A_0-to-A_s comparison isolates risk sharing rather than capital deepening. This differs deliberately from the institutional experiment in Section 4, which changes the wage-setting regime and re-solves the capital stock.

To measure the welfare importance of the employment lottery, let $e(V)$ denote certainty-equivalent consumption delivering utility V. For an asset menu A, let $V(A)$ be equilibrium utility and let $V^F(A)$ evaluate the same financial allocation after the resource-balanced transfer in (10) removes the retained–displaced consumption difference within each aggregate state. The implemented per-period burden fraction is

$$d(A) = \frac{e\big(V^F(A)\big) - e(V(A))}{e(V(A_0))}. \tag{17}$$

The A_0 certainty equivalent is the common numeraire for all menus within a given economy. With the model period equal to three years, we report

$$\mathcal{D}(A) = 10^4 \left\{ \big[1 + d(A)\big]^{1/3} - 1 \right\} \quad \text{annual basis points.} \tag{18}$$

Thus $\mathcal{D}(A)$ is not the total welfare cost of technological change. It is the menu-indexed consumption-equivalent burden associated with the employment-contingent displacement lottery. Appendix B provides the exact construction and the self-normalized companion used for comparisons across wage-setting regimes.

A complementary statistic asks how much of the welfare-relevant displacement object is measurable with respect to the aggregate state before changing the asset menu. For any random variable Y, define the correlation ratio

$$\eta_s^2(Y) = \frac{\mathrm{Var}(\mathbb{E}[Y \mid s])}{\mathrm{Var}(Y)}. \tag{19}$$

Section 2 showed why the informative target is $Y = u'(c)d$, not the raw balanced transfer d. Hence $\eta_s^2(u'(c)d)$ is a security-menu-independent ceiling on the variance of the welfare-relevant displacement object that any payoff measurable with respect to s can span. Section 5 reports this statistic alongside the equilibrium change in $\mathcal{D}(A)$ when the menu is expanded from A_0 to A_s.

All headline welfare comparisons use the same one-period model horizon; the consumption-equivalent difference is formed first and then reported in annualized basis points. Persistence is economically important but is not estimated in the benchmark, so Section 6 treats the horizon as a separate specification sensitivity rather than embedding an uncalibrated continuation parameter in the headline result.

3.2 Empirical discipline and parameter accounting

The quantitative exercise uses three different forms of discipline, and the paper keeps them distinct. Six parameters are jointly fitted to eight moments. The wage-adjustment parameter ν_θ is instead disciplined using wage and employment responses to an earlier large labor-demand shock. The equity loading on capability news is disciplined by announcement returns around AI capability events. Preference curvature, the scale of the capability shock, the outside-option specification, and the horizon are maintained calibrations rather than objects identified by the eight targeted moments. Table 1 makes this accounting explicit.

The historical evidence does not identify a mature-AI wage regime directly. Instead, we map wage and employment responses to the China import-competition shock (Autor et al., 2013) into a scenario range for the relative importance of wage and employment adjustment. The primary mapping places the point estimate near $\nu_\theta = 0.021$. We use $[0, 0.13]$ as a historically informed U.S. scenario range rather than as a confidence interval for a technology-invariant structural parameter. The mapping is based on wage changes among workers who remain directly exposed and employed; broader local-labor-market estimates mix that margin with reemployment wages and spillovers and therefore do not measure the object entering the firm's retention problem. The Internet Appendix gives the underlying mapping and source audit.

The financial block requires the response of AI-sector equity to capability news. We summarize that response by β_R. Announcement-window estimates are noisy: the pooled response is close to zero and event-level estimates vary in sign. We therefore use the event evidence to delimit an admissible region, not to claim a precisely estimated return loading or to infer worker exposure from stock returns. The benchmark constrained economy lies at $\beta_R = -0.0559$ on the boundary of

Table 1: Parameter discipline

Parameter	Economic role	Benchmark	Treatment/source
Panel A. Parameters fitted jointly to the eight targeted moments			
$\bar{\kappa}$	Mean participation cost	0.6217	Fitted
σ_κ	Dispersion of participation costs	2.3279	Fitted
a	Automation intensity	0.8173	Fitted
A_T	Traditional-sector productivity	2.2524	Fitted
A_{AI}	AI-sector productivity scale	2.7095	Fitted
$\bar{w}_0$	AI-sector wage-schedule scale	1.8590	Fitted
Panel B. Labor and financial exposure			
ν_θ	Wage response to automation	0	Institutional scenario; mapped point 0.021, paper range $[0, 0.13]$
β_R	Equity loading on capability news	-0.0559	Constrained by event-study magnitude restriction $\lvert\beta_R\rvert \leq 0.056$; profiled in Figure 1
Panel C. External welfare and outside-option inputs			
$\mathrm{Var}(\theta)$	Capability-shock variance	0.04	Anchored externally; varied in Section 6
γ	Relative risk aversion	5	Externally fixed; varied in Section 6
ϕ_R	Reallocation cost after displacement	0.2	Externally fixed
b_0	Floor on displaced outside option	0.3	Externally fixed
β_T	Traditional-sector curvature	0.65	Externally fixed
$\bar{L}_T$	Traditional-sector labor stock	1	Normalization
Horizon	Welfare horizon	one period	CE difference formed first, then annualized

The six Panel-A parameters are fitted jointly. Panel B contains objects that are not estimated by the eight-moment criterion: ν_θ is mapped from historical wage and employment responses and β_R is constrained using AI announcement-return evidence. Panel C lists external inputs with first-order importance for the welfare calculation. The Internet Appendix reports the complete parameter inventory, including the remaining technology parameters, wealth-distribution parameters, shock-grid definitions, derived quantities, and numerical settings.

the event-study magnitude restriction $|\beta_R| \leq 0.056$. Because this constraint is empirically rather than economically central, the relevant test is whether the paper's financial conclusion survives movement along the admissible β_R ridge. Section 3.3 performs that test.

The six fitted parameters minimize

$$L_8(\vartheta) = \sum_{j=1}^{8} \left(\frac{m_j(\vartheta) - \widehat{m}_j}{\widehat{\sigma}_j} \right)^2, \tag{20}$$

conditional on the maintained inputs in Table 1. The eight moments cover market participation, its wealth gradient, the capitalist financing share of AI capital, average automation and displacement, the AI-to-traditional wage ratio, the AI equity premium, and relative equity volatility. Because these moments come from different sources and their joint sampling covariance matrix is unavailable, the criterion uses the reported standard errors as diagonal weights. We therefore use L_8 as a

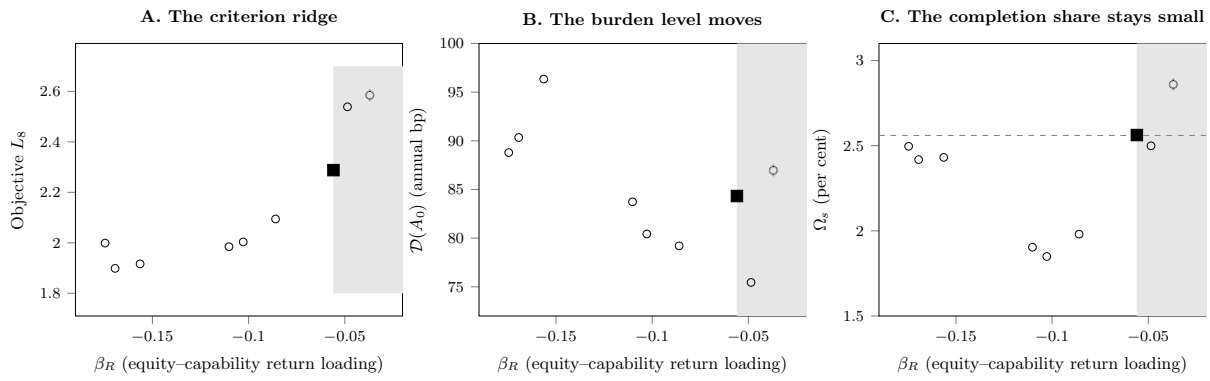


Figure 1: Identification along the equity–capability return-loading ridge.

Notes. Panel A reports the eight-moment criterion L_8 across stored β_R economies. Panel B reports the benchmark displacement burden $\mathcal{D}(A_0)$ in annualized basis points. Panel C reports $\Omega_s = [\mathcal{D}(A_0) - \mathcal{D}(A_s)]/\mathcal{D}(A_0)$. The shaded region is the event-study magnitude restriction $|\beta_R| \leq 0.056$ on the negative-loading side shown in the figure. The filled square is the constrained benchmark $\beta_R = -0.0559$; the gray diamond marks the raw confidence-interval edge $\beta_R = -0.037$. The horizontal dashed line in Panel C marks the benchmark completion share.

transparent minimum-distance fit criterion; we do not interpret its minimized value as a formal efficient-GMM J test. Search, continuation, and numerical-accuracy details are reported in the Internet Appendix.

3.3 Identification of the financial conclusion

The main identification concern is not whether every primitive has a tight standard error. It is whether parameter combinations that fit nearly as well can reverse the economic answer the paper ultimately gives. The clearest weak direction is the equity–capability loading β_R. Figure 1 therefore profiles both the criterion and the economic objects of interest along this direction.

The figure separates parameter uncertainty from uncertainty about the paper's financial conclusion. Panel A shows that the criterion is shallow in β_R and continues to improve outside the event-study admissible region. Thus the data used in the model do not sharply pin down this loading. Panel B shows why that matters: even within the admissible region, the level of $\mathcal{D}(A_0)$ moves appreciably, from roughly 75 to 87 annual basis points across the displayed admissible nodes. An absolute welfare level should therefore not be presented as tightly identified by this exercise.

Panel C gives the comparison that matters for the financial conclusion. Over the same admissible nodes, completing the modeled aggregate market removes only about 2.5–2.9 percent of the displacement burden. The figure therefore supports a narrower statement: the qualitative conclusion that aggregate-market completion has limited reach is stable along this weakly disciplined

Table 2: Targeted fit and an untargeted displacement validation

Moment	Data	S.E.	Model	Std. deviation
Panel A. Targeted moments				
Equity-market participation rate	0.5200	0.0300	0.5286	+0.29
Participation gradient in wealth	0.3500	0.0800	0.3531	+0.04
Capitalist financing share of AI capital	0.1000	0.0600	0.1172	+0.29
Expected automation rate $\mathbb{E}[z]$	0.2500	0.0800	0.3617	+1.40
Expected displacement $\mathbb{E}[\lambda]$	0.0700	0.0200	0.0667	−0.16
Wage ratio w_A/w_T	1.3000	0.1500	1.2955	−0.03
AI equity premium, annual	0.0622	0.0694	0.0820	+0.29
Relative equity volatility σ_{AI}/σ_m	1.1434	0.0993	1.1680	+0.25
			$L_8 = \sum_j z_j^2$	2.288
Panel B. Untargeted validation				
Displaced-worker earnings loss: model				0.3292
Documented empirical range				$[0.30, 0.40]$

Panel A reports the data moments used in estimation and the benchmark model values. The standardized deviation is (model − data)/S.E., so the squared deviations sum to $L_8 = 2.288$. The equity-premium and volatility moments use the adopter-portfolio convention employed in estimation. Panel B is untargeted: the empirical earnings-loss range does not enter L_8. The comparison therefore evaluates the severity of the modeled displacement process, not the welfare level, which also depends on the outside option, wealth, portfolios, and preferences.

return- loading direction, even though the level of the burden is not. It does not establish that every structural parameter or every welfare magnitude is sharply identified.

A local Jacobian reveals another weak direction involving the mean and dispersion of participation costs $(\bar{\kappa}, \sigma_\kappa)$. We report the full local diagnostics in the Internet Appendix rather than assigning each parameter a single "identifying moment." Those diagnostics help explain the fit, but they do not replace the economically relevant profile exercise in Figure 1. The remaining question is whether the benchmark economy fits the targeted labor and financial moments and, importantly, whether its displacement process reproduces evidence that was not used to fit it.

3.4 Targeted fit and untargeted validation

Table 2 answers two distinct questions. Panel A asks how closely the benchmark economy matches the eight moments used in estimation. Panel B then asks whether the same economy produces a plausible earnings loss for a displaced worker, an outcome that is not part of the fitting criterion. The second comparison is particularly important here because a model could fit asset-market moments while making displacement mechanically inexpensive.

The targeted fit is uneven in an informative way. Seven of the eight moments lie within 0.3

standard deviations of their targets. The exception is the average automation rate, for which the model is 1.40 standard deviations above the data value. Its squared deviation is 1.96, about 86 percent of the total criterion $L_8 = 2.288$. Thus the low criterion should not be described as eight uniformly close matches; most of the remaining tension is concentrated in the automation moment. This is also the moment most directly connected to the technology-to-labor mapping, so the discrepancy is economically relevant rather than a numerical footnote.

Panel B supplies a different kind of evidence. Empirical displacement studies place the earnings loss of displaced workers in a range of roughly 30–40 percent. The model produces 32.9 percent, inside that range, even though the range is not targeted in estimation. This comparison supports the claim that the model generates an economically plausible severity of displacement. It does not validate the absolute welfare burden: translating earnings losses into welfare also depends on the outside option, wealth, preferences, and portfolios. With the displacement mechanism and its severity disciplined in this sense, the next section can ask how the amount of worker-specific risk changes when the wage-setting regime changes.

3.5 What the quantitative discipline does not establish

Three limitations should be kept separate. First, the ranges reported for $\mathcal{D}(A_0)$ and for the welfare effect of completing aggregate markets are criterion-set and calibration ranges, not sampling confidence intervals. They describe how the objects vary across the economies the evidence admits, not how they would vary across repeated samples from a fixed economy.

Second, Figure 1 and the near-minimum criterion exercises measure numerical or profile sensitivity, not sampling uncertainty. They ask whether alternative parameter combinations that the criterion has difficulty distinguishing imply a different economic answer. Section 6 reports the broader criterion-set exercises; the Internet Appendix gives the full local diagnostics.

Third, several economically important inputs are maintained calibrations. Preference curvature, the capability-shock variance, the displaced-worker outside option, and the horizon can materially move the level of the welfare burden. Section 6 varies these objects explicitly rather than folding them into an artificial confidence interval. For this reason, the quantitative results below place more weight on comparisons and orderings than on interpreting any one annual-basis-point level as a precisely estimated constant.

The discipline in this section is sufficient for the next question but does not answer it. Section 4 now changes the wage-setting regime while holding the benchmark technology and preference

primitives fixed and re-solving equilibrium capital. That experiment asks how much worker-specific displacement risk the labor market produces. Section 5 then asks the separate financial question: how much of the risk that remains can existing equity, the aggregate state, and complete aggregate financial markets actually reach?

4 How Labor-Market Adjustment Changes the Risk to Be Insured

Section 2.4 showed that a more responsive retained wage weakens firms' direct incentive to adjust to a capability shock through displacement. That local result does not determine how much worker-specific risk survives after capital, retention, and the outside option adjust in equilibrium. We now measure that response quantitatively. The purpose of this section is to answer two questions in sequence. First, how much does the wage-setting regime change the *amount* of worker-specific displacement risk? Second, when adjustment shifts away from employment, where does the technology risk go? The financial question—how much of the remaining displacement loss aggregate securities can reach—is left to Section 5.

4.1 The institutional experiment

The preferred comparison changes one institutional object. Starting from the benchmark economy disciplined in Section 3, we hold every fitted and externally calibrated primitive fixed except the wage–automation elasticity ν_θ. At each value of ν_θ we re-solve the equilibrium capital stock K^* and then recompute firms' state-contingent retention decisions, household allocations, and the displacement burden. Thus the exercise changes the wage-setting regime while allowing the equilibrium quantities that respond to that regime to adjust. Because the wage-scale parameter $\bar{w}_0$ is held fixed, varying ν_θ changes the wage schedule in levels as well as its responsiveness to capability. The exercise should therefore be interpreted as a fixed-primitives comparison of wage-setting regimes, not as a compensated elasticity experiment that holds the mean or reference-state retained wage fixed.

This fixed-calibration experiment is deliberately different from re-estimating the six free parameters at every value of ν_θ. Re-estimation asks how the other primitives must move to keep matching the same U.S. moments when the wage regime is changed. That is useful for understanding what U.S. moment matching absorbs, and we report it in Appendix IA.E. It is not the clean institutional comparative static. Here a higher value of the minimum-distance criterion against the U.S. targets

Table 3: Labor-market adjustment and worker-specific displacement risk

ν_θ	$\mathbb{E}[\lambda]$ (per cent)	$\mathcal{I}(\nu_\theta)/\mathcal{I}(0)$	$\mathcal{D}(A_0)$ (annual bp)	K^*
0.00	6.670	1.000	84.33	0.635
0.05	5.622	0.686	60.01	0.682
0.10	4.585	0.447	39.85	0.702
0.13	3.942	0.330	29.76	0.709

Notes. The table varies only the wage–automation elasticity ν_θ. All fitted and externally calibrated primitives are held at their benchmark values, while equilibrium capital is re-solved at every node. The four rows lie in the historically informed U.S. range. λ is the incumbent worker's displacement probability; $\mathcal{I}(\nu_\theta) = \mathbb{E}[\lambda(1-\lambda)(c_R - c_D)^2]$ is the within-state identity component; and $\mathcal{D}(A_0)$ is the menu-indexed displacement burden defined in Section 3.1. The rows are institutional comparative statics, not separate re-estimations to the U.S. target moments.

simply says that the counterfactual economy no longer resembles the United States; it is not a criterion for discarding the counterfactual.

The empirical mapping in Section 3.2 places the historically informed U.S. range at $\nu_\theta \in [0, 0.13]$. We therefore use $\{0, 0.05, 0.10, 0.13\}$ for the main numerical comparison. Values above that range are useful for tracing the mechanism, but we do not attach country labels to them.

4.2 Wage flexibility sharply reduces worker-specific displacement risk

Does the wage-setting regime matter enough to change the economic size of the risk studied in the paper? Table 3 gives a clear answer. Moving from $\nu_\theta = 0$ to 0.13 lowers expected displacement from 6.7% to 3.9%, a decline of about 41%. The conditional identity component $\mathcal{I}(\nu_\theta)$ falls more sharply, to about one third of its benchmark value. The menu-A_0 displacement burden falls from 84.33 to 29.76 annual basis points, a decline of about 65%. The intermediate nodes show the same ordering rather than an endpoint comparison driving the result.

The economic interpretation is narrower than "wage flexibility raises welfare." A higher ν_θ reduces displacement partly because retained wages absorb more of the capability shock. Table 3 therefore establishes that the institution changes the *incidence and size of the worker-specific lottery*; it does not rank total welfare across wage regimes. Nor does the table show that financial markets can insure a larger fraction of the remaining lottery. Those are separate questions.

The last column also raises an immediate alternative explanation. Equilibrium capital rises by about 11.6% between the first and last rows. If capital deepening, rather than wage adjustment,

produced the fall in displacement, the institutional interpretation would be misleading. We return to this concern in Section 4.5. Before doing so, Figure 2 asks what replaces the worker-specific risk as wages become more responsive.

4.3 Wage adjustment reroutes technology risk

Figure 2 makes the change in the composition of risk visible. Panel A normalizes expected displacement and the conditional identity component to 100 in the benchmark economy so that their rates of decline can be compared on a common scale; Table 3 provides the economically interpretable levels. Both objects fall throughout the U.S.-mapped region and continue to fall as the experiment moves toward more wage-responsive regimes. The identity component falls faster because wage adjustment affects both the probability of displacement and the retained–displaced consumption gap.

Panel B asks where some of the technology risk goes. For nonparticipants, the best linear payoff on the capability shock θ explains 5.8% of total-consumption variation at $\nu_\theta = 0$ and 10.4% at $\nu_\theta = 0.13$. Thus, as firms absorb more of the capability shock through retained wages, total household consumption becomes more exposed to an aggregate capability-linked payoff. This is the quantitative counterpart of the wage channel in Proposition 1: the institution changes the *form* in which technology risk reaches households.

The rise in the capability-linked consumption span should not be attributed entirely to the wage institution. Because K^* rises with ν_θ, greater installed technology capital also amplifies household exposure to capability. Pinning K at its benchmark value leaves about 78% of the increase in the nonparticipant θ-linked consumption span at $\nu_\theta = 0.13$. Hence most, but not all, of Panel B's movement is a direct wage-regime effect; roughly one quarter is associated with capital deepening.

Just as important is what Figure 2 does *not* show. Panel B concerns *total consumption*, not the marginal-utility-weighted displacement loss. A more wage-responsive economy can make capability-linked aggregate insurance more useful for consumption while simultaneously making the remaining employment lottery smaller. Whether the remaining displacement loss itself becomes more aggregate-state measurable is the question taken up in Section 5. Appendix IA.E also reports a secondary implication for asset prices: the equilibrium loading of AI-sector equity on capability news changes sign as ν_θ rises. We do not use that untested cross-economy prediction to support the main institutional result.

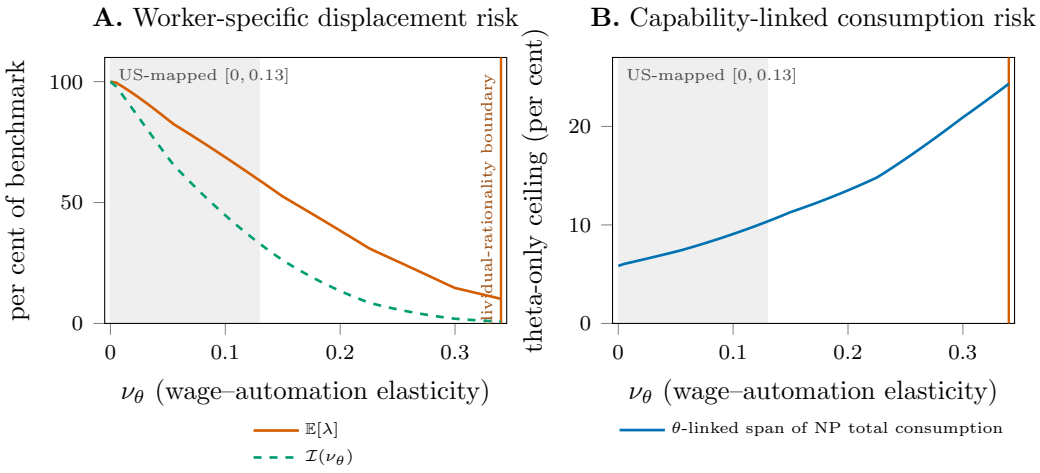


Figure 2: Labor-market adjustment changes the amount and form of technology risk

Notes. The figure uses the fixed-calibration institutional experiment of Section 4.1. The shaded region is the historically informed U.S. range $\nu_\theta \in [0, 0.13]$. Panel A reports expected displacement and the conditional identity component as percentages of their values at $\nu_\theta = 0$; the corresponding levels are reported in Table 3. Panel B reports the fraction of nonparticipant total-consumption variance spanned by the best linear payoff on θ. The vertical line at $\nu_\theta = 0.34$ marks the individual-rationality boundary described in Section 4.4. No main-text inference uses values to the right of that boundary.

4.4 The economically interpretable domain

How far can this comparative static be pushed before the maintained labor-market model stops representing a coherent worker choice? The answer is considerably below the mechanical "full flexibility" limit. At approximately $\nu_\theta = 0.34$, the retained wage first falls below the displaced worker's outside option in a high-capability state that still carries displacement. A worker in that state would prefer to leave, but the model contains no quit margin: employment is chosen by the firm and the worker is treated as accepting retention. Beyond this point the allocation therefore violates individual rationality, so 0.34 is the upper bound of the economically interpretable comparative static.

This boundary is distinct from the later point at which displacement mechanically disappears. The model drives expected displacement below 0.1% only around $\nu_\theta = 0.449$. That later extinction point is useful for understanding the model's mechanics but is outside the domain in which the labor allocation should be interpreted economically. Figure 2 therefore ends at the first boundary. All four main rows in Table 3 lie well inside it. Appendix IA.E reports the fine-grid boundary search, the later extinction point, and the diagnostic continuation beyond the main-paper domain.

The domain restriction also clarifies what a richer labor model could change. Allowing workers to quit would endogenously prevent retained wages from falling below the outside option and would

distinguish voluntary from involuntary separations. Doing so would add a new selection margin rather than provide a small robustness perturbation. We therefore leave that extension outside the current paper and interpret the comparative static only on the individually rational domain.

4.5 Mechanism checks

The first check isolates the capital channel suggested by Table 3. We repeat selected ν_θ changes while pinning K at its benchmark value. These cells are a partial-equilibrium decomposition—the capital market does not clear away from the benchmark—and we do not describe them as counterfactual equilibria. At $\nu_\theta = 0.13$, however, the exercise is highly informative about attribution: approximately 105% of the free-K movement in expected displacement, 99% of the movement in $\mathcal{I}(\nu_\theta)$, and 98% of the movement in $\mathcal{D}(A_0)$ survive when capital is pinned. Values slightly above 100 percent mean that capital deepening modestly offsets rather than creates the direct wage-regime effect. The central fall in worker-specific displacement risk is therefore not a capital-accumulation confound.

A second check separates displacement incidence from its consequences. The outside option affects displaced-worker consumption and therefore the welfare cost of the lottery, but it does not enter the firm's retention first-order condition in the maintained model. Consistent with that structure, varying the pre-existing outside-option grid changes the burden while leaving expected displacement exactly unchanged. The two exercises together yield a useful economic decomposition: ν_θ primarily determines how much adjustment occurs through displacement, whereas the outside option determines how costly a given displacement outcome is. The full capital-pinning and outside-option calculations are reported in Appendix IA.E.

The evidence in this section therefore changes the interpretation of the paper in a specific way. Labor-market institutions do not merely change a parameter in the worker's income process; they determine how much of a common capability shock is converted into worker-specific employment risk and how much remains on an aggregate wage margin. What this section has not established is that the *fraction* of the remaining displacement loss reachable by financial markets changes with the institution. Section 5 now turns to that separate question.

5 The Limits of Financial Insurance

Section 4 showed that the labor-market adjustment regime can change the size of the displacement problem dramatically. Over the historically informed range, expected displacement falls by about forty percent and the menu-A_0 displacement burden falls by about sixty-five percent. That result does not answer the financial question. A smaller lottery could become easier to hedge, harder to hedge, or remain mostly outside the payoff space of aggregate securities. This section asks which of those possibilities the model delivers.

We use two deliberately different tests. The first is a statistical span test: at a given allocation, what fraction of the welfare-relevant displacement object can be represented by existing equity or, more generously, by an arbitrary payoff on the full aggregate state? The second is an equilibrium welfare test: how much does the displacement burden change when households can trade a complete menu of aggregate-state claims? The first locates where the risk sits; the second asks what market completion actually accomplishes. Their magnitudes need not coincide, but they should point in the same economic direction if the paper's interpretation is correct.

5.1 Existing equity spans consumption, not displacement

Does owning AI-exposed equity provide the natural hedge suggested by the motivating intuition? A projection of total consumption initially makes the case look strong. At the benchmark, equity explains 99.6% of lottery-averaged consumption variation for participants and 94.8% for nonparticipants. Panel A of Table 4 reports these numbers. They are nearly complete statistical spans, but they do not measure displacement insurance.

The reason is the distinction between the two aggregate shocks introduced in Section 2. AI-sector equity is primarily exposed to the productivity shock ε, whereas displacement is driven primarily by the capability shock θ and then by the worker-specific retention draw. At the benchmark, the correlation between the AI-equity return and displacement is about -0.57. Almost all of that correlation runs through the productivity shock, which lowers profits and raises separations together, as it would for any firm; the capability shock accounts for about four percent of it, and the sign of equity's loading on capability is not determined by the evidence (Section 6.2). The high total-consumption R^2 instead reflects the common ε exposure of equity and household consumption. For participants the comparison is additionally partly mechanical because the equity return already enters their consumption bundle.

A projection R^2 also says nothing about the sign of the required position. For total consumption the variance-minimizing equity position is a short position, so a long-only worker could not implement the statistical hedge. That fact weakens the simple "own the firms that benefit" argument, but it is not the paper's main financial limitation. When the target is changed to the welfare-relevant displacement object $u'(c)d$, the desired equity hedge is long and lies inside the model's no-leverage constraint. The constraint is therefore slack for the object that matters. Yet equity explains only 1.3% of $u'(c)d$ for participants and 0.6% for nonparticipants. The problem is covariance, not the short-sale constraint.

This comparison changes the reader's interpretation of the apparent natural hedge. Equity is very good at spanning aggregate consumption risk that happens to share its productivity exposure. It is a poor hedge for the retained-versus-displaced loss. The remaining question is whether this failure is specific to the particular equity payoff or whether displacement risk lies mostly outside the space of *all* aggregate financial payoffs.

5.2 Even the best aggregate payoff reaches little displacement risk

The full-state correlation ratio defined in Section 3.1 answers that question without committing to a particular security design. At the benchmark,

$$\eta_s^2\big(u'(c)d\big) = 4.10\% \quad \text{for participants,} \qquad \eta_s^2\big(u'(c)d\big) = 4.04\% \quad \text{for nonparticipants.}$$

Thus even an arbitrary payoff measurable with respect to the complete modeled state $s = (\varepsilon, \theta)$ can span only about four percent of the variance of the welfare-relevant displacement object.

It is useful to interpret these numbers as two distinct gaps. For participants, about 95.9% of the displacement object is outside the aggregate-state payoff space altogether; for nonparticipants the corresponding share is about 96.0%. This is the *labor gap*: it reflects the worker-specific outcome that remains after the aggregate state has been realized. Conditional on getting into the aggregate-state payoff space, there is a smaller *menu gap*: the difference between the roughly four-percent full-state ceiling and the part already captured by equity. At the benchmark that gap is about 2.8 percentage points for participants and 3.5 points for nonparticipants.

This accounting establishes which friction carries the quantitative burden. A new aggregate security can work only on the menu gap. It cannot reach the roughly ninety-six percent labor gap because, by Lemma 1, every aggregate-state payoff is common to retained and displaced workers

Table 4: The financial reach of displacement risk

Panel A. Benchmark statistical span

Object and projection	Participants	Nonparticipants
Equity R^2: lottery-averaged total consumption	99.62%	94.82%
Equity R^2: $u'(c)d$	1.304%	0.572%
Full-state ceiling $\eta_s^2(u'(c)d)$	4.105%	4.042%

Panel B. Statistical reach across the wage-setting experiment

ν_θ	η_s^2, P	η_s^2, NP	Equity R^2, P	Equity R^2, NP
			target is $u'(c)d$; entries are per cent	
0.00	4.105	4.042	1.304	0.572
0.05	3.693	3.966	0.970	0.443
0.10	3.074	3.360	0.648	0.291
0.13	2.785	2.986	0.479	0.219

Panel C. Equilibrium aggregate-market completion

ν_θ	$\mathcal{D}(A_0)$ (annual bp)	$\mathcal{D}(A_0) - \mathcal{D}(A_s)$ (annual bp)	raw Ω_s (per cent)	Ω_s^{self} (per cent)
0.00	84.326	2.161	2.562	2.205
0.05	60.007	1.909	3.181	2.203
0.10	39.847	1.408	3.533	1.986
0.13	29.764	1.146	3.850	1.980

Notes. Panel A is evaluated at the benchmark allocation. The first row projects each group's lottery-averaged total consumption on the excess return to AI-sector equity. The second row changes the target to the marginal-utility-weighted balanced displacement transfer $u'(c)d$. The third row is the correlation-ratio ceiling over all payoffs measurable with respect to the full aggregate state $s = (\varepsilon, \theta)$. Panel B repeats the displacement span calculations along the fixed-calibration wage-setting experiment of Section 4; technology and preference primitives are held fixed and K^* is re-solved at every node. Panel C solves the financial counterfactual at fixed K within each node. A_0 is the benchmark financial menu and A_s is a complete set of zero-net-supply Arrow claims on all eighteen aggregate states. The absolute basis-point change is the primary level companion. The raw $\Omega_s = [\mathcal{D}(A_0) - \mathcal{D}(A_s)]/\mathcal{D}(A_0)$ uses the original fixed-A_0 normalization. Ω_s^{self} is the companion measure in which each menu is normalized by its own certainty equivalent; it is the preferred share for comparisons across ν_θ. Appendix B gives the normalization construction and the full asset-menu ladder.

in the same state. Security design can therefore improve the mapping from aggregate states to payoffs, but it cannot make the payoff contingent on worker identity.

5.3 Completing aggregate financial markets

The span calculation is an upper bound at a fixed allocation. Does the same limit appear when households actually trade a richer security menu and equilibrium prices adjust? We answer this with the A_0-to-A_s experiment defined in Section 3.1. The A_s menu is intentionally generous: households may take two-sided positions in zero-net-supply Arrow claims on all eighteen modeled aggregate states. Installed capital is held fixed so that the comparison isolates financial risk sharing from the production and labor-allocation changes studied in Section 4.

At the benchmark, the displacement burden is 84.326 annual basis points under A_0 and 82.165 basis points under A_s. Complete aggregate markets therefore reduce the measured burden by 2.161 basis points. Using the paper's original benchmark normalization,

$$\Omega_s = \frac{\mathcal{D}(A_0) - \mathcal{D}(A_s)}{\mathcal{D}(A_0)} = 2.562\%. \tag{21}$$

Thus the deliberately complete aggregate market leaves about 97.4% of the benchmark displacement burden in place.

The roughly four-percent span ceiling and the 2.6% equilibrium welfare effect are not competing estimates of the same parameter. The span statistic is a baseline variance decomposition: it asks how much of $u'(c)d$ could in principle be represented by an arbitrary function of the aggregate state. The A_0-to-A_s comparison instead re-solves portfolios and security prices, changes consumption across states, and then recomputes a nonlinear certainty-equivalent gap. There is no identity equating the two. Their common implication is what matters: both put the financially reachable part of displacement risk in the low single digits.

This interpretation also clarifies why $\mathcal{D}(A_0) - \mathcal{D}(A_s)$ is not "direct insurance" of worker identity. Section 2.3 showed that the balanced transfer needed to equalize retained and displaced workers has zero conditional mean within every aggregate state. Aggregate claims cannot replicate it. What trading can do is redistribute consumption across aggregate states and thereby change the marginal-utility weights applied to the within-state lottery. The positive 2.161-basis-point effect is that equilibrium reweighting channel.

The intermediate asset menus are useful diagnostics but not structural contributions. In the benchmark ladder, universal equity access raises the burden slightly, a complete capability-state menu lowers it, and adding the remaining aggregate-state claims raises it slightly again. Those increments are order dependent because $\mathcal{D}(A)$ is menu indexed and prices clear at each step. For that reason the main text emphasizes the first-to-last A_0 versus A_s comparison; the complete ladder is reported in Appendix IA.F.

5.4 The financial boundary remains small as the wage regime changes

Section 4 produced a large change in the *amount* of risk. Does that same wage-setting experiment create a regime in which aggregate finance can reach a large fraction of the remaining displacement loss? Panels B and C of Table 4 show that it does not.

The statistical upper bound becomes smaller, not larger. For participants the full-state ceiling falls from 4.1% to 2.8% between $\nu_\theta = 0$ and 0.13; for nonparticipants it falls from 4.0% to 3.0%. Existing equity reaches still less, with the displacement R^2 falling from 1.3% to 0.5% for participants and from 0.6% to 0.2% for nonparticipants. Greater wage responsiveness therefore does not turn the surviving displacement lottery into an aggregate financial risk. It mainly makes the lottery smaller, as Section 4 documented.

The equilibrium experiment reaches the same conclusion, but it also reveals an important reporting issue. The absolute reduction in the social burden falls from 2.2 to 1.1 annual basis points across the U.S.-mapped range. The comparable self-normalized completion share is essentially flat and slightly declining, 2.2% to 2.0%. By contrast, the raw Ω_s rises from 2.6% to 3.9%. Reading that last column alone would suggest that wage flexibility raises financial reach by roughly one half. That interpretation is false: the decomposition in Appendix B shows that the rise is generated by the level-normalization term as the baseline burden collapses. For cross-regime comparisons we therefore report the raw measure only beside the absolute basis-point effect and Ω_s^{self}.

This is the paper's main interaction between labor institutions and finance. The wage-setting regime has first-order consequences for how much worker-specific employment risk is produced. Conditional on that risk existing, however, the aggregate financial boundary remains in the low single digits. Over the empirically relevant range there is no economy in which the model simultaneously generates a large displacement lottery and an aggregate financial market capable of insuring a large share of it.

The re-estimated wage-regime family reported in Appendix IA.E provides a different check on this conclusion. There the remaining parameters are allowed to move so that each economy continues to match the U.S. moments. That exercise answers a different question from the fixed-calibration comparison above, but it likewise does not generate a large aggregate-state displacement span. We therefore do not interpret the small boundary as a peculiarity of the single rigid-wage benchmark.

5.5 Participation changes incidence, not the boundary

The limited reach of aggregate securities does not imply that financial participation is irrelevant for household welfare. It implies the narrower result that access to aggregate assets is not a substitute for employment-contingent insurance.

At the benchmark allocation the displacement burden is 77.82 annual basis points for partici-

pants and 88.79 basis points for nonparticipants. That cross-sectional difference should not be read as the causal insurance value of stock ownership. Participants are selected on wealth and participation costs, and greater wealth buffers a displacement loss even if it is held outside equity. The access counterfactual separates this selection from the act of opening the market: giving previously excluded households equity access can raise ordinary household welfare through access to the equity premium, but the social displacement burden moves from 84.33 to 85.19 basis points rather than falling. The gain from access is therefore not displacement insurance.

Equilibrium prices also make incidence distinct from average welfare. New equity demand compresses the premium and can make incumbent participants worse off even when newly admitted households gain. Detailed group-level certainty equivalents and price-clearing diagnostics belong in Appendix IA.F because they answer who gains from access, not how much displacement risk aggregate finance can span.

Finally, price, position, and welfare should not be conflated. A reservation price for a capability-linked claim is a marginal valuation; the optimal quantity depends on the market price relative to that valuation; and the welfare gain depends on the resulting equilibrium allocation. A large reservation price per unit can therefore coexist with a small total welfare effect. The detailed reservation-price frontier and portfolio-position calculations are useful for security design, but they do not alter the conclusion from Table 4: the dominant limitation arises before the choice of aggregate security, because most displacement risk is not measurable with respect to the aggregate state.

The paper has therefore answered the two questions posed at the end of Section 4. Labor adjustment strongly changes the *size* of the worker-specific displacement problem. Expanding the aggregate financial payoff space changes only a small share of the welfare burden that remains. The next section asks what maintained assumptions or calibrations could overturn either side of that conclusion, with particular attention to the aggregate state space, the outside option, preferences, and the horizon.

6 Sensitivity, Scope, and What Could Change the Conclusion

Sections 4 and 5 establish two separate quantitative conclusions. The wage-setting regime has a large effect on how much worker-specific displacement risk the economy produces. Conditional on that risk remaining, aggregate financial markets reach only a small fraction of its welfare-relevant

component. A useful sensitivity exercise must therefore say which of these conclusions it threatens. Some assumptions mainly rescale the welfare cost of displacement; others can change the amount of displacement; and a smaller set can change the boundary between worker-specific and aggregate risk. Treating all of them as generic "robustness checks" would obscure those distinctions.

This section asks what would have to change for the paper's conclusions to change. We begin with near-minimum economies and external calibrations, where the existing computations are informative about quantitative stability. We then turn to the more consequential scope questions: whether the aggregate state omits information that could make displacement substantially more contractible, and whether richer labor contracts could condition directly on employment outcomes. The section ends by separating these model and calibration exercises from sampling uncertainty, which remains a distinct statistical question.

6.1 Near-minimum economies: the financial conclusion is not tied to one calibration

Could the small financial boundary be an artifact of choosing one point from a weakly identified region? Section 3.3 showed that the answer cannot be read from the criterion alone: economically important levels must be recomputed along directions that the moments have difficulty distinguishing. Figure 1 already gives one version of that exercise. Within the event-study magnitude restriction on β_R, the baseline displacement burden moves from roughly 75.45 to 86.96 annual basis points, while the raw A_0-to-A_s completion share remains between 2.50% and 2.86%. The weak return-loading direction therefore matters more for the level of the burden than for the conclusion that aggregate completion has limited reach.

A broader near-minimum exercise gives the same ordering. Define the calibrated criterion set

$$\mathcal{C}(\Delta) = \{\vartheta : L_8(\vartheta) \leq L_{8,\min} + \Delta,\ |\beta_R(\vartheta)| \leq 0.056\},$$

conditional on the benchmark external calibrations. For the tight $\Delta = 0.5$ set, the exhibited feasible endpoints place $\mathcal{D}(A_0)$ between 50.87 and 100.73 annual basis points. Yet the nonparticipant aggregate-state ceiling for $u'(c)d$ remains between 3.46% and 4.33%, and the raw completion share remains between 1.94% and 3.24%. Thus the criterion set allows almost a twofold range in the welfare level without producing an economy in which aggregate securities reach a large share of the displacement object.

Table 5: Sensitivity of the welfare level and aggregate-market completion

Exercise	$\mathcal{D}(A_0)$ (annual bp)	Raw Ω_s (%)	What the exercise changes
Benchmark	84.33	2.56	Benchmark calibration
Panel A. Near-minimum fitted economies			
$\Delta L_8 \leq 0.5$ criterion set	50.87–100.73	1.94–3.24	Six internally fitted parameters, conditional on benchmark external calibrations
β_R magnitude restriction	75.45–86.96	2.50–2.86	Weakly disciplined return-loading direction
Panel B. External calibrations, with the six fitted parameters re-estimated			
$\mathrm{Var}(\theta) \in [0.01, 0.12]$	80.62–95.72	−2.10–9.40	Scale of aggregate capability risk
$\gamma \in [2, 10]$	18.75–269.05	−1.09–10.54	Curvature of household marginal utility

Notes. $\mathcal{D}(A_0)$ is the menu-indexed displacement burden under the benchmark asset menu. $\Omega_s = [\mathcal{D}(A_0) - \mathcal{D}(A_s)]/\mathcal{D}(A_0)$ uses the original node-specific normalization and is reported here because the older external-calibration grids were solved under that convention. It should not be compared across wage-setting regimes without the level and self-normalized companions reported in Table 4. A negative value means that complete aggregate-state trading raises the measured burden slightly through equilibrium reweighting; it does not imply negative direct insurance of worker identity. The $\Delta L_8 \leq 0.5$ row is an inner envelope over exhibited feasible near-minimum economies and is not a confidence interval. The outside-option and persistence exercises are omitted because the complete A_s menu was not recomputed at every node of those older grids.

The qualification is statistical as well as economic. $\mathcal{C}(0.5)$ is a *calibrated criterion set*, not a confidence set. The eight target moments come from different sources and the diagonal weighting matrix is not accompanied by a joint sampling covariance matrix. The exercise therefore establishes stability across near-minimum fitted economies; it does not attach repeated-sample coverage to the reported ranges. Wider Δ sets are reported in Appendix IA.C because they answer how far one can move away from the minimum, not how uncertain an estimator is in repeated samples.

Table 5 makes the main comparison visible. The welfare level is much more sensitive than the financial-completion conclusion. That distinction is not evidence that the benchmark 2.6% is a precisely identified constant. It is not. Rather, the exercises for which the complete A_s menu has actually been re-solved do not produce an environment in which aggregate financial completion becomes first-order displacement insurance.

6.2 External calibrations mainly change the size of the welfare loss

The strongest level sensitivity comes from objects that the eight target moments do not estimate: the scale of the capability shock, preference curvature, the displaced worker's outside option, and the horizon. These exercises are useful precisely because they show where the paper should and

should not claim numerical precision.

Capability risk. The benchmark sets $\mathrm{Var}(\theta) = 0.04$. Re-estimating the six free parameters over $\mathrm{Var}(\theta) \in [0.01, 0.12]$ moves the baseline burden only moderately, from 80.62 to 95.72 basis points, but the equilibrium completion effect changes sign: $\mathcal{D}(A_0) - \mathcal{D}(A_s)$ ranges from about -1.70 to $+9.00$ basis points. The corresponding raw share ranges from -2.10% to 9.40%. The sign of a small reweighting effect is therefore not a robust object. More importantly, even a twelvefold change in the capability variance does not make aggregate completion comparable in size to the displacement burden itself.

The fit gives little reason to privilege the benchmark variance. The criterion falls from about 2.91 at $\mathrm{Var}(\theta) = 0.01$ to 2.29 at the benchmark and 1.98 at 0.12, with no targeted moment missing by more than two standard errors over the grid. The capability variance should therefore remain an external calibration, not be described as identified by the eight moments.

Preference curvature. Risk aversion has a much larger effect on the level. Re-estimating the six free parameters over $\gamma \in [2, 10]$ moves $\mathcal{D}(A_0)$ from 18.75 to 269.05 annual basis points. The complete-market difference ranges from roughly -0.20 to $+28.35$ basis points and the raw completion share from -1.09% to 10.54%. This exercise rules out presenting 84.33 basis points as a model-free estimate of the social cost of displacement. It also shows why the financial result should be stated as a boundary rather than as an exact percentage: the benchmark share moves materially, but even the wide curvature grid does not turn aggregate-state trading into insurance for most of the displacement burden.

The targeted moments do not select $\gamma = 5$. In fact, after re-estimation the criterion is $L_8 = 1.57$ at $\gamma = 2$, compared with 2.29 at the benchmark. The paper therefore treats preference curvature as an external welfare input and lets the sensitivity speak for itself.

Outside option and horizon. The displaced worker's outside option is even more important for the welfare level. Across the existing safety-net grid, $\mathcal{D}(A_0)$ ranges from about 16.55 to 279.23 basis points. Section 4.5 showed, however, that these outside-option changes leave expected displacement unchanged in the current model: the outside option enters the welfare calculation but not the firm's retention first-order condition. The exercise therefore identifies severity conditional on displacement, not the incidence of displacement.

The complete A_s menu was not re-solved over that older outside-option grid, so we do not

convert those level results into claims about financial completion. The same discipline applies to persistence. The benchmark welfare object is one-period and then annualized. Existing persistence exercises show that welfare costs can grow substantially with the duration of displacement, but they do not supply a persistence-adjusted A_0-to-A_s comparison. The horizon is therefore a scope parameter for the level of the loss until the complete financial counterfactual is recomputed over a disciplined dynamic specification.

These results sharpen rather than weaken the paper's main division of labor. The wage-setting regime determines how much adjustment occurs through separations; the outside option and horizon determine how costly those separations are; preference curvature determines how strongly low-consumption outcomes enter welfare. None of those statements is equivalent to saying that an aggregate security can condition on worker identity.

Does the result depend on AI creating no value for the firm? The benchmark gives AI capability no direct productivity benefit: capability raises the automated share of tasks and nothing else. That is a deliberate modelling choice, and it is the choice that makes the benchmark equity loading on capability negative. It also sits uneasily with the evidence that generative-AI exposure has been positively priced in firm values since late 2022 (Eisfeldt et al., 2026). Table 6 therefore lets capability raise the productivity of installed capital directly, through a multiplicative factor with elasticity φ_θ, holding the six fitted parameters at their benchmark values and re-solving equilibrium capital at each point. A small benefit is enough to reverse the loading: it crosses zero at $\varphi_\theta \approx 0.03$ and is $+0.11$ at 0.10. The eight targeted moments barely register the change, so the sign of the loading is an assumption the data do not test rather than an estimate.

What matters for the paper is the second half of the table. As the loading turns positive, equity's projection on the welfare-relevant displacement loss falls rather than rises, from 1.3 to 0.9 percent for participants and from 0.6 to 0.3 percent for nonparticipants, and the aggregate-state ceiling remains between three and four percent. Giving the natural hedge the sign the hedging argument presumes makes it reach less of displacement, not more, because the additional return variation that capability now supplies is common to retained and displaced workers alike. The entries are fixed-parameter sensitivities, not re-estimated economies; a re-estimated cell in this region is the natural next addition to the criterion-set analysis of Section 6.1.

Table 6: The capability productivity channel and the reach of the hedge

φ_θ	β_R	L_8	η_s^2, P	η_s^2, NP	Equity R^2, P	Equity R^2, NP
			target is $u'(c)d$; entries are per cent			
0 (benchmark)	-0.056	2.29	4.10	4.04	1.30	0.57
0.05	$+0.032$	2.54	3.84	3.80	1.13	0.45
0.10	$+0.113$	2.71	3.59	3.58	1.02	0.37
0.20	$+0.262$	2.90	3.18	3.19	0.89	0.26

Notes. Capability enters production through the factor $\exp\{\varphi_\theta(\theta - 1)\}$ multiplying installed capital; the benchmark sets $\varphi_\theta = 0$. The six fitted parameters are held at their benchmark values and equilibrium capital is re-solved at each row, so the entries are fixed-parameter sensitivities and L_8 is reported, not minimized. β_R crosses zero at $\varphi_\theta \approx 0.031$. η_s^2 is the ceiling on any payoff measurable with respect to the full aggregate state; the equity R^2 is the projection of the marginal-utility-weighted displacement loss on the AI-sector equity return.

6.3 The aggregate state is the most important threat to the financial boundary

What change could genuinely make the conclusion in Section 5.2 fail? The most direct route is a richer information set. The benchmark ceiling of roughly 4% is a ceiling relative to the modeled aggregate state $s = (\varepsilon, \theta)$. It is not a theorem that every observable common shock in the economy is uninformative about displacement.

Suppose the relevant contractible state were instead

$$S^* = (\varepsilon, \theta, q),$$

where q collects additional common information such as sectoral demand, occupation-level technology exposure, financing conditions, or other systematic forces omitted from the benchmark. The corresponding ceiling would be

$$\eta^2_{S^*} = \frac{\mathrm{Var}(E[u'(c)d \mid S^*])}{\mathrm{Var}(u'(c)d)}.$$

Additional state variables overturn the financial-boundary result only to the extent that they raise this numerator materially. It is not enough that a new variable forecast aggregate consumption, equity returns, or automation in general; it must predict the welfare-relevant retained-versus-displaced loss that remains unexplained by the current aggregate state.

This is the paper's most important specification caveat because it attacks the measurability result itself rather than merely rescaling welfare. Earlier drafts used an auxiliary experiment in which extra orthogonal aggregate income risk equal to roughly 0.43 times the existing aggregate variance made a welfare proxy assign half of the measured burden to an aggregate component. That calculation remains a useful scale diagnostic, but it is not a market-completion threshold: it changes

the underlying risk environment rather than solving the zero-net-supply A_s trading problem. The detailed exercise therefore belongs in Appendix IA.F.

The main-text claim is consequently conditional and falsifiable. Given the aggregate state represented in the quantitative model, completing trade in that state has limited value for displacement insurance. Evidence that additional contractible common variables explain a large share of $u'(c)d$ would directly weaken that conclusion and would motivate rebuilding the financial menu around the richer state.

6.4 Labor contracts can bypass the financial boundary

A second way to change the conclusion is not to enlarge the aggregate state but to change the information on which contracts can condition. The limitation in Section 5 applies to aggregate securities. It does not say that worker-specific displacement is uninsurable in a broader contractual sense. A payment contingent on individual employment status can transfer resources between retained and displaced workers inside the same aggregate state; an s-measurable security cannot.

This distinction locates severance, unemployment insurance, wage insurance, and employment-contingent contracts outside the financial payoff space studied here. Such arrangements may be able to reach the worker-identity margin directly, but they raise moral-hazard, adverse-selection, verification, bargaining, and fiscal incidence questions that the present model does not solve. The paper therefore does not rank those instruments or claim that they implement full insurance. It locates the point at which another aggregate financial security stops addressing the relevant incompleteness.

The same scope discipline applies to the wage institution itself. The model takes the retained-wage schedule as an institutional object and lets the firm choose retention conditional on it. It does not endogenize bargaining between workers and firms, and it contains no voluntary quit margin. Section 4.4 showed why that omission matters: around $\nu_\theta \simeq 0.34$, the retained wage first falls below the displaced worker's outside option in a live state. Beyond that point a retained worker would quit, but the model has no margin through which to represent the response. The model should therefore not be used as a literal calibration of still more flexible labor institutions beyond that boundary.

Adding quits would require more than changing one first-order condition. The model would have to distinguish voluntary from involuntary separation, specify worker participation and selection, and likely tie feasible retained wages to outside options or bargaining. Those extensions could

change the amount of displacement risk generated by the labor market. They would not, by themselves, overturn the measurability lemma: if two otherwise comparable workers still receive different employment outcomes conditional on the full contractible aggregate state, an aggregate payoff remains unable to condition on which worker received which outcome.

6.5 What remains uncertain

The sensitivity analysis does not convert the quantitative model into a fully sample-based inference exercise. Four uncertainties should remain explicit.

First, the ranges reported for $\mathcal{D}(A_0)$, $\eta_s^2(u'(c)d)$, and the complete-market counterfactual are criterion-set and calibration ranges. Six parameters are fitted to eight noisy empirical moments, so sampling variation in those objects exists; the criterion profiles and calibration grids answer a different question and are not relabeled as confidence intervals.

Second, the benchmark range for ν_θ is a historically motivated scenario range, not a confidence interval and not a mapping of named labor-market institutions. AI may induce a wage-versus-employment response that differs from the earlier labor-demand episode used to discipline that range. Such evidence would change the amount-of-risk conclusion in Section 4 before it changed anything about finance.

Third, external welfare inputs remain quantitatively consequential. Preference curvature, the displaced outside option, and the persistence of earnings losses can move the level of the welfare burden substantially. The paper can therefore make a stronger statement about the location of the financial boundary than about a single universal number of annual basis points.

Fourth, the aggregate-state specification is inherently incomplete. The paper's most important empirical extension would be to ask whether richer common information materially predicts the welfare-relevant displacement loss conditional on (ε, θ). That test targets the assumption most capable of raising the financial ceiling rather than adding another security to an unchanged state space.

Taken together, these exercises narrow the claim rather than turn it into an omnibus robustness statement. The absolute welfare cost of displacement is calibration dependent, and the exact benchmark completion percentage is not a universal constant. What survives the exercises actually computed is the economic ordering: labor-market adjustment can change the amount of worker-specific risk by a large amount, while completing trade in the modeled aggregate state changes a much smaller share of the remaining burden. The routes that could overturn that ordering

are economically specific—a substantially different labor-adjustment process, omitted contractible aggregate information that predicts displacement, or contracts that condition directly on individual employment outcomes. Those are the boundaries the conclusion should carry forward.

7 Conclusion

Can workers hedge technological displacement by owning the firms that benefit from the technology? The answer turns on a distinction between the aggregate shock and the worker-level risk it ultimately creates. A common capability shock moves wages, employment, and asset values for everyone. But when firms respond through separations, the same shock also creates a lottery over *which* worker is retained and which is displaced, and a payoff written only on the aggregate state cannot condition on that identity. The size of the lottery is not a primitive. It is set by how firms adjust headcount to an institutionally determined wage.

The first result is therefore about the labor market rather than the security menu. Holding technology and preferences fixed and varying only the wage response to capability across the historically informed range, greater wage responsiveness shifts adjustment away from separations: the conditional identity component of worker risk falls to about one third of its benchmark value and the displacement burden under the baseline asset menu falls by roughly two thirds, from 84.33 to 29.76 annual basis points. Labor-market adjustment determines the amount and form of worker risk that arrives at financial markets. It does not merely alter the payoff on an otherwise fixed insurable claim.

The second result is that the remaining displacement risk is poorly aligned with aggregate financial payoffs, and that this does not change as the labor market changes. Equity appears at first to be an almost complete hedge: at the benchmark it explains 99.6% of the variation in lottery-averaged consumption for participants. But that reflects common exposure to AI-sector productivity, not insurance of the retained-versus-displaced outcome, and it does not depend on the sign of equity's loading on capability, which the evidence does not determine. When the target is the marginal-utility-weighted displacement loss, equity explains only 1.3% of its variance, and even an arbitrary payoff on the full modeled aggregate state reaches only a few percent. The main financial limitation is not a short-sale constraint, limited participation, or the absence of a particular AI-linked security. Most of the welfare-relevant displacement risk lies outside the aggregate-state payoff space.

Completing aggregate financial markets leads to the same conclusion in equilibrium. Giving households a complete set of zero-net-supply Arrow claims on all eighteen aggregate states lowers the benchmark displacement burden by 2.16 annual basis points, a completion share of about 2.2%. More importantly, the completion share stays near 2% as wage responsiveness rises through the historically informed range, even while the level of the burden falls by two thirds. Greater wage responsiveness does not turn a large worker-specific displacement lottery into an aggregate insurable risk; it makes the lottery smaller.

These magnitudes are conditional on the quantitative environment. The absolute welfare cost of displacement is sensitive to preference curvature, the displaced worker's outside option, the scale and persistence of capability risk, and other external inputs, and the ranges we report for it are criterion-set and calibration ranges rather than sampling intervals. The aggregate-state ceiling is a statement about the state variables in the model, not a theorem that no additional common, contractible information could predict displacement. The wage-versus-employment mapping is disciplined by an earlier large labor-demand shock and need not describe the adjustment margin under future AI adoption. These qualifications limit the numerical claims, but they also identify what evidence would matter most for overturning the economic conclusion: a substantially different labor-adjustment process, or additional contractible aggregate information that strongly predicts the welfare-relevant displacement loss.

The broader implication is a boundary between aggregate financial risk sharing and employment-contingent insurance. Broader participation and richer securities can improve household welfare and redistribute aggregate risk. They cannot transfer resources between two workers who face the same aggregate state but realize different employment outcomes. Insurance at that margin must condition on displacement itself, through employment contracts, severance, wage insurance, unemployment insurance, or related arrangements. Owning the firms that benefit from AI can hedge the aggregate exposures AI creates. It is not, by itself, a substitute for insuring the worker-specific risk of losing a job, and the labor market, not the security market, decides how large that risk is.

A Proofs and the Domain of the Wage-Adjustment Exercise

A.1 Aggregate-state measurability

The idea behind Lemma 1 is simply conditioning. Once the aggregate state is fixed, every payoff written only on that state is a constant across workers. Such a payoff can move consumption across aggregate states, but it cannot make one transfer to a retained worker and a different transfer to a displaced worker who faces the same state.

Proof of Lemma 1. Fix a worker type and portfolio holdings. Let $X_1(s), \ldots, X_J(s)$ be aggregate-state payoffs and let

$$P(s) = \sum_{j=1}^{J} h_j X_j(s)$$

be the resulting portfolio payoff. Conditional on s, $P(s)$ is a constant. Therefore

$$\mathrm{Var}\big(Y^L + P(s) \mid s\big) = \mathrm{Var}\big(Y^L \mid s\big),$$

which proves (12).

The same conditioning argument shows why the portfolio cannot reproduce the balanced employment contingent transfer. When $0 < \lambda(s) < 1$ and $\Delta c(s) \neq 0$, (10) requires

$$d_R(s) = -\lambda(s)\Delta c(s) \quad \text{and} \quad d_D(s) = \big(1 - \lambda(s)\big)\Delta c(s).$$

These two transfers are different, whereas $P(s)$ is the same for both employment outcomes. Hence no portfolio measurable only with respect to s can implement the within-state full-insurance transfer. □

The result is about the conditional component only. Aggregate-state securities can alter consumption across states and therefore can change $\mathrm{Var}(\mathbb{E}[Y^L \mid s])$, marginal utility, and equilibrium prices. Appendix B formalizes this distinction for the welfare-relevant object $u'(c)d$.

A.2 Wage adjustment and the displacement response

The proof of Proposition 1 has two steps. First, the capability shock changes the marginal product of labor both through the labor share $1 - z$ and through the Cobb–Douglas exponents. Second, the same shock changes the retained wage by an amount governed by ν_θ. Log differentiating the firm's retention first-order condition puts these two responses on the same scale and isolates the substitution between wage adjustment and headcount adjustment.

For the mechanism in Section 2.1, suppress terms that do not depend on N and write the firm's state-s retention payoff as

$$\Pi(N; K, s) = P_Y Y_A(N, K; s) - \bar{w}(s)N - \frac{\psi \bar{w}(s)}{2}(1 - N)^2, \qquad 0 \leq N \leq 1, \tag{22}$$

where

$$Y_A = A_{AI}\varepsilon \left(\frac{K}{K_0}\right)^{z} \left(\frac{N}{N_0}\right)^{1-z}, \qquad z = z(K, \theta). \tag{23}$$

For a fixed state, $0 < z < 1$. Hence production is strictly concave in N, and the negative quadratic separation-cost term is concave. The retention problem is therefore strictly concave.

Proof of Proposition 1. Differentiating (22) with respect to N gives

$$\Pi_N = \mathrm{MPL}(N, K; s) - \bar{w} + \psi\bar{w}(1 - N).$$

At an interior solution, $\Pi_N = 0$, so with $\lambda = 1 - N$,

$$\mathrm{MPL}(N, K; s) = \bar{w}(1 - \psi\lambda). \tag{24}$$

Because $\mathrm{MPL} > 0$ and $\bar{w} > 0$, an interior solution also satisfies $1 - \psi\lambda > 0$, so taking logs of (24) is valid.

From (23),

$$\mathrm{MPL} = P_Y(1 - z)\frac{Y_A}{N}.$$

Thus

$$\log \mathrm{MPL} = \text{constant} + \log(1 - z) + z \log\left(\frac{K}{K_0}\right) + (1 - z)\log\left(\frac{N}{N_0}\right) - \log N. \tag{25}$$

Holding K and ε fixed and allowing optimal retention to respond to θ, total differentiation gives

$$\frac{d \log \mathrm{MPL}}{d\theta} = z_\theta \left[\ell_N - \frac{1}{1 - z}\right] - \frac{z}{N}\frac{\partial N}{\partial \theta}, \qquad \ell_N \equiv \log\left(\frac{KN_0}{K_0 N}\right). \tag{26}$$

The automation rule $z = \bar{z}(1 - e^{-aK\theta})$ implies

$$z_\theta = aK\bar{z}e^{-aK\theta} = aK(\bar{z} - z) > 0. \tag{27}$$

The wage schedule is $\bar{w} = \bar{w}_0\varepsilon^{\varpi}(1 - z)^{\nu_\theta}$. Therefore

$$\frac{d \log \bar{w}}{d\theta} = -\nu_\theta \frac{z_\theta}{1 - z}. \tag{28}$$

Taking logs of (24), differentiating, and using $\partial N/\partial\theta = -\partial\lambda/\partial\theta$ yields

$$z_\theta \left[\ell_N - \frac{1}{1 - z}\right] + \frac{z}{N}\frac{\partial \lambda}{\partial \theta} = -\nu_\theta \frac{z_\theta}{1 - z} - \frac{\psi}{1 - \psi\lambda}\frac{\partial \lambda}{\partial \theta}.$$

Collecting the terms multiplying $\partial\lambda/\partial\theta$ gives

$$\left[\frac{z}{N} + \frac{\psi}{1 - \psi\lambda}\right]\frac{\partial \lambda}{\partial \theta} = z_\theta \left[\frac{1 - \nu_\theta}{1 - z} - \ell_N\right].$$

Defining

$$D = \frac{z}{N} + \frac{\psi}{1 - \psi\lambda} > 0$$

produces

$$\frac{\partial \lambda}{\partial \theta} = \frac{z_\theta}{D}\left[\frac{1 - \nu_\theta}{1 - z} - \ell_N\right],$$

which is (14).

For a fixed state and local retention point, z, z_θ, D, and ℓ_N are held fixed when comparing the

direct wage-margin effect. The derivative of the bracketed term with respect to ν_θ is

$$-\frac{1}{1-z} < 0.$$

Hence greater wage responsiveness weakens the direct displacement response to a capability shock. Setting the bracket equal to zero gives the state-dependent threshold

$$\nu^*(s) = 1 - \ell_N(1-z),$$

as stated in (16). □

The last step is intentionally local. It does not assert that $\partial^2\lambda/(\partial\theta\,\partial\nu_\theta) < 0$ along a general-equilibrium path on which N, K, and the other equilibrium objects also change. This is why Section 4 solves the fixed-calibration economies rather than extrapolating the derivative.

A.3 The full-retention corner

The interior derivative is not valid when the firm is already retaining its entire incumbent workforce. The corner is easy to characterize because the retention problem is strictly concave: the sign of the marginal payoff at the upper bound tells us whether the unconstrained optimum lies above that bound.

Define

$$F(N; K, s) \equiv \mathrm{MPL}(N, K; s) - \bar{w}\{1 - \psi(1-N)\} = \Pi_N(N; K, s).$$

Strict concavity implies $F_N = \Pi_{NN} < 0$. Therefore, if

$$F(1; K, s) = \mathrm{MPL}(1, K; s) - \bar{w}(s) \geq 0, \tag{29}$$

the constrained optimum is $N^* = 1$, so $\lambda = 0$. If the inequality is strict, continuity implies that the corner remains binding for sufficiently small changes in θ, and consequently

$$\frac{\partial\lambda}{\partial\theta} = 0$$

locally. At a transition into or out of the corner the policy rule can kink, so the interior formula in Proposition 1 should not be continued through the boundary. The numerical solution therefore solves the retention problem state by state.

A.4 The economically interpretable domain

The upper end of the wage experiment is governed by worker participation, not by the algebraic limit $\nu_\theta = 1$. The model lets the firm choose retention but does not give a retained worker a quit option. Economically, the no-quit allocation can therefore be interpreted only while a worker weakly prefers retention to taking the displaced-worker outside option.

Because financial wealth is the same across the two employment realizations within a state, this participation condition is

$$c_R(s) - c_D(s) = \bar{w}(s) - b(s, K; \lambda(s)) = \Delta c(s) \geq 0 \tag{30}$$

for every state that carries positive surviving displacement risk. If $\Delta c(s) < 0$, a worker who could obtain $b(s, K; \lambda)$ by leaving would prefer the outside option; the maintained no-quit allocation

would then violate individual rationality.

This is a domain criterion, not an analytical restriction used to solve the model. Applying it to the fixed-calibration family of Section 4, the fine-grid calculation has $\Delta c(s) > 0$ at $\nu_\theta = 0.340$ for every surviving displacement state and finds the first negative gap immediately above that point, in the interval $(0.340, 0.345]$. We therefore use

$$\nu_\theta \simeq 0.34$$

as the upper boundary of the economically interpretable comparative static. The historically informed range $[0, 0.13]$ lies strictly inside this domain.

A different numerical threshold occurs later. The expected displacement rate falls below 0.1% at approximately

$$\nu_\theta = 0.44881.$$

This is a near-extinction threshold for the displacement lottery; it is not the economic domain boundary. The worker-participation condition fails first. The diagnostic continuation beyond 0.34, including the later near-extinction point, is therefore reported only in Appendix IA.E and is not used for main-text economic inference.

B Welfare Construction and Aggregate Financial Markets

This appendix formalizes the welfare objects used in Sections 3 and 5. The key distinction is between direct insurance of the retained-versus-displaced outcome and aggregate risk sharing. An employment-contingent transfer can move resources between the two worker outcomes inside a given aggregate state. Aggregate securities cannot make that transfer directly, but they can change consumption and marginal utility across aggregate states and thereby change the welfare cost assigned to the within-state lottery.

B.1 The menu-indexed displacement burden

Fix an economy—including its wage-setting regime—and let A denote an asset menu. Solve the actual menu-A allocation first. Let $V(A)$ be expected utility at that allocation. Let $V^F(A)$ evaluate the *same financial allocation* after the resource-balanced transfer defined below removes the retained-versus-displaced lottery state by state. Thus $V^F(A)$ is a measurement counterfactual: portfolios are not re-optimized after the employment lottery is removed.

Let

$$C(A) \equiv e(V(A)), \qquad C^F(A) \equiv e(V^F(A))$$

denote CRRA certainty-equivalent consumption, and let $C_0 \equiv C(A_0)$. The implemented per-period burden fraction is

$$d(A) = \frac{C^F(A) - C(A)}{C_0}. \tag{31}$$

All menus within the same economy therefore use the same A_0 numeraire. With the model period equal to $\tau = 3$ years, define the annualization map

$$\mathcal{A}_\tau(x) = (1 + x)^{1/\tau} - 1. \tag{32}$$

The burden reported in annual basis points is

$$\mathcal{D}(A) = 10^4 \, \mathcal{A}_3(d(A)) \, . \tag{33}$$

The annualization is applied only after the certainty-equivalent difference has been formed.

The social object is computed ex ante, before the household-type draw: expected utilities are aggregated across types first and then converted once to a certainty equivalent. Participant and nonparticipant burdens use the same construction within the corresponding group. The social burden is therefore not a mass-weighted average of the two group certainty equivalents.

Menu indexing is essential. Both $V(A)$ and $V^F(A)$ are evaluated with the same financial opportunities. Hence $\mathcal{D}(A) - \mathcal{D}(A')$ does not mechanically relabel the ordinary utility gain from a richer asset menu as displacement insurance. It asks how the richer menu changes the welfare gap associated with removing the employment lottery.

B.2 The balanced employment-contingent transfer

The transfer is easiest to understand before introducing utility. Within an aggregate state, full insurance must do two things at once: equalize consumption across the retained and displaced outcomes and leave aggregate resources unchanged. Those two requirements uniquely determine the transfer.

Claim B.1 (balanced transfer). Fix a household type and aggregate state s. Suppose a worker is retained with probability $1 - \lambda(s)$ and displaced with probability $\lambda(s)$, with consumption $c_R(s)$ and $c_D(s)$, respectively. For $0 < \lambda(s) < 1$, the unique transfers $d_R(s)$ and $d_D(s)$ that (i) equalize post-transfer consumption and (ii) satisfy conditional resource balance are

$$\begin{aligned} d_R(s) &= -\lambda(s)\Delta c(s), \\ d_D(s) &= \big(1 - \lambda(s)\big)\Delta c(s), \qquad \Delta c(s) \equiv c_R(s) - c_D(s). \end{aligned} \tag{34}$$

They imply

$$\mathbb{E}[d \mid s] = 0. \tag{35}$$

Proof. Let the common post-transfer consumption level be $\widetilde{c}(s)$. Then

$$d_R = \widetilde{c} - c_R, \qquad d_D = \widetilde{c} - c_D.$$

Conditional resource balance requires

$$(1 - \lambda)d_R + \lambda d_D = 0.$$

Substituting the two expressions above gives

$$\widetilde{c} = (1 - \lambda)c_R + \lambda c_D \equiv c^F.$$

Therefore

$$d_R = c^F - c_R = -\lambda(c_R - c_D)$$

and

$$d_D = c^F - c_D = (1 - \lambda)(c_R - c_D),$$

which proves (34). Substitution immediately yields

$$(1 - \lambda)d_R + \lambda d_D = 0,$$

which is (35). □

A useful implication follows without any asset-pricing assumptions. For every square-integrable payoff $X(s)$ measurable with respect to the aggregate state,

$$\mathrm{Cov}(d, X(s)) = 0. \tag{36}$$

Indeed, $\mathbb{E}[d] = 0$ by iterated expectations and

$$\mathbb{E}[dX(s)] = \mathbb{E}[X(s)\mathbb{E}[d \mid s]] = 0.$$

This is the welfare-side counterpart of Lemma 1: the raw resource transfer required to insure worker identity has no aggregate-state component.

B.3 Why marginal-utility weighting can be aggregate

Conditional resource balance does not imply that the transfer has zero welfare weight. A dollar transferred to a displaced worker is valued at a different marginal utility than a dollar taken from a retained worker. The next calculation is the reason the paper projects $u'(c)d$, rather than d, on aggregate payoffs.

Claim B.2 (marginal-utility-weighted transfer). If $0 < \lambda(s) < 1$, $c_D(s) < c_R(s)$, and utility is strictly concave, then

$$\mathbb{E}\big[u'(c)d \mid s\big] = \lambda(s)\big(1 - \lambda(s)\big)\Delta c(s)\,\big[u'(c_D(s)) - u'(c_R(s))\big] > 0. \tag{37}$$

Proof. Using Claim B.1,

$$\begin{aligned}\mathbb{E}[u'(c)d \mid s] &= (1-\lambda)u'(c_R)\big[-\lambda\Delta c\big] + \lambda u'(c_D)\big[(1-\lambda)\Delta c\big] \\ &= \lambda(1-\lambda)\Delta c\big[u'(c_D) - u'(c_R)\big].\end{aligned}$$

Strict concavity implies that u' is strictly decreasing. Since $c_D < c_R$, we have $u'(c_D) > u'(c_R)$. All remaining factors are positive, which proves the strict inequality. □

If $\lambda \in \{0, 1\}$ or $c_R = c_D$, the expression is zero. Thus the aggregate component created by marginal-utility weighting exists only where there is a genuine within-state employment lottery and the two outcomes have different consumption.

B.4 The aggregate-state projection bound

The statistical upper bound used in Section 5.2 does not depend on a particular security. The idea is the standard L^2 projection result: among all functions of the aggregate state, the conditional expectation is the best mean-square predictor. The correlation ratio is therefore the largest variance share that any aggregate-state payoff could possibly explain.

Claim B.3 (aggregate-state projection bound). Let Y be square integrable with $\mathrm{Var}(Y) > 0$, and let $\mathcal{H}_s$ denote the square-integrable random variables measurable with respect to s. Then the orthogonal projection of Y on $\mathcal{H}_s$ is $\mathbb{E}[Y \mid s]$, and

$$\sup_{g \in \mathcal{H}_s} \left\{1 - \frac{\mathbb{E}[(Y-g)^2]}{\mathrm{Var}(Y)}\right\} = \frac{\mathrm{Var}(\mathbb{E}[Y \mid s])}{\mathrm{Var}(Y)} \equiv \eta_s^2(Y), \tag{38}$$

where the predictor includes a constant.

Proof. Write

$$m(s) = \mathbb{E}[Y \mid s].$$

Then

$$Y - \mathbb{E}[Y] = \big(m(s) - \mathbb{E}[Y]\big) + \big(Y - m(s)\big).$$

For every $g(s) \in \mathcal{H}_s$,

$$\mathbb{E}[(Y - m(s))g(s)] = \mathbb{E}[g(s)\mathbb{E}[Y - m(s) \mid s]] = 0.$$

Hence $m(s)$ is the orthogonal projection of Y on $\mathcal{H}_s$. The Pythagorean identity gives

$$\mathrm{Var}(Y) = \mathrm{Var}(m(s)) + \mathbb{E}[\mathrm{Var}(Y \mid s)]\,.$$

No other s-measurable predictor can have a smaller mean-square residual. Dividing the explained variance $\mathrm{Var}(m(s))$ by $\mathrm{Var}(Y)$ gives (38). □

For the paper's target $Y = u'(c)d$, the benchmark values are approximately 4.10% for participants and 4.04% for nonparticipants. This is a deliberately generous statistical ceiling. It does not impose market-clearing, zero-net-supply, or portfolio constraints, and it is not an equilibrium welfare calculation.

B.5 Aggregate financial-market completion

The equilibrium experiment asks a different question from Claim B.3. Instead of projecting a fixed target, it lets households trade new securities and clears their prices. Installed capital is held fixed so that the experiment changes financial risk sharing without changing the production and labor allocation that generated the displacement risk.

The central menus are:

A_0. The benchmark menu, with the observed participation friction and existing AI-sector equity.

A_θ. Universal equity access plus zero-net-supply claims spanning the nine capability states.

A_s. Universal equity access plus two-sided, zero-net-supply Arrow claims spanning all eighteen states $s = (\varepsilon, \theta)$.

Two-sided Arrow positions are necessary for nontrivial trade in zero net supply. Security prices clear under each menu. The capital stock, wage schedule, retention rule, and displacement probabilities are held fixed within the financial counterfactual.

At the benchmark,

$$\mathcal{D}(A_0) = 84.3262 \text{ bp}, \qquad \mathcal{D}(A_s) = 82.1654 \text{ bp},$$

so the absolute first-to-last change is 2.1607 annual basis points. The intermediate access and A_θ menus are order-dependent diagnostics and are reported in Appendix IA.F rather than interpreted as structural components.

B.6 Why the span statistic and the welfare comparison differ

The roughly 4% span statistic and the roughly 2.6% benchmark completion share are not two estimators of the same object. The former holds the benchmark allocation fixed and asks how

much of $u'(c)d$ is measurable with respect to s. The latter changes the financial opportunity set, re-solves portfolios and security prices, and then recomputes a nonlinear certainty-equivalent burden.

Formally, Claim B.3 is a property of the joint distribution of (Y, s) at one allocation. By contrast, moving from A_0 to A_s changes $V(A)$, $V^F(A)$, $C(A)$, and the state-by-state marginal-utility weights. No projection identity equates $\eta_s^2(u'(c)d)$ with

$$\frac{\mathcal{D}(A_0) - \mathcal{D}(A_s)}{\mathcal{D}(A_0)}.$$

Their usefulness comes from approaching the same economic boundary from different directions: one is a statistical measurability bound and the other is an equilibrium welfare experiment.

B.7 Raw and self-normalized completion across wage regimes

Cross-regime comparisons require one additional accounting distinction. The raw burden in (31) uses $C_\nu(A_0)$ as a common numeraire for all menus at wage regime ν. This is appropriate for a within-economy menu ladder, but a financial menu can itself change the level of certainty-equivalent consumption. An absolute certainty-equivalent lottery gap can therefore move partly because the consumption level on which that gap is measured changes.

To make the source of the normalization transparent, write

$$G_\nu(A) \equiv C_\nu^F(A) - C_\nu(A), \qquad C_{\nu 0} \equiv C_\nu(A_0).$$

The raw burden fraction is

$$d_\nu^{\text{raw}}(A) = \frac{G_\nu(A)}{C_{\nu 0}}, \tag{39}$$

whereas the self-normalized burden fraction is

$$d_\nu^{\text{self}}(A) = \frac{G_\nu(A)}{C_\nu(A)}. \tag{40}$$

Both are converted to annual basis points using the same map $\mathcal{A}_3$ in (32). Thus

$$\mathcal{D}_\nu^{\text{raw}}(A) = 10^4 \mathcal{A}_3(d_\nu^{\text{raw}}(A)), \qquad \mathcal{D}_\nu^{\text{self}}(A) = 10^4 \mathcal{A}_3(d_\nu^{\text{self}}(A)).$$

At A_0 the two burdens are identical because $C_{\nu 0} = C_\nu(A_0)$. For A_s, the *only* difference between the two constructions is whether the same absolute gap $G_\nu(A_s)$ is divided by $C_\nu(A_0)$ or by $C_\nu(A_s)$.

The reported shares are therefore

$$\Omega_{\nu,s}^{\text{raw}} = 100\, \frac{\mathcal{D}_\nu^{\text{raw}}(A_0) - \mathcal{D}_\nu^{\text{raw}}(A_s)}{\mathcal{D}_\nu^{\text{raw}}(A_0)}, \tag{41}$$

$$\Omega_{\nu,s}^{\text{self}} = 100\, \frac{\mathcal{D}_\nu^{\text{self}}(A_0) - \mathcal{D}_\nu^{\text{self}}(A_s)}{\mathcal{D}_\nu^{\text{self}}(A_0)}. \tag{42}$$

This makes the interpretation exact: disagreement between the two shares is entirely attributable to the menu-level normalization of the A_s certainty-equivalent gap, because $G_\nu(A_s)$, $G_\nu(A_0)$, and the annualization rule are otherwise the same inputs.

The wage-regime results of Section 4 illustrate why the distinction matters. Over $\nu_\theta \in [0, 0.13]$,

the absolute reduction

$$\mathcal{D}(A_0) - \mathcal{D}(A_s)$$

falls from 2.161 to 1.146 annual basis points. The self-normalized share moves only from 2.2% to 2.0%. The raw share instead rises from 2.6% to 3.9%. Over the same comparison, $C_\nu(A_s)/C_\nu(A_0)$ falls from approximately 0.9963 to 0.9809. The apparent increase in the raw share is therefore a level-normalization effect, not evidence that aggregate securities reach a larger fraction of the remaining worker-specific displacement risk.

For this reason the paper uses three objects together across wage regimes: the absolute basis-point difference, the self-normalized completion share, and the statistical ceiling $\eta_s^2(u'(c)d)$. The raw share remains useful for the original within-economy menu comparison, but it is not interpreted alone as an institutional comparative static.

B.8 Interpretation and scope

The full-insurance transfer is a measurement device, not a proposed contract. It is assumed to be verifiable, resource balanced, and free of incentive or administrative costs. Employment-contingent insurance in practice may be limited by moral hazard, adverse selection, verification, search incentives, or fiscal incidence. Conversely, a policy that changes the wage-setting regime or firms' retention choices changes the size of the lottery itself and is not the financial counterfactual studied here.

Likewise, A_s is complete only with respect to the modeled aggregate state. Additional contractible aggregate variables could raise the projection bound if they predict the welfare-relevant displacement loss conditional on (ε, θ). What A_s cannot do, by construction, is condition on which worker is displaced after that aggregate state is known.

References

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Athanasoulis, S. G. and Shiller, R. J. (2001). World income components: Measuring and exploiting risk-sharing opportunities. *American Economic Review*, 91(4):1031–1054.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–2168.

Benzoni, L., Collin-Dufresne, P., and Goldstein, R. S. (2007). Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *Journal of Finance*, 62(5):2123–2167.

Catherine, S. (2022). Countercyclical labor income risk and portfolio choices over the life cycle. *Review of Financial Studies*, 35(9):4016–4054.

Cocco, J. F., Gomes, F. J., and Maenhout, P. J. (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies*, 18(2):491–533.

Constantinides, G. M. and Duffie, D. (1996). Asset pricing with heterogeneous consumers. *Journal of Political Economy*, 104(2):219–240.

Davis, S. J. and Willen, P. (2013). Occupation-level income shocks and asset returns: Their covariance and implications for portfolio choice. *Quarterly Journal of Finance*, 3(3–4):1350011.

Duffie, D. and Rahi, R. (1995). Financial market innovation and security design: An introduction. *Journal of Economic Theory*, 65(1):1–42.

Eisfeldt, A. L., Schubert, G., Zhang, M. B., and Taska, B. (2026). Generative AI and firm values. *Journal of Finance*. Forthcoming.

Gârleanu, N., Kogan, L., and Panageas, S. (2012). Displacement risk and asset returns. *Journal of Financial Economics*, 105(3):491–510.

Guiso, L., Pistaferri, L., and Schivardi, F. (2005). Insurance within the firm. *Journal of Political Economy*, 113(5):1054–1087.

Hampole, M., Papanikolaou, D., Schmidt, L. D. W., and Seegmiller, B. (2025). Artificial intelligence and the labor market. NBER Working Paper 33509, National Bureau of Economic Research. Revised September 2025.

Huang, X., Martinez, J. V., Pokojovy, M., and Qian, M. (2026). Hedging ai shocks to human capital. Technical report, Working paper. April 2026.

Kogan, L., Papanikolaou, D., Schmidt, L. D. W., and Seegmiller, B. (2023). Technology and labor displacement: Evidence from linking patents with worker-level data. NBER Working Paper 31846, National Bureau of Economic Research. Revised November 2025.

Kogan, L., Papanikolaou, D., and Stoffman, N. (2020). Left behind: Creative destruction, inequality, and the stock market. *Journal of Political Economy*, 128(3):855–906.

Mankiw, N. G. and Zeldes, S. P. (1991). The consumption of stockholders and nonstockholders. *Journal of Financial Economics*, 29(1):97–112.

Meeuwis, M., Papanikolaou, D., and Schmidt, L. D. W. (2026). Risk premia, limited firm insurance, and heterogeneous earnings risk. Technical report, Working paper. Updated July 2026.

Rockall, E. J., Mendes Tavares, M., and Pizzinelli, C. (2025). Ai adoption and inequality. IMF Working Paper 2025/068, International Monetary Fund.

Shiller, R. J. (1993). *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks*. Oxford University Press.

Internet Appendix to: Can Workers Hedge AI by Owning It? Labor Adjustment and the Limits of Financial Insurance*

Murray Z. Frank Vojislav Maksimovic

August 26, 2026

This Internet Appendix is not intended to be read independently of the paper.

This appendix contains the model audit, empirical mappings, identification diagnostics, validation and numerical checks, wage-regime diagnostics, and additional financial experiments referenced in the paper. Section, table, figure, and equation numbers beginning with "IA." refer to this document. References to numbered sections, tables, figures, or appendices of the paper are written explicitly in the text.

Contents

IA.A Full Quantitative Environment and Parameter Audit **3**
IA.A.1 Timing, aggregate uncertainty, and household heterogeneity 3
IA.A.2 Technology, production, and the labor-market allocation 4
IA.A.3 Equity payoffs 6
IA.A.4 Households, participation, and portfolio choice 7
IA.A.5 Capital-market equilibrium 8
IA.A.6 Financial menus used in the counterfactuals 9
IA.A.7 Numerical solution and calibration 9
IA.A.8 Calibration rationale and empirical provenance 10
IA.A.9 Complete run-of-record parameter inventory 16
IA.A.10 Run-of-record and reproduction conventions 18

IA.B Empirical Mapping of the Wage-Adjustment Parameter **20**
IA.B.1 The empirical object 20
IA.B.2 From the observed share to the structural elasticity 20
IA.B.3 Sampling uncertainty and the covariance problem 22
IA.B.4 Why local-labor-market wage responses are a different object 23
IA.B.5 What the historical range does and does not establish 23

*Frank is from University of Minnesota. E-mail: murra280@umn.edu. Maksimovic is from University of Maryland. E-mail: max@rhsmith.umd.edu. Generative AI assistants were used during the preparation of this manuscript to search for errors, to improve the writing, and to help with the code. We reviewed and edited every suggestion, and the substantive scientific decisions are ours alone. We take full responsibility for the final manuscript, including any remaining mistakes or misinterpretations.

IA.C Identification and Criterion-Set Diagnostics **26**
IA.C.1 Criterion and search design . . . 26
IA.C.2 The return-loading ridge and the event-study boundary . . . 27
IA.C.3 Local Jacobian and moment sensitivity . . . 28
IA.C.4 Finite-difference and quadrature stability . . . 29
IA.C.5 Criterion-set envelopes for the economic objects . . . 30
IA.C.6 What the diagnostics establish—and what they do not . . . 31

IA.D Fit, Validation, Persistence, and Numerical Checks **33**
IA.D.1 Construction of the eight target moments . . . 33
IA.D.2 What the benchmark fits—and where it does not . . . 34
IA.D.3 Untargeted validation: the severity of displacement . . . 35
IA.D.4 Horizon and persistence . . . 36
IA.D.5 Exact welfare versus quadratic risk accounting . . . 36
IA.D.6 Numerical accuracy and internal consistency . . . 36
IA.D.7 What these checks establish . . . 37

IA.E Wage-Regime Diagnostics **39**
IA.E.1 Two distinct wage-regime exercises . . . 39
IA.E.2 Asset-menu ladders . . . 39
IA.E.3 Span ceilings across wage regimes . . . 40
IA.E.4 Capital deepening and the outside option . . . 40
IA.E.5 The economically interpretable domain . . . 41
IA.E.6 Why the raw completion share rises . . . 42
IA.E.7 A secondary asset-pricing implication . . . 43
IA.E.8 Summary . . . 44

IA.F Additional Financial Experiments and Specification Sensitivity **45**
IA.F.1 Equity access: selection is not insurance . . . 45
IA.F.2 Portfolio constraints, security design, and reservation prices . . . 46
IA.F.3 External calibration: capability risk and risk aversion . . . 46
IA.F.4 Outside options and the horizon . . . 48
IA.F.5 State-space scope: what an omitted aggregate shock would have to do . . . 48
IA.F.6 Summary of the financial sensitivity exercises . . . 49

IA.A Full Quantitative Environment and Parameter Audit

The paper states only the pieces of the quantitative environment needed to understand the economic argument. This section supplies the audit trail. It states the complete benchmark timing, shock construction, firm and household problems, market-clearing condition, financial-menu conventions, solution method, and run-of-record parameter structure. The goal is not to introduce a second model. It is to make clear exactly what the numerical economy solves and which inputs are estimated, externally fixed, normalized, derived, or purely numerical.

The organization follows the order in which the quantitative economy is solved. We first construct the aggregate and household states. We then map an aggregate technology state into production, wages, retention, and worker income. Next we describe household participation and portfolio choice and the capital-market equilibrium. We finish with the numerical algorithm and a complete 47-field parameter audit. The wage-regime mapping, identification profiles, and convergence exercises are intentionally deferred to IA.B–IA.E, where they can be evaluated without interrupting the description of the benchmark environment.

IA.A.1 Timing, aggregate uncertainty, and household heterogeneity

The first step is to say what is known when each decision is made. This matters because the paper's central risk is created *after* the aggregate state is known: households choose participation and portfolios before the state is realized, while the firm chooses retention after observing the technology state. The model therefore separates uncertainty about the aggregate state from the retained-versus-displaced realization conditional on that state.

There are two dates. At date 0, the AI-adopting sector installs capital K and households choose financial-market participation and portfolios. The wage schedule for retained workers is predetermined by the wage-setting regime. At date 1, the aggregate state $s = (\varepsilon, \theta)$ is realized, the firm chooses retained headcount, production and labor income are realized, financial claims pay off, and households consume. One model period is $\tau = 3$ years. The parameter field `Rf` is the gross *annual* risk-free return, which is $R_f^a = 1.02$ at the benchmark. The gross risk-free return used over one model period is therefore

$$R_f^{(\tau)} = (R_f^a)^\tau = 1.02^3 = 1.061208. \tag{IA.1}$$

The two aggregate shocks represent different economic objects. The productivity shock is a symmetric two-point shock,

$$\varepsilon \in \{1 + \sigma_\varepsilon,\, 1 - \sigma_\varepsilon\}, \qquad \Pr(\varepsilon = 1 + \sigma_\varepsilon) = \frac{1}{2}. \tag{IA.2}$$

AI capability is positive and has unit mean. The benchmark assumes a lognormal capability distribution with

$$\sigma^2_{\ln\theta} = \ln(1 + \mathrm{Var}(\theta)), \qquad \mu_{\ln\theta} = -\frac{1}{2}\sigma^2_{\ln\theta}. \tag{IA.3}$$

Let x_j and ω_j denote Gauss–Hermite nodes and weights. Before the final mean correction, define

$$\widetilde{\theta}_j = \exp\Big(\mu_{\ln\theta} + \sqrt{2}\,\sigma_{\ln\theta} x_j\Big), \qquad \widetilde{\pi}_j = \frac{\omega_j}{\sqrt{\pi}}. \tag{IA.4}$$

The probabilities are normalized to sum to one and the nodes are rescaled by their quadrature mean,

$$\pi_j = \frac{\widetilde{\pi}_j}{\sum_k \widetilde{\pi}_k}, \qquad \theta_j = \frac{\widetilde{\theta}_j}{\sum_k \pi_k \widetilde{\theta}_k}, \qquad \sum_j \pi_j \theta_j = 1. \tag{IA.5}$$

The shocks ε and θ are independent. The run of record uses nine capability nodes and two productivity realizations, for eighteen aggregate states. The disaster branch contained in the parameter structure is switched off at the benchmark, so the active state space is the pure tensor product of these two shocks.

Households differ in initial wealth and in the fixed cost of entering the equity market. These dimensions are needed to reproduce both the level and the wealth gradient of stock-market participation; they are not additional aggregate shocks. Initial wealth satisfies

$$\log w_0 \sim N(\mu_w, \sigma_w^2), \tag{IA.6}$$

and the participation cost is independently lognormal,

$$\log \kappa \sim N(\mu_\kappa, \sigma_\kappa^2), \qquad \mu_\kappa = \log \bar{\kappa} - \frac{1}{2}\sigma_\kappa^2, \qquad \mathbb{E}[\kappa] = \bar{\kappa}. \tag{IA.7}$$

The benchmark uses 30 Gauss–Hermite nodes for the wealth integral. For participation, the preferred equilibrium implementation solves directly for the cost threshold $\kappa^*(w_0; K)$ and integrates the lognormal cost density up to that threshold. This removes artificial jumps that arise when a moving participation cutoff is evaluated only at a fixed set of cost nodes. Some projection diagnostics use a fixed household tensor; whenever that convention is used, IA.D reports it explicitly rather than mixing its participation mass with the mass used in estimation and welfare.

IA.A.2 Technology, production, and the labor-market allocation

The next step maps a realized technology state into worker outcomes. Capability first determines how many tasks can be automated. Given that task allocation, the wage-setting regime determines the retained wage and the firm chooses how many incumbent workers to keep. Displacement is therefore an equilibrium outcome of the production and wage environment rather than an exogenous shock to household income.

Installed AI capital K and capability θ determine the automated task share

$$z(K, \theta) = \bar{z}\left(1 - e^{-aK\theta}\right), \qquad 0 \le z < \bar{z} < 1. \tag{IA.8}$$

For a fixed capability realization, z is increasing and concave in installed capital; for a fixed capital stock it is increasing and concave in capability. The parameter a determines how quickly the frontier approaches its ceiling $\bar{z}$.

The parameter structure also allows two extensions that are inactive at the benchmark. A capability-dependent machine-productivity term is

$$\Gamma_\theta(\theta) = \exp(\phi_\theta(\theta - 1)), \tag{IA.9}$$

and effective human-task capacity per retained worker is

$$h(z) = 1 + \chi z. \tag{IA.10}$$

The benchmark sets $\phi_\theta = 0$ and $\chi = 0$, so $\Gamma_\theta \equiv 1$ and $h \equiv 1$. Keeping these fields in the parameter structure makes alternative production specifications transparent without changing the benchmark definition.

With the output price normalized to one, AI-sector output is

$$Y_A(s, K, N) = A_{AI}\varepsilon \left(\frac{\Gamma_\theta(\theta)K}{K_0}\right)^{z(K,\theta)} \left(\frac{Nh(z(K,\theta))}{N_0}\right)^{1-z(K,\theta)}. \tag{IA.11}$$

The reference scales K_0 and N_0 keep the production function written in dimensionless capital and headcount ratios. A common proportional rescaling of both is a unit change when accompanied by the corresponding productivity rescaling. Their relative scale is economically more consequential because the exponent on each input changes with z. We therefore treat them as frozen run-of-record production anchors rather than as additional estimated productivity parameters or as two innocuous free normalizations.

The retained wage is taken as given by the firm,

$$\bar{w}_A(s, K) = \bar{w}_0 \varepsilon^{\varpi} \left[1 - z(K, \theta)\right]^{\nu_\theta}. \tag{IA.12}$$

The exponent ν_θ is the paper's wage-regime parameter. It is an elasticity with respect to the unautomated task share $1 - z$, not a coefficient on the level of θ. The benchmark has $\nu_\theta = 0$; Section 4 of the paper and IA.E vary it.

Separations cost

$$C_{\text{sep}}(s, K, N) = \frac{\psi}{2} \bar{w}_A(s, K)(1 - N)^2. \tag{IA.13}$$

After observing s, the firm chooses $N \in [0, 1]$ to maximize

$$\Pi_A(s, K, N) = Y_A(s, K, N) - \bar{w}_A(s, K)N - C_{\text{sep}}(s, K, N). \tag{IA.14}$$

For fixed (s, K), the second derivative with respect to retained headcount is

$$\frac{\partial^2 \Pi_A}{\partial N^2} = -z(1 - z)\frac{Y_A}{N^2} - \psi \bar{w}_A < 0, \tag{IA.15}$$

so the retention problem is strictly concave. An interior solution satisfies

$$(1 - z)\frac{Y_A}{N} - \bar{w}_A + \psi \bar{w}_A(1 - N) = 0. \tag{IA.16}$$

The numerical routine evaluates the interior root and the boundaries. In particular, when the marginal profit from retention is nonnegative at $N = 1$, the firm retains the full incumbent workforce. Appendix A of the paper treats this corner formally.

The displacement probability of an incumbent worker is the residual

$$\lambda(s, K) = 1 - N^*(s, K). \tag{IA.17}$$

Workers who are displaced move to a traditional sector over the model horizon. Traditional-sector output and the marginal wage are

$$Y_T(s, K) = A_T \left(\bar{L}_T + \lambda(s, K)\right)^{\beta_T}, \tag{IA.18}$$

$$w_T(s, K) = \beta_T A_T \left(\bar{L}_T + \lambda(s, K)\right)^{\beta_T - 1}. \tag{IA.19}$$

The displaced-worker outside option is

$$b(s, K) = \max\left\{b_0,\ w_T(s, K) - \phi_R\right\}. \tag{IA.20}$$

This formula is worth stating exactly. The reallocation cost ϕ_R is a *level subtraction* from the traditional-sector wage. It is not a proportional haircut and it is not subtracted from the retained AI wage. Because $\beta_T < 1$, additional displaced workers reduce the traditional-sector marginal wage, so the outside option can be lower in states with more separations.

Let $e \in \{0, 1\}$ be the employment realization, with $e = 1$ denoting retention. Then

$$\Pr(e = 0 \mid s, K) = \lambda(s, K), \qquad \Pr(e = 1 \mid s, K) = 1 - \lambda(s, K), \tag{IA.21}$$

and labor income is

$$\ell(s, e; K) = e\, \bar{w}_A(s, K) + (1 - e) b(s, K). \tag{IA.22}$$

Conditional on s, employment outcomes are integrated analytically. For any function f of consumption,

$$\mathbb{E}[f(c) \mid s] = (1 - \lambda(s)) f(c_R(s)) + \lambda(s) f(c_D(s)). \tag{IA.23}$$

This is the integration used in portfolio choice, participation, and welfare.

IA.A.3 Equity payoffs

The financial side needs a payoff to the capital installed at date 0. There are two natural profit-per-capital objects, and separating them avoids a common source of confusion. Average optimized profit per unit of installed capital is useful descriptively, but a household that supplies one additional unit of capital receives the *marginal* value of installed capital. Portfolio choice and the equity-market equilibrium use this marginal payoff.

Define optimized profit by

$$\Pi_A^*(s, K) = \Pi_A(s, K, N^*(s, K)). \tag{IA.24}$$

The descriptive average payoff is

$$\overline{R}_E(s, K) = \frac{\Pi_A^*(s, K)}{K}. \tag{IA.25}$$

The traded payoff used by households is instead

$$R_E(s, K) = \frac{d\Pi_A^*(s, K)}{dK}. \tag{IA.26}$$

Because $N^*(s, K)$ solves the retention problem, the envelope theorem gives

$$R_E(s, K) = \left. \frac{\partial \Pi_A(s, K, N)}{\partial K} \right|_{N = N^*(s,K)}. \tag{IA.27}$$

Equivalently,

$$\begin{aligned} R_E(s, K) = \frac{\partial Y_A(s, K, N^*)}{\partial K} - N^* \frac{\partial \bar{w}_A(s, K)}{\partial K} \\ - \frac{\psi}{2}(1 - N^*)^2 \frac{\partial \bar{w}_A(s, K)}{\partial K}. \end{aligned} \tag{IA.28}$$

The production derivative includes both the direct effect of installed capital and the effect of capital on the automation share, which changes the two production exponents. The implementation evaluates this derivative analytically; IA.D reports the finite-difference check against optimized profits.

A useful reporting object is the equilibrium loading of equity on capability,

$$\beta_R = \frac{\mathrm{Cov}(R_E, \theta)}{\mathrm{Var}(\theta)}. \tag{IA.29}$$

This is a *derived equilibrium object*, not one of the 47 fields in the parameter structure. The eight-moment fit is solved subject to the external magnitude restriction on this derived loading described in Section 3 of the paper. IA.C reports the corresponding profile and explains why the admissibility fence is not a sampling confidence interval.

IA.A.4 Households, participation, and portfolio choice

The household block turns the firm-side payoffs into consumption and asset demand. The economic distinction is straightforward. A nonparticipant saves only in the risk-free asset. A participant pays a fixed cost and then chooses how much of the remaining wealth to invest in AI-sector equity. Wealth heterogeneity and participation costs therefore affect who owns the hedge, but they do not change the worker's retained-versus-displaced identity within a state.

All households have CRRA utility,

$$u(c) = \begin{cases} \dfrac{c^{1-\gamma}}{1-\gamma}, & \gamma \neq 1, \\ \log c, & \gamma = 1. \end{cases} \tag{IA.30}$$

The exact certainty-equivalent operator is

$$e(V) = \big[(1-\gamma)V\big]^{1/(1-\gamma)} \qquad (\gamma \neq 1). \tag{IA.31}$$

No quadratic approximation is used for the headline welfare objects. The welfare construction and annualization convention are given in Appendix B of the paper rather than repeated here.

A worker with initial wealth w_0 and participation cost κ has investable wealth

$$\widetilde{w} = \begin{cases} w_0 - \kappa, & \text{if the worker participates,} \\ w_0, & \text{if the worker does not participate.} \end{cases} \tag{IA.32}$$

A participant chooses an equity share $\alpha_w \in [0,1]$. The interval imposes the benchmark long-only, unlevered restriction. Conditional on state s and employment outcome e, participant consumption is

$$c^P(s,e) = \widetilde{w}\left[\alpha_w R_E(s,K) + (1-\alpha_w)R_f^{(\tau)}\right] + \ell(s,e;K). \tag{IA.33}$$

A nonparticipant has

$$c^{NP}(s,e) = w_0 R_f^{(\tau)} + \ell(s,e;K). \tag{IA.34}$$

Expected utility integrates over both aggregate states and the employment lottery in (IA.23). A participating worker's optimal share is

$$\alpha_w^*(K,w_0,\kappa) \in \arg\max_{\alpha\in[0,1]} \mathbb{E}\big[u(c^P(s,e;\alpha))\big]. \tag{IA.35}$$

At an interior solution its first-order condition is

$$\mathbb{E}\Big[u'(c^P)\,\widetilde{w}\,\big(R_E - R_f^{(\tau)}\big)\Big] = 0. \tag{IA.36}$$

The numerical routine also evaluates both portfolio boundaries, so an interior first-order condition is not imposed when a corner is optimal.

Let $V^P(K,w_0,\kappa)$ denote maximized participant utility and $V^{NP}(K,w_0)$ nonparticipant utility. Participation requires both

$$w_0 > \kappa \qquad \text{and} \qquad V^P(K,w_0,\kappa) \geq V^{NP}(K,w_0). \tag{IA.37}$$

Over the calibrated region, the participation gain is monotone in κ at a fixed wealth level, so the participation set can be represented by a scalar threshold $\kappa^*(w_0;K)$. The smooth implementation integrates participant demand over $[0,\kappa^*(w_0;K)]$ against the lognormal cost density.

A mass μ_C of capitalists has wealth W and no labor income. A capitalist chooses $\alpha_C \in [0,1]$ to maximize expected CRRA utility of

$$c_C(s) = W\left[\alpha_C R_E(s,K) + (1-\alpha_C)R_f^{(\tau)}\right]. \tag{IA.38}$$

The capitalist first-order condition has the same form as (IA.36), without labor income or a participation cost.

Consumption is required to be positive at every priced state. The code contains a very small numerical floor to prevent floating-point overflow while rejecting infeasible parameter vectors; accepted benchmark and sensitivity allocations are checked not to rely on that floor.

IA.A.5 Capital-market equilibrium

The benchmark equilibrium closes the model by requiring households collectively to finance the installed AI capital. This step is important because capital is not fixed when the benchmark or the labor-institution counterfactual is solved: changes in household demand can move K, and that in turn changes automation and retention.

Capitalist demand for AI-sector equity is

$$D_C(K) = \mu_C W \alpha_C^*(K). \tag{IA.39}$$

Worker demand is

$$\begin{aligned} D_W(K) = \int\int_{(w_0,\kappa)\in\mathcal{P}(K)} & \alpha_w^*(K,w_0,\kappa)(w_0-\kappa) \\ & \times dF_\kappa(\kappa)\, dG_w(w_0), \end{aligned} \tag{IA.40}$$

where $\mathcal{P}(K)$ is the participation set defined by (IA.37). The date-0 price of one unit of installed AI capital is normalized to one, so market clearing requires

$$K^* = D_C(K^*) + D_W(K^*). \tag{IA.41}$$

Equivalently, define excess demand

$$\mathcal{E}(K) = D_C(K) + D_W(K) - K. \tag{IA.42}$$

An equilibrium is a root of $\mathcal{E}(K) = 0$. The numerical implementation uses the stable root satisfying

$$\mathcal{E}'(K) < 0 \tag{IA.43}$$

and connected continuously to the calibrated benchmark. At the run of record, $K^* \simeq 0.63545$.

Thus a benchmark equilibrium consists of the capital stock, the state-by-state retention rule, worker participation and portfolio rules, capitalist portfolio choice, labor income, and equity payoffs satisfying household optimality, firm optimality, and (IA.41). Section 4 of the paper changes the wage-setting parameter and re-solves this equilibrium. That is why the institutional experiment can change both labor adjustment and installed capital.

IA.A.6 Financial menus used in the counterfactuals

The paper also asks a different question: holding the real allocation that generates displacement fixed, how much can a richer financial market change risk sharing? To isolate that question, the financial counterfactuals hold installed capital fixed at the equilibrium value for the economy being studied. The wage schedule, production allocation, retention rule, and displacement probabilities therefore remain fixed while household portfolios and security prices adjust.

The opportunity sets are:

A_0. The benchmark allocation with the observed participation friction, the risk-free asset, and existing AI-sector equity.

A_{access}. All households receive access to existing AI-sector equity. This is an access experiment, not aggregate-market completion.

A_θ. Universal equity access plus a complete set of zero-net-supply claims spanning the nine capability states. This menu is complete with respect to θ but not with respect to the productivity dimension ε.

A_s. Universal equity access plus zero-net-supply Arrow claims spanning all eighteen aggregate states $s = (\varepsilon, \theta)$. Positions in the new Arrow claims are two-sided; imposing nonnegative positions on every zero-net-supply claim would eliminate trade by construction. Claim prices clear in equilibrium.

The A_s menu is deliberately generous. It is complete with respect to the *modeled aggregate state*, but not with respect to the worker's employment identity conditional on that state. This is why the A_0–A_s experiment can change state-by-state consumption and marginal utility without directly replicating the retained-versus-displaced transfer. Appendix B of the paper states the exact welfare comparison; IA.F reports the full intermediate menu ladder and the additional security-design diagnostics.

For some diagnostics, the model also uses the standardized linear capability claim

$$X_\theta = \frac{\theta - \mathbb{E}[\theta]}{\sqrt{\text{Var}(\theta)}}. \tag{IA.44}$$

This claim is useful for interpreting exposure to the capability dimension, but it is not the financial upper bound used in the paper. The complete A_θ and A_s menus provide those upper benchmarks.

IA.A.7 Numerical solution and calibration

The numerical problem is simple in dimension but contains economically real kinks. Retention and portfolio choices are one-dimensional, yet participation is discrete and portfolios can optimally sit at their bounds. The solution method therefore combines exact one-dimensional optimization with a global-to- local search over the six fitted parameters rather than relying on a single smooth first-order system.

Aggregate productivity is integrated exactly over its two realizations. Capability and initial wealth are integrated using Gaussian quadrature after the lognormal transformations above. Employment outcomes are integrated analytically conditional on the aggregate state using (IA.23). The firm solves the strictly concave retention problem at each state. Worker and capitalist portfolio problems compare any interior optimum with the two portfolio boundaries.

For participation, a fixed-node indicator can make aggregate equity demand jump when a threshold moves across a cost node. The preferred implementation therefore solves for $\kappa^*(w_0; K)$ and integrates the participant portion of the lognormal cost distribution directly. This smooths the numerical integration without smoothing the underlying participation decision.

For a given parameter vector, the equilibrium solver evaluates $\mathcal{E}(K)$ on a log-spaced capital grid, uses sign changes to bracket roots, and refines candidate roots with a safeguarded scalar method. A root must clear the capital market and satisfy the stability condition (IA.43). Sensitivity exercises use continuation from the preceding parameter value so that a switch between equilibrium branches is not mistaken for an economic comparative static.

The six internally fitted parameters minimize the eight-moment criterion stated in Section 3 of the paper. The calibration search begins globally over the admissible parameter box, improves promising candidates with a derivative- free local routine, and uses a bounded local polish once a sufficiently smooth neighborhood is reached. Every objective evaluation solves the full equilibrium and returns the eight model moments. Parameter vectors are rejected when markets fail to clear, consumption or portfolio restrictions are violated, or the solution lies on an inadmissible branch. IA.C reports the search and profile diagnostics; IA.D reports quadrature, derivative, and numerical-accuracy checks. Keeping those diagnostics separate from the present section lets IA.A describe the economy without turning the model audit into an optimizer log.

IA.A.8 Calibration rationale and empirical provenance

The parameter inventory is useful only if the reader can also see why each economic input takes the value it does. We therefore distinguish four cases. First, six parameters are jointly fitted to the eight moments in Section 3 of the paper; their justification is the moment fit rather than an external point estimate. Second, several inputs are tied to outside data or standard quantitative benchmarks. Third, some economically consequential primitives are maintained structural anchors: the available data do not separately identify them, so the paper fixes transparent benchmark values and interprets the quantitative results conditional on those choices. Fourth, a small number of objects are genuine numerical or unit conventions. We do not call a fixed parameter a "normalization" merely because it is set rather than estimated.

This distinction is especially important in the labor block. The output price is a genuine numeraire normalization. By contrast, the separation-cost curvature, the size of the traditional sector, the automation ceiling, and the relative production reference scales can affect real responses. Those objects are therefore described below as maintained structural anchors, not disguised as innocuous normalizations.

Table IA.1 records this distinction. The "principal discipline" shown for a fitted parameter is not an exclusion restriction and should not be read as a one-parameter/one-moment identification claim. All six fitted parameters move jointly, and IA.C reports the Jacobian and profile diagnostics. The table instead tells the reader what economic information makes each choice reasonable and, equally importantly, which values should be viewed as maintained benchmarks rather than estimates.

Table IA.1: Calibration rationale and empirical provenance

Object	Benchmark	Status	Reason for choice / empirical discipline
Parameters fitted jointly to the eight target moments			

Continued on next page

Table IA.1 continued

Object	Benchmark	Status	Reason for choice / empirical discipline
$\bar{\kappa}$	0.6217	Fitted jointly	Location of the participation-cost distribution. The level of stock-market participation is its principal discipline, with the wealth gradient helping to separate location from dispersion. Fixed participation costs are standard in quantitative household-finance models; see, for example, Vissing-Jørgensen (2002). The numerical value here is generated by the paper's moment fit, not imported from that study.
σ_κ	2.3279	Fitted jointly	Dispersion of participation costs. The participation gradient across wealth is especially informative about this parameter, jointly with the participation level. IA.C shows that $(\bar{\kappa}, \sigma_\kappa)$ form the weakest local direction, so neither is presented as a sharply estimated stand-alone primitive.
a	0.8173	Fitted jointly	Curvature of the automation frontier. The average automation and displacement moments are its main quantitative discipline. The task-allocation interpretation follows the automation literature, including Acemoglu and Restrepo (2018); the numerical curvature is chosen by the model fit rather than copied from that literature.
A_T	2.2524	Fitted jointly	Traditional-sector productivity. Relative wages and the displaced-worker outside option provide the principal information for this scale. Its units are model units; only ratios and implied labor-income outcomes have empirical content.
A_{AI}	2.7095	Fitted jointly	AI-sector productivity scale. It affects automation incentives, installed capital, and equity payoffs and is therefore jointly disciplined by the automation, financing-share, and return moments. It is a model scale parameter, not a measured TFP level.
$\bar{w}_0$	1.8590	Fitted jointly	Scale of the retained-wage schedule. The AI-to-traditional wage ratio and displacement moment are its main discipline. This parameter is distinct from household wealth; its value is in the model's wage units.
External data and literature benchmarks			
γ	5	Literature benchmark	A coefficient of relative risk aversion of five is a conventional intermediate value in quantitative household portfolio models. Benzoni et al. (2007), for example, use $\gamma = 5$ as their baseline and report sensitivity around it. Because welfare levels are sensitive to curvature, the paper re-estimates the fitted block over $\gamma \in [2, 10]$ rather than treating five as identified.

Continued on next page

Table IA.1 continued

Object	Benchmark	Status	Reason for choice / empirical discipline
R_f^a	1.02	Literature benchmark	A 2% annual safe return is standard in life-cycle portfolio calibrations; Cocco et al. (2005) use the same annual risk-free rate. The model compounds this value over the three-year horizon according to (IA.1).
(μ_w, σ_w)	$(-0.8, 1.1)$	SCF-based	These lognormal parameters are chosen to approximate the dispersion of household wealth in the 2022 Survey of Consumer Finances after expressing wealth in the model's units. The SCF is the appropriate source for U.S. household balance-sheet heterogeneity; see Aladangady et al. (2023). The pair $(-0.8, 1.1)$ is not a pair of published SCF moments and should not be read as such. The present calibration archive supports the SCF as the source of the intended wealth heterogeneity, but not a claim that these two transformed parameters are independently estimated from the SCF.
β_T	0.65	Literature benchmark	In the traditional-sector Cobb–Douglas technology, β_T is the labor elasticity and therefore the labor share under competitive pricing. Gollin (2002) reports adjusted labor shares for most countries in the range 0.65–0.80. The benchmark uses the lower end of that conventional range; A_T is then fitted to the wage data.
τ	3 years	Data-horizon alignment	The displacement incidence is interpreted over a three-year horizon. This lines up the model's job-loss window with the three-year recall window used by the modern U.S. Displaced Workers Survey; see Farber (2017). Certainty equivalents are computed over the model period and only then converted to annual reporting units.
Externally anchored but not directly estimated objects			
$\mathrm{Var}(\theta)$	0.04	External anchor	There is no independently observed time series of AI capability innovations that identifies this variance without using the same short return sample that is used to discipline the financial block. The benchmark therefore sets $\mathrm{sd}(\theta) = 0.20$ as a moderate reference dispersion and does not claim that 0.04 is a data estimate. Section 6 re-estimates the model over $\mathrm{Var}(\theta) \in [0.01, 0.12]$; the target criterion does not select 0.04, so the sensitivity rather than the benchmark point carries the inferential burden.

Continued on next page

Table IA.1 continued

Object	Benchmark	Status	Reason for choice / empirical discipline
σ_ε	0.12	Maintained shock scale	The two-point productivity shock places productivity at 1 ± 0.12. This is a maintained scale for aggregate AI-sector productivity risk, not a direct estimate of a structural TFP innovation and not a pure normalization. Its implications are disciplined indirectly by the targeted AI-equity premium and relative equity volatility. The choice matters for the quantitative decomposition of aggregate consumption and returns, so quantities such as the near-complete equity span of lottery-averaged consumption should be read as benchmark-model outcomes rather than universal constants. We keep this shock distinct from θ so that the capability state is not inferred mechanically from the same returns the model later asks securities to span.
(ϕ_R, b_0)	$(0.20, 0.30)$	Outside-option benchmark	These two parameters summarize reallocation costs and a lower bound on displaced-worker income. They are not separately estimated from a single empirical replacement-rate statistic. Their joint implication is instead checked out of sample: the run-of-record model produces a 32.9% earnings loss conditional on displacement, close to the large initial losses documented in administrative displacement studies such as Couch and Placzek (2010) and to the persistent losses in Jacobson et al. (1993). Section 6 varies the outside option because these inputs matter substantially for the level of the welfare burden.
(μ_C, W)	$(0.04, 2)$	Stylized ownership block	The model compresses the highly concentrated upper tail of U.S. wealth into a small capitalist group. The SCF motivates the need for a concentrated ownership group (Aladangady et al., 2023), but the 4% mass and wealth scale of two are not literal SCF percentile estimates. For the capitalist financing block, CRRA homotheticity implies that the optimal portfolio share is independent of the wealth scale W, so the key supply scale is the product $\mu_C W = 0.08$. The separate split into $\mu_C = 0.04$ and $W = 2$ is therefore a stylized representation; its economically relevant financing implication is disciplined by the targeted capitalist share of installed AI capital.

Continued on next page

Table IA.1 continued

Object	Benchmark	Status	Reason for choice / empirical discipline
ϖ	0.5	Partial pass-through benchmark	The productivity elasticity of retained wages is fixed at an intermediate value so that productivity shocks are neither fully absorbed by wages nor completely excluded from wages. Evidence that firms partially smooth worker earnings against firm-level shocks, such as Guiso et al. (2005), motivates partial pass-through, but does not identify the value 0.5 for an aggregate AI-sector shock. We therefore treat 0.5 as a maintained benchmark rather than an empirical estimate. This parameter affects the aggregate wage exposure to ε and hence some portfolio and total-consumption span statistics; the paper does not interpret the particular value 0.5 as a separately identified structural elasticity.
Maintained structural anchors and deliberately stripped benchmark restrictions			
$\bar{L}_T$	1	Relative-size benchmark	The traditional sector contains one unit of incumbent labor relative to the unit mass of exposed workers. Because displaced workers enter as $\bar{L}_T + \lambda$, this ratio affects how strongly displacement depresses the traditional-sector marginal wage; it is therefore not a pure scale normalization. A_T is fitted conditional on this relative-size benchmark, and the paper does not claim to estimate $\bar{L}_T$ separately.
$\bar{z}$	0.90	Maintained technology ceiling	The ceiling keeps a strictly positive residual human-task share and helps determine the shape of the automation frontier jointly with the fitted parameter a. It should not be read as a claim that 90% of U.S. jobs or tasks are empirically automatable. The moment set disciplines average automation and therefore the composite frontier generated by $(\bar{z}, a)$; it does not separately identify the ceiling.
ψ	1	Maintained adjustment-cost curvature	The separation-cost coefficient is economically consequential: it enters the firm's retention first-order condition and the denominator of the displacement response in Proposition 1. We set it to one as a simple benchmark curvature for the quadratic adjustment cost, not as an estimate of statutory firing costs and not as a normalization. The six fitted parameters, including the automation and wage scales, are estimated conditional on this value.
χ	0	Restricted baseline	Setting human-task complementarity to zero removes an additional productivity channel and makes the benchmark mechanism operate through task automation and retention. Positive complementarity is an extension, not something the benchmark data are claimed to reject.

Continued on next page

Table IA.1 continued

Object	Benchmark	Status	Reason for choice / empirical discipline
ν_θ	0	Reference regime	Zero is the rigid-with-respect-to-automation reference node used to define the institutional comparative static. It is not the paper's estimate of the U.S. wage regime. IA.B maps historical automation evidence to a point near 0.021 and a scenario range $[0, 0.13]$.
ϕ_θ	0	Restricted baseline	The benchmark turns off a direct capability effect on machine productivity. Capability therefore operates through the automation frontier. This restriction isolates the paper's labor-adjustment mechanism rather than claiming that direct machine-productivity effects are empirically zero.
(K_0, N_0)	$(0.6933, 0.9629)$	Frozen production reference anchors	These values are inherited from the run-of-record production normalization and are not empirical estimates. A common proportional rescaling of both anchors is a change of units when the productivity scale is adjusted correspondingly. Their *relative* scale is more than notation, however, because the state-dependent exponent reallocation depends on $\log(KN_0/(K_0N))$. The archive establishes these as the frozen run-of-record reference scales; it does not provide an independent empirical estimate of their ratio. We therefore treat that ratio as a maintained production anchor rather than claiming full normalization invariance.
Numerical choices			
$n_\theta = 9$, $n_q = 30$, brackets, tolerances	—	Numerical	These values trade computation time against accuracy and have no direct economic interpretation. IA.D checks higher quadrature orders, alternative root grids, the analytic equity derivative, and the smooth participation limit. The consumption floor 10^{-15} is an overflow guard and does not bind in accepted reported allocations.

Four rules summarize the calibration logic. An internally fitted parameter is defended by the moments and identification diagnostics that discipline it. A data-based input is defended by a transparent mapping to the outside evidence. A maintained structural anchor is defended by stating its economic role and limiting the interpretation of results conditional on it. A genuine normalization is defended by invariance to the associated change of units. This classification is more informative than attaching a citation to every fixed number.

The capitalist block illustrates why the economically relevant combination can be simpler than the primitive parameterization suggests. With no capitalist labor income and CRRA utility, capitalist consumption can be written $c_C = WR_p(\alpha_C)$. For $\gamma \neq 1$,

$$\mathbb{E}[u(c_C)] = W^{1-\gamma} \frac{\mathbb{E}[R_p(\alpha_C)^{1-\gamma}]}{1-\gamma},$$

so W does not affect the optimal portfolio share. Capitalist demand for installed AI capital is

therefore

$$D_C = \mu_C W \alpha_C^*,$$

and the financing block depends on the wealth scale primarily through $\mu_C W = 0.08$. This is why the paper treats the separate values $(\mu_C, W) = (0.04, 2)$ as a stylized ownership decomposition while disciplining their main equilibrium implication with the capitalist financing-share moment.

The outside-option block provides another example of this discipline. The values $\phi_R = 0.20$ and $b_0 = 0.30$ are not calibrated by forcing the model to match the displaced-worker earnings loss. Nevertheless, after the model is fit to displacement incidence and the other seven target moments, it implies a conditional earnings loss of 32.9%. Administrative studies find losses above 30% immediately following mass layoffs (Couch and Placzek, 2010), while the classic evidence of Jacobson et al. (1993) documents persistent losses of roughly one-quarter even years later. The untargeted match therefore supports the economic severity of the outside-option block without pretending that either ϕ_R or b_0 is separately identified.

Finally, the two aggregate-risk parameters σ_ε and ϖ deserve a narrow interpretation. They help determine how the productivity shock appears in equity payoffs, wages, and total consumption. They therefore matter for quantitative statements such as the very high equity span of lottery-averaged consumption. They do not create the paper's logical measurability result: conditional on a given aggregate state, an aggregate-state payoff still cannot distinguish a retained worker from a displaced worker. The reported quantitative span and welfare magnitudes, however, remain conditional on the maintained aggregate-risk specification. The paper accordingly treats those magnitudes as properties of the calibrated benchmark environment rather than structural constants.

IA.A.9 Complete run-of-record parameter inventory

Having explained the economic rationale for the substantive choices, the final step is complete parameter accounting. A reader should be able to tell not only why the benchmark values are used but also exactly which fields the code receives. Table IA.2 therefore lists every field in the run-of-record parameter structure, including derived and numerical fields. It is an audit table rather than a second calibration table; the justifications for the economically substantive inputs are in Table IA.1.

Only six fields are fitted jointly to the eight targeted moments: $(\bar{\kappa}, \sigma_\kappa, a, A_T, A_{AI}, \bar{w}_0)$. The capability variance is anchored at 0.04, while risk aversion and the remaining preference, technology, labor-market, and wealth-distribution inputs are externally fixed, maintained structural anchors, genuine unit conventions, or scenario parameters. The inventory also records fields that are numerically necessary but economically inert at the run of record. This distinction matters because the parameter constructor's defaults are not identical to the values used by the benchmark loader.

Table IA.2: Complete run-of-record parameter inventory

Parameter (code field)	Economic or numerical role	Value	Treatment
Fitted and anchored quantitative parameters			
$\bar{\kappa}$ (`bar_kappa`)	Mean of the lognormal participation-cost distribution.	0.6217441735	Fitted
σ_κ (`sigma_kappa`)	Log dispersion of the participation-cost distribution.	2.327881312	Fitted
Var(θ) (`var_theta`)	Variance of the mean-one capability shock.	0.04	External anchor; sensitivity

Continued on next page

Table IA.2 continued

Parameter (code field)	Economic or numerical role	Value	Treatment
a (`a`)	Curvature/intensity of the automation frontier.	0.8172872943	Fitted
A_T (`A_T`)	Traditional-sector productivity scale.	2.25240693	Fitted
A_{AI} (`A_AI`)	AI-sector productivity scale.	2.709462151	Fitted
$\bar{w}_0$ (`bar_w0`)	Scale of the retained AI-wage schedule.	1.85902044	Fitted
Externally fixed preferences and asset-market inputs			
γ (`gamma`)	CRRA coefficient of relative risk aversion.	5	Literature benchmark
R_f^a (`Rf`)	Gross annual risk-free return.	1.02	Literature benchmark
σ_ε (`sigma_eps`)	Half-width of the symmetric two-point AI-productivity shock.	0.12	Maintained shock scale
ϕ_R (`phi_R`)	Level reallocation cost subtracted from the traditional-sector wage.	0.2	Outside-option benchmark
b_0 (`b0`)	Floor on the displaced-worker outside option.	0.3	Outside-option benchmark
μ_w (`mu_w`)	Mean of log initial worker wealth.	-0.8	SCF-based
σ_w (`sig_w`)	Standard deviation of log initial worker wealth.	1.1	SCF-based
μ_C (`mu_C`)	Population mass of capitalists.	0.04	Stylized ownership block
W (`W`)	Investable wealth per capitalist.	2	Stylized ownership block
Technology and labor-market inputs			
β_T (`beta_T`)	Diminishing-returns exponent in the traditional sector.	0.65	Literature benchmark
$\bar{L}_T$ (`Lbar_T`)	Incumbent traditional-sector labor stock.	1	Relative-size benchmark
$\bar{z}$ (`zbar`)	Upper bound on the automated task share.	0.9	Maintained technology ceiling
ψ (`psi`)	Scale of the quadratic separation cost.	1	Maintained curvature; loader override
ϖ (`varpi`)	Elasticity of the retained AI wage to ε.	0.5	Partial pass-through benchmark
χ (`chi`)	Complementarity term in $h(z) = 1 + \chi z$.	0	Restricted baseline
ν_θ (`nu_theta`)	Elasticity of retained wage to the unautomated share $1 - z$.	0	Reference regime; IA.B mapping
ϕ_θ (`phi_theta`)	Capability loading in machine productivity Γ_θ.	0	Restricted baseline
Shock-grid specification			
n_θ (`n_theta`)	Number of capability quadrature nodes.	9	As run; loader override
π_d (`pi_disaster`)	Probability assigned to the optional disaster branch.	0	As run; loader override
θ_d (`theta_disaster`)	Capability realization in the optional disaster branch.	0.05	Inert at benchmark
Derived state-space objects			
ε_H (`eps_H`)	High productivity realization $1 + \sigma_\varepsilon$.	1.12	Derived
ε_L (`eps_L`)	Low productivity realization $1 - \sigma_\varepsilon$.	0.88	Derived
π_ε (`pi_eps`)	Probability of the high productivity realization.	0.5	Derived
$\{\theta_j\}$ (`theta_grid`)	Mean-one capability grid.	`vector[9]`	Derived
$\{\pi_{\theta,j}\}$ (`pi_theta`)	Probabilities on the capability grid.	`vector[9]`	Derived
n_s (`n_states`)	Number of active aggregate states, $2n_\theta$.	18	Derived
$R_f^{(\tau)}$ (`Rf_model`)	Gross risk-free return over the three-year model period.	1.061208	Derived
Production reference anchors			

Continued on next page

Table IA.2 continued

Parameter (code field)	Economic or numerical role	Value	Treatment
K_0 (`K0`)	Frozen capital reference scale in production.	0.6933081746	Maintained reference; loader override
N_0 (`N0`)	Frozen retained-headcount reference scale in production.	0.9628764434	Maintained reference; loader override
Scenario guard			
s_k (`persist_sk`)	Persistence-scenario scale, equal to one at the benchmark.	1	Inert scenario guard
Numerical settings			
n_q (`n_quad`)	Quadrature order used for the worker-wealth integral and fixed-grid diagnostics.	30	Numerical
$K_{\min}$ (`K_lo`)	Lower bracket for the capital-equilibrium search.	0.01	Numerical
$K_{\max}$ (`K_hi`)	Upper bracket for the capital-equilibrium search.	10	Numerical
$r_{E,\min}$ (`rE_lo`)	Lower bracket on the diagnostic return-based solver path.	-0.05	Numerical
$r_{E,\max}$ (`rE_hi`)	Upper bracket on the diagnostic return-based solver path.	2	Numerical
$n_{r_E}^{\text{fast}}$ (`n_rE_fast`)	Coarse return-grid size on the diagnostic solver path.	80	Numerical
$n_{r_E}^{\text{full}}$ (`n_rE_full`)	Fine return-grid size on the diagnostic solver path.	600	Numerical
tol (`tol`)	Equilibrium convergence tolerance.	10^{-10}	Numerical
τ (`period`)	Length of one model period, in years.	3	DWS-aligned horizon
$c_{\min}$ (`c_floor`)	Numerical consumption floor used only as an overflow guard.	10^{-15}	Numerical

There are 47 fields in total. Five deserve special attention because the run-of-record loader overrides the default constructor: $\psi = 1$ rather than 2, $n_\theta = 9$ rather than 11, $\pi_d = 0$ rather than 0.02, and $(K_0, N_0) = (0.6933081746, 0.9628764434)$ rather than $(1, 1)$. None is part of the six-parameter fitted block. The value of ψ is a maintained adjustment-cost curvature; the next two fields determine the active numerical shock grid; and (K_0, N_0) are frozen production reference anchors. The last pair should not be interpreted as two estimated physical quantities, but neither should their relative scale be described as economically irrelevant: because the production exponent varies with automation, the ratio K_0/N_0 enters the exponent-reweighting channel.

Two distinctions also prevent common bookkeeping mistakes. First, $\bar{w}_0 = 1.85902044$ is the *wage-schedule scale*; household wealth is governed separately by $(\mu_w, \sigma_w) = (-0.8, 1.1)$. Second, the capability-return loading β_R in Section 3 of the paper is not a 48th parameter. It is computed from the equilibrium return vector by (IA.29); the admissible region for that derived object constrains the fitted economy.

IA.A.10 Run-of-record and reproduction conventions

The audit closes by fixing the numerical vintage. This is necessary because the project contains both constructor defaults and scenario outputs; quoting a parameter without its run-of-record loader can reproduce the wrong economy even when the code itself is unchanged.

The benchmark parameter inventory in Table IA.2 is the Job-9 audit of `ror-constrained-v1`, dated August 5, 2026. The canonical facts record for that benchmark was created at commit

`82bef6c`; the parameter audit that generated the 47-field inventory was produced at commit `7d30d5f`. Later wage-regime calculations use this benchmark as the $\nu_\theta = 0$ anchor and then apply the counterfactual rules stated in Section 4 of the paper and IA.E.

The reporting pipeline keeps economic objects tied to their numerical source. Generated tables and macros are promoted from recorded outputs rather than retyped from the prose. The facts registry records, for each promoted object, its source file, model vintage, units, annualization convention, reporting group, and welfare convention. Numerical validation in IA.D checks the stable capital root, the smooth participation limit, the analytic equity derivative, and grid refinement. IA.C separately reports identification and criterion-set diagnostics. This division of labor is deliberate: IA.A tells the reader what economy is being solved; the later sections ask how accurately it is solved and how strongly its parameters and conclusions are disciplined.

IA.B Empirical Mapping of the Wage-Adjustment Parameter

The wage parameter in the model is not directly observed in existing data. What the historical evidence provides is a comparison of how much an adverse technology shock moves wages and how much it moves employment. This section constructs that observable adjustment-share statistic, maps it into the model's wage elasticity, and states exactly what can and cannot be inferred from the exercise.

The distinction matters for interpretation. The paper does not estimate a mature-AI wage-setting regime from an AI displacement episode; no such episode yet exists. Instead, it uses an earlier U.S. automation and trade-displacement episode to discipline a range of wage-versus-employment adjustment that is plausible for the quantitative exercise. The resulting interval $[0, 0.13]$ is therefore a historically informed U.S. scenario range. It is not a confidence interval for a technology-invariant structural parameter, and it does not locate any foreign country on the model's wage-setting axis.

IA.B.1 The empirical object

The first step is to define the statistic that can be measured in historical data. Suppose a common adverse shock produces a wage coefficient β_w and an employment coefficient β_E for the same exposed population, measured on a common base. The adjustment share is

$$\widehat{\nu} = \frac{|\beta_w|}{|\beta_w| + |\beta_E|}. \tag{IA.45}$$

It equals zero when all measured adjustment occurs through employment and approaches one when all measured adjustment occurs through wages. Absolute values are used because the object is the split across adjustment margins, rather than the sign convention used to code the underlying shock.

The primary evidence comes from the within-manufacturing specifications in Table 7 of Autor et al. (2013). Panel A reports the change in log manufacturing employment and Panel B reports the change in average log weekly manufacturing wages for the same broad manufacturing population and specification. The published coefficients used in the corrected audit are

$$\widehat{\beta}_E = -4.231 \quad (1.047), \qquad \widehat{\beta}_w = 0.150 \quad (0.482), \tag{IA.46}$$

where parentheses contain reported standard errors. Hence

$$\widehat{\nu} = \frac{0.150}{0.150 + 4.231} = 0.0342. \tag{IA.47}$$

The use of this pair is intentional. An earlier internal construction combined the wage coefficient in Table 7 with an employment-to-population coefficient from Table 5 and converted the latter using an assumed 10% manufacturing share. That construction implied an employment response of -5.96 and an adjustment share of 0.0245. It is superseded. Table 7 already provides the directly comparable log-manufacturing-employment coefficient, so no employment-base imputation is required. The old calculation is retained in the source audit only so that the revision can be reconstructed; no result in the paper uses it.

IA.B.2 From the observed share to the structural elasticity

The empirical share in (IA.45) and the model parameter ν_θ are not the same object. In the model, ν_θ is the elasticity of the retained wage with respect to the unautomated task share,

$$\nu_\theta = \frac{\partial \ln \bar{w}_A}{\partial \ln(1 - z)}.$$

The same value of this elasticity can generate different observed wage shares because employment also responds endogenously. We therefore map the empirical share through the model's retention system rather than equating $\widehat{\nu}$ with ν_θ.

The mapping can be seen directly from the local responses. At an interior retention solution and holding installed capital fixed,

$$g_w(s;\nu_\theta) \equiv \frac{\partial \ln \bar{w}_A(s)}{\partial \theta} = -\nu_\theta \frac{z_\theta}{1-z}, \tag{IA.48}$$

while Proposition 1 of the paper implies

$$g_N(s;\nu_\theta) \equiv \frac{\partial \ln N(s)}{\partial \theta} = \frac{z_\theta}{N} \frac{\ell_N - (1-\nu_\theta)/(1-z)}{z/N + \psi/(1-\psi+\psi N)}, \qquad \ell_N = \ln\left(\frac{KN_0}{K_0 N}\right). \tag{IA.49}$$

Both responses contain the common factor z_θ. The empirical adjustment share is therefore informative about the relative wage and employment margins without requiring a separate estimate of the scale of the capability innovation.

The published statistic is a ratio of two regression coefficients. Its closest model analogue is consequently a ratio of aggregate responses, rather than an average of state-by-state ratios. Our primary mapping is

$$m_{\text{agg}}(\nu_\theta) = \frac{\left|\sum_s \pi_s g_w(s;\nu_\theta)\right|}{\left|\sum_s \pi_s g_w(s;\nu_\theta)\right| + \left|\sum_s \pi_s g_N(s;\nu_\theta)\right|}. \tag{IA.50}$$

All eighteen aggregate states enter. At a full-retention corner, $N^*(s) = 1$, the employment response is locally zero and the wage response continues to enter. This convention is important because the empirical regression coefficient is an aggregate response and does not condition on the observation being at an interior firm solution.

For transparency, we also report the secondary object

$$m_{\text{int}}(\nu_\theta) = \frac{\sum_{s\in\mathcal{S}_I} \pi_s \dfrac{|g_w(s;\nu_\theta)|}{|g_w(s;\nu_\theta)| + |g_N(s;\nu_\theta)|}}{\sum_{s\in\mathcal{S}_I} \pi_s}, \tag{IA.51}$$

where $\mathcal{S}_I$ is the set of states with an interior retention solution. This is the probability-weighted mean of state-specific adjustment shares among states in which employment actually contracts. It is a useful diagnostic, but it is not the primary mapping because an average of ratios is not the same estimand as a ratio of two aggregate regression responses.

The mapping used to construct the historical range fixes capital at the run-of-record value. This isolates the wage-versus-retention margin from a second general-equilibrium response through installed capital. Section 4 of the paper instead changes ν_θ and re-solves K; IA.E reports the corresponding forward mapping for that family. The distinction is quantitative rather than conceptual. For example, at $\nu_\theta = 0.13$ the fixed-K primary share is 0.242, while the same fixed-calibration economy with capital re-solved has a share of about 0.261.

Inverting the primary map gives

$$\widehat{\nu} = 0.0342 \quad \Longrightarrow \quad \nu_\theta = 0.0214, \tag{IA.52}$$

reported in the paper as approximately 0.021. Under the secondary interior-state convention the same empirical point maps to 0.0258. The difference is small at the point estimate but becomes larger toward the upper end of the range because more states reach the full-retention corner.

The discrete state grid also makes m_{agg} a stepwise-smooth function of ν_θ. When a state first reaches $N^* = 1$, its employment derivative drops to zero. This is a feature of combining a corner

Table IA.3: Structural wage elasticity and the observable adjustment share

ν_θ	$m_{\text{agg}}(\nu_\theta)$	$m_{\text{int}}(\nu_\theta)$	Probability at full-retention corner
0.00	0.000	0.000	0.175
0.05	0.103	0.066	0.378
0.10	0.193	0.129	0.378
0.13	0.242	0.165	0.378
0.20	0.401	0.249	0.500

The table uses the corrected Job-12R reference mapping: all fitted primitives and K are held at the run-of-record economy while ν_θ varies. The primary column is the ratio of aggregate responses in (IA.50); the secondary column is the mean of interior-state ratios in (IA.51). Corner probabilities refer to states with $N^* = 1$, at which the local employment response is zero.

solution with a nine-node capability grid, not an optimization failure. Repeating the mapping with 9, 15, and 21 capability nodes changes the primary inversion of the upper endpoint by only 0.004 in ν_θ, so the reported mapping is not materially driven by the chosen quadrature order.

IA.B.3 Sampling uncertainty and the covariance problem

The most important statistical qualification concerns the joint sampling uncertainty of the wage and employment coefficients. Their standard errors are published, but their sampling covariance is not. Treating the two coefficients as independent is therefore an assumption, not a fact. This matters more for the upper endpoint than for the point estimate.

The clean way to see the issue is to work first with the signed ratio

$$r = \frac{\beta_w}{\beta_E}, \qquad \widehat{\nu} = \frac{|r|}{1 + |r|}. \tag{IA.53}$$

Let s_w and s_E be the reported standard errors and let ρ denote the unknown sampling correlation between the two coefficient estimates. A Fieller confidence set for r contains the values satisfying

$$\left(\widehat{\beta}_w - r\widehat{\beta}_E\right)^2 \leq c^2 \left[s_w^2 + r^2 s_E^2 - 2r\rho s_w s_E\right], \tag{IA.54}$$

with c the relevant normal critical value; see Fieller (1954). The resulting set for the nonnegative adjustment share is obtained by applying $|r|/(1 + |r|)$ to the Fieller set.

This calculation corrects an imprecise statement in an earlier draft. The ratio is not problematic because its denominator is weak. The employment coefficient in (IA.46) has an absolute t-ratio of about 4.0, so the Fieller set is bounded. The issue is instead that the unreported covariance changes the shape and endpoints of the joint confidence region.

At $\rho = 0$, the exact 95% Fieller set for the signed ratio is approximately

$$[-0.303,\ 0.210].$$

Folding this set through (IA.53) gives an adjustment-share set with upper endpoint 0.232. The primary structural map then gives

$$\nu_\theta \leq 0.124. \tag{IA.55}$$

A simpler one-sided delta construction gives an adjustment-share upper bound of 0.187, which maps to $\nu_\theta \leq 0.097$. We use the Fieller construction as the more conservative benchmark and round its zero-correlation structural endpoint upward when defining the paper's $[0, 0.13]$ scenario range.

The unknown covariance prevents this range from being described as an unconditional 95% confidence set. Letting ρ vary over its full admissible interval $[-1, 1]$, the Fieller adjustment-share upper endpoint ranges from 0.204 to 0.334. Under the primary structural map, these endpoints correspond to

$$\nu_\theta \in [0.107,\ 0.156] \quad \text{for the upper endpoint alone.} \tag{IA.56}$$

The fixed region $[0, 0.13]$ contains the full Fieller set only for a subset of possible correlations, approximately $\rho \in [-0.75, 0.09]$. This is why the main paper calls $[0, 0.13]$ a historically informed benchmark range rather than a confidence interval.

The covariance uncertainty is also why we do not attach a conventional standard error to the mapped structural elasticity. The point estimate 0.021 is useful as a reference location, but the paper's economic results are reported over a range rather than built around that point. Estimating the joint covariance from the underlying microdata would sharpen this mapping, but that is additional empirical work and is not required for the claims made here.

IA.B.4 Why local-labor-market wage responses are a different object

Broader local-labor-market studies often report substantially larger wage responses to automation or trade shocks. Those estimates are important for the total incidence of technological change, but they do not directly measure the wage elasticity in the firm's retention problem. The reason is composition.

The retained wage in the model is $\bar{w}_A(s)$. A broader average earnings measure for the initially exposed cohort is instead

$$\omega_L(s) = [1 - \lambda(s)]\bar{w}_A(s) + \lambda(s)b(s). \tag{IA.57}$$

Even when $\nu_\theta = 0$ and the wage of a retained worker is perfectly rigid with respect to automation, ω_L moves because the fraction displaced changes and displaced workers receive the outside option b. A commuting-zone wage measure may additionally include workers outside the directly exposed sector.

This distinction is quantitatively important. At $\nu_\theta = 0$, the retained-wage adjustment share in (IA.50) is exactly zero, but the model's composition-inclusive analogue is about 0.249. Thus a positive wage share measured for a broad local population is entirely consistent with a zero retained-wage elasticity. Mapping such a composition-inclusive estimate through the retained-wage bridge would mechanically attribute separation and reemployment effects to wage flexibility and would therefore answer a different question.

Table IA.4 records how the historical sources are used. Only the within-manufacturing pair from Autor et al. (2013) is used to construct the benchmark region. The commuting-zone estimates in Acemoglu and Restrepo (2020) and the broader Autor–Dorn–Hanson results are reported as context for composition-inclusive adjustment, not as alternative measurements of ν_θ.

IA.B.5 What the historical range does and does not establish

The evidence supports a narrow use of the mapping. It says that earlier U.S. adjustment episodes with directly measured incumbent manufacturing wages were employment-intensive, and it provides a disciplined scale for asking how the model changes as more adjustment is moved onto wages. The point share 0.0342 maps near $\nu_\theta = 0.021$, while the conservative zero-covariance Fieller construction motivates an upper benchmark near 0.13.

Three stronger statements are not warranted. First, $[0, 0.13]$ is not a sampling confidence interval because the wage–employment covariance is not reported. Second, it is not technology invariant: AI

may affect different tasks, bargaining arrangements, and occupations from the historical episodes used here. Third, it is not a cross-country institutional ranking. In particular, the paper does not assign Germany, France, Europe, or any other named economy to a value of ν_θ without a same-shock, same-sample wage/employment pair satisfying the audit requirements above.

The quantitative fit provides a separate consistency check, not a statistical validation of the historical range. Holding ν_θ fixed and re-estimating the remaining fitted block, the eight-moment criterion rises from about 2.29 at $\nu_\theta = 0$ to 2.68 at 0.13, 5.65 at 0.20, and 9.01 at 0.35. We describe this as deterioration in fit relative to the U.S. targets, not as rejection based on a known sampling distribution. IA.C reports the criterion and identification diagnostics in full.

The role of IA.B is therefore limited but useful. It converts a transparent historical adjustment statistic into the structural axis used by the model, shows the sensitivity of that conversion to aggregation and sampling covariance, and makes clear which broader automation evidence does not measure the same object. IA.E then takes the structural axis as given and asks how the amount and form of displacement risk change along the fixed-calibration wage-regime family.

Table IA.4: Source audit for the wage–employment mapping

Source	Population / margin	Treatment	Reason
Autor et al. (2013), Table 7, Panels A–B	Manufacturing workers in exposed U.S. commuting zones	Primary	Wage and employment coefficients are reported on the same within-manufacturing base. This is the closest published pair to the retained-worker margin in the model.
Autor–Dorn–Hanson Table-5-derived pair	Manufacturing employment constructed from an employment/population coefficient	Superseded	An earlier internal calculation divided the Table-5 coefficient by an assumed manufacturing share. Table 7 makes that conversion unnecessary. The row is retained only to document the correction.
Acemoglu and Restrepo (2020), commuting-zone and aggregate robot estimates	All local or aggregate workers	Context only	Average wages include composition and reallocation effects. They do not isolate the wage of workers retained at the directly exposed employer.
Dauth et al. (2021)	German robot exposure	Classified out	The study is informative about institutional adjustment, including retention and entry hiring, but the archived evidence does not supply the same-shock, same-sample retained-wage/employment pair required by (IA.45). It therefore does not locate Germany on the ν_θ axis.
Graetz and Michaels (2018)	Industry robot adoption across 17 countries	Classified out	The estimates do not provide the adverse same-population wage/employment pair needed for the adjustment-share mapping; aggregate employment is not estimated to contract in the required way.

The audit requires a wage and employment response to the same shock, for the same population, on a common base, and with the wage concept identified as retained-worker or composition-inclusive earnings. A source that is useful for the economics of automation need not satisfy these requirements for the specific mapping in this paper.

IA.C Identification and Criterion-Set Diagnostics

The quantitative model is fitted with more moments than free fitted parameters, but that fact alone does not establish that the economically important objects are well pinned down. This section asks a more useful question: can parameter combinations that fit nearly as well as the benchmark produce a materially different conclusion about the financial reach of displacement risk? We answer that question with three complementary diagnostics—a multi-start and profile search, a local moment Jacobian, and direct criterion-set envelopes for the economic objects.

These calculations have a deliberately limited interpretation. They describe the numerical and local identification of the calibrated model conditional on the eight target moments, the maintained external calibration, and the event-study restriction on the return loading. They do not provide a theorem of global identification, and they are not sampling confidence sets. The distinction is important because the target moments come from different data sources and the paper does not have their joint sampling covariance matrix.

IA.C.1 Criterion and search design

The first task is to make clear what the optimizer is minimizing. Let ϑ denote the six internally fitted parameters,

$$\vartheta = (\bar{\kappa}, \sigma_\kappa, a, A_T, A_{AI}, \bar{w}_0)',$$

and let $m_i(\vartheta)$, $\widehat{m}_i$, and s_i denote the model moment, target, and assigned uncertainty scale for target i. The criterion is

$$L_8(\vartheta) = \sum_{i=1}^{8} \left[\frac{m_i(\vartheta) - \widehat{m}_i}{s_i} \right]^2 = r(\vartheta)' r(\vartheta). \tag{IA.58}$$

Equivalently, in unstandardized units, $W = \mathrm{diag}(s_1^{-2}, \ldots, s_8^{-2})$. The run-of-record minimum is

$$L_{8,\min} = 2.28797. \tag{IA.59}$$

Equation (IA.58) is a diagonal minimum-distance criterion. It is not an efficient-GMM objective based on an estimated joint covariance matrix. Consequently, neither $L_{8,\min}$ nor differences $L_8 - L_{8,\min}$ are given a chi-squared sampling interpretation in this paper. The eight moments and their construction are reported in IA.D.

The capability-return loading

$$\beta_R = \frac{\mathrm{Cov}(R_E, \theta)}{\mathrm{Var}(\theta)} \tag{IA.60}$$

is also not a ninth target inside L_8. Instead, the event evidence defines the maintained magnitude restriction

$$|\beta_R(\vartheta)| \leq 0.056. \tag{IA.61}$$

Parameter vectors outside this region are rejected rather than penalized. Risk aversion, the capability variance, the wage-regime parameter, and the other externally fixed inputs in IA.A are held at their benchmark values throughout the identification exercises in this section.

The search is designed to reduce the risk that a non-smooth objective is being read from a single local optimum. The run of record uses six polishing starts, two warm and four cold. All six return feasible solutions. Two finish within $\Delta L_8 \leq 0.05$ of the best value; the other four are materially worse and are classified as dominated rather than treated as sampling dispersion. The two competitive solutions imply reservation prices within 0.28 annual basis points of one another.

Table IA.5: Selected economies along the equity–capability return-loading ridge

Economy	β_R	L_8	$\mathcal{D}(A_0)$ (annual bp)	Raw Ω_s (%)	Reservation price (annual bp)
Run of record	-0.05594	2.288	84.33	2.56	218.99
Committed reference	-0.04861	2.539	75.45	2.50	206.76
Raw-CI lower edge	-0.03700	2.585	86.96	2.86	201.20
Best stored unrestricted point	-0.16938	1.899	90.34	2.42	340.57

The first three rows satisfy the maintained magnitude fence $|\beta_R| \leq 0.056$; the last row is a diagnostic outside that region. $\Omega_s = [\mathcal{D}(A_0) - \mathcal{D}(A_s)]/\mathcal{D}(A_0)$ uses the benchmark cross-menu normalization described in Appendix B of the paper. The reservation price is included only to show that economically meaningful price objects move along the ridge; it is not a welfare gain.

This agreement is useful but should not be overstated. Multiple starts can show that the reported basin is repeatedly found; they cannot prove that no other minimum exists elsewhere in the parameter box. For the same reason, the equilibrium solver follows the stable capital-market root continuously from the benchmark rather than mechanically choosing the first crossing. A change in equilibrium branch would otherwise look like a parameter response even when it is only a root-selection discontinuity.

IA.C.2 The return-loading ridge and the event-study boundary

The simplest weak direction is visible in the return loading. The criterion would improve if the model were allowed to choose a more negative β_R, while the event evidence limits its magnitude. The useful diagnostic is therefore not merely whether the constraint binds, but whether moving along this weakly disciplined direction changes the paper's economic answer.

The boundary check verifies that the restriction is active. Relaxing (IA.61) from 0.056 to 0.062 lowers the criterion from 2.288 to 2.192; tightening it to 0.050 raises the criterion to 2.376. The fit therefore pulls outward through the benchmark boundary. This is why the run-of-record value

$$\beta_R = -0.05594$$

is described as a constrained calibration rather than as an unconstrained estimate.

Two empirical regions should not be conflated. The pooled event estimate implies the raw asymmetric interval

$$[-0.037,\ 0.056],$$

whereas the maintained magnitude fence is

$$[-0.056,\ 0.056].$$

The benchmark lies inside the magnitude fence but outside the raw interval on its negative side. Re-estimating with the raw lower endpoint $\beta_R = -0.037$ raises L_8 to 2.585. Dropping the magnitude restriction altogether allows the criterion to continue down the ridge, with the best stored unrestricted point at $\beta_R = -0.16938$ and $L_8 = 1.899$.

Table IA.5 records the economic quantities at representative points. The point of the table is not that the burden level is invariant—it is not. The point is that even where the criterion has difficulty distinguishing the return loading, the fraction of the displacement burden changed by complete aggregate-state trading remains small.

Across the three stored economies inside the magnitude fence, $\mathcal{D}(A_0)$ ranges from 75.45 to 86.96 annual basis points, while Ω_s ranges only from 2.50% to 2.86%. Across the broader nine-point stored ridge, which deliberately extends outside the event fence, the burden ranges from 75.45 to 96.34 basis points and Ω_s from 1.85% to 2.86%. Figure 1 of the paper plots the criterion, burden, and completion share over this same ridge.

This distinction is the central identification lesson from the profile. The absolute welfare level is less tightly disciplined than the financial boundary. A flat or sloping criterion profile does not by itself establish either stability or instability of an economic conclusion; the economic object must be recomputed along the weak direction.

IA.C.3 Local Jacobian and moment sensitivity

The ridge is a global numerical diagnostic. The local Jacobian answers a different question: which combinations of moments move when each fitted parameter is perturbed near the benchmark? This is useful for understanding where the local weak directions come from, but it should be read as a local linearization rather than as a substitute for the profiles above.

Let

$$J_{ij} = \frac{\partial m_i(\vartheta)}{\partial \vartheta_j} \tag{IA.62}$$

denote the 8×6 moment Jacobian and let

$$G = W^{1/2} J \tag{IA.63}$$

be the Jacobian in standardized-moment units. The reported derivative uses central differences on the smooth analytic-threshold participation path,

$$J_{\cdot j}(h) \approx \frac{m(\vartheta + h_j e_j) - m(\vartheta - h_j e_j)}{2h_j}, \qquad h_j = h|\vartheta_j|, \tag{IA.64}$$

with $h \in \{10^{-3}, 2 \times 10^{-3}\}$.

At the benchmark, G is numerically full column rank. Its condition number is approximately 763.9, while

$$\operatorname{cond}(J'WJ) = \operatorname{cond}(G'G) \simeq 5.81 \times 10^5.$$

The second number is essentially the square of the first, as it should be for the corresponding normal matrix. The weakest singular direction loads mainly on the two participation-cost parameters: in absolute value, approximately 0.823 on $\bar{\kappa}$ and 0.568 on σ_κ. Thus the local weakness is not a general failure of rank; it is concentrated in the location–dispersion tradeoff of the participation cost distribution.

For scale, define the curvature quantity

$$s_j^{\mathrm{curv}} = \sqrt{[(J'WJ)^{-1}]_{jj}}, \qquad q_j = \frac{|\widehat{\vartheta}_j|}{s_j^{\mathrm{curv}}}. \tag{IA.65}$$

Because W is a diagonal matrix of assigned uncertainty scales rather than an estimated joint sampling covariance matrix, q_j is a *local curvature diagnostic*, not a sampling t-statistic. The values in Table IA.6 make the same point as the singular vector: $\bar{\kappa}$ is locally much less sharply curved than the production and wage parameters.

Two complementary moment rankings are also useful. The "direct response" column asks which standardized target moment moves most when a parameter is perturbed. The "estimator

Table IA.6: Local identification diagnostics for the six fitted parameters

Parameter	Estimate	Curvature q_j	Largest direct-response moment	Largest estimator-sensitivity moment
$\bar{\kappa}$	0.6217	0.71	Participation rate	Participation gradient
σ_κ	2.3279	3.89	Participation rate	Participation gradient
a	0.8173	11.79	Displacement rate	Displacement rate
A_T	2.2524	7.33	AI/traditional wage ratio	AI/traditional wage ratio
A_{AI}	2.7095	17.85	Participation rate[a]	Capitalist financing share
$\bar{w}_0$	1.8590	13.76	Participation rate	Capitalist financing share

The curvature quantity is defined in (IA.65) and is not a sampling t-statistic. The direct-response ranking uses the largest absolute standardized Jacobian entry for each parameter. The estimator-sensitivity ranking uses the largest absolute scaled entry of Λ in (IA.66). For the latter ranking, the top moment exceeds the runner-up by factors of 4.80, 11.26, 3.23, 1.64, 5.54, and 6.03, respectively, in the order shown.

[a] The direct-response ranking for A_{AI} is the one unstable case in the finite-difference battery: its leading moment switches between the participation rate and the equity-volatility ratio across the two step sizes.

sensitivity" column instead asks which one-standard uncertainty-unit perturbation of a target would move the jointly fitted parameter most after all six parameters adjust. With the sign convention of the implementation, the latter uses

$$\Lambda = -(J'WJ)^{-1}J'W, \tag{IA.66}$$

and ranks moments by the absolute scaled elements of Λ. This is the local sensitivity logic of Andrews et al. (2017). The sign convention is immaterial for the rankings reported here.

The two rankings agree for a and A_T but differ for four of the six parameters. There is no contradiction. A moment can respond strongly to one parameter while still have little effect on that parameter estimate once the other five parameters are allowed to adjust jointly. This is why the paper does not attach a one-to-one statement such as "moment i identifies parameter j" to the calibration.

A further useful fact is that β_R, although excluded from the eight-moment objective, is highly responsive to several fitted parameters. If it is admitted as a ninth diagnostic moment, it is the most responsive object for five of the seven calibrated quantities in the broader audit. This helps explain why the external event-study fence materially shrinks the admissible parameter region even though β_R is not included in L_8.

IA.C.4 Finite-difference and quadrature stability

A local Jacobian in a model with participation thresholds and portfolio corners is only useful if its qualitative message survives reasonable numerical choices. The finite-difference battery therefore changes both the derivative step and the household integration grid. The purpose is not to select the step size that gives the most attractive ranking; it is to identify which rankings are genuinely stable.

The primary calculation uses $n_q = 30$ household quadrature nodes and relative central-difference steps $h = 10^{-3}$ and 2×10^{-3}. The grid check compares the primary result with calculations averaged over $n_q \in \{36, 40, 44, 48\}$. The estimator-sensitivity ranking in the final column of Table IA.6 is unchanged for all six parameters across these checks. The direct-response ranking is unchanged for five of six; only A_{AI} switches between the participation rate and the equity-volatility ratio.

Table IA.7: Calibrated criterion-set envelopes

Criterion slack	Candidates	$\mathcal{D}(A_0)$ (annual bp)	Raw Ω_s (%)	$\eta^2_{s,P}(u'd)$ (%)	$\eta^2_{s,NP}(u'd)$ (%)
Benchmark	1	84.33	2.56	4.10	4.04
$\Delta = 0.50$	32	50.87–100.73	1.94–3.24	3.37–4.33	3.46–4.33
$\Delta = 1.00$	66	37.39–122.47	1.65–3.64	3.00–4.34	3.16–4.38
$\Delta = 2.00$	95	21.89–129.85	0.52–4.86	2.54–5.06	2.83–5.16
$\Delta = 3.84$	123	15.62–141.33	0.44–5.54	2.34–5.12	2.66–5.26
$\Delta = 6.00$	159	13.25–154.36	-1.35–5.91	2.12–5.12	2.40–5.28

The six fitted parameters vary jointly; the external calibration is held fixed and the magnitude fence $|\beta_R| \leq 0.056$ is imposed at every candidate. The reported candidate count is the pooled feasible set at each slack. The labels 3.84 and 6.00 are scale references only. Although they are numerically close to familiar chi-squared critical values, no chi-squared coverage interpretation is attached to them. The ranges are attained inner approximations, not confidence intervals or analytical outer bounds.

The derivative magnitudes themselves are less stable than the ordinal rankings. The largest relative step-size sensitivity is about 0.72, for the equity-premium response to A_{AI}. The largest quadrature-grid sensitivity is about 0.14, for the participation-gradient response to $\bar{w}_0$. These movements are large enough that the paper does not use the Jacobian to make precise elasticity claims. The defensible use is ordinal: the participation-cost block contains the weak local direction, the production and wage parameters have substantially more local curvature, and the estimator-sensitivity rankings are stable.

Using the analytic participation threshold is important for this exercise. It removes the artificial jumps generated when a moving participation cutoff crosses a fixed cost node. It does not remove economically real kinks from portfolio corners, and it does not eliminate finite-difference sensitivity. The remaining sensitivity is therefore reported rather than smoothed away.

IA.C.5 Criterion-set envelopes for the economic objects

The local Jacobian can identify a weak direction without telling us whether that direction matters economically. The criterion-set exercise therefore reverses the question. Instead of asking how parameters move, it directly searches for the smallest and largest values of the paper's economic objects among parameter vectors that remain close to the minimum.

Conditional on the benchmark external calibration, define

$$\mathcal{C}(\Delta) = \{\vartheta : L_8(\vartheta) \leq L_{8,\min} + \Delta,\ |\beta_R(\vartheta)| \leq 0.056\}\,. \tag{IA.67}$$

For each value of Δ, the search separately targets the minimum and maximum of the baseline displacement burden, the aggregate-market completion share, and the participant and nonparticipant full-state measurability ceilings. Candidate trajectories from all these searches are then pooled so that the reported envelopes are nested as Δ increases.

The resulting ranges are *attained inner approximations*. Every reported endpoint is an actual feasible parameter vector found by the search. That makes the endpoints constructive—the reported economic values can occur inside the stated criterion set—but it does not prove that no more extreme feasible vector exists. The exercise is therefore stronger than evaluating a small collection of stored calibrations and weaker than an analytical outer bound.

The tight set gives the most relevant read. At $\Delta = 0.5$, the burden ranges from 50.87 to 100.73 annual basis points—almost a factor of two—while the nonparticipant full-state ceiling remains

between 3.46% and 4.33% and the raw completion share between 1.94% and 3.24%. Thus the same near-minimum region that permits substantial uncertainty about the welfare *level* does not contain an economy in which aggregate finance reaches a large fraction of the worker-specific displacement object.

The conclusion remains similar when the criterion set is widened substantially. At $\Delta = 2$, the largest full-state ceiling is 5.16% and the largest completion share is 4.86%. These are no longer especially tight fits to the minimum, but they still do not turn aggregate securities into a first-order insurance mechanism for displacement.

The $\Delta = 6$ row is intentionally much looser and should not be used as an inferential range. Here the burden level spans more than an order of magnitude and the completion statistic can be slightly negative. A negative Ω_s does not mean that an aggregate security directly creates worker-identity risk. It means that, after equilibrium portfolio and marginal-utility reweighting, the menu-indexed measured burden under A_s can slightly exceed the burden under A_0. The row is useful mainly for showing where sign stability eventually gives way as one permits large departures from the minimum.

Several checks verify that the displayed endpoints are genuine solved economies rather than optimizer artifacts. The benchmark welfare objects reproduce to numerical precision. No pooled candidate inside the parameter box and magnitude fence improves on $L_{8,\min}$. The envelope collapses back to the benchmark as Δ approaches zero. Forty-four pass-1 optima were re-solved cold and reproduce their criterion values. No reported candidate fails Arrow-market clearing; the worst residual is approximately 4.6×10^{-10}, against a 10^{-8} tolerance. Of the forty min/max endpoints in the full table, thirty-eight have neither an individual standardized target residual exceeding two in absolute value nor a fitted parameter on its box boundary; the two exceptions occur only at $\Delta = 6$.

Equilibrium multiplicity is also tracked rather than hidden. Of the 159 pooled candidates at the widest slack, 133 have a unique stable capital-market crossing. Six reported endpoints have two stable crossings; in those cases the continuation rule selects the branch connected to the benchmark, with selected capital stocks in the neighborhood of 0.636 to 0.665. The envelope therefore compares like equilibrium branches rather than mixing branch switches with parameter variation.

IA.C.6 What the diagnostics establish—and what they do not

The identification evidence supports a deliberately asymmetric conclusion. The calibration does not tightly pin down every primitive or every welfare level. The participation-cost location is locally weak, the return-loading criterion continues beyond the event-study fence, and even a tight near-minimum criterion set permits substantial movement in $\mathcal{D}(A_0)$. Reporting one parameter vector without these diagnostics would therefore overstate precision.

At the same time, the economic object central to the finance claim is more stable than the level. Along the return-loading ridge and throughout the tight near-minimum criterion set, both the full-state measurability ceiling and the equilibrium effect of completing aggregate markets remain in the low single digits. This does not make the financial boundary "point identified" in a formal population sense. It establishes the narrower fact the paper needs: the numerical weak directions found by the current calibration do not contain a near-minimum economy in which aggregate financial markets insure a large share of displacement risk.

No repeated-sample confidence interval for the welfare objects is reported. The legacy bootstrap in the project was built around a superseded soft-penalty estimator, did not recompute the current asset-menu hierarchy, and had a high cap rate. Treating that exercise as uncertainty for $\mathcal{D}(A_0)$, $\mathcal{D}(A_s)$, Ω_s, or the measurability ceilings would therefore be incorrect. A valid sampling exercise

would have to respect the source structure of the target moments, re-estimate the current hard-constrained model, and recompute the complete welfare chain in every successful replication.

The criterion-set envelopes answer a different question and are useful on their own terms: how much can the economic conclusions move while the calibrated fit is allowed to deteriorate by a stated amount? External calibration uncertainty—including risk aversion, the capability variance, the outside option, and the horizon—is separate again. Section 6 of the paper and IA.F report those exercises rather than combining fundamentally different sources of uncertainty into one synthetic interval.

IA.D Fit, Validation, Persistence, and Numerical Checks

IA.C asks whether parameter vectors that fit nearly as well can change the paper's economic conclusion. The present section asks a different question: are the benchmark moments and welfare objects constructed as intended, and are the reported results numerically reliable? We first document the eight target moments, then examine an untargeted labor-market outcome, the horizon and persistence convention, the use of exact rather than quadratic welfare, and the principal numerical checks.

The organizing principle is simple. A fit statistic is informative only when the reader knows what each moment measures; a validation exercise is useful only when it was not used to choose the fitted parameters; and a numerical comparison is useful only when it tests an approximation that could matter for the result. The checks below are therefore reported together with their limitations rather than as a generic list of robustness exercises.

IA.D.1 Construction of the eight target moments

The calibration uses eight moments to discipline six fitted parameters, conditional on the external inputs in IA.A. The moments deliberately cover separate parts of the model: participation and ownership, the technology and labor response, relative wages, and the return process. This breadth is useful because a model that fits only asset returns could still make displacement too rare or too inexpensive to be relevant for the paper's welfare question.

Let K^* denote equilibrium installed AI capital and let $\kappa^*(w; K^*)$ be the largest participation cost at which a worker with wealth w enters the equity market. With F_κ and G_w denoting the participation-cost and wealth distributions, define

$$p(w) = F_\kappa(\kappa^*(w; K^*)) \,. \tag{IA.68}$$

The first target, the equity-market participation rate, is

$$m_1 = \int p(w)\, dG_w(w). \tag{IA.69}$$

The second target is the participation gradient across wealth. If w_{50} is the median of the model wealth distribution,

$$m_2 = \mathbb{E}[p(w) \mid w > w_{50}] - \mathbb{E}[p(w) \mid w \le w_{50}]. \tag{IA.70}$$

These moments discipline the level and dispersion of fixed participation costs. The participation target is motivated by the U.S. household-finance evidence on limited stock-market participation; see, for example, Vissing-Jørgensen (2002). The model's smooth threshold implementation is used for both moments, so movement of a participation cutoff is not converted into a jump merely because a fixed quadrature node happens to cross it.

The third target measures ownership concentration in the financing of installed AI capital. Capitalist demand is $\mu_C W \alpha_C^*$, so the model statistic is

$$m_3 = \frac{\mu_C W \alpha_C^*(K^*)}{K^*}. \tag{IA.71}$$

The target is a deliberately coarse representation of concentrated U.S. equity ownership. Its assigned width therefore includes uncertainty in mapping observed household balance sheets into the paper's two-group ownership block; it is not treated as a sampling standard error for a directly observed model primitive.

The fourth and fifth targets discipline the physical technology-to-labor mapping. With π_s denoting the calibrated probabilities of the eighteen aggregate states,

$$m_4 = \mathbb{E}[z] = \sum_s \pi_s z(K^*, \theta_s), \tag{IA.72}$$

$$m_5 = \mathbb{E}[\lambda] = \sum_s \pi_s \lambda(s; K^*). \tag{IA.73}$$

The automation target summarizes the task share that the calibration asks the model to place on the AI technology; its uncertainty scale includes the translation from task-level evidence into the model's single automation share. The displacement target is a three-year worker-separation probability. Its construction synthesizes the automation and exposure evidence represented in Acemoglu and Restrepo (2020), Webb (2020), Felten et al. (2021), Eloundou et al. (2023), and Autor (2024). The three-year model period also matches the modern Displaced Worker Survey recall window discussed by Farber (2017).

The sixth target is the AI-to-traditional wage ratio. The implementation constructs the retained AI wage and the traditional-sector wage at the solved state allocation and applies the calibration's state weights to form the reported ratio $m_6 = w_A/w_T$. The target is intentionally broad: its width allows for the fact that observed sector wage comparisons need not coincide exactly with the model's retained-worker wage concept. For this reason the 0.150 width is a mapping scale rather than a literal standard error on one published regression coefficient.

The last two targets discipline the payoff available through AI-sector equity. They are the annualized adopter-equity premium and the adopter-to-market volatility ratio. The adopter portfolio is used because the model's equity claim is the payoff to installed capital in the adopting sector. The equity premium is measured imprecisely in the short return sample and therefore receives a much wider sampling scale than the relative-volatility target. Both targets use the same adopter-vintage convention used in estimation; neither is silently replaced by a producer or market-wide return series.

Table IA.8 reports the complete target vector. We label the third column an *uncertainty scale*, rather than a standard error, because the first, second, fifth, and two return scales are sampling-based, whereas the financing-share, automation, and wage-ratio scales also include structural uncertainty in mapping outside evidence into the model.

IA.D.2 What the benchmark fits—and where it does not

The table is more informative when read as a pattern rather than as a single criterion value. Seven of the eight moments are within 0.3 assigned uncertainty units of their targets. The exception is expected automation: the model value 0.3617 is 1.40 units above the target 0.2500. Its squared deviation is about 1.96, or roughly 86% of the total criterion

$$L_8 = 2.288.$$

The remaining fit tension is therefore concentrated rather than evenly spread across the eight moments.

That concentration is economically meaningful. A lower automation share would improve the largest individual residual, but automation also affects installed capital, displacement, wages, and equity payoffs. The fitted model chooses a higher automation share while matching those other margins closely. The paper therefore does not describe the benchmark as matching every moment nearly exactly. It reports the automation miss visibly and treats the scale of capability risk and the return loading as separate sensitivity objects.

Table IA.8: Construction and fit of the eight targeted moments

Moment	Target	Scale	Model	Std. dev.	Scale type
Equity participation rate	0.5200	0.0300	0.5286	+0.29	Sampling
Participation gradient in wealth	0.3500	0.0800	0.3531	+0.04	Sampling
Capitalist financing share	0.1000	0.0600	0.1172	+0.29	Mapping
Expected automation, $\mathbb{E}[z]$	0.2500	0.0800	0.3617	+1.40	Mapping
Expected displacement, $\mathbb{E}[\lambda]$	0.0700	0.0200	0.0667	−0.16	Sampling / synthesis
AI/traditional wage ratio	1.3000	0.1500	1.2955	−0.03	Mapping
AI equity premium, annual	0.0622	0.0694	0.0820	+0.29	Sampling
Relative equity volatility	1.1434	0.0993	1.1680	+0.25	Sampling
			$L_8 = \sum_j z_j^2$	2.288	

The standardized deviation is (model − target)/scale. The labels in the last column describe how the scale is used in the calibration; they are not claims that the sampling-based moments share a common data set or sampling covariance matrix. The criterion therefore measures standardized fit but is not treated as an efficient-GMM overidentification statistic. The return moments use the adopter-portfolio convention employed in estimation.

The table also should not be read as a formal specification test. Eight targets and six fitted parameters provide two additional restrictions in an accounting sense, but the targets are assembled from several sources and three uncertainty scales explicitly include model-to-data translation uncertainty. Without a joint sampling covariance matrix, a chi-squared interpretation of the minimized criterion would give the number more statistical content than the construction supports. IA.C accordingly uses changes in L_8 as calibrated-distance measures, not as confidence-set critical values.

IA.D.3 Untargeted validation: the severity of displacement

The most useful validation exercise asks the model to reproduce an outcome that matters for welfare but was not placed in the fitting criterion. The retained versus displaced consumption gap is such an outcome: a model could fit the asset-market moments while making job loss almost costless, in which case the paper's insurance question would be mechanically small.

The run-of-record economy implies an earnings loss conditional on displacement of

$$\text{earnings loss} = 0.3292. \tag{IA.74}$$

This number is not a ninth target and does not enter L_8. It lies inside the roughly 30–40% range documented for displaced workers in administrative and survey evidence. In particular, Couch and Placzek (2010) report large initial losses in administrative data, while Jacobson et al. (1993) document that substantial losses persist well beyond the displacement event. The comparison therefore provides an opportunity for the labor-income block to fail that is separate from fitting the displacement probability itself.

The validation is deliberately narrower than a validation of the full earnings *path*. The only archived multi-horizon path in the project was generated at a superseded capital stock. It is retained in the project history to show what was computed, but no numerical value from that stale path is used in the paper or in this Internet Appendix. More fundamentally, the maintained model has two dates. It can discipline the size of the loss over its three-year reference period; it cannot reproduce the slow many-year recovery dynamics in the displacement literature. Thus the 32.9% comparison supports the *severity* of the modeled job-loss state, not the model's dynamic recovery process and not the absolute welfare burden.

IA.D.4 Horizon and persistence

Persistence is economically important because a job-loss state that lasts for many years is more costly than one that lasts for a single reporting period. It is also an area in which it would be easy to create false precision. The paper therefore fixes one horizon convention for the headline calculations and keeps the older continuation exercise as an explicitly uncalibrated diagnostic.

One model period is three years. The benchmark first computes expected utility and certainty equivalents over that model period. A model-period consumption change x is converted to an annual rate only after the welfare comparison, using

$$\mathcal{A}_3(x) = (1+x)^{1/3} - 1. \tag{IA.75}$$

Thus "one-period" in the benchmark means one three-year model period; it does not mean one calendar year and it does not mean that a three-year welfare number is divided mechanically by three.

The archived persistence ladder applies a common continuation horizon to the entire displacement comparison. This convention matters: persistence cannot be attached only to whichever residual component one wishes to make large while leaving the other menu increments at the one-period horizon. The existing audit classifies the calculation as the common-horizon case: every increment is recomputed under the same persistence assumption.

No node of that continuation ladder is promoted as a calibrated persistent horizon. The stored scale values are scenario parameters, not estimates from the earnings-recovery data. Selecting one of them as "the" persistent case would therefore invent empirical discipline that the model does not have. Persistence and alternative horizons are consequently treated as specification sensitivity in Section 6 of the paper and IA.F, rather than folded into the benchmark or into a sampling interval.

IA.D.5 Exact welfare versus quadratic risk accounting

The welfare calculations use exact CRRA expected utility. This choice is not cosmetic. A displacement lottery puts substantial probability on two consumption outcomes that are far apart, precisely the situation in which a second-order approximation around mean consumption can be inaccurate.

For a worker type in state s, let $C_R(s)$ and $C_D(s)$ be consumption if the worker is retained or displaced. The exact conditional certainty equivalent of the employment lottery is

$$C^e(s) = u^{-1}([1-\lambda(s)]u(C_R(s)) + \lambda(s)u(C_D(s)))\,. \tag{IA.76}$$

The corresponding quadratic approximation expands expected utility around mean consumption and produces the familiar variance correction. The two methods are compared in Table IA.9 using the archived benchmark within-state lottery diagnostic.

The approximation understates this diagnostic by roughly one-half or more. We therefore use the quadratic representation only where additivity is the object of interest—for example, descriptive variance accounting. Every promoted welfare comparison, including $\mathcal{D}(A_0)$ and $\mathcal{D}(A_s)$, is computed by evaluating exact expected utility under the relevant allocation and then converting the resulting certainty-equivalent difference to reporting units. The paper never substitutes the quadratic values in Table IA.9 for an exact market-completion welfare calculation.

IA.D.6 Numerical accuracy and internal consistency

The remaining checks ask whether discretization, root selection, or numerical approximations could be creating the reported patterns. No single check proves that the entire computation is exact, so

Table IA.9: Exact and quadratic valuation of the benchmark employment lottery

Group	Exact (annual bp)	Quadratic	Exact / quadratic
Nonparticipants	63.10	31.08	2.03
Participants	44.00	16.94	2.60

This is an auxiliary group-level *within-state employment-lottery* certainty-equivalent diagnostic. It is not the menu-indexed social burden $\mathcal{D}(A)$ defined in Appendix B of the paper. In particular, it should not be compared numerically with the benchmark social value $\mathcal{D}(A_0) = 84.33$ annual basis points as though the objects were two estimators of the same quantity. The diagnostic is retained only to assess the accuracy of the quadratic approximation for the worker-specific lottery.

Table IA.10 reports the specific failure mode each diagnostic is designed to detect.

Two numerical conventions deserve emphasis. First, the capital-market root is selected by stability and continuation, not by mechanically taking the lowest root. This prevents a newly appearing low-capital branch from being reported as an economic response to a nearby parameter change. Second, the facts registry keeps the reported number, source vintage, units, annualization rule, reporting group, and welfare convention together. The purpose is mundane but important: a numerically accurate calculation can still be wrong in the paper if a value from a superseded model vintage or a different welfare normalization is inserted into the current table.

IA.D.7 What these checks establish

Taken together, the checks support four conclusions. First, the benchmark matches seven targeted moments closely while leaving a visible and economically meaningful miss on automation. Second, the model generates a plausible one-horizon earnings loss without having been fitted to that outcome. Third, exact utility matters for the displacement lottery; the quadratic approximation is too inaccurate to carry the paper's welfare claims. Fourth, the tested quadrature, continuation, clearing, and derivative choices do not generate the headline benchmark mechanically.

They do not establish more than that. The two-date model does not validate the multi-year recovery path after job loss. The persistence ladder is not an estimated duration model. The numerical checks do not create sampling confidence intervals, and grid agreement cannot rule out every omitted state or model specification. Those limitations are kept separate from numerical accuracy because they require economic or empirical information, not tighter solver tolerances.

Table IA.10: Numerical audit of the run-of-record calculations

Check	Recorded comparison	What the check establishes
Participation integration	Analytic cost-threshold integration; benchmark wealth grid $n_q = 30$	Removes artificial jumps from fixed participation-cost nodes. IA.C repeats the local derivative rankings over $n_q \in \{36, 40, 44, 48\}$; the estimator-sensitivity ranking is unchanged for all six fitted parameters.
Capability-state grid	Baseline $n_\theta = 9$; wage-mapping check also uses 15 and 21 nodes	The primary upper-endpoint inversion in IA.B changes by only about 0.004 in ν_θ, so the historical wage-regime mapping is not materially generated by the nine-node capability grid.
Retention derivative	Analytic $\partial\lambda/\partial\theta = 0.3351$; central finite difference with step 10^{-6} gives 0.3351	Verifies the derivative used in the labor-adjustment proposition against the state-by-state numerical retention solver at the benchmark.
Benchmark continuation checkpoint	Independent zero-ν_θ continuation solve changes K^* by 1.18×10^{-10}; promoted objects agree to better than 10^{-6}	Confirms that the wage-regime machinery reproduces the run-of-record economy rather than starting from a numerically different benchmark.
Aggregate-claim clearing	Benchmark maximum residuals: 4.23×10^{-12} for A_θ and 2.10×10^{-12} for A_s	The new claims clear to well below the tolerances used for reporting the asset-menu experiments.
Balanced-transfer identity	$\max_s \lvert\mathbb{E}[d \mid s]\rvert = 2.64 \times 10^{-16}$ at the benchmark	Numerically verifies the state-by-state resource-balance identity used to define the displacement transfer.
Consumption floor	Numerical floor $c_{\min} = 10^{-15}$; slack at accepted benchmark allocations	The reported benchmark utility and marginal-utility objects are not created by truncating consumption at the numerical safeguard.

The table summarizes existing audit outputs; it does not report new numerical work. The capital-market solver scans for stable crossings and uses continuation when a parameter is varied. IA.C separately documents branch multiplicity in the criterion-envelope searches and the finite-difference sensitivity of the local Jacobian.

IA.E Wage-Regime Diagnostics

Section 4 of the paper varies the wage–automation elasticity, ν_θ, to ask how labor-market adjustment changes the displacement risk that financial markets are asked to insure. This appendix documents that exercise and two checks that are important for interpreting it. First, we distinguish the paper's institutional comparative static, which holds the benchmark primitives fixed, from an alternative exercise that re-estimates the six fitted parameters at every value of ν_θ. Second, we trace the comparison only over the part of the model in which retained employment is individually rational. The historically informed U.S. range, $\nu_\theta \in [0, 0.13]$, is well inside that domain.

The main result is simple. Changing the wage regime has a large effect on the *amount* of displacement risk. It has much less effect on the fraction of the remaining welfare-relevant displacement loss that aggregate financial markets can reach. The additional results below show that this conclusion is not driven by re-estimation, capital deepening, the outside-option calibration, or the normalization of the completion statistic.

IA.E.1 Two distinct wage-regime exercises

The fixed-calibration family is the comparative static used in the paper. All technology, preference, and fitted primitives are held at their benchmark values; only ν_θ changes, and equilibrium capital K^* is re-solved at each node. The rows therefore represent different economies with different wage-setting regimes. Their distance from the eight U.S. target moments is informative about how unlike the benchmark U.S. calibration they become, but it is not a criterion for excluding the counterfactual.

The re-estimated family asks a different question. At each value of ν_θ, the six fitted parameters are re-estimated against the same eight U.S. moments. This exercise shows how much of the direct wage-regime effect can be absorbed by changing other primitives while continuing to match the U.S. data. It is therefore a useful identification diagnostic, but not a clean institutional comparative static.

Table IA.11 places the two exercises side by side. Within the U.S.-mapped range, the fixed-calibration burden falls monotonically from 84.33 to 29.76 annual basis points. The re-estimated burden instead moves nonmonotonically between 71.48 and 106.79 basis points as other parameters adjust. The distinction is equally visible in the equity–capability loading: re-estimation keeps β_R near the event-evidence region, while varying the wage regime alone moves it sharply. Re-estimation is therefore absorbing an economically important part of the wage-regime change rather than measuring the same comparative static more flexibly.

One implication is worth making explicit. In the re-estimated family the full-state measurability ceiling for the marginal-utility-weighted displacement loss remains between 3.55 and 4.22 percent for nonparticipants and between 3.50 and 4.31 percent for participants over $[0, 0.13]$. Thus re-estimation can absorb large movements in the burden level without producing an economy in which a large share of the displacement loss becomes aggregate-state measurable. This is the same distinction between the level of the problem and the financial boundary emphasized in IA.C.

IA.E.2 Asset-menu ladders

The menu calculations provide a second way to see what re-estimation absorbs. For each economy, A_0 is the benchmark opportunity set, A_{access} removes the equity-access barrier, A_θ adds the capability claim, and A_s completes claims on the full aggregate state. Table IA.12 reports the social burden at each rung. The central paper uses the nested comparison from A_0 to A_s; the intermediate rungs are reported here to make the accounting transparent.

Table IA.11: Fixed-calibration and re-estimated wage-regime families

		Re-estimated			Fixed calibration		
ν_θ	Region	$\mathcal{D}(A_0)$ (bp)	Ω_s^{self} (%)	β_R	$\mathcal{D}(A_0)$ (bp)	Ω_s^{self} (%)	β_R
0.00	U.S.	84.33	2.20	−0.056	84.33	2.20	−0.056
0.05	U.S.	106.79	2.94	−0.024	60.01	2.20	+0.067
0.10	U.S.	71.48	1.85	+0.029	39.85	1.99	+0.140
0.13	U.S.	90.97	2.58	+0.056	29.76	1.98	+0.177
0.20	diagnostic	54.21	1.35	+0.048	12.86	0.57	+0.252
0.35	diagnostic	23.86	−1.80	+0.041	0.61	2.78	+0.390

The U.S.-mapped range is $[0, 0.13]$. The fixed-calibration family holds all benchmark primitives fixed except ν_θ and re-solves K^*. The re-estimated family re-fits the six internally disciplined parameters at each node. Ω_s^{self} is the self-normalized completion statistic defined in Appendix B of the paper. The 0.20 and 0.35 rows are diagnostics rather than empirically mapped U.S. regimes; 0.35 is also just outside the individual-rationality domain of the fixed-calibration model, as documented below.

The intermediate rungs also clarify why access to equity is not itself "displacement insurance." In the re-estimated family, opening access raises the measured burden at every reported node. At the benchmark this reflects premium capture and the resulting marginal-utility reweighting, not a direct increase in worker-identity risk. In the fixed-calibration family the access rung becomes small and eventually changes sign. These details matter for the menu accounting but not for the central result that even the full A_s menu changes only a small share of the displacement burden over the U.S.-mapped range.

IA.E.3 Span ceilings across wage regimes

Table IA.13 reports the corresponding projection bounds. Three columns answer different questions. The full-state ceiling $\eta_s^2(u'd)$ asks how much of the welfare-relevant displacement loss is measurable with respect to the entire aggregate state. The equity $R^2(u'd)$ asks how much is linearly projected onto the existing AI-equity payoff. The θ-only total-consumption ceiling instead asks how much of ordinary total consumption variation can be associated with capability. The last object can rise even while the first two fall, because a more responsive wage routes more capability risk into common wages while simultaneously shrinking the retained-versus-displaced lottery.

Within $[0, 0.13]$, the fixed-calibration nonparticipant full-state ceiling falls from 4.04 to 2.99 percent while the θ-only ceiling for total consumption rises from 5.84 to 10.36 percent. The distinction is useful: greater wage responsiveness makes capability-linked aggregate payoffs more relevant for the worker's paycheck, but it does not turn the remaining worker-specific employment lottery into aggregate risk. Re-estimation partly offsets the fall in the displacement ceiling, yet that ceiling still remains in a narrow 3.55–4.22 percent range for nonparticipants over the mapped U.S. region.

IA.E.4 Capital deepening and the outside option

Because equilibrium capital rises with wage responsiveness in the fixed-calibration family, the main comparative static combines a direct wage-regime effect with endogenous capital deepening. To measure the size of this feedback, we repeat selected nodes with capital pinned at its benchmark value, $K = 0.635449674$. These cells are a partial-equilibrium decomposition, not alternative

Table IA.12: Asset-menu ladders across wage regimes

ν_θ	$\mathcal{D}(A_0)$	$\mathcal{D}(A_{\text{access}})$	$\mathcal{D}(A_\theta)$	$\mathcal{D}(A_s)$	$\mathcal{D}(A_0) - \mathcal{D}(A_s)$	Ω_s^{self} (%)
			annual basis points			
Panel A. Fixed-calibration family						
0.00	84.326	85.194	81.335	82.165	2.161	2.205
0.05	60.007	60.263	57.114	58.098	1.909	2.203
0.10	39.847	39.809	37.532	38.439	1.408	1.986
0.13	29.764	29.661	27.838	28.618	1.146	1.980
0.20	12.864	12.841	11.929	12.453	0.411	0.574
0.35	0.606	0.603	0.526	0.566	0.041	2.781
Panel B. Re-estimated family						
0.00	84.326	85.194	81.335	82.165	2.161	2.205
0.05	106.789	108.377	102.446	103.361	3.428	2.941
0.10	71.485	72.469	69.025	69.924	1.561	1.855
0.13	90.972	92.802	87.387	88.354	2.618	2.578
0.20	54.209	56.623	52.219	53.331	0.878	1.354
0.35	23.862	25.442	23.591	24.236	−0.374	−1.802

All burden measures use the common conventions in Appendix B of the paper. The intermediate access and capability rungs are diagnostic decompositions, not additive welfare components. In the re-estimated 0.35 row, aggregate completion raises the measured burden slightly. The sign comes from the intermediate rungs: the negative access and final Arrow rungs together exceed the positive capability-claim rung. The row is outside the U.S.-mapped range and is not used for the paper's central financial conclusion.

equilibria: away from $\nu_\theta = 0$ the capital market does not clear.

At $\nu_\theta = 0.13$, pinning capital leaves $E[\lambda] = 3.813\%$, $\mathcal{I} = 0.007089$, and $\mathcal{D}(A_0) = 30.61$ basis points, compared with 3.942 percent, 0.006985, and 29.76 basis points when capital is re-solved. In terms of movement away from the benchmark, 105 percent of the fall in expected displacement, 99 percent of the fall in $\mathcal{I}$, and 98 percent of the fall in $\mathcal{D}(A_0)$ survive the capital pin. Values slightly above one mean that capital deepening modestly offsets rather than amplifies the direct wage response. By contrast, only about 78 percent of the increase in the nonparticipant θ-only total-consumption ceiling survives. Capital therefore matters for the rerouting of aggregate consumption risk, but is not the main source of the decline in displacement risk.

The outside option separates another margin cleanly. The replacement-rate, home-value, severance, tax, and reemployment-wage variations used in the outside-option grid enter displaced income in the welfare layer; they do not enter the firm's retention condition or the equilibrium capital-clearing condition. Direct evaluation confirms that expected displacement is bit-identical across those variations even though $\mathcal{D}(A_0)$ moves substantially. Thus the wage regime governs the incidence and amount of displacement, while the safety net governs the cost conditional on displacement. This is why the outside-option exercises are reported as burden-level sensitivity rather than as alternative wage-regime calibrations.

IA.E.5 The economically interpretable domain

The high-ν_θ continuation has two distinct boundaries. The first is economic. As ν_θ rises, the retained wage falls in high-capability states while the displaced outside option does not. Just above $\nu_\theta = 0.34$, a state with positive displacement first satisfies $\bar{w}_A < b$. A retained worker would then prefer the displaced outside option, but the model has no quit margin. The allocation is therefore no longer an individual-rationality description of employment choice. The fine grid places the first crossing in $(0.340, 0.345]$, and the paper uses approximately 0.34 as the upper bound of the

Table IA.13: Projection ceilings across wage regimes

ν_θ	$\eta_s^2(u'd)$ P	$\eta_s^2(u'd)$ NP	Equity $R^2(u'd)$ P	Equity $R^2(u'd)$ NP	θ-only total cons. NP
			percent		
Panel A. Fixed-calibration family					
0.00	4.105	4.042	1.304	0.572	5.840
0.05	3.693	3.966	0.970	0.443	7.286
0.10	3.074	3.360	0.648	0.291	9.079
0.13	2.785	2.986	0.480	0.219	10.357
0.20	1.951	2.241	0.301	0.176	13.512
0.35	0.397	0.472	0.022	0.019	25.224
Panel B. Re-estimated family					
0.00	4.105	4.042	1.304	0.572	5.840
0.05	4.312	4.222	1.576	0.755	8.718
0.10	3.496	3.547	1.253	0.587	9.903
0.13	3.922	3.993	1.507	0.822	11.657
0.20	3.751	3.832	1.187	0.912	10.856
0.35	2.607	2.690	0.783	0.717	13.574

P and NP denote participants and nonparticipants. The θ-only column is the linear projection ceiling for nonparticipant total consumption. It is not a ceiling for the displacement transfer. The fixed-calibration 0.35 row is just outside the model's economically interpretable domain and is shown only to indicate the direction of the continuation; it is not used as an institutional benchmark.

economically interpreted comparative static.

The second boundary is mechanical and comes later. Expected displacement falls below 0.1 percent at $\nu_\theta \simeq 0.44881$ when the benchmark separation-cost parameter is used. This is a near-extinction threshold, not the individual-rationality boundary. The model reaches exact zero displacement only still farther out; a $\nu_\theta = 0.70$ calculation is retained solely as a numerical diagnostic of that limiting case.

Two numerical checks confirm that the mechanical disappearance is not a coarse state-grid artifact. Refining the capability quadrature from 9 to 15 and 21 nodes leaves $E[\lambda]$ at $\nu_\theta = 0.50$ in the range 2.7–3.1 $\times 10^{-4}$ and leaves the 0.70 value between zero and 2.8×10^{-9}. Changing the separation-cost parameter from $\psi = 1$ to $\psi = 2$ moves the 0.1-percent extinction threshold from 0.44881 to 0.41494, whereas the outside-option variations do not move it at all. These checks support the qualitative conclusion that displacement disappears well before a polar value such as $\nu_\theta = 1$, while also reinforcing that the relevant model domain ends earlier for individual-rationality reasons.

IA.E.6 Why the raw completion share rises

Cross-regime comparisons require care with normalization. The raw statistic uses the fixed A_0 certainty-equivalent numeraire embedded in $\mathcal{D}(A)$,

$$\Omega_s = 100\frac{\mathcal{D}(A_0) - \mathcal{D}(A_s)}{\mathcal{D}(A_0)}.$$

The self-normalized companion instead normalizes the lottery-removal comparison by each menu's own certainty equivalent before annualization; Appendix B of the paper gives the exact construction.

Across the U.S.-mapped range, the raw statistic rises from 2.562 to 3.850 percent. Read alone, that would suggest that wage responsiveness raises the share of displacement risk reachable by aggregate finance by roughly one-half. The self-normalized statistic instead moves from 2.205 to 1.980 percent, while the absolute burden reduction falls from 2.161 to 1.146 annual basis points.

Table IA.14: Domain diagnostics for the fixed-calibration continuation

ν_θ	$E[\lambda]$ (%)	Live states	Lottery mass with $\bar{w}_A < b$ (%)	Share of $\mathcal{I}$ there (%)	Mean $\bar{w}_A - b$
0.13	3.942	10	0.000	0.000	+0.436
0.20	2.560	9	0.000	0.000	+0.343
0.34	0.678	7	0.000	0.000	+0.154
0.35	0.602	7	0.038	0.000	+0.140
0.40	0.264	6	6.028	0.158	+0.067
0.42	0.139	5	10.466	13.583	+0.023
0.44	0.112	5	11.652	84.851	−0.003
0.45	0.098	5	97.868	96.351	−0.016

A live state has positive displacement. The two share columns weight states by the surviving employment lottery; $\mathcal{I} = E[\lambda(1-\lambda)(\bar{w}_A - b)^2]$. The first negative gap occurs immediately above 0.34, so the 0.35 row is already outside the formal domain even though its violating state has negligible weight. The 0.45 row is also near the separate mechanical extinction threshold $\nu_\theta \simeq 0.44881$.

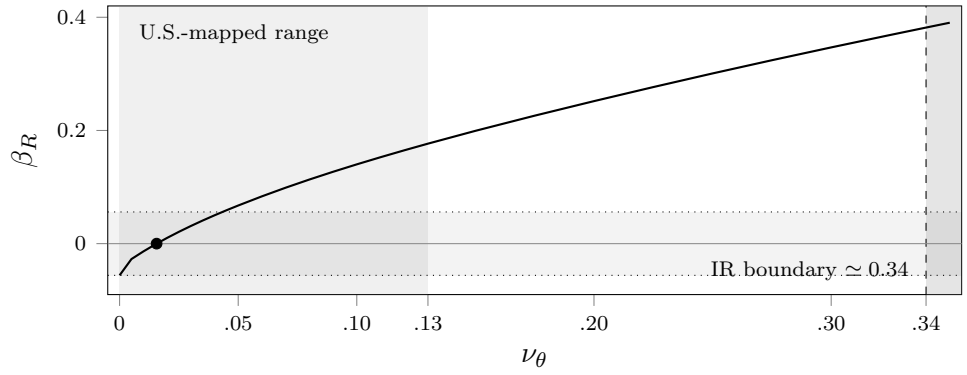


Figure IA.1: Wage responsiveness and the equity–capability loading. The solid line is the derived equilibrium β_R along the fixed-calibration family. The lightly shaded interval $[0, 0.13]$ is the historically informed U.S. range; the horizontal band marks $|\beta_R| \leq 0.056$ for reference. The vertical line at approximately 0.34 is the individual-rationality boundary. The zero crossing at $\nu_\theta \simeq 0.0156$ is a model implication, not a separately estimated cross-country fact.

The apparent rise in the raw share is therefore a normalization effect, not an increase in financial reach. This is why the paper reports the absolute basis- point difference and Ω_s^{self} whenever wage regimes are compared.

IA.E.7 A secondary asset-pricing implication

The fixed-calibration exercise also generates a testable secondary prediction. The equity–capability loading is an equilibrium object, not a parameter held fixed in this experiment. As wages respond more strongly to capability, the capability shock is routed differently through the wage bill and profits, and β_R changes sign. The zero crossing occurs at approximately $\nu_\theta = 0.0156$, inside the U.S.-mapped range. Figure IA.1 shows this result only over the economically relevant part of the continuation. The event-study magnitude band is drawn for reference; leaving that band on the fixed-calibration family is a model prediction, not a violation of a restriction imposed on these counterfactual economies.

Finally, the structural parameter can be translated back into the observable wage-versus-

employment adjustment share used in IA.B. At $\nu_\theta = (0.05, 0.10, 0.13, 0.20)$, the primary ratio-of-aggregate-responses map is

$$\widehat{\nu}_{K\,\text{fixed}} = (0.103, 0.193, 0.242, 0.401), \qquad \widehat{\nu}_{K\,\text{free}} = (0.083, 0.207, 0.261, 0.431).$$

The difference is modest for the purpose of locating the historically informed range, but it is another reason not to interpret ν_θ itself as a literal fraction of adjustment occurring through wages. IA.B provides the empirical inversion and the source audit; no named foreign economy is assigned a value of ν_θ here.

IA.E.8 Summary

The diagnostics leave a narrower conclusion than the earlier "endogenous insurable share" interpretation, but a stronger one. Over the historically informed range, wage responsiveness substantially reduces displacement incidence, the within-state identity component, and the exact displacement burden. Yet the full-state measurability ceiling remains in the low single digits, and the self-normalized effect of completing aggregate markets remains about two percent. Re-estimating the other primitives changes the burden level substantially but does not produce a large aggregate-state share. Capital deepening explains little of the decline in displacement risk, and the outside-option calibration changes severity rather than incidence. Thus the wage regime primarily changes the size of the worker-specific lottery; it does not convert a large displacement lottery into one that aggregate financial markets can insure.

IA.F Additional Financial Experiments and Specification Sensitivity

The main paper uses the most demanding financial comparison—the baseline allocation versus complete trading in the modeled aggregate state—and leaves several useful diagnostics to this appendix. They address three natural questions. Does limited stock-market participation explain the result? Could a better-designed aggregate security or a different portfolio constraint change it? And how dependent are the quantitative magnitudes on external calibration and on the low-dimensional aggregate state?

The answer differs across these questions. Market access and security design matter for household welfare and for the prices and positions of aggregate claims, but they do not create a payoff conditional on which worker is displaced. External preference and shock calibrations can move the *level* of the welfare burden substantially. The most consequential remaining scope issue is instead the dimension of the aggregate state: omitted common shocks that predict displacement could increase the component reachable by financial markets. We keep these statements separate because they place different limits on the paper's conclusion.

IA.F.1 Equity access: selection is not insurance

The first diagnostic separates the correlation between participation and worker outcomes from the causal effect of opening market access. Participants have higher wealth and lower participation costs, so their lower displacement burden cannot by itself be interpreted as evidence that owning equity insures job loss. At the benchmark allocation the group-level burdens are 77.82 annual basis points for participants and 88.79 for nonparticipants. Wealth smooths consumption directly, whether or not it is invested in equity.

The treatment experiment holds installed capital fixed and gives previously excluded households access to the existing equity claim, allowing the equity price to clear. The population displacement burden moves from

$$\mathcal{D}(A_0) = 84.3262 \quad \text{to} \quad \mathcal{D}(A_{\text{access}}) = 85.1940 \quad \text{annual bp.}$$

Thus the access rung lowers the burden by -0.8678 basis points—that is, it raises the measured burden slightly. We do not attach economic significance to the sign of this small difference. The important result is that broadening access does not turn existing equity into meaningful displacement insurance.

This statement does not imply that access has no ordinary welfare value. A separate halt-condition calculation removes displacement exposure entirely. The displacement-insurance effect of opening equity access is then exactly 0.0000 basis points, while the ordinary certainty-equivalent value of access remains 33.06 basis points. The first number must vanish if the menu-indexed burden is isolating displacement risk; the second can remain positive because a new participant obtains access to an asset with a positive expected excess return. The distinction is premium capture versus displacement insurance.

The incidence is also not one-sided. New equity demand compresses the premium, and incumbent participants lose from the price change; their fixed-K certainty equivalent falls by about 51 basis points in the benchmark access experiment. Universal participation can therefore benefit entrants while redistributing away from incumbent asset holders. This is why neither the cross-sectional participant–nonparticipant gap nor the ordinary welfare gain from opening access is used as the paper's measure of displacement insurance.

For reference, the complete benchmark menu ladder is

$$84.3262 \xrightarrow{A_{\text{access}}} 85.1940 \xrightarrow{A_\theta} 81.3350 \xrightarrow{A_s} 82.1654 \quad \text{annual bp.} \tag{IA.77}$$

The intermediate increments are order dependent and are not interpreted as an additive decomposition into access, capability, and productivity "effects." The first-to-last comparison remains the financial counterfactual of interest.

IA.F.2 Portfolio constraints, security design, and reservation prices

A second concern is that existing equity may simply be the wrong instrument or that the household portfolio constraints prevent a useful hedge from being held. The diagnostics show that these are different issues for total consumption and for the welfare-relevant displacement object.

For lottery-averaged total consumption, equity has an enormous projection R^2: 99.62% for participants and 94.82% for nonparticipants. But the variance-minimizing equity positions are shorts, approximately -1.36 and -0.76 times investable wealth, respectively. Under the benchmark long-only restriction the attainable total-consumption variance reduction is therefore zero. This is a useful warning that a high projection R^2 need not describe an implementable hedge.

The result that matters for the paper is different. When the target is $u'(c)d$, the variance-minimizing equity position is *long* and the no-leverage constraint is slack by more than two orders of magnitude. The attainable and unconstrained spans consequently coincide:

$$R^2_{E,P}\big(u'd\big) = 1.304\%, \qquad R^2_{E,NP}\big(u'd\big) = 0.572\%. \tag{IA.78}$$

Thus the low displacement span is not produced by a short-sale or leverage constraint. Even when the correctly signed equity position is feasible, equity has little covariance with the welfare-relevant displacement loss.

A purpose-built capability claim provides a further check on security design. For nonparticipants, the linear θ payoff explains 5.84% of lottery-averaged total-consumption variance; allowing an arbitrary θ-measurable payoff raises that ceiling only to 6.06%. The linear claim is therefore already close to the full θ-measurable frontier for that diagnostic. In the archived price-position exercise, nonparticipants would put essentially their full investable portfolio in the capability claim at prices below about 198 annual basis points; the claim share ranges from about 0.91 to 1.00 before the budget ceases to bind. The participant budget does not bind over the same diagnostic grid, with the claim share between about 0.59 and 0.97. These positions demonstrate that a long capability hedge can be feasible. They do not turn the total-consumption projection into a welfare result, and the claim in this projection exercise is not an equilibrium traded security.

Reservation prices reinforce the same point. A reservation price is a marginal valuation, not a certainty-equivalent gain. Table IA.15 reports the ranges already generated by the project. The price moves substantially when the return-loading direction, the scale of the capability shock, or risk aversion changes. The displacement welfare conclusion should therefore not be organized around a particular claim price.

IA.F.3 External calibration: capability risk and risk aversion

The most complete external-calibration exercises are the grids for $\mathrm{Var}(\theta)$ and γ. In both cases the six fitted parameters are re-estimated at every grid point and the full A_s asset-menu experiment is re-solved. These are therefore genuine sensitivities of the current complete-market calculation, unlike the older outside-option and persistence ladders discussed below.

Table IA.15: Reservation-price diagnostics for the linear capability claim

Variation	Range/exercise	Reservation price (annual bp)
Broad stored β_R ridge	return-loading diagnostic	about 135–219
Inside event-study magnitude fence	fitted economies	201.20–218.99
Capability variance	$\mathrm{Var}(\theta) = 0.01$–$0.12$	118.80–341.93
Risk aversion	$\gamma = 2$–10	86.91–403.65

The first two rows vary the equilibrium return-loading direction while holding the benchmark external calibration fixed. The last two rows instead re-estimate the six fitted parameters at each value of the named external calibration. Reservation prices are conditional marginal valuations. They are not welfare gains, and ranges from different rows answer different sensitivity questions.

Table IA.16: External-calibration sensitivity of welfare and market completion

Exercise	L_8	$\mathcal{D}(A_0)$ (annual bp)	Ω_s (%)	Claim price (annual bp)
Benchmark	2.288	84.33	2.56	218.99
$\mathrm{Var}(\theta) = 0.01$–$0.12$	1.984–2.913	80.62–95.72	−2.10–9.40	118.80–341.93
$\gamma = 2$–10	1.570–3.670	18.75–269.05	−1.09–10.54	86.91–403.65

Each range is over the stored re-estimated grid for the named external calibration. At every grid point the complete eighteen-state A_s menu is re-solved. No targeted moment exceeds two assigned uncertainty units anywhere on either grid. A negative Ω_s means that complete aggregate-state trading slightly raises the menu-indexed displacement burden. The sign of a small completion effect is therefore not treated as a robust object.

The capability-variance exercise moves the burden much less than it moves the claim price. Across a twelve-fold increase in $\mathrm{Var}(\theta)$, $\mathcal{D}(A_0)$ remains between 80.62 and 95.72 basis points, whereas the claim price rises from about 119 to 342 basis points. The target moments do not select the benchmark variance: the criterion falls from 2.913 at 0.01 to 2.288 at the benchmark 0.04 and to 1.984 at 0.12. This is why IA.A treats the capability variance as an external anchor rather than an internally estimated primitive.

Risk aversion has a much larger effect on the welfare level. Raising γ from 2 to 10 moves the burden from 18.75 to 269.05 annual basis points. The benchmark $\gamma = 5$ is likewise not chosen by the eight target moments:

$$L_8(\gamma = 2) = 1.570 < L_8(\gamma = 5) = 2.288.$$

Accordingly, the paper does not claim that 84.33 basis points is a preference- free estimate of the cost of displacement.

The completion statistic is not point-stable on either wide grid. It can be slightly negative and can rise to about 9–11%. This weakens any claim that 2.56% is a precisely identified structural constant. It does not produce an economy in these grids in which completing the modeled aggregate market removes a large fraction of the displacement burden. The appropriate conclusion is therefore qualitative: external calibrations can move welfare levels sharply, while aggregate-state completion remains a secondary rather than first-order insurance channel over the stored exercises.

IA.F.4 Outside options and the horizon

The outside option and persistence change the severity of displacement rather than the information available to aggregate securities. They are nonetheless important because the paper reports welfare in annual consumption units.

The existing outside-option grid spans alternative replacement rates, nonemployment values, and reemployment wage penalties. Across that grid the baseline burden ranges from 16.55 to 279.23 annual basis points—more than an order of magnitude. This is strong evidence against interpreting the benchmark 84.33 basis points as a model-free estimate of the welfare cost of job loss. The exercise predates the current complete A_s calculation, however, and does not re-solve the Arrow menu at every node. We therefore use it only to document sensitivity of the *level*; we do not report its older residual- share statistic as robustness evidence for the current Ω_s.

The same rule applies to persistence. The benchmark is one three-year model period, with annualization applied after the welfare comparison. Older continuation exercises show that displacement welfare costs can increase more than proportionally with duration and that the value of equity access can change sign at alternative fitted return loadings. At one admissible economy, a longer-horizon calculation produces a positive access increment of about 22.3 basis points. That observation is useful because it rules out an overly strong statement such as "market access can never help." But the full A_s menu was not recomputed along that persistence ladder, so 22.3 basis points is not a persistence-adjusted estimate of the paper's complete-market welfare effect.

The horizon exercise also explains why all pieces must be put on a common time basis. Persisting only the worker-specific residual while leaving the financial increment at the benchmark horizon would mechanically make the residual share larger. Every promoted comparison in the paper therefore uses one common horizon and the exact annualization rule in Appendix B of the paper.

IA.F.5 State-space scope: what an omitted aggregate shock would have to do

The most important specification limitation is not a portfolio constraint or a parameter value. It is the dimension of the aggregate state. The benchmark asks how much of $u'(c)d$ is measurable with respect to $s = (\varepsilon, \theta)$. The approximately four-percent ceiling is therefore a ceiling relative to the modeled state, not a theorem about all potentially observable common shocks in the economy.

An omitted state variable could overturn the quantitative conclusion if it adds substantial common variation in the welfare-relevant displacement loss. For example, sector demand, occupation-level exposure, credit conditions, or other common shocks could predict which employment states are especially painful or could make displacement probabilities move systematically in dimensions not captured by (ε, θ). If those variables are observable and contractible, securities written on them could expand the aggregate payoff space.

The project contains an older scale diagnostic intended to quantify this concern. It adds orthogonal, displacement-relevant aggregate variation outside the two- shock state and asks how much would be required before an older aggregate-risk proxy assigns one half of the burden to the aggregate component. The answer is about 0.25 times current aggregate-income variance under variance accounting and about 0.43 times under the associated welfare proxy.

These numbers should *not* be called reversal thresholds for the current market-completion result. The calculation adds or removes aggregate income risk; it does not solve a new zero-net-supply Arrow economy with the enlarged state, and its welfare proxy predates the current $u'(c)d$ and menu-indexed $\mathcal{D}(A)$ construction. The correct interpretation is more modest and, for the paper, more useful: the amount of omitted systematic risk needed to make a richer aggregate state economically important is not obviously enormous. The current model cannot rule out such missing state

variables.

This scope limitation also clarifies what would *not* overturn the boundary. Adding another security whose payoff is only a different function of the same (ε, θ) state cannot exceed the full-state ceiling. Nor can a richer aggregate state directly insure worker identity if, conditional on that enlarged state, otherwise identical workers still realize different employment outcomes. To alter the conclusion materially, new common information must make a sizable part of the welfare-relevant displacement loss aggregate-measurable; the remaining within-state identity lottery still requires an employment- contingent contract or transfer.

IA.F.6 Summary of the financial sensitivity exercises

The diagnostics support a narrower conclusion than a claim of quantitative invariance. Stock-market access can have ordinary welfare value, reservation prices can be large, and external calibrations can move the level of the measured displacement burden by an order of magnitude. The exact 2.56% benchmark completion effect is therefore not a universal constant.

What is more stable within the exercises that solve the current financial problem is the ordering of magnitudes. Existing equity reaches only a small share of $u'(c)d$; the relevant equity hedge is feasible rather than ruled out by the portfolio constraint; and even complete trading in the modeled aggregate state remains a secondary insurance channel. The main route by which this conclusion could change is a richer, contractible aggregate state containing substantial additional information about displacement. That limitation is a scope condition of the paper, not something the existing two-shock calibration can resolve.

References

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Aladangady, A., Bricker, J., Chang, A. C., Goodman, S., Krimmel, J., Moore, K. B., Reber, S., Volz, A. H., and Windle, R. A. (2023). Changes in U.S. family finances from 2019 to 2022: Evidence from the survey of consumer finances. Technical report, Board of Governors of the Federal Reserve System, Washington, DC.

Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4):1553–1592.

Autor, D. (2024). Applying AI to rebuild middle class jobs. NBER Working Paper 32140, National Bureau of Economic Research.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–2168.

Benzoni, L., Collin-Dufresne, P., and Goldstein, R. S. (2007). Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *Journal of Finance*, 62(5):2123–2167.

Cocco, J. F., Gomes, F. J., and Maenhout, P. J. (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies*, 18(2):491–533.

Couch, K. A. and Placzek, D. W. (2010). Earnings losses of displaced workers revisited. *American Economic Review*, 100(1):572–589.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6):3104–3153.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. Technical report, arXiv:2303.10130.

Farber, H. S. (2017). Employment, hours, and earnings consequences of job loss: US evidence from the displaced workers survey. *Journal of Labor Economics*, 35(S1):S235–S272.

Felten, E. W., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.

Fieller, E. C. (1954). Some problems in interval estimation. *Journal of the Royal Statistical Society: Series B (Methodological)*, 16(2):175–185.

Gollin, D. (2002). Getting income shares right. *Journal of Political Economy*, 110(2):458–474.

Graetz, G. and Michaels, G. (2018). Robots at work. *The Review of Economics and Statistics*, 100(5):753–768.

Guiso, L., Pistaferri, L., and Schivardi, F. (2005). Insurance within the firm. *Journal of Political Economy*, 113(5):1054–1087.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *American Economic Review*, 83(4):685–709.

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy*, 110(4):825–853.

Webb, M. (2020). The impact of artificial intelligence on the labor market. Technical report, Stanford University. Working paper.