Taxing Entrepreneurs and Workers

A Linear Optimization Approach for Multidimensional Screening*

Nicolò Ceneri†, Giuseppe Lopomo‡, Alessandro T. Villa§, Nicolas Werquin¶

October 2026

Abstract

We study optimal taxation in economies with general equilibrium market clearing, where agents with privately known labor skills and entrepreneurial abilities choose between deterministic labor income and risky firm operation. The government observes labor income and realized dividends but not effort or technology shocks. We formulate the multidimensional screening problem as a lottery-based linear optimization, accounting for global incentive constraints, fixed costs and other non-convexities. Optimal policies exhibit tax breaks, which can render net taxes negative, for agents with intermediate entrepreneurial abilities and labor skills above a threshold. General equilibrium strengthens this effect under decreasing returns, as labor-market clearing requires sufficient entry into entrepreneurship, further increasing subsidies for agents with high worker options. In a calibrated U.S. economy, optimal taxes are lower and can be negative for low-profit realizations. Subsidies rise when risk declines and when the frequency of high-ability entrepreneurs in the population diminishes. Global incentive constraints bind only near the occupational frontier, amplifying tax breaks for marginal entrepreneurs.

JEL Codes: D82, E60, H21, H24, H25, L26.
Keywords: Optimal taxation, multidimensional screening, occupational choice, entrepreneurial risk, linear optimization.

*We thank Anmol Bhandari, Marco Bassetto, Job Boerma, Joao Guerreiro, Christian Hellwig, Ellen McGrattan, and Yucheng Yang, as well as participants at NASMES 2026, Barcelona Summer Forum 2026, and SED 2026 for helpful comments and suggestions. The views in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of Chicago, or any person associated with the Federal Reserve System.

†Department of Economics, New York University. E-mail: nicolo.ceneri@nyu.edu.
‡Fuqua School of Business, Duke University. E-mail: giuseppe.lopomo@duke.edu.
§Federal Reserve Bank of Chicago. E-mail: alessandro.villa@chi.frb.org.
¶Federal Reserve Bank of Chicago. E-mail: nicolas.werquin@chi.frb.org.

1 Introduction

Entrepreneurial income is both risky and highly concentrated, with a small share of founders earning large rewards. At the same time, entrepreneurial choices have significant general-equilibrium effects. Firm entry and expansion raise labor demand, affect equilibrium wages and aggregate output. This creates a policy tension that is muted in partial equilibrium. A government that values redistribution and insurance may wish to tax high business incomes aggressively, implement transfers financed by the proceeds, and provide substantial protection against entrepreneurial risk. Yet such policies can also weaken entry incentives for high-productivity entrepreneurs, reducing labor demand and — through equilibrium wage adjustments — lowering aggregate labor income. Because labor skill, entrepreneurial ability, and effort are privately known, and fiscal instruments can condition only on realized labor income and realized firm profits, optimal tax and transfer design must balance redistribution and entry incentives, not only to affect entrepreneurs' behavior, but because the occupational allocation of talent is itself a determinant of economy-wide earnings and welfare.

This paper studies optimal taxation in economies where individuals differ along two dimensions—labor skill and entrepreneurial ability—choose ex ante between paid work and operating a firm, and entrepreneurs face idiosyncratic productivity risk. The government observes labor income and realized dividends, but not effort, entrepreneurial skills, or shock realizations. Optimal policies can be characterized only after solving mechanism-design problems with multidimensional private information, occupational choice, and risk-contingent observables. A key implication is that redistribution and insurance cannot be designed independently of the allocation of talent across occupations: taxes affect entry and firm operation, which in turn affect labor demand, wages, and output in general equilibrium.

Our main results shed light on reasons why a utilitarian planner may tolerate—and even engineer—pockets of high entrepreneurial payoffs. First, the optimal allocation exhibits sharp sorting in the two-dimensional type space. For each labor skill level in the population, there is a threshold in entrepreneurial ability above which agents choose to run firms. This threshold increases with the labor skill, because being capable of generating more labor income requires larger informational rents to be induced to run a firm. As a result, optimal taxes make entrepreneurship locally attractive at the extensive margin. The planner grants information rents to marginal entrepreneurs to prevent high-ability agents from selecting into labor and thus sustain socially valuable entry into entrepreneurship.

Second, these entry incentives have significant general-equilibrium effects. With decreasing returns, even the most productive firms, run by the most skilled entrepreneurs, cannot absorb the entire aggregate labor supply on their own. The labor market clearing must generate a sufficiently large extensive margin of entrepreneurship, in order to stimulate entry by the most talented agents. This amplifies the entry-margin tax break for agents who are close to indifferent between working and operating a firm.

Third, the distribution of entrepreneurial talent and labor skills in the population determines

who receives these subsidies. The largest entry subsidies accrue to agents with the highest options as workers. Although the planner is utilitarian, the presence of private information requires granting rents to these marginal entrepreneurs, which must then be financed by higher net taxes elsewhere. Put differently, in order to support job creation and stimulate output the planner must treat agents with intermediate entrepreneurial skills more favorably than others.

Our analysis is organized in four steps. First, we introduce the model. Agents differ along two dimensions, labor skill and entrepreneurial ability, face fixed participation costs in both occupations, and choose between paid work and operating a firm. Entrepreneurs hire labor at the market clearing wage and are subject to idiosyncratic productivity shocks that are realized after all occupational choices are made. The government observes only realized labor income and realized dividends. Therefore, feasible tax policy cannot condition directly on productivity shocks, agents' types, worker effort, entrepreneurial effort, or firms' labor demand. The planner's problem is a two-dimensional screening problem with risk-contingent observables, where the wage is determined endogenously by the labor market clearing. We compute optimal tax policies using a discrete-lottery linear optimization framework that accommodates a two-dimensional type space and entrepreneurial risk. Unlike approaches that rely on smoothness, convexity, or local first-order conditions, our linear optimization framework accommodates non-convexities—such as fixed entry costs and discrete occupational choices—and delivers a characterization of allocations and tax schedules in environments where local methods are unreliable. Using our framework we can quantify the impact that features of the environment – such as the curvature of the production function, the joint distribution of abilities, and the magnitude of entrepreneurial risk – have on optimal tax policies.

Second, to build intuition we analyze a tractable special case that highlights the mechanism behind the optimal entry subsidy. We show that, with only two entrepreneurial skill levels and no heterogeneity on the labor side, when the high-talent incentive constraint binds, the planner must distort the occupational menu in favor of entrepreneurship. Productive entrepreneurs receive positive information rents, consuming more than workers. These rents are not merely redistributive; they also weaken entrepreneurial effort. In our general-equilibrium setting, this reduces productivity, output, and wages and requires higher labor supply from workers. The analysis of this special case thus delivers a transparent trade-off: relaxing screening constraints requires shifting resources toward marginal entrepreneurs, but the same shifts propagate through labor market clearing and can lower aggregates. It also yields sharp comparative statics. Screening distortions are smaller when production is more labor-intensive, and are larger when high-talent entrepreneurs become scarce.

Third, we calibrate the model to U.S. moments on labor force participation, entrepreneurship, firm size, and the dispersion of business income, and solve for optimal taxes and transfers in general equilibrium. The calibrated optimum features sharp occupational sorting and a highly localized entry-margin subsidy: tax liabilities fall discretely—and often turn negative—right where individuals are close to switching into entrepreneurship. Entrepreneurial risk matters both because

the planner values insurance and because state-contingent taxes relax incentive constraints. Taxes are higher in high-profit states, while subsidies expand in low-profit states where the temptation to select into labor is strongest. As highlighted in the analysis of the special case, when high-ability entrepreneurs are scarcer, subsidies extend deeper into entrepreneurship (even in high-profit states) to sustain labor demand. We also find that when labor and entrepreneurial skills are positively correlated (the "superstars" scenario), the optimal tax policy changes little relative to the baseline. By contrast, under negative correlation (the "specialized talents" scenario), tax incentives are reduced for low-s, high-z entrepreneurs, whose worker outside options are weak, while larger subsidies are directed toward the relatively scarce superstars.

Comparing the full solution with a locally incentive-compatible benchmark reveals that global incentive constraints matter only near the entrepreneurship frontier. The binding non-local constraints involve types close to switching occupations and amplify localized tax breaks for marginal entrepreneurs.

To provide a simple benchmark, we compare the tax schedules that implement the planner's allocation with those in the calibrated U.S. economy. We then ask whether the shape of the optimum can be summarized by a simple parametric tax schedule. To this end, we fit the log-linear schedule augmented with a lump-sum transfer studied by Ferriere et al. (2023), which allows for a uniform transfer and hence for net subsidies at low incomes. We find that, away from the entrepreneurship entry region, this parametric family tracks the optimum closely, but only by reversing the sign of progressivity: it pairs a sizable transfer with negative progressivity, so that redistribution is delivered through the transfer while marginal tax rates decline with income. By contrast, the calibrated U.S. labor-income schedule features mild positive progressivity and no transfer. What the parametric family cannot reproduce is the localized subsidy to low-dividend entrepreneurs near the entrepreneurship entry margin; including that region sharply deteriorates the fit. This entry-margin subsidy is therefore precisely the feature of the optimum that lies outside the parametric class ordinarily used to describe observed tax systems.

Finally, we extend the model to distinguish own-account self-employment from employer firms. Self-employed agents decide after the productivity shock whether to operate alone or hire workers, so that only the employer margin generates labor demand. Calibrating the extension to the U.S. self-employment and employer shares documented by Bhandari et al. (2026), we find that the main mechanism is preserved: optimal taxes fall sharply for agents close to entering employer entrepreneurship. At the same time, the additional production option makes global incentive compatibility substantially more important: the set of types with profitable non-local deviations expands monotonically in the productivity of solo production, reflecting the fact that agents have multiple technologies for generating the same observable business income.

Related literature. This paper contributes to the literature on optimal taxation and mechanism design in environments with multidimensional private information and endogenous occupational choice. Multidimensional screening problems have been studied first by Armstrong (1996), Rochet

and Choné (1998), and Manelli and Vincent (2007). These papers highlight that in multidimensional screening problems without any specific structure, the optimal mechanism tends to be highly sensitive to the belief of the designer. More recent work in public finance includes Kleven et al. (2009), Scheuer (2014), Rothschild and Scheuer (2013, 2014), Bergstrom and Dodds (2021), Boerma et al. (2022), Dodds (2024), Bierbrauer et al. (2024), Spiritus et al. (2025), Golosov and Krasikov (2025), and Shen et al. (2026).[1] These papers examine different economic problems (e.g., taxation of couples or multi-sector abilities), and identify analytically tractable solutions. The main distinguishing feature of our work is the numerical analysis of the general model, which relies on a linear formulation that can account for global incentive constraints as well as non-convexities generated by the economic primitives. Most closely related to our work are Prescott and Townsend (1984), Scheuer (2014), Boerma et al. (2022), and Shen et al. (2026).

It is well known that optimal allocations in multidimensional environments may exhibit global phenomena such as bunching and binding non-local incentive constraints which cannot be handled by local first-order methods. Solving large-scale problems of this kind requires identifying all binding incentive constraints, which is analytically challenging and computationally demanding, especially so in settings with non-convexities. Prescott and Townsend (1984) introduced pioneering techniques to represent nonlinear incentive problems as linear optimization problems using lottery-based transformations (see, e.g., Phelan and Townsend 1991, Doepke and Townsend 2006, and Myerson 2013). However, this approach requires introducing a large number of control variables. Our paper combines this discrete-lottery formulation with constraint generation techniques, and leverages modern large-scale linear programming solvers (such as Gurobi or CPLEX), to solve the multidimensional screening problem by accounting for all incentive constraints numerically. Related large-scale linear programming methods have been applied to dynamic finance and procurement auctions; see, for example, Nikolov et al. (2021); Lopomo et al. (2023). As such, our analysis provides a tractable way to enforce full global incentive compatibility in multidimensional type spaces, while allowing for non-smooth primitives and non-convexities that naturally arise in economic applications, such as endogenous occupational choices with entry costs.

Carlier et al. (2024) also develop a solution algorithm applicable to non-convex multidimensional screening problems but restrict preferences to be quasi-linear, which rules out a direct redistributive or insurance motive. By contrast, our lottery-based formulation allows for generic curvature in preferences. As we show in Subsection 5.3.1, curvature in consumption utility can affect the level and composition of entrepreneurial activity, as well as the set of binding global incentive constraints, both qualitatively and quantitatively. From a methodological standpoint, their approach relies on a primal-dual algorithm, while ours relies instead on a more direct linear representation of the mechanism-design problem using the lottery-based techniques pioneered by Prescott and Townsend (1984) and related work. While our method is theoretically straightforward, it is more computationally intensive, because the number of controls and incentive constraints grows rapidly

[1] Beyond these closest connections, our paper relates to the broader literature on taxing entrepreneurial income under one-dimensional private information; see, e.g., Shourideh (2012); Piketty et al. (2014); Ales et al. (2017).

with the size of the type and allocation grids. In practice, however, modern large-scale LP solvers, combined with barrier methods and constraint generation, allow us to solve the resulting discretized problem with full global incentive compatibility at empirically relevant grid sizes.

Closely related in method is Shen et al. (2026), who also study nonconvex multidimensional screening problems, including optimal Mirrleesian taxation, and use lotteries to enforce global incentive compatibility. Their approach reformulates the lottery problem through a Lagrangian dual and solves it by subgradient descent. This strategy sidesteps the dimensionality of the linear program and is faster in their applications. By contrast, we work directly with the primal linear formulation. This requires more control variables and simplex constraints, but it has two advantages for our purposes. First, linear programs are especially transparent and straightforward to implement. Second, it allows us to exploit modern large-scale LP solvers, such as Gurobi, which deliver robust global solutions efficiently, especially when combined with constraint-generation techniques; this robustness is useful in environments with non-smooth features, such as fixed occupational costs and discrete occupational choices, where the nonlinear subproblems associated with the dual approach may be harder to solve reliably. Our contribution is also economic: we apply this formulation to optimal taxation with occupational choice between paid work and risky entrepreneurship, occupation-specific resource costs, and general-equilibrium labor-market clearing. In this environment, binding global constraints concentrate near the occupational frontier, where agents are tempted to switch occupations, and induce localized tax breaks for marginal entrepreneurs. These rents come at a distributional cost, but they are optimal because sustaining entrepreneurial entry is valuable when production features decreasing returns and firms generate labor demand.

Our approach is also related to Boerma et al. (2022), who develop an alternative LP representation for multidimensional optimal taxation. Their method exploits convex duality: incentive compatibility implies that the indirect utility function satisfies a convexity restriction, and Legendre-type transformations allow the planner's problem to be written in a dual linear form. On the one hand, our discrete-lottery formulation does not require constructing a dual convex representation of the indirect utility function. Instead, we work directly with lotteries over feasible allocation rules and impose incentive compatibility through the expected utilities generated by those rules. This makes the approach well suited to environments with discrete choices, fixed costs, or other non-smooth features for which a dual convex representation is less immediate. On the other hand, the cost of this generality is computational: the number of variables in our formulation grows with the size of the type, shock, and allocation grids, whereas their dual representation is designed to exploit the convex structure of the multidimensional screening problem more directly.

Finally, Scheuer (2014) and Rothschild and Scheuer (2013, 2014) study optimal taxation in general-equilibrium environments with endogenous occupational choice by agents who are privately informed about multiple characteristics: their labor productivity in all sectors and/or their fixed utility cost of entering entrepreneurship. Our theoretical framework builds on the seminal model of Scheuer (2014). However, in these papers, multidimensional skills can be summarized by an endoge-

nous one-dimensional variable that is sufficient for preferences over consumption-income bundles.[2] By contrast, in our model where labor income and dividends are distinct observables with effort costs that depend on different privately known abilities and heterogeneous fixed resource costs of entering each occupation, there is no analogous one-dimensional sufficient statistic, and our incentive problem remains multidimensional. As is common in the heterogeneous firm macroeconomic literature (e.g. Khan and Thomas 2013), we also introduce idiosyncratic productivity shocks in the production function.

Our extension is motivated by Bhandari et al. (2026), who use Internal Revenue Service and Social Security administrative records to study entrepreneurial choice among U.S. business owners and document that only about a third of the self-employed have employees. We therefore extend our baseline model to distinguish own-account self-employment from employer entrepreneurship. This distinction matters in our environment because only employer entrepreneurship generates labor demand, while the hiring decision remains a private action that the planner cannot directly condition on.

Against this background, our aim in this paper is to build a primarily quantitative analysis of the optimal tax system in the presence of entrepreneurs.

Structure. The remainder of the paper is organized as follows. Section 2 presents the full model. Section 3 develops the discrete-lottery linear formulation. Section 4 analyzes a tractable special case of the model to clarify the mechanism with analytical results. Section 5 describes the calibration and reports the quantitative results. Section 6 presents an extension of the baseline model with own-account self-employment and an endogenous hiring margin. Section 7 concludes.

2 Model

In this section we describe the model. We allow for heterogeneity in labor skills and entrepreneurial abilities, fixed participation costs in both occupational choices, and idiosyncratic firm-level productivity shocks.

2.1 Overview and notation

Types. There is a unit mass of agents with privately known type

$$(s, z) \in \mathcal{S} \times \mathcal{Z} \subset \mathbb{R}^2_+,$$

distributed with joint density $f(s, z)$, constructed from marginal distributions F_s and F_z via a copula. We interpret s as idiosyncratic labor productivity ("skill") and z as idiosyncratic entrepreneurial productivity. In the quantitative analysis, each marginal follows a Fréchet (Type II

[2] In particular, Scheuer (2014)'s privately-known fixed cost of entering entrepreneurship is a utility cost that enters the payoff function separably.

extreme value) distribution and the joint is constructed via a Gaussian copula with correlation parameter ρ.

Labor market and wage. There is a competitive labor market with wage $w > 0$ per efficiency unit of labor. Labor supplied by workers and labor hired by firms are measured in efficiency units.

Occupational choice and private actions (exclusive occupations). Each agent chooses *one* occupation *ex ante*, before the realization of entrepreneurial risk. The agent either (i) works as a worker by choosing labor effort $\ell \geq 0$, or (ii) operates a firm by choosing entrepreneurial effort $e \geq 0$ and labor demand $n \geq 0$ (efficiency units) to hire in the competitive labor market. Occupations are mutually exclusive:

$$\ell > 0 \;\Rightarrow\; (n, e) = (0, 0), \qquad \text{and} \qquad n > 0 \;\Rightarrow\; \ell = 0.$$

Labor effort ℓ (if working), entrepreneurial effort e (if operating), and labor demand n (if operating a firm) are privately chosen and unobservable to the government.

Fixed participation costs and home production. Working requires a fixed participation cost $\kappa_1 \geq 0$ whenever $\ell > 0$. Operating a firm requires a fixed cost $\kappa_2 \geq 0$ whenever $n > 0$. Agents who choose neither occupation—the *inactive*—consume a home-production level $c_0 > 0$ at zero effort cost, representing home production and means-tested transfers.

Preferences. Preferences are separable in consumption and effort:

$$u(c, \ell, e) = U(c) - V(\ell) - V(e), \tag{1}$$

where $U' > 0$, $U'' < 0$ and $V' > 0$, $V'' > 0$. The concavity of U implies risk aversion, so insurance against entrepreneurial risk is valued.

Observables and policy instruments. The government observes realized labor income y and realized entrepreneurial dividend/profit d, but does *not* observe (s, z), ℓ, e, n, nor the entrepreneurial shock separately. A tax/transfer schedule $T(y, d)$ is feasible. For the mechanism-design analysis below, we formulate the planner's problem as a direct mechanism in which agents report their private type and the planner assigns allocations contingent on those reports.

2.2 Technology, income, and entrepreneurial risk

Labor income. If the agent works (i.e. $\ell > 0$), observed labor income is

$$y \;=\; w \cdot s \cdot \ell \;-\; \kappa_1 \cdot \mathbf{1}[\ell > 0]. \tag{2}$$

If the agent does not work ($\ell = 0$), then $y = 0$.

Entrepreneurial technology. If the agent operates a firm ($n > 0$), output is

$$q \;=\; z \cdot \varepsilon \cdot e^{\alpha} \cdot n^{1-\alpha}, \qquad \alpha \in (0,1), \tag{3}$$

where $e \geq 0$ is entrepreneurial effort and $\varepsilon > 0$ is an idiosyncratic productivity shock. Dividends/profits (observable) are

$$d \;=\; q - w \cdot n - \kappa_2 \cdot \mathbf{1}[n > 0] \;=\; z \cdot \varepsilon \cdot e^{\alpha} \cdot n^{1-\alpha} - w \cdot n - \kappa_2 \cdot \mathbf{1}[n > 0]. \tag{4}$$

If the agent does not operate ($n = 0$), then $d = 0$.

Shock timing and information. The entrepreneurial shock ε is realized after the occupational choice and reporting. The government does not observe ε separately; it only observes its implications through realized dividends d.

Shock distribution. Entrepreneurial risk is i.i.d. across agents and independent of (s, z). We assume ε is lognormal:

$$\log \varepsilon \sim \mathcal{N}(\mu_\varepsilon, \sigma_\varepsilon^2), \qquad \varepsilon > 0, \qquad \mathbb{E}[\varepsilon] = 1,$$

so that $\mu_\varepsilon = -\frac{1}{2}\sigma_\varepsilon^2$.

Mutual exclusivity in observables. Exclusive occupations imply that an agent cannot simultaneously have positive labor income and positive entrepreneurial dividends. Hence implementable outcomes satisfy

$$y > 0 \;\Rightarrow\; d = 0, \qquad \text{and} \qquad d > 0 \;\Rightarrow\; y = 0. \tag{5}$$

We restrict attention to allocations satisfying (5).

Risk and insurance. Because d is observed, the tax/transfer schedule can condition transfers on realized dividends. Hence the planner can provide insurance against entrepreneurial risk, subject to incentive constraints.

2.3 Implementable allocations and constraints

Direct mechanism. A (deterministic) direct mechanism specifies, for each report $(\hat{s}, \hat{z})$, an intended occupational outcome and a mapping from realized observables to consumption:

$$c = c(y, d).$$

Under truthful reporting, a type (s, z) receives the allocation intended for (s, z); under a misreport $(\hat{s}, \hat{z})$ it receives the allocation intended for $(\hat{s}, \hat{z})$.

Implied labor effort from a labor-income assignment. If the intended outcome assigns positive labor income $y(\hat{s}, \hat{z}) > 0$ (hence, by (5), $d = 0$), then a true type (s, z) must supply effort

$$\ell \;=\; \frac{y(\hat{s}, \hat{z}) + \kappa_1 \mathbf{1}[y(\hat{s}, \hat{z}) > 0]}{w \cdot s}. \tag{6}$$

If $y(\hat{s}, \hat{z}) = 0$, then $\ell = 0$.

Implied entrepreneurial effort from a dividend realization. If the intended outcome assigns entrepreneurial activity (i.e. $d > 0$ and hence $y = 0$), then for a true type (s, z) and realized shock ε the observed dividend d must be generated by some pair of private inputs (e, n) satisfying (4), or equivalently

$$z \cdot \varepsilon \cdot e^{\alpha} n^{1-\alpha} \;=\; d + wn + \kappa_2, \qquad e \geq 0,\; n \geq 0. \tag{7}$$

For any given $n > 0$, (7) uniquely pins down the entrepreneurial effort required to deliver dividend d:

$$e(n; d, z, \varepsilon, w) \;=\; \left(\frac{d + \kappa_2 + wn}{z \cdot \varepsilon \cdot n^{1-\alpha}} \right)^{1/\alpha}. \tag{8}$$

Hence many (e, n) pairs can rationalize the same observed d. To obtain a unique "implied effort" mapping from dividends to disutility, we impose a selection rule: among all feasible (e, n) satisfying (7), the agent chooses the pair that minimizes entrepreneurial effort e (and thus minimizes the effort disutility $V(e)$).[3] Minimizing (8) over $n > 0$ yields the least-effort labor demand

$$n^*(d; w) \;:=\; \frac{1-\alpha}{\alpha} \cdot \frac{d + \kappa_2}{w} \quad \text{for } d > 0, \tag{9}$$

and substituting $n^*(d; w)$ into (7) gives the associated implied entrepreneurial effort

$$e^*(d; z, \varepsilon, w) \;:=\; \left(\frac{d + \kappa_2 + w \cdot n^*(d; w)}{z \cdot \varepsilon \cdot \big(n^*(d; w)\big)^{1-\alpha}} \right)^{1/\alpha}. \tag{10}$$

We set $e^*(0; z, \varepsilon, w) = 0$ and $n^*(0; w) = 0$.

Implied effort costs. Define implied worker effort from labor income y as

$$\ell(y; s, w) := \begin{cases} \dfrac{y + \kappa_1}{ws}, & y > 0, \\ 0, & y = 0, \end{cases}$$

and implied entrepreneurial effort from dividends d as

$$e(d; z, \varepsilon, w) := \begin{cases} e^*(d; z, \varepsilon, w), & d > 0, \\ 0, & d = 0. \end{cases}$$

[3] Because V is strictly increasing, minimizing $V(e)$ is equivalent to minimizing e.

Government spending and the aggregate resource constraint. The government must finance an exogenous level of public expenditure $G \geq 0$. As typical in the literature, G is a pure claim on resources: it is levied and consumed by the government, and is neither rebated to agents nor valued in their utility. With a continuum of agents, cross-sectional averages coincide with expectations. In what follows, we therefore use expectation notation over the entrepreneurial shock interchangeably with cross-sectional averages. The aggregate resource constraint is

$$\int_{\mathcal{S}\times\mathcal{Z}} \mathbb{E}_\varepsilon\big[c\big(y(s,z),d(s,z,\varepsilon)\big)\big]\; f(s,z)\,ds\,dz \;+\; G \;\leq\; \int_{\mathcal{S}\times\mathcal{Z}} \mathbb{E}_\varepsilon[y(s,z)+d(s,z,\varepsilon)]\; f(s,z)\,ds\,dz. \quad (11)$$

where $d(s,z,\varepsilon)=0$ for agents who do not operate a firm. In the benchmark economy we set G equal to the revenue raised by the U.S. tax system in the calibrated positive economy, so that the optimal policy is revenue-neutral relative to the data.

Labor market clearing. Total effective labor supplied equals total labor hired:

$$\int_{\mathcal{S}\times\mathcal{Z}} s\cdot\ell(s,z)\cdot f(s,z)\,ds\,dz = \int_{\mathcal{S}\times\mathcal{Z}} \mathbb{E}_\varepsilon\big[n^*(d(s,z,\varepsilon);w)\big]\cdot f(s,z)\,ds\,dz. \quad (12)$$

Incentive compatibility (ex ante). Truthful reporting must be optimal. Fix a true type (s,z) and consider any report $(\hat{s},\hat{z})$. Under a report $(\hat{s},\hat{z})$, the induced observables are $(y(\hat{s},\hat{z}),d(\hat{s},\hat{z},\varepsilon))$, and the true type incurs the corresponding effort costs implied by $\ell(\cdot;s,w)$ and $e(\cdot;z,\varepsilon,w)$.

The incentive constraints require that for all (s,z) and all reports $(\hat{s},\hat{z})$,

$$\begin{aligned}&\mathbb{E}_\varepsilon\big[U\big(c(y(s,z),d(s,z,\varepsilon))\big)-V\big(e(d(s,z,\varepsilon);z,\varepsilon,w)\big)\big]-V\big(\ell(y(s,z);s,w)\big)\;\geq\\ &\mathbb{E}_\varepsilon\big[U\big(c(y(\hat{s},\hat{z}),d(\hat{s},\hat{z},\varepsilon))\big)-V\big(e(d(\hat{s},\hat{z},\varepsilon);z,\varepsilon,w)\big)\big]-V\big(\ell(y(\hat{s},\hat{z});s,w)\big).\end{aligned} \quad (13)$$

It is useful to contrast these constraints with the one-dimensional reductions that arise in related occupational-choice models, such as Scheuer (2014) and Rothschild and Scheuer (2013, 2014). In a deterministic environment with equal operating costs, entrepreneurial production can be reduced to an effective wage for entrepreneurial effort, and each agent's type can be summarized by the best effective wage across occupations. In our baseline environment the occupational costs κ_1 and κ_2 differ across occupations, and enter as resource costs in the production of labor income and entrepreneurial dividends, rather than as additive utility costs as in Scheuer (2014). The occupational costs affect the implied effort required to generate income and can make the payoff from deviations depend on more than a single effective wage. Moreover, entrepreneurial risk introduces a further state-contingent component to entrepreneurial payoffs. Hence the incentive constraints generally remain multidimensional.

Participation constraint (individual rationality). Agents have access to a home-production technology that delivers consumption $c_0 > 0$ at zero effort. Participation in the mechanism is voluntary: the planner cannot assign any type an expected utility below the outside option. For

all (s, z),

$$\mathbb{E}_\varepsilon\big[U\big(c(y(s,z), d(s,z,\varepsilon))\big) - V\big(e(d(s,z,\varepsilon); z, \varepsilon, w)\big)\big] - V\big(\ell(y(s,z); s, w)\big) \;\geq\; U(c_0). \tag{14}$$

2.4 Planner's problem

Let $W(s,z) \geq 0$ be welfare weights. The planner chooses an allocation through the direct mechanism described above, anticipating that the wage w adjusts to clear the labor market.

$$\begin{aligned}\max_{\{c(\cdot),y(\cdot),d(\cdot)\},w} \quad \int W(s,z) \cdot \Big[\, \mathbb{E}_\varepsilon\big[U\big(c(y(s,z), d(s,z,\varepsilon))\big) - V\big(e(d(s,z,\varepsilon); z, \varepsilon, w)\big)\big] \\ - V\big(\ell(y(s,z); s, w)\big) \,\Big] \cdot f(s,z)\, ds\, dz\end{aligned} \tag{15}$$

s.t. (5), (11), (12), (13), (14) $\quad \forall (s,z), (\hat{s}, \hat{z})$.

3 Lottery Formulation and Discrete LP Characterization

In this section, we describe our approach for quantitative analysis. We proceed in two steps. First, we rewrite the planner's problem of Section 2 using a lottery representation in the spirit of Prescott and Townsend (1984). Second, we discretize the type, shock, and allocation spaces. On the resulting finite grid, the lottery representation becomes a finite-dimensional linear problem (LP): the planner optimizes over lotteries with support contained in the grid, and expected utilities, resource constraints, labor-market clearing, participation constraints, and incentive constraints are all linear in the lottery weights.

3.1 Lottery representation

We first rewrite the planner's problem in Section 2 in a lottery form, following the approach pioneered by Prescott and Townsend (1984) and used in subsequent work. The economic environment is unchanged from Section 2. The lottery representation provides an alternative formulation of the planner's problem, in which allocations are described by probability weights over feasible allocation rules.

Let $\mathcal{A}$ denote the set of feasible individual allocation rules. An element $a \in \mathcal{A}$ specifies observable labor income y_a, shock-contingent dividends $d_a(\varepsilon)$, and shock-contingent consumption $c_a(\varepsilon)$, subject to the occupational exclusivity restriction in (5). Given an allocation rule a, a true type (s, z) incurs the implied effort costs described in Section 2: worker effort is given by $\ell(y_a; s, w)$, and entrepreneurial effort is given by $e(d_a(\varepsilon); z, \varepsilon, w)$.

A randomized direct mechanism assigns to each report $(\hat{s}, \hat{z})$ a probability distribution over feasible allocation rules. We denote this distribution by $\pi(\cdot \mid \hat{s}, \hat{z})$. Thus, instead of choosing a single allocation rule for each reported type, the planner chooses probability weights over feasible

allocation rules. A deterministic mechanism is the special case in which, for every report $(\hat{s}, \hat{z})$, the measure $\pi(\cdot \mid \hat{s}, \hat{z})$ places all mass on a single allocation rule.

Because there is a continuum of agents of each type, this representation also has a deterministic interpretation. The measure $\pi(a \mid \hat{s}, \hat{z})$ can be read as the fraction of agents reporting $(\hat{s}, \hat{z})$ who are assigned allocation rule a. Hence the lottery formulation need not be interpreted as idiosyncratic randomization at the individual level. It can instead be viewed as assigning different fractions of the mass of agents with a given report to different feasible allocation rules.

For a true type (s, z), the expected utility from reporting $(\hat{s}, \hat{z})$ is

$$\mathcal{U}(s, z; \hat{s}, \hat{z}) := \int_{\mathcal{A}} \left\{ \mathbb{E}_{\varepsilon}\big[U\big(c_a(\varepsilon)\big) - V\big(e(d_a(\varepsilon); z, \varepsilon, w)\big)\big] - V\big(\ell(y_a; s, w)\big) \right\} d\pi(a \mid \hat{s}, \hat{z}). \tag{16}$$

The incentive-compatibility and participation constraints in (13) and (14) can therefore be written compactly as

$$\mathcal{U}(s, z; s, z) \geq \mathcal{U}(s, z; \hat{s}, \hat{z}), \qquad \forall (s, z), (\hat{s}, \hat{z}), \tag{17}$$

and

$$\mathcal{U}(s, z; s, z) \geq U(c_0), \qquad \forall (s, z). \tag{18}$$

Similarly, the resource constraint (11) becomes

$$\int_{\mathcal{S}\times\mathcal{Z}} \int_{\mathcal{A}} \mathbb{E}_{\varepsilon}[c_a(\varepsilon)] \; d\pi(a \mid s, z)\, f(s, z)\, ds\, dz \; + \; G \leq \int_{\mathcal{S}\times\mathcal{Z}} \int_{\mathcal{A}} \mathbb{E}_{\varepsilon}[y_a + d_a(\varepsilon)] \; d\pi(a \mid s, z)\, f(s, z)\, ds\, dz, \tag{19}$$

and labor-market clearing (12) becomes

$$\int_{\mathcal{S}\times\mathcal{Z}} \int_{\mathcal{A}} s\, \ell(y_a; s, w)\, d\pi(a \mid s, z)\, f(s, z)\, ds\, dz = \int_{\mathcal{S}\times\mathcal{Z}} \int_{\mathcal{A}} \mathbb{E}_{\varepsilon}[n^*(d_a(\varepsilon); w)] \; d\pi(a \mid s, z)\, f(s, z)\, ds\, dz. \tag{20}$$

Let $\Delta(\mathcal{A})$ denote the set of probability measures over $\mathcal{A}$. The planner's lottery problem is

$$\begin{aligned} \max_{\pi, w} \quad & \int_{\mathcal{S}\times\mathcal{Z}} W(s, z)\, \mathcal{U}(s, z; s, z)\, f(s, z)\, ds\, dz \\ \text{s.t.} \quad & (17),\ (18),\ (19),\ (20), \\ & \pi(\cdot \mid s, z) \in \Delta(\mathcal{A}) \qquad \forall (s, z) \in \mathcal{S} \times \mathcal{Z}. \end{aligned} \tag{21}$$

3.2 LP representation

Given the planner's problem (21), we define a finite grid on the continuum of feasible allocations, and optimize over the set of all lotteries with support contained in the grid. On a finite allocation grid, the probability measures $\pi(\cdot \mid s, z)$ become finite-dimensional probability weights, and the objective, resource constraint, labor-market clearing condition, participation constraints, and incentive constraints are linear in these weights. This construction yields a finite LP which approximates the original planner's problem, with the approximation improving as the grids are refined.

The wage. The wage w is a choice variable of the planner subject to labor-market clearing. A worker assigned labor income y must exert effort $\ell = (y+\kappa_1)/(sw)$ and therefore supplies $(y+\kappa_1)/w$ efficiency units, while least-effort labor demand at dividend d is $n^*(d;w) = \frac{1-\alpha}{\alpha}(d+\kappa_2)/w$. Both sides of (12) are homogeneous of degree -1 in w, so clearing can be written equivalently as

$$\int (y+\kappa_1)\, d\Phi \;=\; \int \tfrac{1-\alpha}{\alpha}(d+\kappa_2)\, d\Phi,$$

which contains no w at all: the condition restricts the allocation, not the wage. Labor remains responsive to the wage. Hence wage is welfare-relevant, and in opposite directions for the two occupations. Raising w lowers the effort a worker needs to deliver a given labor income, since $\ell \propto w^{-1}$; but it raises the effort an entrepreneur needs to deliver a given dividend, since a higher wage bill must be covered out of the same output and $e \propto w^{(1-\alpha)/\alpha}$. Welfare is therefore single-peaked in w, and the planner chooses it. We solve a linear program on a grid of wages and pick the one that delivers the highest welfare.

3.2.1 Finite sets

In order to set up the LP formulation we need to define: (i) type and shock spaces, and (ii) grids for allocations.

Types. Let N denote the number of possible types indexed by $i \in \{1,\dots,N\}$. Each type (s_i, z_i) is given weight $\omega_i \geq 0$ with $\sum_{i=1}^{N} \omega_i = 1$.

Shock states. Discretize entrepreneurial risk into R shock states $r = 1,\dots,R$:

$$\varepsilon_r > 0, \qquad p_r \geq 0, \qquad \sum_{r=1}^{R} p_r = 1,$$

with (ε_r, p_r) common across types.

Grids for allocations. Let

$$\mathcal{C} = \{\bar{c}_k\}_{k=1}^{K}, \quad \mathcal{Y} = \{\bar{y}_a\}_{a=1}^{A_y}, \quad \mathcal{D} = \{\bar{d}_b\}_{b=1}^{B_d},$$

with $\bar{c}_k > 0$ and with $\bar{y}_a, \bar{d}_b$ including 0. Precompute $U_k := U(\bar{c}_k)$.

Exclusive occupations (discrete admissible set). At the discrete level, restrict attention to admissible pairs

$$\mathcal{J} \;\subset\; \{1,\dots,A_y\} \times \{1,\dots,B_d\},$$

defined by

$$(a,b) \in \mathcal{J} \quad\Longleftrightarrow\quad \Big(\bar{y}_a > 0 \Rightarrow \bar{d}_b = 0\Big) \text{ and } \Big(\bar{d}_b > 0 \Rightarrow \bar{y}_a = 0\Big). \tag{22}$$

Equivalently, $\mathcal{J}$ contains (i) worker bundles ($\bar{y}_a > 0, \bar{d}_b = 0$), (ii) entrepreneur bundles ($\bar{y}_a = 0, \bar{d}_b > 0$), and (iii) the inactive point $(0, 0)$, but excludes all ($\bar{y}_a > 0, \bar{d}_b > 0$).

3.2.2 Precomputations for the LP for any given wage w

(i) Worker effort disutility. Given a labor-income gridpoint $\bar{y}_a$, a type-i worker must supply effort

$$\ell_{i,a} := \frac{\bar{y}_a + \kappa_1 \mathbf{1}[\bar{y}_a > 0]}{s_i\, w},$$

and we precompute

$$V^{\ell}_{i,a} := V(\ell_{i,a}). \tag{23}$$

Under exclusivity, this term is relevant only for worker bundles $(a, b) \in \mathcal{J}$ with $\bar{d}_b = 0$; for entrepreneur bundles with $\bar{y}_a = 0$, $\ell_{i,a} = 0$.

(ii) Implied inputs and entrepreneurial effort disutility at a dividend node. Under the least-effort selection rule in the continuous model, a dividend realization pins down a least-effort labor demand and the associated entrepreneurial effort. For each dividend node $\bar{d}_b > 0$ define

$$n^*_b \;:=\; \frac{1-\alpha}{\alpha} \cdot \frac{\bar{d}_b + \kappa_2}{w}, \qquad e^*_{i,b,r} \;:=\; \left(\frac{\bar{d}_b + \kappa_2 + w\, n^*_b}{z_i \cdot \varepsilon_r\, (n^*_b)^{1-\alpha}} \right)^{1/\alpha}. \tag{24}$$

Both mappings are evaluated at the wage chosen by the planner. If $\bar{d}_b = 0$, set $n^*_b = 0$ and $e^*_{i,b,r} = 0$ for all (i, r). We precompute entrepreneurial effort disutility

$$V^{e}_{i,b,r} := V(e^*_{i,b,r}). \tag{25}$$

(iii) Total effort disutility coefficient. For each type i, admissible bundle $(a, b) \in \mathcal{J}$, and shock state r, define

$$\mathcal{V}_{i,a,b,r} \;:=\; V^{\ell}_{i,a} + V^{e}_{i,b,r}. \tag{26}$$

By exclusivity, exactly one of the two components is nonzero for worker or entrepreneur bundles, and both are zero at $(0, 0)$. Then $U_k - \mathcal{V}_{i,a,b,r}$ is the per-state utility coefficient used in the LP.

3.2.3 Decision variables

The numerical implementation uses a lottery representation. For each reported type i and shock state r, the mechanism chooses:

- a lottery over consumption nodes $\mathcal{C}$:

$$\lambda_{i,k,r} \geq 0, \qquad k = 1, \ldots, K;$$

- a lottery over admissible income/dividend pairs $(a, b) \in \mathcal{J}$:

$$\mu_{i,a,b,r} \geq 0, \qquad (a, b) \in \mathcal{J}.$$

The planner's objective and constraints depend only on these marginals because preferences are additively separable ($U(c) - \mathcal{V}$) and the feasibility constraints in the implementation depend on c only through its expectation, and on (y, d) only through their expectations and pre-computed coefficients. Any pair of marginals (λ, μ) can be coupled into a joint lottery over (c, y, d) conditional on (i, r), so this marginal representation is without loss for the discretized direct mechanism, in which allocations condition on the reported type.

Lottery interpretation. We interpret lotteries over continuous allocation variables as a piecewise-linear interpolation across grid nodes. Thus, the allocation space for consumption, labor income, and dividend nodes is continuous in lottery space. As the grids are refined, the piecewise-linear interpolation becomes finer, allowing a better approximation of nonlinearities.

A lottery over discrete choices, such as occupational choices, can be interpreted either as individual randomization over occupations, driven by idiosyncratic shocks, or, equivalently in an economy with a continuum of agents, as a deterministic split of a reported type into fractions assigned to each occupation.

3.2.4 LP representation

Fix $(s_i, z_i, \omega_i, W_i)_{i=1}^N$, grids $(\mathcal{C}, \mathcal{Y}, \mathcal{D})$, the admissible set $\mathcal{J}$ in (22), the discretized shock system $(\varepsilon_r, p_r)_{r=1}^R$, and the precomputations (23)–(26) and (24). For any given wage w, the objective and constraints below are linear in the decision variables (λ, μ). We solve this linear problem several times until we find the welfare-maximizing wage.

LP representation for a given wage w

The planner maximizes weighted expected utility:

$$\max_{\lambda,\mu} \quad \sum_{i=1}^{N} \omega_i \cdot W_i \cdot \left[\sum_{r=1}^{R} p_r \cdot \left(\sum_{k=1}^{K} \lambda_{i,k,r} \cdot U_k - \sum_{(a,b)\in\mathcal{J}} \mu_{i,a,b,r} \cdot \mathcal{V}_{i,a,b,r} \right) \right]. \tag{27}$$

subject to:

(i) Simplex constraints. For every reported type i and shock state r,

$$\sum_{k=1}^{K} \lambda_{i,k,r} = 1, \qquad \lambda_{i,k,r} \geq 0, \tag{28}$$

$$\sum_{(a,b)\in\mathcal{J}} \mu_{i,a,b,r} = 1, \qquad \mu_{i,a,b,r} \geq 0. \tag{29}$$

(ii) Resource feasibility. Expected consumption plus government spending cannot exceed expected total (observable) income:

$$\sum_{i=1}^{N} \omega_i \sum_{r=1}^{R} p_r \sum_{k=1}^{K} \lambda_{i,k,r} \cdot \bar{c}_k \; + \; G \; \leq \; \sum_{i=1}^{N} \omega_i \sum_{r=1}^{R} p_r \sum_{(a,b)\in\mathcal{J}} \mu_{i,a,b,r} \cdot (\bar{y}_a + \bar{d}_b). \tag{30}$$

Because G is a constant, (30) remains linear in the lottery weights, so the revenue requirement leaves the linear-programming structure of the problem intact.

(iii) Labor market clearing. Define effective labor supplied at a labor-income node as

$$L_a^s := \frac{\bar{y}_a + \kappa_1 \cdot \mathbf{1}[\bar{y}_a > 0]}{w}.$$

A worker assigned $\bar{y}_a$ exerts effort $(\bar{y}_a + \kappa_1)/(s_i w)$ and therefore supplies $(\bar{y}_a + \kappa_1)/w$ efficiency units, independently of skill. Labor demanded at dividend node $\bar{d}_b$ (under the least-effort rule) is n_b^* from (24), which carries the same factor $1/w$; both sides of (31) are therefore homogeneous of degree -1 in the wage. Labor market clearing is:

$$\sum_{i=1}^{N} \omega_i \sum_{r=1}^{R} p_r \sum_{(a,b)\in\mathcal{J}} \mu_{i,a,b,r} \cdot L_a^s = \sum_{i=1}^{N} \omega_i \sum_{r=1}^{R} p_r \sum_{(a,b)\in\mathcal{J}} \mu_{i,a,b,r} \cdot n_b^*. \tag{31}$$

(iv) Incentive compatibility. For any true type i and any report j, truthful reporting must yield weakly higher expected utility. Under misreport j, the agent receives the lottery designed for report j, but effort disutility is evaluated using the true (s_i, z_i) (via $\mathcal{V}_{i,a,b,r}$). Thus, for all i, j,

$$\sum_{r=1}^{R} p_r \left(\sum_{k=1}^{K} \lambda_{i,k,r} \cdot U_k - \sum_{(a,b)\in\mathcal{J}} \mu_{i,a,b,r} \cdot \mathcal{V}_{i,a,b,r} \right) \geq \sum_{r=1}^{R} p_r \left(\sum_{k=1}^{K} \lambda_{j,k,r} \cdot U_k - \sum_{(a,b)\in\mathcal{J}} \mu_{j,a,b,r} \cdot \mathcal{V}_{i,a,b,r} \right). \tag{32}$$

(v) Participation constraint (individual rationality). Each type must receive expected utility at least as large as the outside option $U(c_0)$:

$$\sum_{r=1}^{R} p_r \left(\sum_{k=1}^{K} \lambda_{i,k,r} \cdot U_k - \sum_{(a,b)\in\mathcal{J}} \mu_{i,a,b,r} \cdot \mathcal{V}_{i,a,b,r} \right) \geq U(c_0), \qquad \forall i. \tag{33}$$

(vi) Ex-ante occupational choice (no switching across shocks). The implementation imposes that occupational status does not depend on the entrepreneurial shock realization. Let

$$\mathcal{J}^W := \{(a,b) \in \mathcal{J} : \bar{y}_a > 0\}, \qquad \mathcal{J}^0 := \{(a,b) \in \mathcal{J} : \bar{y}_a = 0,\ \bar{d}_b = 0\}.$$

Then for each reported type i and each $r \geq 2$,

$$\sum_{(a,b)\in\mathcal{J}^W} \mu_{i,a,b,r} = \sum_{(a,b)\in\mathcal{J}^W} \mu_{i,a,b,1}, \tag{34}$$

$$\sum_{(a,b)\in\mathcal{J}^0} \mu_{i,a,b,r} = \sum_{(a,b)\in\mathcal{J}^0} \mu_{i,a,b,1}. \tag{35}$$

Since $\mathcal{J}$ excludes $(\bar{y}_a > 0, \bar{d}_b > 0)$, these equalities imply that the remaining probability mass on entrepreneur bundles $(\bar{y}_a = 0, \bar{d}_b > 0)$ is also constant across r.

4 A Tractable Special Case

Before turning to the quantitative analysis in Section 5, we analyze a tractable special case of the model in Section 2 with common labor skill and two entrepreneurial types. We compare a "first-best" benchmark, in which there is no private information and thus taxes can be contingent

on entrepreneurs' characteristics, with a "second-best" setting, in which entrepreneurial talent is privately known and thus taxes can depend only on observed dividends.

The analysis delivers three main insights. First, with private information the planner must make the entrepreneurship option sufficiently attractive to deter high-talent agents from becoming workers. Specifically, when the high-type incentive constraint binds, entrepreneurs must be given positive information rents: they consume more than workers and, relative to the full-information benchmark, exert less entrepreneurial effort. In our general-equilibrium setting, this reduces productivity, output, and wages and requires higher labor supply from workers. The planner therefore faces the following trade-off: relaxing the incentive constraint requires shifting resources toward entrepreneurs, but doing so weakens entrepreneurial effort incentives and tends to depress aggregates. Under private information, entrepreneurs are strictly better off both relative to workers and relative to the first-best.

Second, as the production technology becomes more labor-intensive, private information generates smaller distortions in wages, output, and welfare. When the share of high-talent entrepreneurs decreases, the incentive problem becomes more severe, so private information has a larger impact on aggregates and welfare.

Finally, the gap in entrepreneurial talent affects the economy primarily through a scale channel, not through the severity of the incentive problem. The tightness of the incentive constraint—that is, the shadow cost of separation—depends on technology and population composition, rather than on the level (or gap) of entrepreneurial productivity itself. A larger entrepreneurial talent advantage raises equilibrium wages, dividends, and consumption for both occupations while leaving equilibrium effort choices unchanged. As a result, the welfare gains from greater entrepreneurial talent accrue uniformly, benefiting both workers and entrepreneurs at the same rate in both the first-best benchmark and under private information.

4.1 Environment

The setting is as in Section 2, except for the following specializations. Entrepreneurial risk is shut down ($\varepsilon \equiv 1$), participation costs are zero ($\kappa_1 = \kappa_2 = 0$), and the home-production level is $c_0 = 0$ (so that the participation constraint (14) is not binding). A unit mass of agents has common labor productivity $s = 1$ and privately known entrepreneurial productivity $z \in \{z_L, z_H\}$ with $z_H > z_L$ and population shares (ω_L, ω_H), with $\omega_L + \omega_H = 1$. Agents choose ex ante between working (choosing labor effort $\ell \geq 0$) and operating a firm (choosing entrepreneurial effort $e \geq 0$ and labor demand $n \geq 0$). Preferences are $u(c, \ell, e) = \log c - \frac{1}{2}\ell^2 - \frac{1}{2}e^2$. A competitive labor market clears at wage $w > 0$. If working, observable income is $y = w\ell$ and $d = 0$; if operating, output is $q = ze^{\alpha}n^{1-\alpha}$ and observable dividends are $d = q - wn$, with $y = 0$.

As in Section 2, the planner conditions allocations only on observables (y, d) and we restrict attention to exclusive outcomes ($y > 0 \Rightarrow d = 0$ and $d > 0 \Rightarrow y = 0$). Given (w, z) and a target dividend $d \geq 0$, we adopt the least-effort implementation rule: among all (e, n) delivering d, the

entrepreneur chooses the feasible pair minimizing e. For $d > 0$ this yields

$$n^*(d; w) = \frac{1-\alpha}{\alpha}\frac{d}{w}, \qquad q^*(d; w) = \frac{d}{\alpha}, \tag{36}$$

and implied effort

$$e^*(d; z, w) = \Gamma(\alpha)\, w^{\frac{1-\alpha}{\alpha}} \frac{d}{z^{1/\alpha}}, \qquad \Gamma(\alpha) := \frac{1}{\alpha(1-\alpha)^{(1-\alpha)/\alpha}}. \tag{37}$$

4.2 Planner's occupational menu and the separating allocation

Policy instrument. The planner offers an occupational menu with two options: a worker contract $\mathcal{W} = (c^W, y^W)$ and an entrepreneur contract $\mathcal{E} = (c^E, d^E)$. Given wage w, a worker who selects $\mathcal{W}$ supplies $\ell^W = y^W/w$ and attains indirect utility

$$u^W = \log c^W - \tfrac{1}{2}\left(\frac{y^W}{w}\right)^2. \tag{38}$$

A type-z entrepreneur who selects $\mathcal{E}$ chooses the inputs needed to deliver the targeted dividend d^E and attains

$$u^E(z) = \log c^E - \tfrac{1}{2}\big(e^*(d^E; z, w)\big)^2. \tag{39}$$

Sorting. In this environment, any competitive equilibrium with production features *occupational sorting by entrepreneurial productivity*: low-z agents choose to become workers and high-z agents operate firms. The reason is that, for any given dividend target $d > 0$, the optimal entrepreneurial $e^*(d; z, w)$ is strictly decreasing in z; hence higher-z agents can implement any profitable scale at a lower disutility level than lower-z agents. Because workers have no comparative advantage in operating firms and entrepreneurial output is increasing in z, efficiency dictates that entrepreneurial activity is allocated to the high type, while the low type supplies labor.[4] Finally, labor-market clearing requires that both occupations be present: whenever dividends are positive ($d^E > 0$), firm operation entails a positive measure of workers supplying labor and a positive measure of entrepreneurs demanding labor.

Incentive constraints. Any incentive compatible outcome must satisfy:

$$(\mathrm{IC}_H): \quad \log c^E - \tfrac{1}{2}\big(e^*(d^E; z_H, w)\big)^2 \;\geq\; \log c^W - \tfrac{1}{2}\left(\frac{y^W}{w}\right)^2, \tag{40}$$

$$(\mathrm{IC}_L): \quad \log c^W - \tfrac{1}{2}\left(\frac{y^W}{w}\right)^2 \;\geq\; \log c^E - \tfrac{1}{2}\big(e^*(d^E; z_L, w)\big)^2. \tag{41}$$

[4] Formally, among allocations satisfying feasibility and the occupational observability restriction, any allocation with z_L operating firms and z_H working can be improved by swapping occupations, holding fixed the observable outcomes (y, d): the high type can replicate the same dividend with strictly lower effort cost, while the worker contract depends only on y.

Constraint (IC_H) requires that any high-talent agent prefers the entrepreneur option to the worker option. By becoming a worker, she would receive consumption c^W and would have to supply labor effort $\ell = y^W/w$. Constraint (IC_L) requires any low-talent agent prefers to become a worker rather than an entrepreneur. If he selected the entrepreneur option, he would receive consumption c^E and, given his type z_L, would need to exert the higher effort level $e^*(d^E; z_L, w)$ to generate the same observable dividend d^E.

Market clearing conditions. With low types supplying labor and high types demanding labor, labor-market clearing equation can be written as

$$\omega_L \cdot \frac{y^W}{w} \;=\; \omega_H \cdot \overbrace{\frac{1-\alpha}{\alpha} \cdot \frac{d^E}{w}}^{n^*(d^E;w)}. \tag{42}$$

or equivalently

$$y^W = \chi \cdot d^E, \qquad \text{where } \chi := \frac{\omega_H}{\omega_L} \cdot \frac{1-\alpha}{\alpha}. \tag{43}$$

and the aggregate resource feasibility can be written as

$$\omega_L \cdot c^W + \omega_H \cdot c^E \;=\; \omega_H \cdot \overbrace{d^E/\alpha}^{q^*(d^E;w)}. \tag{44}$$

The aggregate amount of labor and the total consumption are thus jointly determined by the dividend target d^E, while the wage w is determined by equilibrium conditions (and remains explicit in the expressions below).

Planner problem. The planner chooses (c^W, y^W, c^E, d^E, w) to maximize utilitarian welfare subject to incentive compatibility and feasibility:

$$\max_{c^W>0,\, y^W\geq 0,\, c^E>0,\, d^E\geq 0,\, w>0} \quad \omega_L \cdot \left[\log c^W - \tfrac{1}{2}\left(\tfrac{y^W}{w}\right)^2\right] + \omega_H \cdot \left[\log c^E - \tfrac{1}{2}\big(e^*(d^E; z_H, w)\big)^2\right] \tag{45}$$

$$\text{s.t.} \quad (\text{IC}_H)\ (40), \qquad (\text{IC}_L)\ (41),$$

$$\omega_L \cdot \frac{y^W}{w} = \omega_H \cdot n^*(d^E; w), \tag{46}$$

$$\omega_L \cdot c^W + \omega_H \cdot c^E = \omega_H \cdot \frac{d^E}{\alpha}. \tag{47}$$

Let $\lambda \geq 0$ denote the multiplier on (IC_H) and $\mu \geq 0$ the multiplier on (IC_L). In the closed-form characterization below we focus on the empirically relevant setting in which

$$\lambda > 0, \qquad \mu = 0,$$

i.e., (IC_H) binds and (IC_L) is slack.

4.3 Closed form when the high-type incentive constraint binds

We now focus on the case where (IC_H) binds and (IC_L) is slack. In words, the planner must deter the high-z agent from wanting to select the worker contract, while the low type finds the entrepreneurship option too costly because doing so would require a substantially higher effort in order to generate the same observable dividend. In this case, the analysis delivers a sharp closed-form characterization of the optimal allocation, and a transparent mapping from the information friction, represented by the magnitude of the multiplier λ, into the differences in consumption and tax burdens the two agent types. Equations (36) and (37) can be rewritten as

$$n^*(d;w) = \beta \cdot \frac{d}{w}, \qquad e_H(d,w) := e^*(d;z_H,w) = \Gamma(\alpha)\cdot w^{\beta}\cdot \frac{d}{z_H^{1/\alpha}},$$

where

$$\beta := \frac{1-\alpha}{\alpha}, \qquad \chi := \frac{\omega_H}{\omega_L}\beta, \qquad \Gamma(\alpha) := \frac{1}{\alpha(1-\alpha)^{(1-\alpha)/\alpha}}$$

and labor-market clearing implies $y^W = \chi \cdot d^E$. The solution to the planner's problem can now be characterized in detail.

Closed-form representation (endogenous wage; (IC_H) binds, (IC_L) slack)

If the solution of (48) below satisfies $\lambda > 0$, then (IC_H) binds, hence the high type is indifferent between the entrepreneur and worker options. The other incentive constraint (IC_L) is automatically satisfied with slack.
If instead (48) yields $\lambda \leq 0$, (IC_H) is slack. By complementary slackness, the relevant multiplier is $\lambda = 0$, and the allocation is obtained by evaluating (49)–(51) at $\lambda = 0$.

Step 1: scalar equation for λ. Imposing (IC_H) as equality implies the following equation in λ

$$\boxed{\log\left(\frac{\omega_L\cdot(\omega_H+\lambda)}{\omega_H\cdot(\omega_L-\lambda)}\right) = \frac{1}{2}\cdot\left[\frac{\alpha}{\omega_H+\lambda} - \frac{1-\alpha}{\omega_L-\lambda}\right].} \tag{48}$$

Step 2: wage $w(\lambda)$. The wage is pinned down by the planner's FOC with respect to w

$$\boxed{w(\lambda) = z_H\cdot\left(\frac{(\omega_L-\lambda)\cdot\left(\frac{\omega_H}{\omega_L}\right)^2\cdot\beta}{(\omega_H+\lambda)\cdot\Gamma(\alpha)^2}\right)^{\alpha/2}.} \tag{49}$$

Step 3: allocations, given $w(\lambda)$. Dividends and labor income are given by

$$\boxed{d^E(\lambda) = \frac{w(\lambda)}{\chi}\cdot\sqrt{\frac{1-\alpha}{\omega_L-\lambda}}, \qquad y^W(\lambda) = \chi\cdot d^E(\lambda).} \tag{50}$$

and the consumption levels are

$$\boxed{c^E(\lambda) \;=\; (\omega_H + \lambda) \cdot \frac{d^E(\lambda)}{\alpha}, \qquad c^W(\lambda) \;=\; \frac{\omega_L - \lambda}{\omega_L} \cdot \omega_H \cdot \frac{d^E(\lambda)}{\alpha}.} \tag{51}$$

Step 4: Effort and Labor, given $w(\lambda)$. The equilibrium labor and entrepreneurial effort simplify to:

$$\ell^W(\lambda) = \frac{y^W(\lambda)}{w(\lambda)} = \sqrt{\frac{1-\alpha}{\omega_L - \lambda}}, \qquad e_H(\lambda) = \sqrt{\frac{\alpha}{\omega_H + \lambda}}. \tag{52}$$

4.4 Information rents and the distributional tilt toward entrepreneurs

If (IC_H) is binding and (IC_L) is slack, the closed-form solution implies the consumption and effort wedges

$$G(\lambda) := c^E(\lambda) - c^W(\lambda) = \frac{d^E(\lambda)}{\alpha} \cdot \lambda \cdot K, \quad N(\lambda) := e_H(\lambda) - \ell^W(\lambda) = \sqrt{\frac{\alpha}{\omega_H + \lambda}} - \sqrt{\frac{1-\alpha}{\omega_L - \lambda}}, \tag{53}$$

where $K := 1 + \omega_H/\omega_L > 0$.

Distribution and welfare implications. Equation (53) summarizes the redistribution induced by private information. When (IC_H) binds, we have $\lambda > 0$ and $d^E(\lambda) > 0$ (by (50)), hence $G(\lambda) > 0$: entrepreneurs consume more than workers, so the separating allocation necessarily tilts resources toward entrepreneurship. This redistribution is especially stark in the two-type economy: since labor income and dividends exhaust output, the resource constraint (44) is equivalent to $\omega_L T^W + \omega_H T^E = 0$ with $T = y + d - c$, so the information rent is financed one-for-one by workers. In particular, workers pay positive net taxes while entrepreneurs receive a net transfer (a subsidy) of equal aggregate size.[5]

To assess the welfare implications, it is useful to compare effort across occupations. There are two cases. If $\alpha \le \omega_H$, then effort gap $N(\lambda) < 0$ for all $\lambda > 0$. Thus entrepreneurs exert less effort than workers while consuming more, so they are strictly better off in utility terms. Instead, if $\alpha > \omega_H$, which is arguably the more relevant case (e.g., the high-talent share is small, say $\omega_H \approx 5\%$–10%). Then, we have

$$N(\lambda) \begin{cases} > 0, & \text{iff } 0 \le \lambda < \alpha - \omega_H, \\ = 0, & \text{iff } \lambda = \alpha - \omega_H, \\ < 0, & \text{iff } \lambda > \alpha - \omega_H. \end{cases} \tag{54}$$

[5] With no private information, the self-selection motive is absent; in the present representation this corresponds to $\lambda = 0$, which collapses the rent wedge to $c^E(0) = c^W(0)$.

For moderate wedges ($\lambda < \alpha - \omega_H$), entrepreneurs still exert more effort than workers, but private information shifts the allocation in their favor: they receive higher consumption and face weaker effort requirements than under the full-information benchmark. As λ increases, the allocation becomes progressively more tilted: entrepreneurial effort falls more relative to worker effort, while the consumption advantage persists. In all cases, entrepreneurs are better off in utility terms relative to the first best.

Information rents and general-equilibrium feedback. A useful way to unpack (53) is to differentiate the rent wedge $G(\cdot)$ with respect to the incentive distortion λ:

$$\frac{dG}{d\lambda} = \frac{K}{\alpha}\Big(\underbrace{d^E(\lambda)}_{\substack{\text{direct effect:}\\ \text{higher rent weight}}} + \underbrace{\lambda \cdot \frac{d}{d\lambda} d^E(\lambda)}_{\substack{\text{general-equilibrium effect:}\\ \text{dividends adjust via } w(\lambda)}} \Big). \tag{55}$$

The first term represents the *direct* screening force: holding aggregate dividends fixed, a higher λ increases the entrepreneurs' consumption to relax (IC_H), thus widening the consumption gap. The second term reflects *general-equilibrium* feedbacks, due to the endogeneity of both dividends (hence GDP) and wages. Using $d^E(\lambda) = \frac{w(\lambda)}{\chi} \cdot \ell^W(\lambda)$, we obtain the decomposition

$$\frac{d}{d\lambda} d^E(\lambda) = \frac{1}{\chi}\left(\underbrace{\frac{dw(\lambda)}{d\lambda} \cdot \ell^W(\lambda)}_{\text{wage channel}} + \underbrace{w(\lambda) \cdot \frac{d\ell^W(\lambda)}{d\lambda}}_{\text{labor-supply channel}} \right). \tag{56}$$

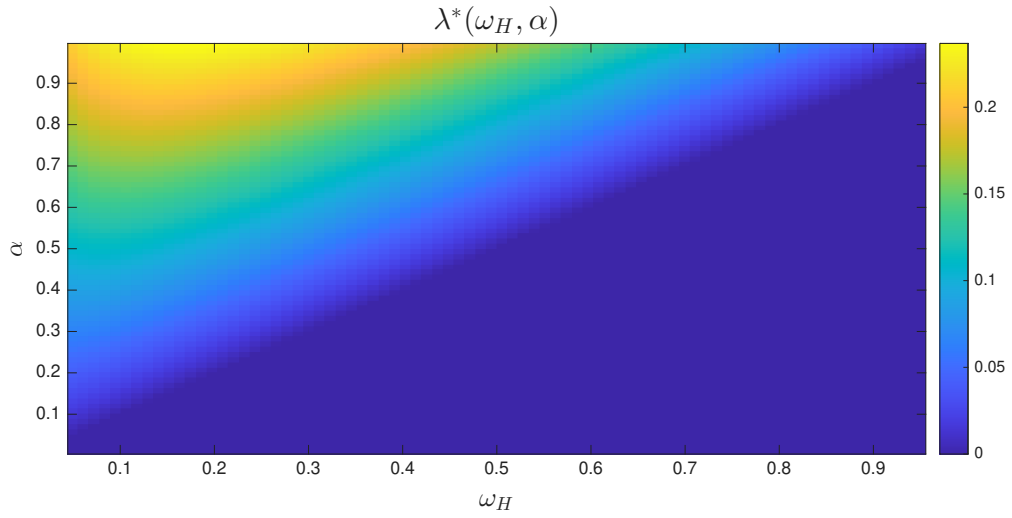
The equations (49) and (52) imply $\frac{dw(\lambda)}{d\lambda} < 0$ and $\frac{d\ell^W(\lambda)}{d\lambda} > 0$. The wage channel pushes dividends *down* (a higher λ is associated with lower entrepreneurial effort and, through the wage formula, a lower equilibrium wage), while the labor-supply channel pushes dividends — and therefore GDP, $\omega_H \alpha^{-1} d^E(\lambda)$ — *up*. Overall, $G(\lambda)$ rises with λ when the direct rent effect and the labor-supply channel dominate the wage decline; if instead $w(\lambda)$ falls sufficiently steeply as λ increases, then $d^E(\lambda)$ and GDP can fall enough so that the consumption gap need not be monotone in λ, even though it remains tilted toward entrepreneurs whenever $\lambda > 0$.

When are rents largest? Figure 1 shows that $\lambda^*(\omega_H, \alpha)$ increases with α, and is largest at an interior and relatively small value of ω_H.[6] Three remarks are in order.

First, holding the population share ω_H fixed, a higher α shifts the production technology toward being more effort-intensive and less labor-intensive, and induces the planner to increase the attractiveness of the entrepreneurship option by increasing their equilibrium payoff relative to the workers. Second, holding technology (α) fixed, a smaller population share ω_H makes it harder for the planner to induce the high type to choose entrepreneurship and thus generate output. As a result, the shadow value of (IC_H) increases, which in turn amplifies information rents.

[6] In Appendix A we show that $\partial\lambda^*/\partial\alpha > 0$.

Figure 1: Comparative statics of the Lagrange multiplier λ



Notes: The figure plots the numerical solution for λ as a function of the high-type share ω_H and the technology parameter α, obtained by solving equation (48) for λ for each parameter pair.

Finally, z_H does not affect the shadow value of the incentive constraint: the multiplier λ is pinned down only by α and ω_H. By contrast, higher entrepreneurial productivity raises equilibrium wages, dividends, and thus consumption levels for both entrepreneurs and workers. The implemented worker labor supply and high-type entrepreneurial effort depend on $(\omega_H, \alpha, \lambda)$ only, not on z_H. As a result, the welfare gains from higher z_H operate entirely through a scale effect on consumption. In particular, welfare for both entrepreneurs and workers increases at rate $1/z_H$ in both the first-best benchmark and under private information (see Appendix A.2 for a formal proof). In this special case, an exogenous increase in entrepreneurial talent therefore raises welfare for both entrepreneurs and workers.

5 Calibration and Quantitative Results

This section describes how we parameterize the model and presents the main quantitative findings. We calibrate nine parameters to match nine empirical moments from U.S. data on labor force participation, entrepreneurship, firm size, income dispersion, effective tax rates, and extensive-margin labor supply. The quantitative analysis delivers three main results: (i) occupational choice exhibits sharp sorting across the two-dimensional type space, with general equilibrium effects playing a central role; (ii) the optimal allocation features distinct tax treatment across occupation-specific margins; and (iii) entrepreneurial risk shapes the mechanism by providing additional screening power to the planner.

5.1 Calibration

The model is static and interpreted as a yearly cross-section. Units are fixed by setting the Fréchet scales to $\sigma_s = 1/\Gamma(1-1/\xi_s)$ and $\sigma_z = 1/\Gamma(1-1/\xi_z)$, so that the untruncated marginals have mean one; labor income, entrepreneurial dividends, consumption and the fixed participation costs are all measured in those units.[7] In the positive economy, we solve for the value of wage that clears the labor market given the observed tax schedule, and obtain $w = 0.835$ at the calibrated parameters. In the planner's problem the wage is chosen optimally by the planner as discussed in Section 5.2. We first solve a decentralized (positive) economy in which agents face observed U.S. tax schedules and make occupational choices, and then calibrate nine structural parameters to match nine empirical moments.

Preferences. We assume CRRA utility over consumption and a standard power disutility of effort:

$$U(c) = \begin{cases} \dfrac{c^{1-\gamma}-1}{1-\gamma}, & \gamma \neq 1, \\ \log c, & \gamma = 1, \end{cases} \qquad V(\ell) = \frac{\ell^{1+\psi}}{1+\psi}.$$

We set $\gamma = 1$ (log utility) and $\psi = 2$, implying a Frisch elasticity of $1/\psi = 0.5$, following Chetty (2012) who reconciles the gap between micro and macro elasticity estimates by accounting for optimization frictions.

Technology and productivity risk. The firm production function is $q = z \cdot \varepsilon \cdot e^{\alpha} n^{1-\alpha}$ with $\alpha \in (0,1)$. The parameter α governs firm scale: lower values imply higher labor demand per unit of entrepreneurial effort, hence larger firms on average. We calibrate α to match the average number of employees per operating firm in the U.S.

Idiosyncratic entrepreneurial risk is multiplicative, i.i.d. across agents, and independent of (s, z). We assume a lognormal shock

$$\varepsilon \sim \log \mathcal{N}\left(-\tfrac{1}{2}\sigma_\varepsilon^2, \sigma_\varepsilon^2\right), \qquad \mathbb{E}[\varepsilon] = 1.$$

We calibrate σ_ε to match the coefficient of variation of business income among entrepreneurs, following DeBacker et al. (2021).

Tax schedules. In the positive economy, labor income y is taxed according to the HSV parametrization of Heathcote et al. (2017):

$$\tilde{y} = \lambda_L \, y^{1-\tau_L},$$

where $\tilde{y}$ denotes after-tax labor income, $\tau_L \in (0,1)$ controls progressivity, and $\lambda_L > 0$ controls the tax level. We calibrate (τ_L, λ_L) to match effective tax rate moments from the Distributional

[7] The discretized grids truncate the marginals at the first and ninety-ninth percentiles and renormalize the resulting probabilities, which discards part of the upper tail: on the grid the realized means are 0.94 and 0.98 rather than exactly one. This affects the interpretation of the unit, not the solution.

National Accounts (Piketty et al., 2022). Entrepreneurial dividends d are taxed according to a stylized version of the U.S. statutory schedule for qualified dividends and long-term capital gains: 0% on the first \$50,000, 15% on \$50,000–\$550,000, and 20% above \$550,000, with thresholds rounded from the single-filer brackets. Dollar thresholds are mapped to model units by normalizing one model unit to \$65,000.

Type distribution. Each type dimension follows a Fréchet (Type II extreme value) distribution, a standard choice in quantitative models of occupational sorting and talent allocation (e.g., Hsieh et al., 2019). The CDF of each marginal is

$$F_x(x) = \exp\left(-\left(\frac{x}{\sigma_x}\right)^{-\xi_x}\right), \qquad x > 0,$$

where $\xi_x > 0$ is the shape parameter governing the thickness of the right tail (lower ξ_x implies a fatter Pareto tail with index ξ_x), and $\sigma_x = 1/\Gamma(1 - 1/\xi_x)$ normalizes the mean to one. The Fréchet distribution combines a smooth, unimodal density with a power-law right tail, making it well suited for capturing both the bulk and the upper tail of the ability distributions. We calibrate ξ_s and ξ_z to match the $\log(p_{90}/p_{50})$ of wages and the Gini coefficient of business income, respectively.

The joint distribution of (s, z) is constructed via a Gaussian copula with correlation parameter ρ. We set $\rho = 0$, so that the two type dimensions are independent conditional on the marginals, following the standard approach when occupations are exclusive and the cross-sectoral ability correlation is not identified from cross-sectional data alone (cf. Hsieh et al., 2019). Subsection 5.3.1 presents comparative statics with respect to ρ.

Extensive margins and home production. Fixed costs κ_1 (workers) and κ_2 (entrepreneurs) govern occupational entry. Inactive agents—those for whom neither occupation yields positive expected utility net of fixed costs—consume a reservation level c_0 representing home production and means-tested transfers. This parameter is calibrated to match the extensive-margin participation elasticity $\eta = 0.25$ estimated by Chetty (2012).

Calibration strategy. We jointly calibrate nine parameters—$(\kappa_1, \kappa_2, \alpha, \sigma_\varepsilon, \xi_s, \xi_z, \tau_L, \lambda_L, c_0)$—to match nine empirical moments using a simulated method of moments procedure that minimizes the sum of squared percentage deviations between model-implied and data moments. Table 1 reports the calibrated parameter values, and Table 2 compares model-implied moments to their data counterparts.

The calibration matches the targeted moments well. The effective tax rates, the top-to-median wage ratio, and the participation elasticity are matched almost exactly, while the remaining moments are within roughly 10% of their targets. The Fréchet shape parameters $\xi_s \approx 2.5$ and $\xi_z \approx 4.6$ imply a fatter right tail for labor skill than for entrepreneurial ability. The span-of-control parameter $\alpha \approx 0.25$ implies labor-intensive entrepreneurial production, generating substantial labor

Table 1: Calibrated parameters

Parameter	Symbol	Value	Target moment	Source
Set externally				
Risk aversion	γ	1.00	Log utility	Standard
Inverse Frisch	ψ	2.00	Frisch = 0.5	Chetty (2012)
Ability correlation	ρ	0.00	Independence	Standard
Calibrated jointly				
Worker fixed cost	κ_1	0.743	Labor participation	Bureau of Labor Statistics (2024b)
Entrep. fixed cost	κ_2	0.110	Entrepreneurship rate	Bureau of Labor Statistics (2024a)
Span of control	α	0.251	Avg. employees/firm	U.S. Census Bureau (2021)
Shock volatility	σ_ε	0.520	CV of business income	DeBacker et al. (2021)
Skill tail (Fréchet)	ξ_s	2.52	$\log(p_{90}/p_{50})$ wage	Piketty et al. (2022)
Ability tail (Fréchet)	ξ_z	4.62	Gini business income	Piketty et al. (2022)
Tax progressivity	τ_L	0.067	ETR bottom 50%	Piketty et al. (2022)
Tax level	λ_L	0.777	ETR middle 40%	Piketty et al. (2022)
Home production	c_0	0.042	Participation elasticity	Chetty (2012)
Statutory business tax (not calibrated)				
Bracket 1	T_B	0%	$d \leq$ \$50k	IRC §1(h)
Bracket 2	T_B	15%	\$50k–\$550k	IRC §1(h)
Bracket 3	T_B	20%	$d >$ \$550k	IRC §1(h)

Notes: The table reports the externally assigned and internally calibrated parameters used in the baseline economy. Parameters in the first block are set externally using standard values or direct empirical counterparts. Parameters in the second block are calibrated jointly to match the corresponding moments listed in the fourth column; the reported target moment indicates the moment most directly associated with each parameter, but identification is joint. The statutory business tax schedule in the last block is imposed directly and is not included in the calibration. Dollar thresholds in the statutory business tax schedule are expressed in annual income units.

demand and significant general equilibrium effects through wage determination. The worker fixed cost ($\kappa_1 \approx 0.74$) exceeds the entrepreneurial fixed cost ($\kappa_2 \approx 0.11$), reflecting the interpretation that entry into paid employment involves search and matching frictions whereas entrepreneurship requires lower upfront investment but faces greater idiosyncratic risk. Home production $c_0 \approx 0.04$ provides a modest outside option, calibrated to reproduce the extensive-margin participation elasticity of 0.25 estimated by Chetty (2012).

Numerical method. We discretize the two-dimensional type space on a 60×60 grid, yielding $N = 3{,}600$ types. Each type is assigned marginal lotteries over $K = 60$ consumption nodes and over admissible income-dividend pairs across $R = 2$ shock states. With $A_y = B_d = 30$, the admissible income-dividend grid contains $|\mathcal{J}| = A_y + B_d - 1 = 59$ nodes. The grids are chosen to approximate the relevant nonlinearities of the allocation rules, with consumption nodes concentrated at low consumption levels where marginal utility is most nonlinear. Since lotteries over nodes provide a piecewise-linear interpolation, the effective choice set is convex in lottery space rather than restricted to deterministic grid points. We verified that further refinements of K, A_y,

Table 2: Targeted moments: model vs. data

Moment	Model	Data	Identifies
Panel A: Targeted moments			
Labor participation	0.645	0.620	κ_1
Entrepreneurship rate	0.093	0.100	κ_2
Avg. employees per firm	9.40	10.00	α
CV of business income	1.23	1.20	σ_ε
$\log(p_{90}/p_{50})$ wage	0.981	0.997	ξ_s
Gini of business income	0.614	0.650	ξ_z
ETR bottom 50%	0.148	0.148	τ_L
ETR middle 40%	0.211	0.214	λ_L
Participation elasticity	0.250	0.250	c_0
Panel B: Untargeted moments (validation)			
Inactivity rate	0.355	0.380	—
Wage Gini	0.363	0.520	—
ETR top 10%	0.150	0.249	—
Avg. ETR labor	0.206	0.234	—
Top 10% wage share	0.289	0.375	—
Top 10% business share	0.389	0.675	—
Business income share	0.422	0.280	—

Notes: The table reports the model-implied moments and their empirical counterparts. Panel A lists the moments targeted in the calibration. The last column reports the parameter most directly associated with each targeted moment; all internally calibrated parameters are chosen jointly. Panel B reports untargeted moments used to assess the model's external fit. Tax moments are model-relevant effective tax rates, including individual income taxes, payroll taxes, and corporate taxes, computed from the distributional national accounts of Piketty et al. (2022). The participation elasticity is the Hicksian extensive-margin elasticity with respect to the net-of-tax rate, following Chetty (2012).

and B_d have only minimal effects on the quantitative results. The planner's problem admits a linear programming formulation over these marginal lotteries, as described in Subsection 3.2. This representation replaces the discretized nonlinear problem, which has $N(1+2R) = 18{,}000$ primitive controls, with a larger but linear problem containing

$$NR\big(K + |\mathcal{J}|\big) = 3{,}600 \cdot 2 \cdot (60 + 59) = 856{,}800$$

primitive lottery controls.

The lottery formulation also introduces auxiliary linear restrictions. For each type and shock state, the consumption and income-dividend lotteries must each sum to one, giving $2NR = 14{,}400$ simplex equalities, together with $NR(K + |\mathcal{J}|) = 856{,}800$ nonnegativity restrictions on the lottery weights. In addition, because occupational status is chosen ex ante and cannot vary with the entrepreneurial shock realization, we impose $2N(R-1) = 7{,}200$ no-switching equalities, which keep the probability of worker and inactive status fixed across shock states. The main computational bottleneck is the set of global incentive-compatibility constraints: imposing global IC directly

requires one restriction for each ordered pair of distinct types, or

$$N(N-1) = 3{,}600 \cdot 3{,}599 = 12{,}956{,}400$$

pairwise inequalities. We handle this dimensionality through constraint generation. Rather than loading all global IC constraints at once, we start from a relaxed LP with a sparse set of local IC constraints, solve it using Gurobi's barrier method, search over type pairs for violated deviations, add the most violated constraints to the active set, and re-solve. The procedure iterates until no global IC violations remain, thereby solving the full global-IC problem while avoiding the memory and numerical burden of imposing all 12.96 million IC constraints simultaneously. Appendix B.1 provides the full algorithmic details. Table 3 summarizes these dimensions, reporting both the general expressions and their baseline numerical values for the lottery LP and the corresponding discretized nonlinear problem.

Table 3: Baseline dimensions of the lottery LP

Object	**Lottery LP**		**Discretized nonlinear problem**	
	Expression	**Count**	**Expression**	**Count**
Types	$N = N_s N_z$	3,600	$N = N_s N_z$	3,600
Income-dividend nodes	$\lvert\mathcal{J}\rvert = A_y + B_d - 1$	59	—	—
Primitive controls	$NR(K + \lvert\mathcal{J}\rvert)$	856,800	$N(1 + 2R)$	18,000
Simplex equalities	$2NR$	14,400	—	—
Nonnegativity restrictions	$NR(K + \lvert\mathcal{J}\rvert)$	856,800	—	—
No-switching equalities	$2N(R-1)$	7,200	—	—
Participation constraints	N	3,600	N	3,600
Resource feasibility	1	1	1	1
Labor-market clearing	1	1	1	1
Global IC constraints	$N(N-1)$	12,956,400	$N(N-1)$	12,956,400

Notes: Baseline values are $N_s = N_z = 60$, $R = 2$, $K = 60$, and $A_y = B_d = 30$. The admissible set $\mathcal{J}$ contains the inactive node, worker nodes, and entrepreneur nodes, excluding bundles with both positive labor income and positive dividends. Primitive lottery controls include consumption lottery weights and income-dividend lottery weights. The discretized nonlinear problem has one labor-income choice and shock-contingent dividend and consumption choices for each type. Global IC constraints count ordered deviations across distinct types.

Government spending. The revenue requirement G is not a free parameter: we set it, as in Boerma et al. (2022), to the revenue that the observed U.S. tax system actually raises in the calibrated positive economy.[8] In the positive economy, aggregate consumption is after-tax income,

[8] Boerma et al. (2022) proceed in the same way, fixing a status-quo labor tax (a flat rate $\tau = 0.3$ in their case) and letting G be the revenue it finances. The difference here is that our status-quo schedule is the calibrated progressive HSV schedule together with the statutory business tax, rather than a flat rate.

so the resources absorbed by the government are

$$G \;=\; \underbrace{\int \mathbb{E}_{\varepsilon}\big[y+d\big]\, f\, ds\, dz}_{\text{output}} \;-\; \underbrace{\int \mathbb{E}_{\varepsilon}\big[c\big]\, f\, ds\, dz}_{\text{consumption}},$$

which equals 17.2% of output in the positive economy at the calibrated parameters, or $G = 0.095$ in levels.[9] Holding G fixed at this level makes the optimal policy *revenue neutral* with respect to the data: the planner must finance the same public expenditure as the current U.S. tax system, and the welfare gains we report therefore come from redesigning the tax schedule, not from shrinking the government.

5.2 The optimal wage

The planner chooses an allocation, and the wage is part of that choice. Since labor-market clearing is homogeneous of degree -1 in w and therefore holds at any wage once the observable allocation (y, d) satisfies it, w is determined by the planner's objective rather than by market clearing, and we optimize over it alongside (c, y, d) subject to the same incentive and participation constraints. We compute the optimum by solving the linear program on a grid of wages, holding the calibrated primitives $(z, \kappa_1, \kappa_2, c_0, G)$ fixed in levels and selecting the wage that delivers the highest welfare. Each grid point is a separate problem solved to full global incentive compatibility, so the comparison is across second-best optima.

The trade-off is visible in the first-order condition. Write

$$\Delta(s,z) \;=\; V\big(\ell(s,z)\big) \;-\; \frac{1-\alpha}{\alpha}\,\mathbb{E}_{\varepsilon}\big[V\big(e(s,z,\varepsilon)\big)\big],$$

the net effort saving a type enjoys from a marginally higher wage: a worker needs less effort to deliver a given labor income, since $\ell \propto w^{-1}$, while an entrepreneur needs more to deliver a given dividend, since a larger wage bill must be covered out of the same output and $e \propto w^{(1-\alpha)/\alpha}$. Let $\eta(\cdot,\cdot) \geq 0$ and $\nu(\cdot) \geq 0$ denote the multipliers on the incentive and participation constraints and $H(\theta) = \int \eta(\theta, \hat{\theta})\, d\hat{\theta}$. Differentiating the Lagrangian and using the envelope theorem — the clearing term drops, because its derivative is proportional to excess demand and hence zero — the optimal wage satisfies

$$\int \big[W(\theta) f(\theta) + \nu(\theta) + H(\theta)\big]\, \Delta(\theta)\, d\theta \;=\; \iint \eta(\theta,\hat{\theta})\, \Delta(\theta \mid \hat{\theta})\, d\hat{\theta}\, d\theta, \tag{57}$$

where $\Delta(\theta \mid \hat{\theta})$ evaluates the same expression at type θ's effort costs but at the bundle assigned to $\hat{\theta}$. Absent private information $\eta = \nu = 0$ and (57) reduces to $\mathbb{E}_W[V(\ell)] = \frac{1-\alpha}{\alpha}\mathbb{E}_W[V(e)]$: the planner equates the average effort cost borne by workers to the average entrepreneurial effort cost,

[9] We set G to the *level* of revenue the observed system raises, not to its share of output. Anchoring the level to a share of the *planner's* output would be circular, since the planner's output depends on G and on the wage the planner chooses; the positive economy is the descriptive benchmark and the planner's economy is the counterfactual.

scaled by the labor share of the wage bill. Screening adds two things. Types whose constraints bind receive extra weight through $\nu + H$, so the wage tilts toward the occupations where incentives are tight. And the right-hand side subtracts the effect of the wage on *deviation* payoffs: a higher wage also lowers the effort a high-ability agent would need in order to mimic a worker, which tightens the constraints that sustain entry into entrepreneurship. The second-best wage therefore differs from the first-best one, and in our calibration it is higher.

Figure B.1 in Appendix B reports welfare across the grid. It is single-peaked, with an optimum at $w^* = 0.93$. The peak is asymmetric: lowering the wage below w^* is more costly than raising it.

It is worth comparing w^* with the wage that prevails in the calibrated positive economy, which is 0.835. The planner raises the wage by about eleven percent relative to the decentralized economy, shifting the terms of trade between the two occupations toward workers.

5.3 Quantitative Results

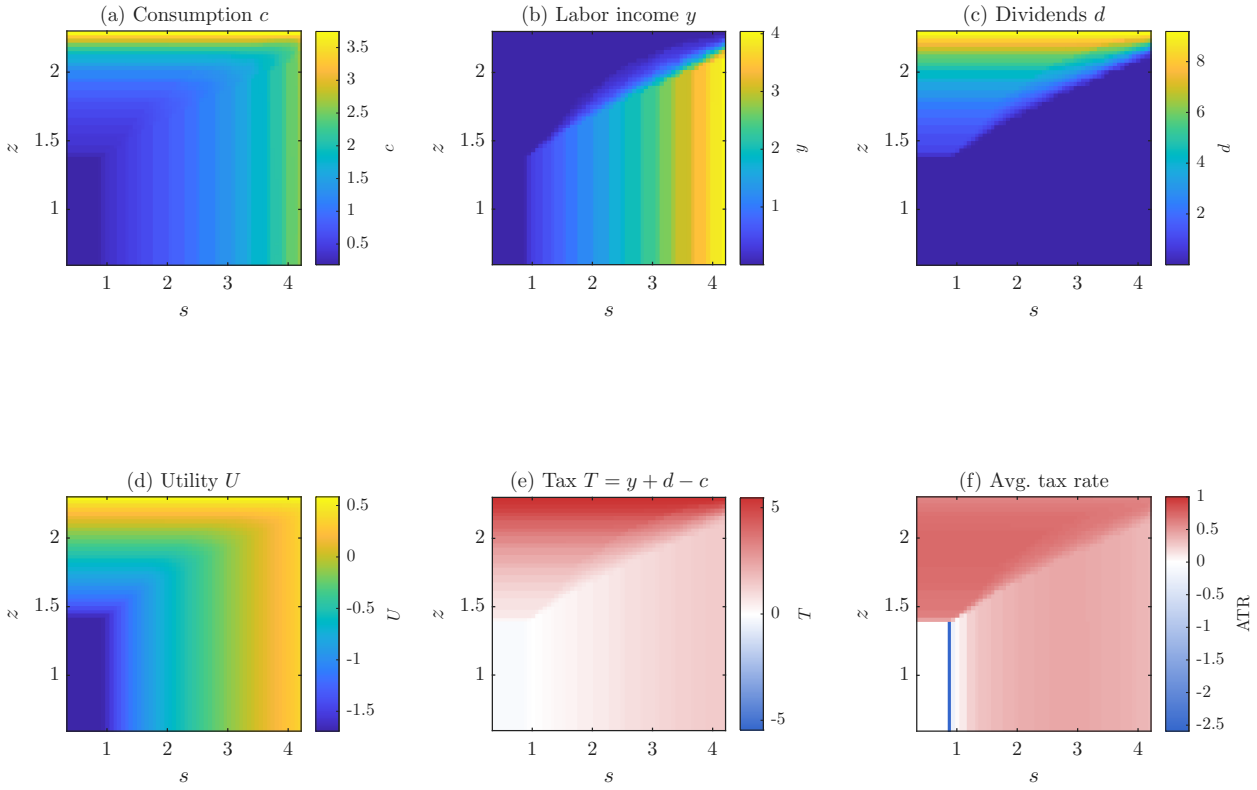
We now describe the structure of the optimal allocation in the calibrated economy and how it changes with (i) curvature in consumption utility, (ii) entrepreneurial risk, (iii) the scarcity of high-z entrepreneurial types, and (iv) affiliation between labor skill s and entrepreneurial ability z. Taxes are reported as net tax liabilities $T(s,z) = y(s,z) + d(s,z) - c(s,z)$, and the average tax rate is $\text{ATR}(s,z) = T(s,z)/\big(y(s,z) + d(s,z)\big)$ (with ATR $= 0$ at the inactive point). Unless otherwise stated, we report *ex ante* (pre-shock) allocations that are constant across entrepreneurial shock realizations; when we display state-by-state taxes, we do so holding fixed the same occupational choice.

Sorting and a tax break at the entry margin. Figure 2 illustrates the main features of the optimal policy and the resulting allocations in the calibrated stochastic economy. First, the equilibrium exhibits sharp occupational sorting. Low entrepreneurial ability types are induced to work, earning positive labor income y and zero dividend d. In contrast, sufficiently high z types operate firms, earning zero labor income y and positive dividend d. Holding z fixed at a sufficiently high value, higher-s agents have a better outside option in paid work and thus require larger entrepreneurial rents to be induced into firm operation.

Second, consumption is increasing in both s and z, while the occupational choice generates a pronounced frontier along the entry boundary. The key redistributional instrument is the tax $T = y + d - c$ in panel (e). Fix a sufficiently high s and increase z. Taxes in the low shock realization are *decreasing* in z around the entrepreneurship entry margin. As z rises into the entrepreneur region, the planner reduces net tax liabilities sharply, creating a localized *entrepreneurial tax break*. This is the quantitative counterpart of the information-rent logic in Section 4: to prevent high-s types with moderate z from mimicking low-z types and selecting a worker contract, the planner makes the entrepreneurial contract more attractive at the point where occupational choice is most elastic.

Finally, due to the presence of decreasing returns in production, these tax breaks are amplified

Figure 2: Optimal policy and resulting allocations



Notes: The figure reports the optimal allocation over the two-dimensional type space (s, z) in the calibrated economy with entrepreneurial risk and baseline affiliation ($\rho = 0$). Entrepreneurial risk has two shock realizations; the plotted objects are *ex ante* allocations (and, when applicable, expectations across shock realizations). Panels display: (a) consumption c, (b) labor income y, (c) dividends d, (d) indirect utility $U(c) - V(\ell) - \mathbb{E}[V(e)]$, (e) net tax liability $T = y + d - c$, and (f) the average tax rate $\text{ATR} = T/(y + d)$ (defined as 0 at $(y, d) = (0, 0)$). Exclusive occupations imply $y > 0 \Rightarrow d = 0$ and $d > 0 \Rightarrow y = 0$.

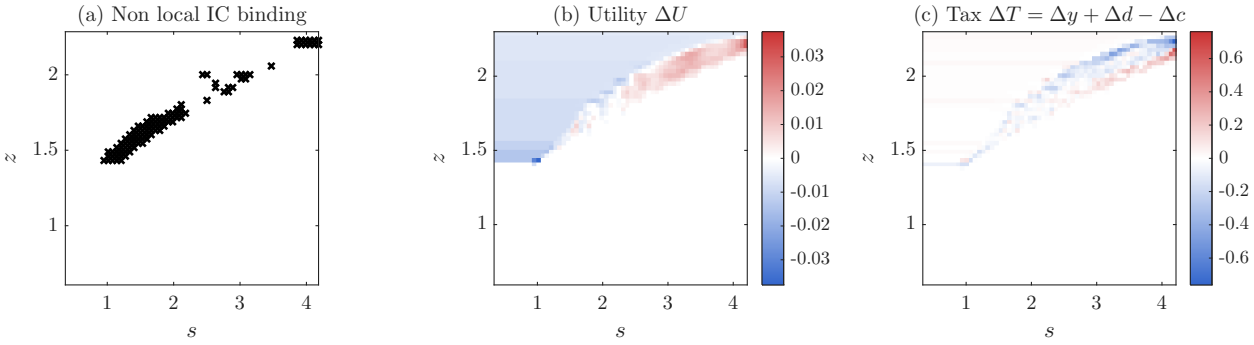
by general equilibrium effects. Because entrepreneurs hire labor, a contraction in entrepreneurship entry reduces labor demand. Market clearing then requires a sufficiently large mass of active entrepreneurs, as decreasing returns discourage the planner from further expanding larger firms. Under the optimal tax schedule, the planner uses tax breaks to sustain entry and stimulate labor demand.

The role of global IC constraints. Figure 3 shows that the relevant global incentive constraints are concentrated almost entirely near the occupational-choice frontier. As shown in panel (a), the binding non-local constraints are clustered near the entrepreneurship entry margin, where agents are close to indifferent between paid work and firm operation. By contrast, for workers away from that boundary, removing non-local IC constraints has no effect on the optimal tax policy and the resulting allocation. When the full set of global constraints is imposed, the optimal policy responds precisely where cross-occupation deviations are most tempting. Panel (c) shows that taxes are reduced at the entrepreneurship entry margin for types just inside the entrepreneur region. These localized tax cuts make entrepreneurship more attractive and deter marginal entrepreneurs—especially those

with relatively strong worker outside options—from selecting the worker contract. Panel (b) shows that imposing global IC reallocates utility toward the entrepreneurship entry margin mainly from types with low s and high z. Utility rises for types near that boundary, where additional rents are needed to sustain separation. Furthermore, in this same region, consumption and dividends fall for most types, as illustrated in panels (a) and (c) of Appendix Figure E.1.

Overall, the figure clarifies that imposing global IC constraints simply sharpens the redistribution required at the extensive margin of entrepreneurship.

Figure 3: Effects of binding non-local IC constraints



Notes: The figure compares the solution in which all global IC constraints are imposed with the solution imposing only local IC constraints. Panel (a) identifies types for which a non-local IC constraint is violated by more than 10^{-2} in the local-IC-only solution. Panel (b) reports the difference in indirect utility, ΔU, between the global and local-IC-only solutions, computed as global minus local. Panel (c) reports the corresponding difference in net tax liabilities, ΔT, again computed as global minus local.

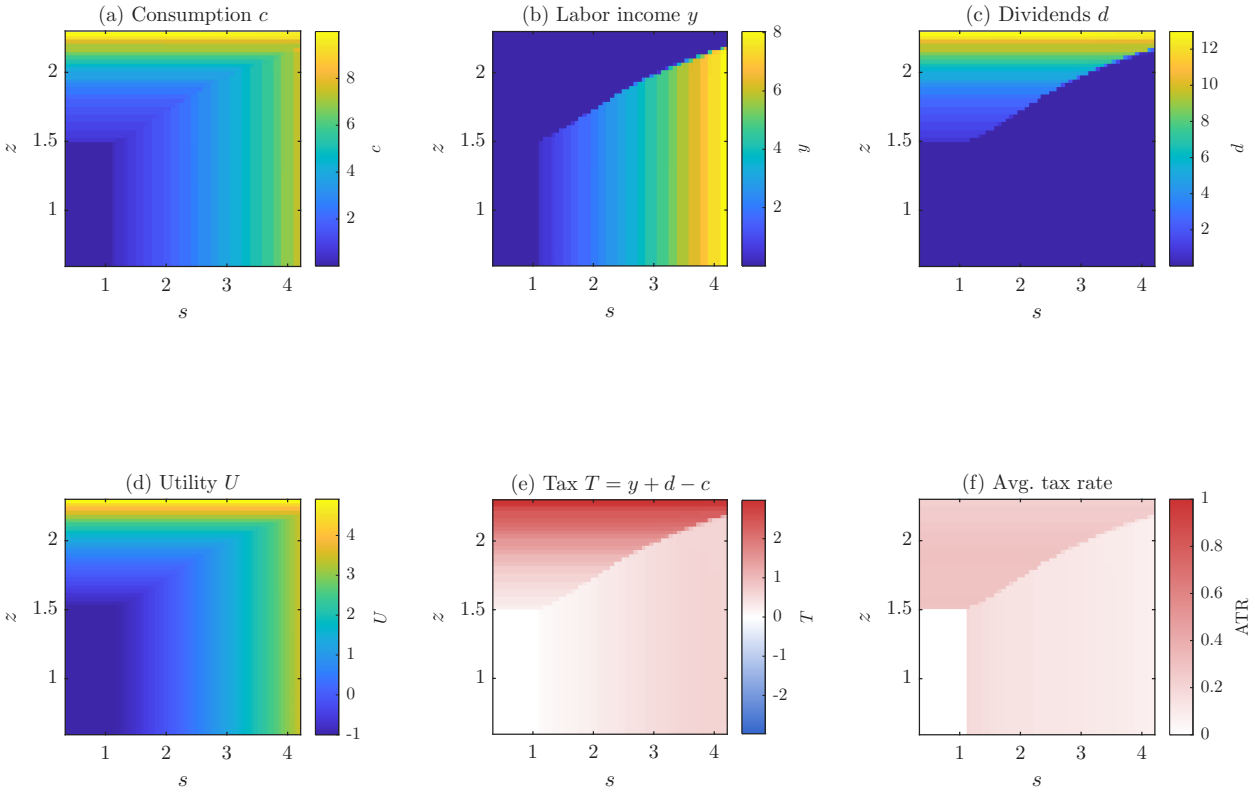
5.3.1 Comparative Statics

We now study how the optimal policy responds to changes in the key economic forces shaping the mechanism: curvature in consumption utility, the scarcity of high-ability entrepreneurs, entrepreneurial risk, and affiliation between labor skill and entrepreneurial ability.

The role of consumption utility curvature. We next isolate the role of curvature in consumption utility by solving an otherwise identical economy with quasi-linear utility, $U(c) = c$. This exercise removes the planner's direct motive to compress consumption inequality. Figure 4 shows that, once this redistributive-insurance motive is shut down, the optimal allocation becomes more tilted toward production through firm scale. The largest entrepreneurial firms become larger and generate substantially higher dividends, especially in the upper region of the z-distribution. By contrast, in the baseline economy with concave consumption utility, the planner is more concerned with inequality and therefore relies more on extensive-margin entry incentives: it fosters entrepreneurship near the occupational frontier, but keeps the largest firms smaller. Thus, curvature in consumption utility changes not only the level of redistribution, but also the composition of entrepreneurial activity, shifting the optimum from larger firms at the top toward broader entry at the margin.

At the same time, high-skill workers also receive higher consumption when utility is quasi-linear, reflecting the increase in aggregate resources generated by a larger entrepreneurial sector. The tax schedule adjusts accordingly: both medium- to high-skill entrepreneurs and high-skill workers face lower net tax liabilities relative to the concave-utility benchmark, reflecting the fact that the planner is less concerned with inequality.

Figure 4: Optimal policy and resulting allocations with quasi-linear utility

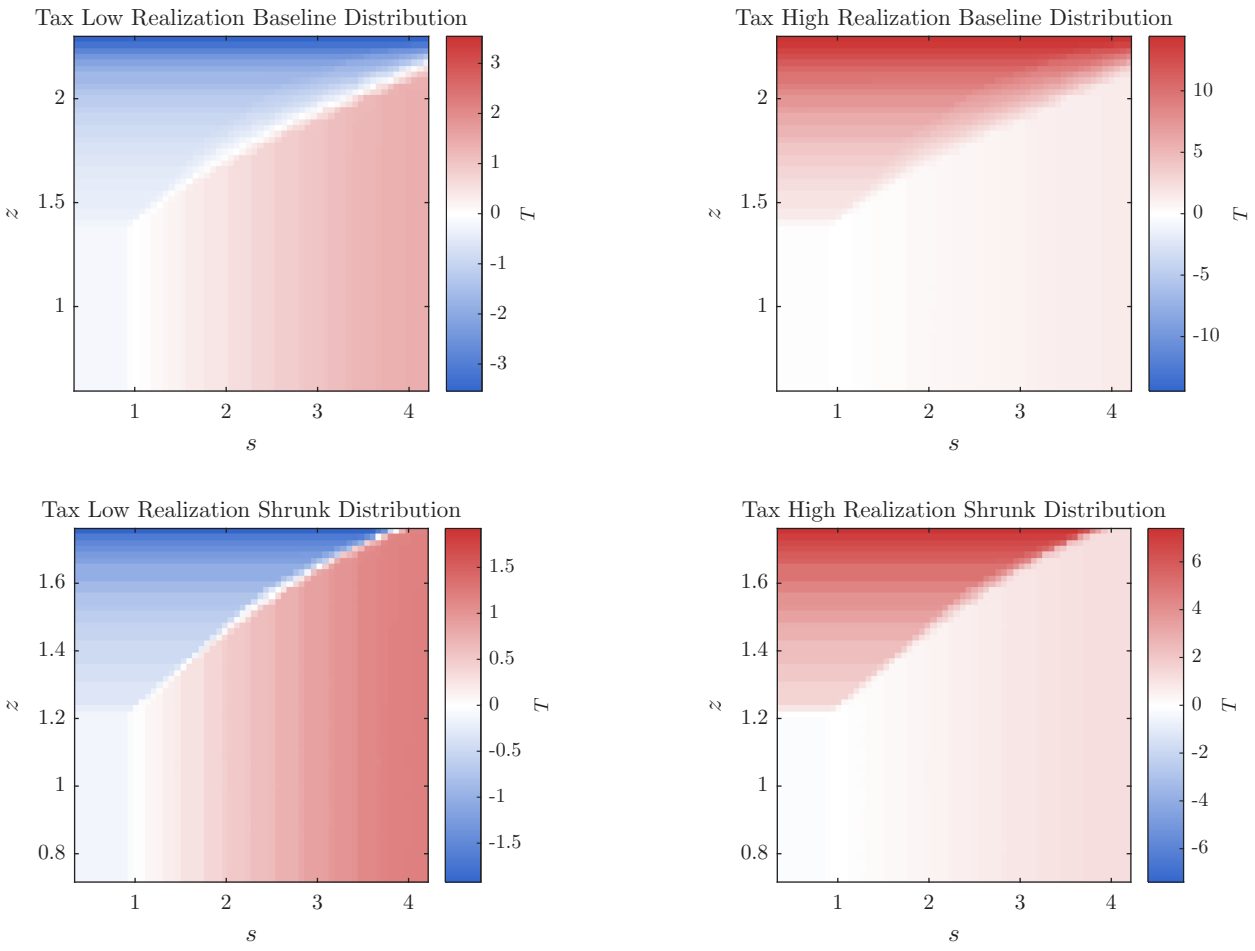


Notes: The figure reports the optimal allocation over the two-dimensional type space (s, z) in the calibrated stochastic economy when consumption utility is quasi-linear, $U(c) = c$. All other parameters are held fixed at their baseline values. Entrepreneurial risk has two shock realizations; the plotted objects are *ex ante* allocations and, when applicable, expectations across shock realizations. Panels display: (a) consumption c, (b) labor income y, (c) dividends d, (d) indirect utility $U(c) - V(\ell) - \mathbb{E}[V(e)]$, (e) net tax liability $T = y + d - c$, and (f) the average tax rate $\text{ATR} = T/(y+d)$, defined as 0 at $(y, d) = (0, 0)$. Exclusive occupations imply $y > 0 \Rightarrow d = 0$ and $d > 0 \Rightarrow y = 0$.

The comparison with the baseline also clarifies how global incentive constraints operate when inequality aversion is absent. Appendix Figure E.2 shows that, in the quasi-linear utility economy, non-local IC constraints bind less extensively, especially along the occupational frontier. The remaining global constraints are concentrated around types who are productive in both occupations and are close to indifferent between working and operating a firm. For these types, the planner reduces the utility assigned to worker contracts just below the frontier and reallocates utility to entrepreneur contracts just above it. This makes entrepreneurship strictly more attractive for marginal high-(s, z) types. Hence, even without a redistributive motive, global incentive compatibility continues to shape the policy locally at the entry margin, but its role is more limited.

Scarcity of entrepreneurial types. We next reduce the dispersion of entrepreneurial ability, compressing the upper tail of z ("worsened entrepreneurial types"). Figure 5 shows how the optimal tax schedule changes in each shock realization relative to the baseline.

Figure 5: Scarcity of high-z types and optimal taxes by realization

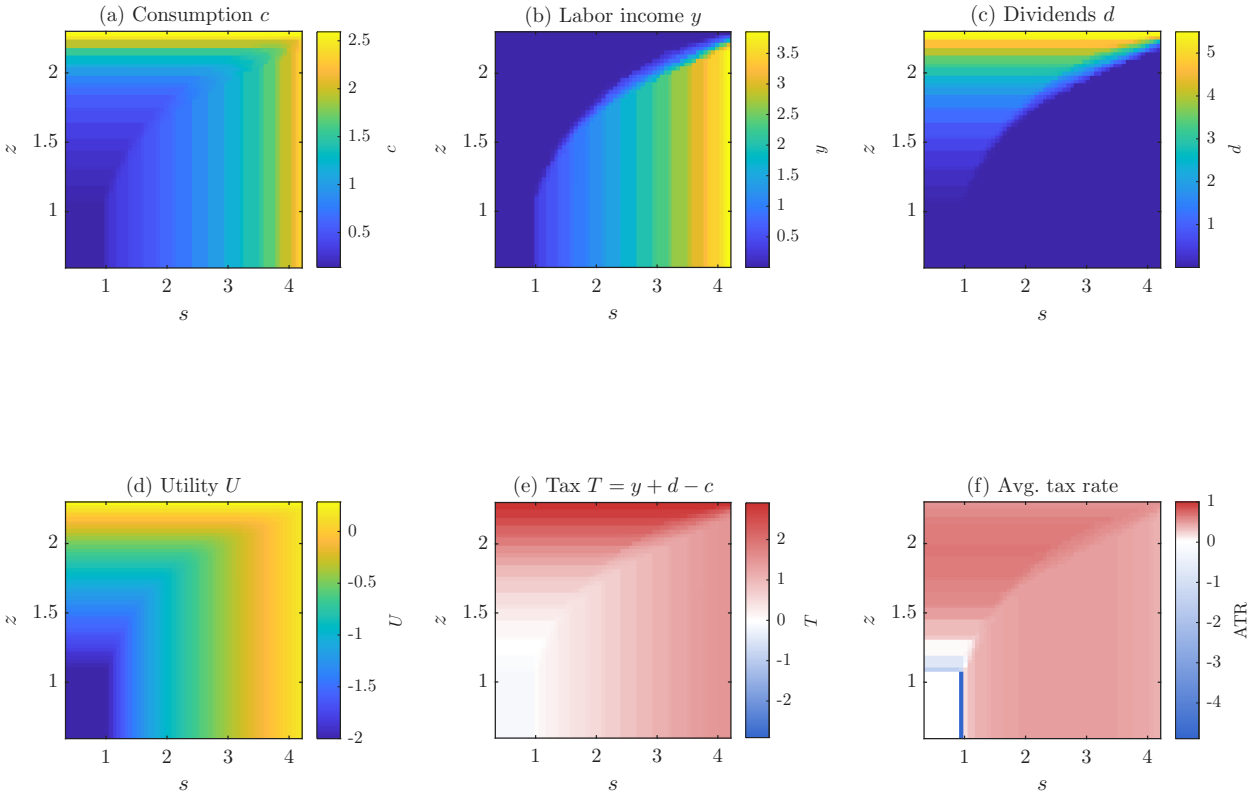


Notes: The figure reports net tax liabilities $T = y + d - c$ across the type space in the stochastic economy, separately by shock realization. The *top row* corresponds to the baseline entrepreneurial-type distribution; the *bottom row* corresponds to an economy with lower dispersion in entrepreneurial ability (a "shrunk" z distribution). The *left column* is the low dividend realization; the *right column* is the high dividend realization. Negative values (blue) indicate net subsidies.

When entrepreneurial types are less dispersed, screening becomes harder and the mass of very productive entrepreneurs falls. Quantitatively, the planner responds by extending tax breaks to lower-z types in the low shock realization. In the baseline economy, given the relative abundance of high z types, taxes are higher in the high-realization state. However, once high z types become scarce, sustaining entry and labor demand requires raising the attractiveness of entrepreneurship across states. This is consistent with the simple two-type logic in Section 4: when it is harder to secure sufficient entrepreneurial participation, the shadow cost of the relevant incentive constraint rises and the optimal policy tilts more resources toward entrepreneurs.

The role of entrepreneurial risk. Entrepreneurial risk affects the optimal policy through two channels: (i) insurance (consumption smoothing across realizations) and (ii) screening (state-contingent distortions that relax ex ante incentive constraints).

Figure 6: Optimal policy and resulting allocations in the economy without entrepreneurial risk



Notes: The figure reports the optimal allocation over the two-dimensional type space (s, z) in the calibrated *deterministic* economy (no entrepreneurial risk) with baseline affiliation ($\rho = 0$). The plotted objects are deterministic allocations. Panels display: (a) consumption c, (b) labor income y, (c) dividends d, (d) indirect utility $U(c) - V(\ell) - V(e)$, (e) net tax liability $T = y + d - c$, and (f) the average tax rate $\text{ATR} = T/(y + d)$ (defined as 0 at $(y, d) = (0, 0)$). Exclusive occupations imply $y > 0 \Rightarrow d = 0$ and $d > 0 \Rightarrow y = 0$.

First, risk induces the planner to *smooth consumption* in the entrepreneur region, raising (lowering) consumption for types that would otherwise be exposed to low (high) realized dividends.

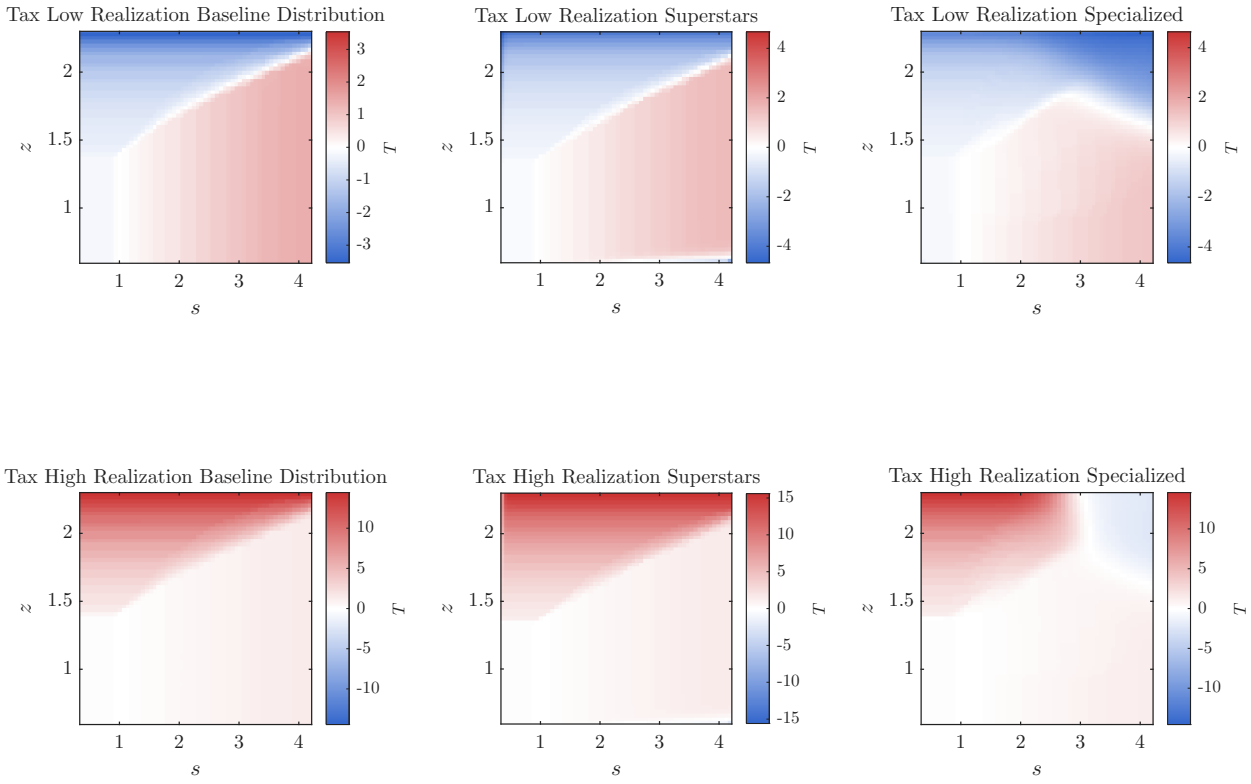
Second, in the stochastic economy, the planner can relax incentive constraints by increasing taxes in high states and decreasing them in low states. As a result, higher z-types have a weaker incentive to underreport their entrepreneurial talent or select into a worker contract. Figure 5 indicates that, on average, the magnitudes of both taxes in high states and subsidies in low states are lower compared to the baseline scenario. This is a composition effect: entry extends to lower-z types, whose dividends and consumption are smaller. For a given type, both magnitudes are instead larger: scarce high-z types are asked to produce higher dividends in the high state and receive higher consumption in both states.

The top panels in Figure 5 show that the net effect of these countervailing planner incentives favors consumption smoothing.

By contrast, in the deterministic economy there is no scope for either insurance or shock-contingent screening. As a result, the optimal policy features larger and more systematic tax breaks for entrepreneurs near the entry boundary. This is shown in Figure 6, which is the deterministic counterpart of Figure 2.

Affiliation between s and z: superstars versus specialized talent. Finally, we vary the correlation between labor skill and entrepreneurial ability. Positive affiliation ("superstars") increases the prevalence of types that are simultaneously high s and high z, while negative affiliation ("specialized talents") generates comparative advantage, with types more likely to be high in one dimension and low in the other. Figure 7 compares the resulting tax schedules for each shock realization.

Figure 7: Affiliation in the joint type distribution and taxes by realization



Notes: The figure reports net tax liabilities $T = y + d - c$ across the type space in the stochastic economy by shock realization, under alternative correlations between log s and log z. The *top row* is the baseline distribution ($\rho = 0$). The *middle row* is the superstar economy (positive affiliation, $\rho > 0$). The *bottom row* is the specialized economy (negative affiliation, $\rho < 0$). The *left column* is the low realization; the *right column* is the high realization. Negative values (blue) indicate net subsidies.

The correlation structure affects both the distribution and the intensity of entrepreneurial rents. The optimal tax policy in the superstar scenario does not change much compared to the baseline case. In both cases, the planner uses sharply state-contingent instruments: in the high realization, taxes are concentrated where realizations are high and incentive constraints are weakest, while subsidies remain targeted to types whose entry would otherwise be most fragile. The only significant

difference is the more favorable treatment of highly specialized agents, i.e. low s and high z entrepreneurs and high s and low z workers.

By contrast, in the specialized economy, many entrepreneurs have relatively low labor skills. For these types, the worker option is less attractive and thus entry into entrepreneurship is easier to sustain. Hence, the planner reduces tax incentives in the entrepreneurial region with low s and high z. In contrast, large subsidies are provided to the relatively scarce superstars.

5.4 Comparing the optimum with the observed tax system

Because the planner is required to finance the same revenue G as the U.S. system, the optimal and observed schedules are directly comparable in levels. Figure 8 plots them side by side. Occupational exclusivity implies that $T(y, d)$ is supported on two rays, so the schedule is displayed as a labor-income schedule $T(y, 0)$ for workers and a dividend schedule $T(0, d)$ for entrepreneurs. The optimal schedule is summarized from the lotteries: for each observable node we average the consumption lottery of every type and shock realization that places mass on it, so the plotted object is the average net tax at each node. Because the optimum conditions consumption on the reported type, this schedule approximates a tax function $T(y, d)$ rather than implementing the optimum exactly.

The comparison is informative about the shape of the optimum, not only its level. Over the range of incomes that the optimal allocation actually populates, the optimum lies well above the observed schedule in both occupations except at the lowest incomes, and its *marginal* rate declines with income: for workers the implied marginal rate falls from about 0.50 at $y \approx 1$ to 0.39 at $y \approx 3.4$, and for entrepreneurs from about 0.94 at $d \approx 2$ to 0.78 at $d \approx 12$. The observed system does the opposite, with marginal rates that are constant or rising. Average rates nonetheless remain high throughout: entrepreneurs face average rates of roughly 81–86% over most of the populated range.

The declining marginal rates are the tax-schedule counterpart of the entry-margin result in Figure 2. Types in this range are those whose occupational choice is still responsive—workers who could plausibly operate a firm, and entrepreneurs who could retreat to paid work—and taxing them harder at the margin would deter the entry that labor-market clearing requires. The same force appears at the bottom of the dividend schedule, where the optimum is negative: close to half of all entrepreneurs operate at very low realized dividends and receive a net subsidy, which is precisely the transfer that sustains entry by marginal operators.

It is useful to ask how the observed functional form would have to change to implement the optimum. Heathcote and Tsujiyama (2021) compare the fully nonlinear Mirrleesian optimum to simple parametric schedules, contrasting a log-linear tax without transfers with an affine tax without progressivity. Ferriere et al. (2023) show that combining the two—a log-linear income tax augmented with a lump-sum transfer—comes systematically close to the Mirrleesian optimum, because the two instruments separate the two objectives: the transfer delivers redistribution through progressive average rates, while low progressivity keeps marginal rates flat and preserves efficiency.

We therefore take that functional form,

$$c = \lambda\, x^{1-\tau} + b, \tag{58}$$

and fit it to the optimal allocation by mass-weighted nonlinear least squares, separately for labor income $x = y$ and dividends $x = d$. [10] Table 4 reports the estimates. Two features stand out. First, the progressivity parameter changes sign: the estimated U.S. schedule has $\tau = 0.067$, whereas the optimum is fitted by $\tau = -0.16$ for labor income and $\tau = -0.69$ for dividends. In the HSV sense the optimal schedule is *regressive*, consistent with the declining marginal rates documented above. Second, redistribution is delivered by the lump-sum term rather than by marginal progressivity: b is positive and sizable in both occupations. The optimum is therefore well described as a large universal transfer financed by high—and mildly regressive—marginal rates, with dividends taxed more heavily than labor income. Both fits are close, with R^2 of 0.998 and 0.995. The optimum therefore lies in the region identified by Ferriere et al. (2023): a sizable demogrant financed by marginal rates that are, if anything, mildly regressive.

The one feature that this family cannot reproduce is the entry subsidy. Fitting (58) to entrepreneurs including the region of near-zero dividends yields a degenerate fit ($R^2 = 0.47$); the estimates in Table 4 are obtained on $d > 1$. The transfer to low-dividend entrepreneurs is thus not representable in the parametric class typically used to describe observed tax systems, even once a demogrant is allowed for.

Table 4: Generalized HSV schedules: estimated U.S. system versus the optimum

	λ	τ	b	R^2
Estimated U.S. schedule, Heathcote et al. (2017)	0.777	+0.067	0 (imposed)	—
Optimum, workers	0.432	−0.157	+0.206	0.998
Optimum, entrepreneurs ($d > 1$)	0.023	−0.693	+0.317	0.995

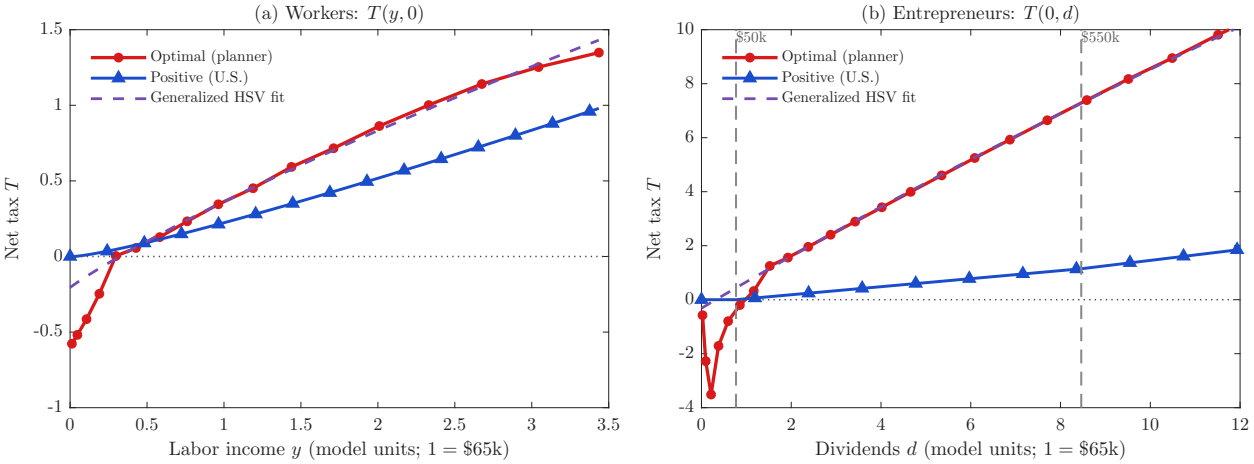
Notes: Mass-weighted nonlinear least-squares fits of $c = \lambda x^{1-\tau} + b$ to the optimal allocation, with $x = y$ for workers and $x = d$ for entrepreneurs. The first row reports the calibrated labor-income schedule of the positive economy, which imposes $b = 0$.

5.5 Positive versus normative economy

Table 5 compares the calibrated economy under the U.S. tax system with the optimum, which raises the same revenue G. The optimum has fewer active agents (the inactive share rises from 36% to 58%), a larger share of entrepreneurs among them (18% against 14%), and a higher wage; GDP falls by about a fifth, but TFP rises by 10% as production concentrates on fewer and more productive entrepreneurs. Figure 9 shows that the optimum nearly halves the Gini of consumption (0.67 to 0.37), while pre-tax income becomes more concentrated (0.71 to 0.81). Welfare rises by the equivalent of a 115% increase in consumption: about two thirds of the gain comes from a more

[10] Setting $b = 0$ recovers the standard Heathcote et al. (2017) schedule, which imposes $T(0) = 0$.

Figure 8: Optimal versus observed tax schedules



Notes: Net tax T against the observables in the calibrated stochastic economy with revenue requirement G. Panel (a): workers, $T(y, 0)$. Panel (b): entrepreneurs, $T(0, d)$. The optimal schedule is recovered from the solution's lotteries: at each observable node (y_a, d_b) we take the mass-weighted average of expected consumption over all types and shock realizations assigning positive probability to that node, $\bar{c}(a, b) = \sum_{i,r} \omega_i p_r \mu_{i,a,b,r} \mathbb{E}[c \mid i, r] / \sum_{i,r} \omega_i p_r \mu_{i,a,b,r}$ with $\mathbb{E}[c \mid i, r] = \sum_k \lambda_{i,k,r} \bar{c}_k$, and set $T = y_a + d_b - \bar{c}(a, b)$. The observed schedule is the calibrated Heathcote et al. (2017) rule for labor income and the statutory three-bracket schedule for dividends, evaluated on the same grid; grey dashed lines mark the bracket thresholds. The dashed purple line is the generalized HSV schedule (58) fitted to the optimum (Table 4). Labor income is shown over the range containing 99.9% of workers' population mass and dividends up to $d = 12$, a range containing 97% of entrepreneurs' population mass.

equal distribution of consumption, notably transfers to inactive agents, who in the positive economy consume only home production c_0, and one third from lower disutility of work.

Table 5: Positive versus normative economy

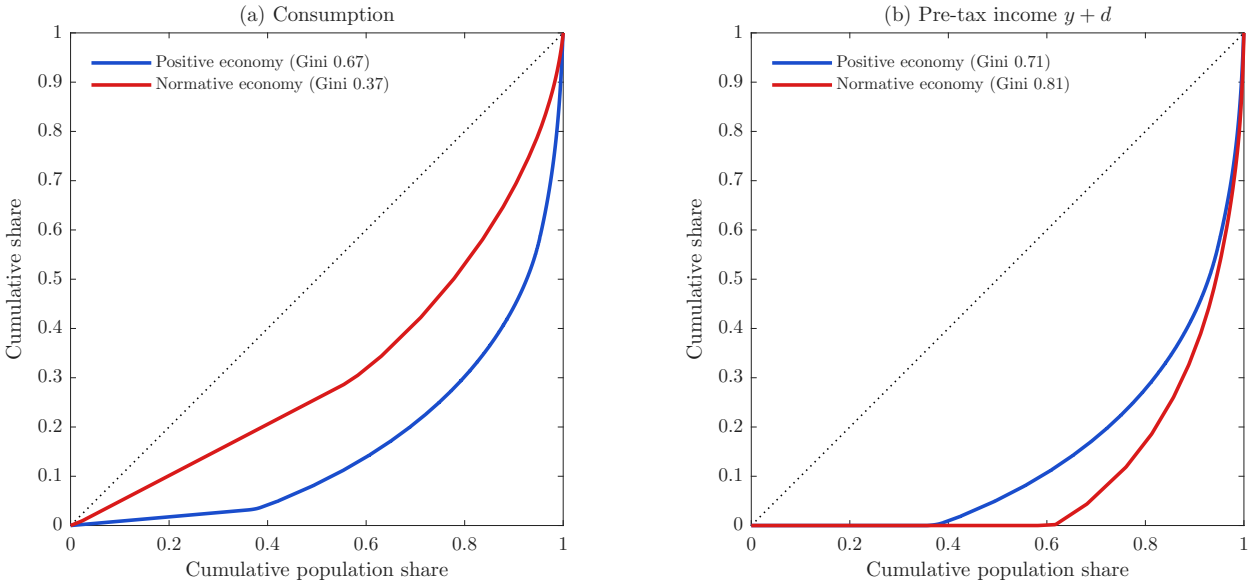
	Positive economy	**Normative economy**
Entrepreneurs (share of active)	0.144	0.176
Workers (share of active)	0.856	0.824
Inactive (share of population)	0.357	0.582
Wage w	0.837	0.933
Consumption (per capita)	0.475	0.357
GDP	0.556	0.453
TFP	1.893	2.089
Gini of consumption	0.670	0.366
Gini of pre-tax income	0.713	0.810
Social welfare	-2.2199	-1.4554

Notes: Both economies use the same types and revenue G, and welfare is the planner's utilitarian objective. TFP is $Y/(E^{\alpha} L^{1-\alpha})$ (entrepreneurial output, effort, and labor). Gini coefficients are of realized consumption and pre-tax income $y + d$. Model units (1 = \$65k).

Figure 9 shows how the optimum reshapes inequality. Before taxes, income is more concentrated in the normative economy: the Gini of pre-tax income rises from 0.71 to 0.81. Two forces drive this. More agents are inactive and earn nothing (58% against 36%), and production concentrates

on fewer and more productive entrepreneurs, so that those who remain active earn more: pre-tax income per active agent is about 25% higher than in the positive economy. After taxes and transfers the ranking reverses, and the Gini of consumption falls from 0.67 to 0.37. The planner separates what agents produce from what they consume: inactive agents, who in the positive economy live on home production alone, receive transfers that raise their consumption to about 0.19, while the consumption of entrepreneurs is compressed. The optimum is therefore more unequal in pre-tax income but substantially more equal in consumption.

Figure 9: Lorenz curves in the positive and normative economies



Notes: Lorenz curves of realized consumption (panel a) and realized pre-tax income $y + d$ (panel b) over the population, in the positive economy (U.S. tax system) and the normative economy (optimal policy). Inactive agents are at c_0 and zero income. Legends report the Gini coefficients.

6 Extension: Self-Employment and the Hiring Margin

The benchmark model treats every entrepreneur as an employer. In the data, this is only true of a minority: roughly one in ten active U.S. workers is self-employed, but only about a third of them hire anyone (Bhandari et al., 2026). The remaining two thirds operate *own-account* — sole proprietors, freelancers, and professionals who sell their own labor without a payroll. The distinction matters for optimal policy, because only employers generate labor demand: it is their entry, not that of own-account operators, that the labor market must clear against. This section extends the model to accommodate both forms of self-employment and shows that the qualitative structure of the optimum survives, while the interpretation of the entry margin sharpens.

The extended technology. Occupational choice is made *ex ante*, as before. What is new is that a self-employed agent decides *ex post*, after the productivity shock ε is realized, whether to

operate alone or to hire. Conditional on paying the entry cost κ_2, the two options are

$$\text{own-account } (n = 0): \quad d_0 = \delta\, s\, \varepsilon\, e^{\alpha} - \kappa_2, \tag{59}$$

$$\text{employer } (n > 0): \quad d_1 = z\, \varepsilon\, e^{\alpha} n^{1-\alpha} - w\, n - \kappa_2. \tag{60}$$

Two features distinguish the branches. Own-account income scales with *labor skill* s: an agent working alone sells the same skill they would sell as an employee. Employer income scales with *entrepreneurial ability* z, the capacity to organize production and manage hired labor.[11] The parameter $\delta \in (0, 1]$ is a solo-production penalty capturing the productivity loss from operating without an organization; it is what makes firms worth forming. The realized dividend is the upper envelope of (59) and (60), so employer status is generated by shock *realizations* within a fixed cross-section of self-employed agents, rather than by permanent heterogeneity. Because the occupational decision is still taken ex ante and only the hiring decision is state-contingent, the planner's problem retains exactly the linear structure of Section 3: the only object that changes is the implied effort at each dividend node, which is now the smaller of the two branch-specific effort requirements, and the implied labor demand, which is positive only when the employer branch is the cheaper way to deliver the dividend.

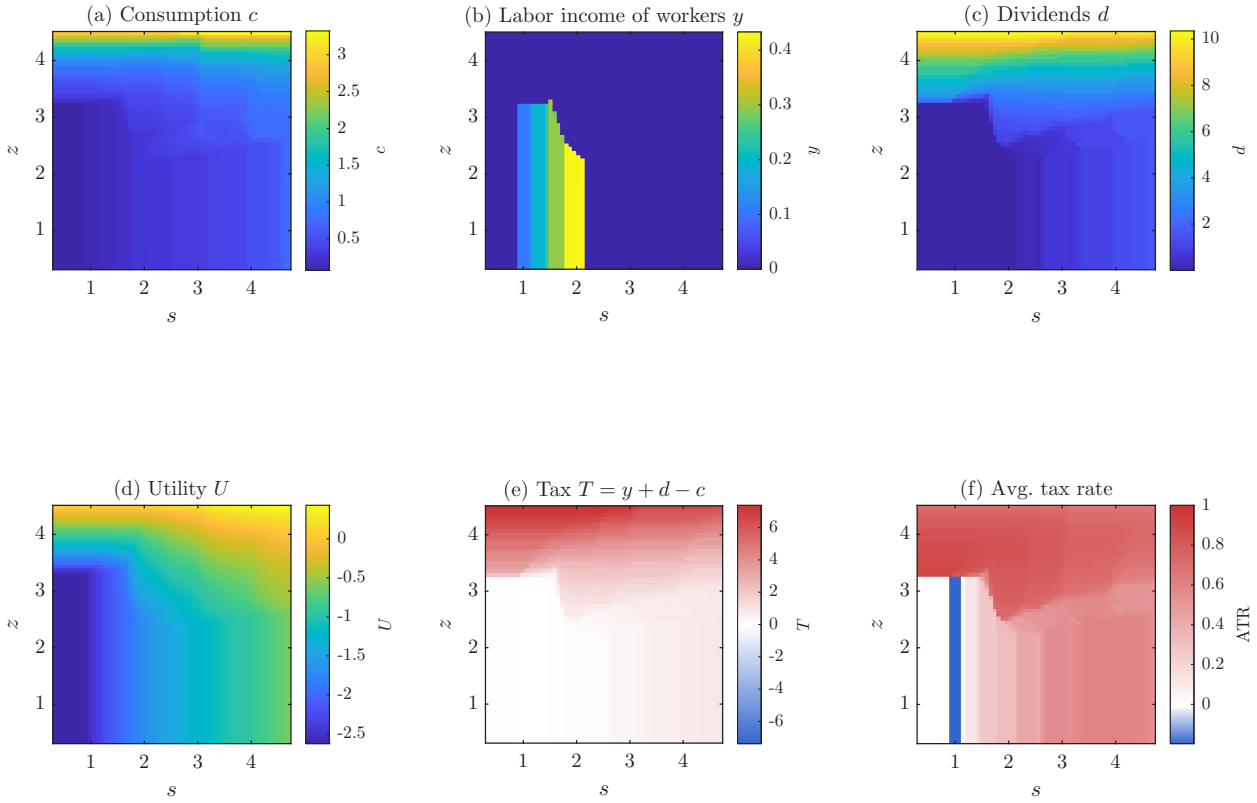
Calibration. The extension adds two parameters, the solo-production penalty δ and the effort-disutility scale η_0, and two occupational targets from Bhandari et al. (2026): the self-employment share among active individuals (10.3%) and the employer share of the self-employed (32.9%). Average hours become a target rather than a normalization, because labor-market clearing ties average firm size to the worker-to-employer ratio and to average effective labor supply. The revenue requirement is pinned exactly as in the benchmark: G equals the revenue raised by the U.S. tax system in the extension's own positive economy, which is 18.1% of output ($G = 0.036$). The full set of eleven parameters and eleven moments is reported in Appendix F, Tables 6 and 7. The calibration matches the occupational block closely — and, importantly, generates the employer share from ex post shock realizations rather than assuming it.

Optimal policy in the extended model. Figure 10 reports the optimal allocation over the type space.

The qualitative structure of Section 5.3 is unchanged: occupational sorting is sharp, and net taxes fall for agents just inside the self-employment region who have strong outside options as workers. What changes is the interpretation. The extensive margin that carries the general-equilibrium externality is now specifically the *employer* margin, since own-account operators hire no one. The planner's entry subsidy is therefore targeted at agents on the verge of operating a firm *with employees*, and the own-account branch acts as an intermediate option that cushions the fall in payoffs for self-employed agents who draw bad shocks. Figure F.6 in Appendix F shows the corresponding

[11] We interpret z as span-of-control managerial ability in the sense of Lucas (1978). Since own-account production does not involve managing employees, we let it depend instead on worker skill s.

Figure 10: Optimal policy with own-account and employer self-employment



Notes: Optimal allocation over the two-dimensional type space (s, z) in the extended economy with an ex post hiring margin, calibrated to the occupational structure of U.S. self-employment. Panels report consumption, labor income, dividends, indirect utility, net tax liability $T = y + d - c$, and the average tax rate. Labor income is shown for types assigned to paid work with probability at least 0.9 and set to zero for the others. As in the benchmark, the extension is solved with a revenue requirement set to the revenue raised by the U.S. tax system in the extension's own calibrated positive economy (18.1% of output, $G = 0.036$).

allocation without entrepreneurial risk: as in the benchmark, removing risk eliminates both the insurance motive and the state-contingent screening instrument, and the optimal policy relies more heavily on systematic entry subsidies.

The extension nonetheless changes how the observed system treats the self-employed. Because own-account income is ordinary self-employment income while employer dividends are business income, the observed tax treatment of a given realized dividend depends on which branch produced it: an own-account operator faces the progressive labor schedule, an employer the flatter statutory business schedule. The observed system thus taxes the two forms of self-employment very differently, whereas the planner conditions only on the realized dividend.

Additional results for the extended model — the comparative statics in the dispersion of entrepreneurial ability and in the affiliation between s and z, the quasi-linear benchmark, and the comparison between the global and local incentive-constrained solutions — are collected in Appendix F.

7 Conclusion

Entrepreneurship is a driver of labor demand and aggregate performance, but also a source of risky and unequal income realizations. In an economy where individuals with concave preferences privately know both their worker skills and entrepreneurial talent, and feasible tax policy can only condition on realized earnings and profits, optimal taxation must balance insurance motives with occupational incentives. Our analysis shows that this interaction generically produces localized tax breaks to entrepreneurs at the extensive margin. Due to the presence of decreasing returns in production, these tax breaks are amplified by general equilibrium effects. A contraction in entrepreneurial entry reduces labor demand. Market clearing then requires a sufficiently large mass of active entrepreneurs, as decreasing returns discourage the planner from further expanding larger firms.

Away from the entrepreneurship entry margin, the optimal tax system is well described by a generalized HSV schedule with a sizable lump-sum transfer and negative progressivity, shifting redistribution from rising marginal tax rates toward transfers. The key departure from this simple form is the localized subsidy to low-dividend entrepreneurs, which the parametric schedule cannot reproduce.

The main methodological contribution of this paper consists of formulating a linear representation of an optimal taxation problem in the presence of two-dimensional private information. We combine this discrete-lottery formulation with constraint generation techniques, and leverage modern large-scale linear programming solvers, such as Gurobi or CPLEX, to solve the multidimensional screening problem by accounting for all incentive constraints numerically. Our analysis provides a tractable way to enforce all incentive compatibility constraints, while allowing for non-smooth primitives and non-convexities that naturally arise in economic applications.

Comparing the full solution with a locally incentive-compatible benchmark reveals that global incentive constraints matter only near the entrepreneurship frontier. The binding non-local constraints involve types close to switching occupations and amplify localized tax breaks for marginal entrepreneurs. In an extension that distinguishes own-account self-employment from employer entrepreneurship, the same mechanism targets the employer margin, while non-local incentive constraints become substantially more important as own-account production becomes more attractive.

References

Ales, L., Bellofatto, A. A., and Wang, J. J. (2017). Taxing atlas: Executive compensation, firm size, and their impact on optimal top income tax rates. *Review of Economic Dynamics*, 26:62–90.

Armstrong, M. (1996). Multiproduct nonlinear pricing. *Econometrica*, 64(1):51–75.

Bergstrom, K. and Dodds, W. (2021). Optimal taxation with multiple dimensions of heterogeneity. *Journal of Public Economics*, 200:104442.

Bhandari, A., Kass, T., May, T. J., McGrattan, E. R., and Schulz, E. (2026). On the nature of entrepreneurship. *Journal of Political Economy*, 134(6):1609–1665.

Bierbrauer, F., Boyer, P. C., Peichl, A., and Weishaar, D. (2024). Pareto-improvements, welfare trade-offs and the taxation of couples. Technical report, ECONtribute Discussion Paper.

Boerma, J., Tsyvinski, A., and Zimin, A. (2022). Bunching and taxing multidimensional skills. Working Paper 30015, National Bureau of Economic Research.

Bureau of Labor Statistics (2024a). Current population survey, table a-9: Selected employment indicators. https://www.bls.gov/cps/cpsaat09.htm. Self-employment (unincorporated) as share of total employment.

Bureau of Labor Statistics (2024b). Employment situation summary, table a-1: Employment status of the civilian population. https://www.bls.gov/news.release/empsit.t01.htm. Labor force participation rate for civilian population 16 years and over.

Carlier, G., Dupuis, X., Rochet, J.-C., and Thanassoulis, J. (2024). A general solution to the quasi linear screening problem. *Journal of Mathematical Economics*, 114:103025.

Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*, 80(3):969–1018.

DeBacker, J., Heim, B. T., Panousi, V., Ramnath, S., and Vidangos, I. (2021). The distribution of business income: Evidence from pass-through entities. *National Tax Journal*, 74(4):911–938.

Dodds, W. (2024). Solving multidimensional screening problems using a generalized single crossing property. *Economic Theory*, 77(4):1025–1084.

Doepke, M. and Townsend, R. M. (2006). Dynamic mechanism design with hidden income and hidden actions. *Journal of Economic Theory*, 126(1):235–285.

Ferriere, A., Grübener, P., Navarro, G., and Vardishvili, O. (2023). On the optimal design of transfers and income tax progressivity. *Journal of Political Economy Macroeconomics*, 1(2):276–333.

Golosov, M. and Krasikov, I. (2025). The optimal taxation of couples*. *The Quarterly Journal of Economics*, 140(3):2163–2211.

Heathcote, J., Storesletten, K., and Violante, G. L. (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review*, 104(7):2075–2126.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics*, 132(4):1693–1754.

Heathcote, J. and Tsujiyama, H. (2021). Optimal income taxation: Mirrlees meets Ramsey. *Journal of Political Economy*, 129(11):3141–3184.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The allocation of talent and U.S. economic growth. *Econometrica*, 87(5):1439–1474.

Khan, A. and Thomas, J. K. (2013). Credit shocks and aggregate fluctuations in an economy with production heterogeneity. *Journal of Political Economy*, 121(6):1055–1107.

Kleven, H. J., Kreiner, C. T., and Saez, E. (2009). The optimal income taxation of couples. *Econometrica*, 77(2):537–560.

Lopomo, G., Persico, N., and Villa, A. T. (2023). Optimal procurement with quality concerns. *American Economic Review*, 113(6):1505–29.

Lucas, R. E. (1978). On the size distribution of business firms. *The Bell Journal of Economics*, 9(2):508–523.

Manelli, A. M. and Vincent, D. R. (2007). Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185.

Myerson, R. B. (2013). *Game Theory: Analysis of Conflict*. Harvard University Press.

Nikolov, B., Schmid, L., and Steri, R. (2021). The sources of financing constraints. *Journal of Financial Economics*, 139(2):478–501.

Phelan, C. and Townsend, R. M. (1991). Computing multi-period, information-constrained optima. *The Review of Economic Studies*, 58(5):853–881.

Piketty, T., Saez, E., and Stantcheva, S. (2014). Optimal taxation of top labor incomes: A tale of three elasticities. *American Economic Journal: Economic Policy*, 6(1):230–271.

Piketty, T., Saez, E., and Zucman, G. (2022). Distributional national accounts. Updated data tables, https://gabriel-zucman.eu/usdina/.

Prescott, E. C. and Townsend, R. M. (1984). Pareto optima and competitive equilibria with adverse selection and moral hazard. *Econometrica*, pages 21–45.

Rochet, J.-C. and Choné, P. (1998). Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(4):783–826.

Rothschild, C. and Scheuer, F. (2013). Redistributive taxation in the Roy model. *The Quarterly Journal of Economics*, 128(2):623–668.

Rothschild, C. and Scheuer, F. (2014). A theory of income taxation under multidimensional skill heterogeneity. Working Paper 19822, National Bureau of Economic Research.

Scheuer, F. (2014). Entrepreneurial taxation with endogenous entry. *American Economic Journal: Economic Policy*, 6(2):126–63.

Shen, C., Kübler, F., Yang, Y., and Zhou, Z. (2026). A Lagrangian approach to optimal randomization. Working paper, first version April 2025.

Shourideh, A. (2012). Optimal taxation of wealthy individuals. Manuscript.

Spiritus, K., Lehmann, E., Renes, S., and Zoutman, F. T. (2025). Optimal taxation with multiple incomes and types. *Theoretical Economics*, 20(2):543–581.

U.S. Census Bureau (2021). Statistics of U.S. businesses (SUSB) 2021 annual data tables. https://www.census.gov/programs-surveys/susb.html. Employment size of employer firms.

A Mathematical Appendix

A.1 Comparative statics of the endogenous-wage multiplier λ

Consider the scalar equation

$$F(\lambda;\omega_H,\alpha)=0, \qquad \lambda\in(-\omega_H,\omega_L), \qquad \omega_L:=1-\omega_H, \tag{61}$$

where

$$F(\lambda;\omega_H,\alpha):=\log\left(\frac{\omega_L\cdot(\omega_H+\lambda)}{\omega_H\cdot(\omega_L-\lambda)}\right)-\frac{1}{2}\left[\frac{\alpha}{\omega_H+\lambda}-\frac{1-\alpha}{\omega_L-\lambda}\right]. \tag{62}$$

Throughout we maintain $\omega_H\in(0,1)$, $\alpha\in(0,1)$, and $\lambda\in(-\omega_H,\omega_L)$ so all denominators are positive.

A.1.1 Step 1: derivatives of F

Write $A:=\omega_H+\lambda$ and $B:=\omega_L-\lambda$. Then $A>0$ and $B>0$ on the admissible set. Differentiating (62) with respect to λ yields

$$F_\lambda(\lambda;\omega_H,\alpha)=\frac{1}{\omega_H+\lambda}+\frac{1}{\omega_L-\lambda}+\frac{1}{2}\left[\frac{\alpha}{(\omega_H+\lambda)^2}+\frac{1-\alpha}{(\omega_L-\lambda)^2}\right]. \tag{63}$$

Hence,

$$F_\lambda(\lambda;\omega_H,\alpha)>0 \qquad \text{for all } \lambda\in(-\omega_H,\omega_L). \tag{64}$$

Differentiating (62) with respect to α gives

$$F_\alpha(\lambda;\omega_H,\alpha)=-\frac{1}{2}\left[\frac{1}{\omega_H+\lambda}+\frac{1}{\omega_L-\lambda}\right]<0. \tag{65}$$

To differentiate with respect to ω_H taking into account $\omega_L=1-\omega_H$, it is convenient to first compute the partials treating ω_H and ω_L as independent, and then apply the chain rule. Holding

λ fixed,

$$F_{\omega_H} = \frac{1}{\omega_H + \lambda} - \frac{1}{\omega_H} + \frac{\alpha}{2(\omega_H + \lambda)^2}, \tag{66}$$

$$F_{\omega_L} = \frac{1}{\omega_L} - \frac{1}{\omega_L - \lambda} - \frac{1 - \alpha}{2(\omega_L - \lambda)^2}. \tag{67}$$

Since $\omega_L = 1 - \omega_H$, we have $d\omega_L/d\omega_H = -1$, so the total derivative of F with respect to ω_H is

$$\frac{d}{d\omega_H} F(\lambda; \omega_H, \alpha) = F_{\omega_H} - F_{\omega_L} = \underbrace{\left[\frac{1}{\omega_H + \lambda} - \frac{1}{\omega_H} - \frac{1}{\omega_L} + \frac{1}{\omega_L - \lambda}\right]}_{\text{log term}} + \underbrace{\left[\frac{\alpha}{2(\omega_H + \lambda)^2} + \frac{1 - \alpha}{2(\omega_L - \lambda)^2}\right]}_{\text{RHS term}}. \tag{68}$$

A.1.2 Step 2: first derivatives via the Implicit Function Theorem

Let $\lambda^* = \lambda^*(\omega_H, \alpha)$ solve (61). By the implicit function theorem,

$$\frac{\partial \lambda^*}{\partial x} = -\left.\frac{F_x}{F_\lambda}\right|_{(\lambda^*, \omega_H, \alpha)} \qquad \text{for } x \in \{\alpha, \omega_H\}. \tag{69}$$

Derivative with respect to α. Combining (64)–(65) with (69) gives

$$\boxed{\frac{\partial \lambda^*}{\partial \alpha} = -\frac{F_\alpha}{F_\lambda} = \frac{\frac{1}{2}\left[\frac{1}{\omega_H + \lambda^*} + \frac{1}{\omega_L - \lambda^*}\right]}{\frac{1}{\omega_H + \lambda^*} + \frac{1}{\omega_L - \lambda^*} + \frac{1}{2}\left[\frac{\alpha}{(\omega_H + \lambda^*)^2} + \frac{1-\alpha}{(\omega_L - \lambda^*)^2}\right]} > 0.} \tag{70}$$

Thus, *the informational wedge indexed by λ^* is increasing in α.*

A.2 Welfare derivative with respect to z_H (endogenous wage).

Consider equilibrium welfare

$$W(z_H) = \omega_L \cdot \left[\log c^W - \tfrac{1}{2} \cdot \left(\frac{y^W}{w}\right)^2\right] + \omega_H \cdot \left[\log c^E - \tfrac{1}{2} \cdot \left(e^*(d^E; z_H, w)\right)^2\right],$$

evaluated at the closed-form allocation in the regime where (IC_H) binds and (IC_L) is slack.

Step 1: λ^* does not depend on z_H. The scalar equation pinning down λ^* is (48), which depends only on $(\omega_H, \omega_L, \alpha)$. Hence, within this regime,

$$\frac{d\lambda^*}{dz_H} = 0.$$

Step 2: the disutility terms are z_H-invariant at the optimum. Using the closed-form

simplifications,

$$\left(\frac{y^W}{w}\right)^2 = \frac{1-\alpha}{\omega_L - \lambda^*}, \qquad \left(e^*(d^E; z_H, w)\right)^2 = \frac{\alpha}{\omega_H + \lambda^*},$$

so both disutility terms are constant with respect to z_H (since λ^* is constant).

Step 3: c^W and c^E scale linearly with z_H. From the wage formula (49),

$$w(\lambda^*) = z_H \cdot \Xi(\lambda^*, \omega_H, \omega_L, \alpha),$$

for some positive factor $\Xi(\cdot)$ independent of z_H. Then (50) implies

$$d^E(\lambda^*) = \frac{w(\lambda^*)}{\chi} \cdot \sqrt{\frac{1-\alpha}{\omega_L - \lambda^*}} = z_H \cdot \kappa(\lambda^*, \omega_H, \omega_L, \alpha),$$

so d^E is linear in z_H. Finally, (51) gives

$$c^E(\lambda^*) = (\omega_H + \lambda^*) \cdot \frac{d^E(\lambda^*)}{\alpha} = z_H \cdot A_E, \qquad c^W(\lambda^*) = \frac{\omega_L - \lambda^*}{\omega_L} \cdot \omega_H \cdot \frac{d^E(\lambda^*)}{\alpha} = z_H \cdot A_W,$$

with constants $A_E, A_W > 0$ independent of z_H.

Step 4: differentiate welfare. Because only the log terms depend on z_H and c^E, c^W are linear in z_H,

$$\frac{d}{dz_H} \log c^E(\lambda^*) = \frac{1}{z_H}, \qquad \frac{d}{dz_H} \log c^W(\lambda^*) = \frac{1}{z_H}.$$

Therefore,

$$\boxed{\frac{dW}{dz_H} = \omega_L \cdot \frac{1}{z_H} + \omega_H \cdot \frac{1}{z_H} = \frac{1}{z_H}.} \tag{71}$$

B Computational Appendix

For any given wage w we solve a linear problem using constraint generation.

B.1 Constraint Generation Algorithm

B.1.1 Notation

$N = N_1 \times N_2$	Number of types on the 2D grid
$\theta_i = (\theta_{1,i}, \theta_{2,i})$	Type i's private information (worker skill, entrepreneur ability)
ω_i	Population weight of type i
W_i	Social welfare weight of type i
K	Number of consumption grid points, $\{c_k\}_{k=1}^{K}$
A	Number of labor income grid points, $\{y_{1,a}\}_{a=1}^{A}$; for notational ease, $A \equiv A_y$
B	Number of entrepreneurial dividend grid points, $\{y_{2,b}\}_{b=1}^{B}$; for notational ease, $B \equiv B_d$
R	Number of shock states, $\{\varepsilon_r\}_{r=1}^{R}$ with probabilities $\{p_r\}_{r=1}^{R}$
$U(c_k)$	CRRA consumption utility at grid point k
$D_{i,a,b,r}$	Total effort disutility for type i at income pair (a, b) in state r
$\lambda_{i,k,r}$	Probability that type i receives consumption c_k in state r
$\mu_{i,a,b,r}$	Probability that type i produces income $(y_{1,a}, y_{2,b})$ in state r
$\mathcal{P}$	Active set of incentive-compatibility (IC) pairs

The disutility coefficients are:

$$D_{i,a,b,r} = V\left(\frac{y_{1,a} + \kappa_1 \cdot \mathbf{1}_{\{y_{1,a}>0\}}}{\theta_{1,i}\, w}\right) + V\big(e^*_{i,b,r}\big)\,,$$

where $V(x) = \frac{x^{1+\psi}}{1+\psi}$ is the power disutility of effort, and $e^*_{i,b,r}$ is the minimum effort required for type i to generate dividend $y_{2,b}$ in shock state r.

B.1.2 Relaxed Master Problem: LP($\mathcal{P}$)

Given an active set of IC pairs $\mathcal{P} \subseteq \{1, \ldots, N\}^2 \setminus \{(i, i)\}$, solve:

$$\max_{\lambda,\mu} \quad \sum_{i=1}^{N} \omega_i \, W_i \, \mathrm{EU}(i) \tag{72}$$

$$\text{s.t.} \quad \sum_{k=1}^{K} \lambda_{i,k,r} = 1, \qquad \forall\, i = 1, \ldots, N,\ r = 1, \ldots, R \tag{73}$$

$$\sum_{(a,b)\in\mathcal{A}} \mu_{i,a,b,r} = 1, \qquad \forall\, i = 1, \ldots, N,\ r = 1, \ldots, R \tag{74}$$

$$\mathrm{EU}(i) \geq \mathrm{DevUtil}(i, j), \qquad \forall\, (i, j) \in \mathcal{P} \tag{75}$$

$$\mathrm{EU}(i) \geq \bar{u}_i, \qquad \forall\, i = 1, \ldots, N \tag{76}$$

$$\sum_{(a,b)\in\mathcal{W}} \mu_{i,a,b,r} = \sum_{(a,b)\in\mathcal{W}} \mu_{i,a,b,1}, \qquad \forall\, i,\ r = 2, \ldots, R \tag{77}$$

$$\sum_{(a,b)\in\mathcal{I}} \mu_{i,a,b,r} = \sum_{(a,b)\in\mathcal{I}} \mu_{i,a,b,1}, \qquad \forall\, i,\ r = 2, \ldots, R \tag{78}$$

$$\sum_{i=1}^{N} \omega_i \sum_{r=1}^{R} p_r \sum_{k=1}^{K} \lambda_{i,k,r} \, c_k \ \leq\ \sum_{i=1}^{N} \omega_i \sum_{r=1}^{R} p_r \sum_{(a,b)\in\mathcal{A}} \mu_{i,a,b,r} (y_{1,a} + y_{2,b}) \ -\ G \tag{79}$$

$$\lambda_{i,k,r} \geq 0, \quad \mu_{i,a,b,r} \geq 0, \qquad \forall\, i, k, a, b, r \tag{80}$$

where the set of allowed income pairs is $\mathcal{A} = \{(a, b) : \text{not}(y_{1,a} > 0 \text{ and } y_{2,b} > 0)\}$ (occupational exclusivity), and the expected utilities are:

$$\mathrm{EU}(i) = \sum_{r=1}^{R} p_r \left[\sum_{k=1}^{K} \lambda_{i,k,r} \, U(c_k) \ -\ \sum_{(a,b)\in\mathcal{A}} \mu_{i,a,b,r} \, D_{i,a,b,r} \right], \tag{81}$$

$$\mathrm{DevUtil}(i, j) = \sum_{r=1}^{R} p_r \left[\sum_{k=1}^{K} \lambda_{j,k,r} \, U(c_k) \ -\ \sum_{(a,b)\in\mathcal{A}} \mu_{j,a,b,r} \, D_{i,a,b,r} \right]. \tag{82}$$

Note: $\mathrm{DevUtil}(i, j)$ uses type j's lottery (λ_j, μ_j) but type i's disutility D_i; it represents the payoff type i would obtain by mimicking type j.

Optionally, a labor market clearing constraint is added (general equilibrium):

$$\sum_{i=1}^{N} \omega_i \sum_{r=1}^{R} p_r \sum_{(a,b)\in\mathcal{A}} \mu_{i,a,b,r}\, L_a^s \;=\; \sum_{i=1}^{N} \omega_i \sum_{r=1}^{R} p_r \sum_{(a,b)\in\mathcal{A}} \mu_{i,a,b,r}\, L_{i,b,r}^d, \tag{83}$$

where $L_a^s = (y_{1,a} + \kappa_1 \mathbf{1}[y_{1,a} > 0])/w$ is effective labor supply, which carries no type index because skill cancels in efficiency units, and $L_{i,b,r}^d = n^*(y_{2,b}; w)$ is equilibrium labor demand from type i's firm at dividend $y_{2,b}$ in state r.

B.1.3 Local IC Pair Initialization

Types are indexed on an $N_1 \times N_2$ grid where $i = (i_2 - 1)N_1 + i_1$ for $i_1 = 1, \ldots, N_1$ and $i_2 = 1, \ldots, N_2$. The local neighbor pairs are:

Algorithm 1 BuildLocalICPairs(N, N_1, N_2)

Input: Grid dimensions N_1, N_2 with $N = N_1 \times N_2$

Output: Set of adjacent-neighbor IC pairs $\mathcal{P}_0$

$\mathcal{P}_0 \leftarrow \emptyset$
for $i = 1, \ldots, N$ **do**
  if $i \bmod N_1 \neq 0$ **then** ▷ Neighbor to the right in θ_1
    $\mathcal{P}_0 \leftarrow \mathcal{P}_0 \cup \{(i, i+1)\}$
  end if
  if $i \bmod N_1 \neq 1$ **then** ▷ Neighbor to the left in θ_1
    $\mathcal{P}_0 \leftarrow \mathcal{P}_0 \cup \{(i, i-1)\}$
  end if
  if $i \leq N - N_1$ **then** ▷ Neighbor above in θ_2
    $\mathcal{P}_0 \leftarrow \mathcal{P}_0 \cup \{(i, i+N_1)\}$
  end if
  if $i > N_1$ **then** ▷ Neighbor below in θ_2
    $\mathcal{P}_0 \leftarrow \mathcal{P}_0 \cup \{(i, i-N_1)\}$
  end if
end for
return $\mathcal{P}_0$

B.1.4 IC Violation Computation

Algorithm 2 ComputeICViolations($\lambda^*, \mu^*, U, D, p, N$)

Input: Optimal marginals λ^*, μ^* from the LP; utilities $U(c_k)$; disutilities $D_{i,a,b,r}$; shock probabilities p_r

Output: Violation matrix $\mathbf{V} \in \mathcal{R}^{N \times N}$, maximum violation, list of violated pairs

for $i = 1, \ldots, N$ **do** ▷ Compute expected consumption utility
  for $r = 1, \ldots, R$ **do**
    $\text{ConsUtil}(i, r) \leftarrow \sum_{k=1}^{K} \lambda^*_{i,k,r} \, U(c_k)$
  end for
  $\text{ECons}(i) \leftarrow \sum_{r=1}^{R} p_r \cdot \text{ConsUtil}(i, r)$
end for

for $i = 1, \ldots, N$ **do** ▷ Compute truthful disutility
  $\text{TruthDis}(i) \leftarrow \sum_{r=1}^{R} p_r \sum_{(a,b) \in \mathcal{A}} \mu^*_{i,a,b,r} \, D_{i,a,b,r}$
end for

for $i = 1, \ldots, N$ **do** ▷ Truthful expected utility
  $\text{EU}(i) \leftarrow \text{ECons}(i) - \text{TruthDis}(i)$
end for

for $i = 1, \ldots, N$ **do** ▷ Compute cross-disutility matrix
  for $j = 1, \ldots, N$ **do**
    $\text{DevDis}(i, j) \leftarrow \sum_{r=1}^{R} p_r \sum_{(a,b) \in \mathcal{A}} \mu^*_{j,a,b,r} \, D_{i,a,b,r}$
  end for
end for

for $i = 1, \ldots, N$ **do** ▷ Compute violation matrix
  for $j = 1, \ldots, N, \; j \neq i$ **do**
    $\mathbf{V}(i, j) \leftarrow \big[\text{ECons}(j) - \text{DevDis}(i, j)\big] - \text{EU}(i)$ ▷ $= \text{DevUtil}(i, j) - \text{EU}(i)$
  end for
end for

$v^{\max} \leftarrow \max_{i \neq j} \mathbf{V}(i, j)$
$\text{ViolPairs} \leftarrow \{(i, j) : \mathbf{V}(i, j) > \tau, \; i \neq j\}$
return $\mathbf{V}$, $v^{\max}$, ViolPairs

B.1.5 Main Algorithm: Constraint Generation

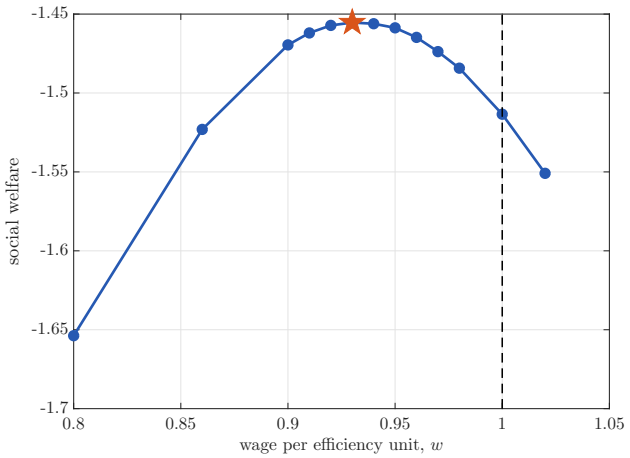
See algorithm 3.

B.1.6 Outer loop: choosing the wage

Algorithm 3 solves the planner's problem at a *given* wage. Because labor-market clearing is homogeneous of degree -1 in w and therefore cannot select among wages, w is chosen by the planner and the algorithm above is nested inside a search over it. The primitives $(z, \kappa_1, \kappa_2, c_0, G)$ and the grids $\{c_k\}, \{y_{1,a}\}, \{y_{2,b}\}$ are held fixed in levels across wages, so no rescaling of welfare is required and $\mathrm{SW}(w)$ is directly comparable from one wage to the next. Only the implied-effort mappings and the labor-market objects are rebuilt at each w.

Each grid point is a separate second-best problem, solved to full global incentive compatibility, so the comparison is across optima rather than across relaxations of a common problem. The wage grid is refined around the maximum until the bracket is narrower than the spacing at which welfare differences fall below solver tolerance.

Figure B.1: Welfare as a function of the wage



Notes: Social welfare at the second-best optimum, computed separately at each wage on a grid with all other primitives held fixed in levels. Each point solves the linear program to full global incentive compatibility. The star marks the optimum w^*; the dashed line marks $w = 1$.

B.1.7 Remarks

1. **Initialization.** The algorithm starts with only $|\mathcal{P}_0| \approx 4N$ local IC constraints (adjacent neighbors on the 2D type grid), rather than the full $N^2 - N$ pairwise constraints. For $N = 3{,}600$ (a 60×60 grid), this reduces the initial constraint count from 12,956,400 to 14,160 directed adjacent-neighbor IC constraints.

2. **Separation oracle.** At each iteration, the full $N \times N$ violation matrix is computed in MATLAB via vectorized matrix operations (the cross-disutility matrix $\mathrm{DevDis}(i,j) = \sum_r p_r \, \mathbf{D}_r \, \mathbf{M}_r^\top$

Algorithm 3 Constraint Generation for the Mirrlees Optimal Tax LP

Input: Type grid $\{\theta_i\}_{i=1}^{N}$ on $N_1 \times N_2$; weights ω_i, W_i; grids $\{c_k\}$, $\{y_{1,a}\}$, $\{y_{2,b}\}$; utility $U(\cdot)$; disutility $D_{i,a,b,r}$; shocks $(\varepsilon_r, p_r)_{r=1}^{R}$; tolerance $\tau > 0$; max iterations T; max pairs per iteration M

Output: Optimal allocations (λ^*, μ^*); social welfare SW*; active IC pair set $\mathcal{P}^*$

Phase 1: Initialization

$\mathcal{P} \leftarrow$ BuildLocalICPairs(N, N_1, N_2) ▷ $|\mathcal{P}| \approx 4N$ local pairs

Phase 2: Constraint Generation Loop

for $t = 1, 2, \ldots, T$ **do**

  Step 2a: Solve Relaxed Master Problem

  $(\lambda^*, \mu^*, \text{SW}) \leftarrow$ Solve LP$(\mathcal{P})$ via AMPL/CPLEX ▷ Eqs. (72)–(80)

  if LP infeasible or solver failure **then**

    return failure

  end if

  Step 2b: Separation Oracle — Check All N^2 IC Constraints

  $(\mathbf{V}, v^{\max}, \text{ViolPairs}) \leftarrow$ ComputeICViolations$(\lambda^*, \mu^*, U, D, p, N)$

  Step 2c: Filter Active-Set Artifacts

  $$\mathbf{V}_{\text{out}}(i,j) \leftarrow \begin{cases} \mathbf{V}(i,j) & \text{if } (i,j) \notin \mathcal{P} \text{ and } i \neq j \\ -\infty & \text{otherwise} \end{cases}$$

  $v_{\text{out}}^{\max} \leftarrow \max_{i,j} \mathbf{V}_{\text{out}}(i,j)$

  NewViol $\leftarrow \{(i,j) \in \text{ViolPairs} : (i,j) \notin \mathcal{P}\}$

  Step 2d: Sort by Violation Magnitude

  Sort NewViol in decreasing order of $\mathbf{V}(i,j)$

  Step 2e: Convergence Check

  if $v_{\text{out}}^{\max} \leq \tau$ **then**

    return $(\lambda^*, \mu^*, \text{SW}, \mathcal{P})$ ▷ Converged: all IC constraints satisfied

  end if

  Step 2f: Augment Active Set

  $n_{\text{add}} \leftarrow \min(|\text{NewViol}|, M)$

  $\mathcal{P} \leftarrow \mathcal{P} \cup \text{NewViol}[1 : n_{\text{add}}]$ ▷ Add most-violated pairs

  Remove duplicates from $\mathcal{P}$

end for

return $(\lambda^*, \mu^*, \text{SW}, \mathcal{P})$ with warning: did not converge in T iterations

Algorithm 4 Optimal wage by grid search

Input: Primitives $(z, \kappa_1, \kappa_2, c_0, G)$ in levels; wage grid $\mathcal{G}_w = \{w_1 < \cdots < w_m\}$; tolerance τ, limits T, M

Output: Optimal wage w^*, allocation (λ^*, μ^*), welfare SW^*

for $w \in \mathcal{G}_w$ **do**

  Step 1: Rebuild the implied-effort mappings at w

  $\ell_{i,a} \leftarrow \big(y_{1,a} + \kappa_1 \mathbf{1}\{y_{1,a} > 0\}\big)/(s_i w)$

  $n_b \leftarrow \frac{1-\alpha}{\alpha}\,(y_{2,b} + \kappa_2)/w$

  $e_{i,b,r} \leftarrow \Big(\dfrac{y_{2,b} + \kappa_2 + w\, n_b}{z_i \varepsilon_r\, n_b^{1-\alpha}}\Big)^{1/\alpha}$

  $D_{i,a,b,r} \leftarrow V(\ell_{i,a}) + V(e_{i,b,r})$

  Step 2: Rebuild the labor-market objects at w

  $L_a^s \leftarrow \big(y_{1,a} + \kappa_1 \mathbf{1}\{y_{1,a} > 0\}\big)/w, \quad L_b^d \leftarrow n_b$

  Step 3: Solve the planner's problem at this wage

  $(\lambda_w, \mu_w, \mathrm{SW}(w)) \leftarrow \textsc{ConstraintGeneration}(D, L^s, L^d; \tau, T, M)$

  ▷ Algorithm 3: full global incentive compatibility at w

end for

$w^* \leftarrow \arg\max_{w \in \mathcal{G}_w} \mathrm{SW}(w)$

return $(w^*, \lambda_{w^*}, \mu_{w^*}, \mathrm{SW}(w^*))$

is computed as a matrix product), making the separation step efficient despite its $O(N^2)$ scaling.

3. **Convergence criterion.** Only violations *outside* the current active set $\mathcal{P}$ determine convergence. Violations on active-set constraints (which should be zero in exact arithmetic) may be nonzero due to LP solver numerical tolerances and are treated as artifacts.

4. **Capping.** At most M new pairs are added per iteration (sorted by violation magnitude in descending order) to prevent the LP size from growing too rapidly.

5. **Marginal lotteries.** The LP uses *marginal* lotteries $\lambda_{i,k,r}$ (consumption) and $\mu_{i,a,b,r}$ (income-dividend pairs) rather than the joint lottery $x_{i,k,a,b,r}$. This is valid because the IC constraints and the objective depend on consumption and income-dividend pairs separately through their marginals, and reduces the number of decision variables from $O(N \cdot K \cdot A_y \cdot B_d \cdot R)$ to $O(N \cdot (K + |\mathcal{J}|) \cdot R)$.

6. **Occupational exclusivity.** The set of admissible income-dividend pairs $\mathcal{J}$ excludes combinations where both $\bar{y}_a > 0$ and $\bar{d}_b > 0$, enforcing that agents either work or run a firm but not both simultaneously.

7. **Solver.** The LP is formulated as an AMPL model and dispatched to CPLEX (barrier method, `lpmethod=4`) or Gurobi (`method=2, crossover=0`) for solution. The basic solution shown in

Figure D.1 is obtained by re-solving the benchmark with crossover (`crossover=1`); its welfare differs from the barrier solution by less than 10^{-4}.

C Data Appendix: Sources for Calibration Targets

This appendix lists every empirical target used in the paper, its numerical value, and its source. All monetary quantities are expressed in the units set by the Fréchet scale normalization of Section 5.1, under which the untruncated ability distributions have mean one; the wage is not normalized, and is an equilibrium price in the positive economy and an instrument in the planner's problem. Dollar amounts are mapped into model units so that one model unit corresponds to \$65,000.

Parameters set externally. Risk aversion is set to $\gamma = 1$ (logarithmic utility). The inverse Frisch elasticity is set to $\psi = 2$, implying a Frisch elasticity of 0.5, the central value in the meta-analysis of Chetty (2012). The correlation between labor skill and entrepreneurial ability is set to $\rho = 0$ in the baseline; comparative statics with respect to ρ are reported in Section 5.3.

Benchmark calibration targets. Table 1 reports the targets of the benchmark economy. Labor-force participation, 0.620, is from Bureau of Labor Statistics (2024b); the entrepreneurship share, 0.100, from Bureau of Labor Statistics (2024a); average employees per firm, 10.0, from U.S. Census Bureau (2021); and the coefficient of variation of business income, 1.20, from DeBacker et al. (2021). The log p_{90}/p_{50} wage ratio, $\log(2.71) \approx 0.997$, the Gini coefficient of business income, 0.65, and the effective tax rates of 0.148 for the bottom 50% and 0.214 for the middle 40% are from the distributional national accounts of Piketty et al. (2022). The extensive-margin participation elasticity, 0.25, is from Chetty (2012).

Extension calibration targets. The occupational moments of the own-account extension are taken from the administrative tax data of Bhandari et al. (2026). The self-employment share among active individuals, 0.1026, and the employer share of the self-employed, 0.329, are computed from the counts in their Table 2; the top-decile share of business income, 0.53, is from the same source. Average hours, 0.333 of the time endowment, follow Heathcote et al. (2014).

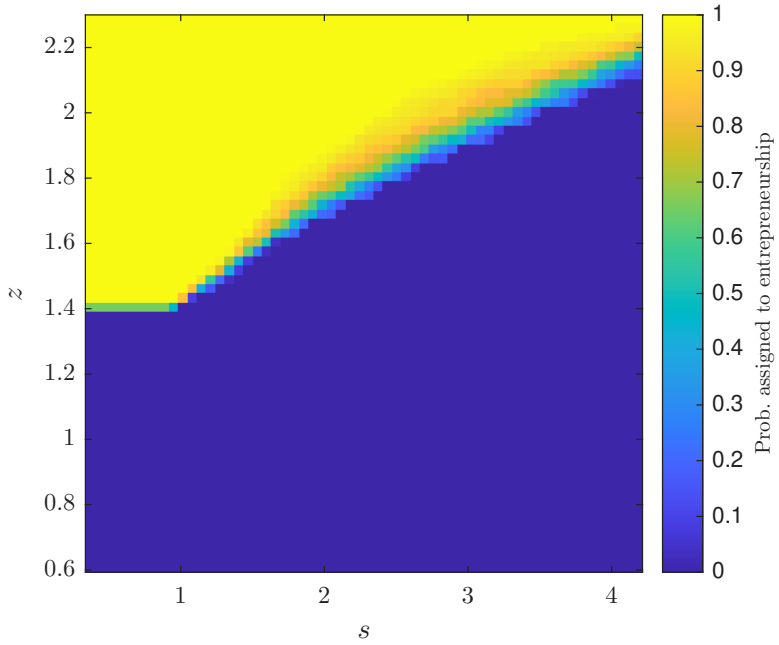
Tax system. Labor income is taxed according to the parametric schedule of Heathcote et al. (2017), $c = \lambda_L y^{1-\tau_L}$, with (τ_L, λ_L) calibrated to the effective tax rates above. Dividends face a stylized version of the statutory schedule for qualified dividends and long-term capital gains under IRC §1(h): a zero rate on the first \$50,000, 15% between \$50,000 and \$550,000, and 20% above \$550,000, applied as marginal rates within brackets. The thresholds round the single-filer brackets (\$48,350 and \$533,400 for 2025) and are applied to dividend income, which is entrepreneurs' only income in the model; we abstract from filing status, from the stacking of dividends on other taxable income, and from the 3.8% net investment income tax. In model units these thresholds are 0.769 and 8.462.

Government spending. The revenue requirement G is not targeted to an external statistic. Following Boerma et al. (2022), the government levies revenue that is not returned to households, and we set G equal to the revenue the calibrated tax system raises in the positive economy. This gives G equal to 17.2% of output in the benchmark and 18.1% in the extension.

D Additional Figures: Comparative Statics

Figure D.1 documents where the optimal mechanism uses randomization. Because occupations are mutually exclusive, the weights that the lottery μ places on bundles with $d > 0$ sum to the probability that a type is assigned to entrepreneurship. That probability is degenerate over most of the type space: it equals zero for 66.2% of types and one for 24.3%. The remaining 9.5% lie in a narrow band along the occupational frontier, where the planner assigns a genuinely random occupation. The band widens with labor skill, tracing the entrepreneurial-ability threshold $z^*(s)$, and reflects the non-convexity created by discrete occupational choice: near the frontier the planner cannot separate types with a deterministic assignment and randomizes instead. Away from the frontier randomization is not used at the occupational margin.

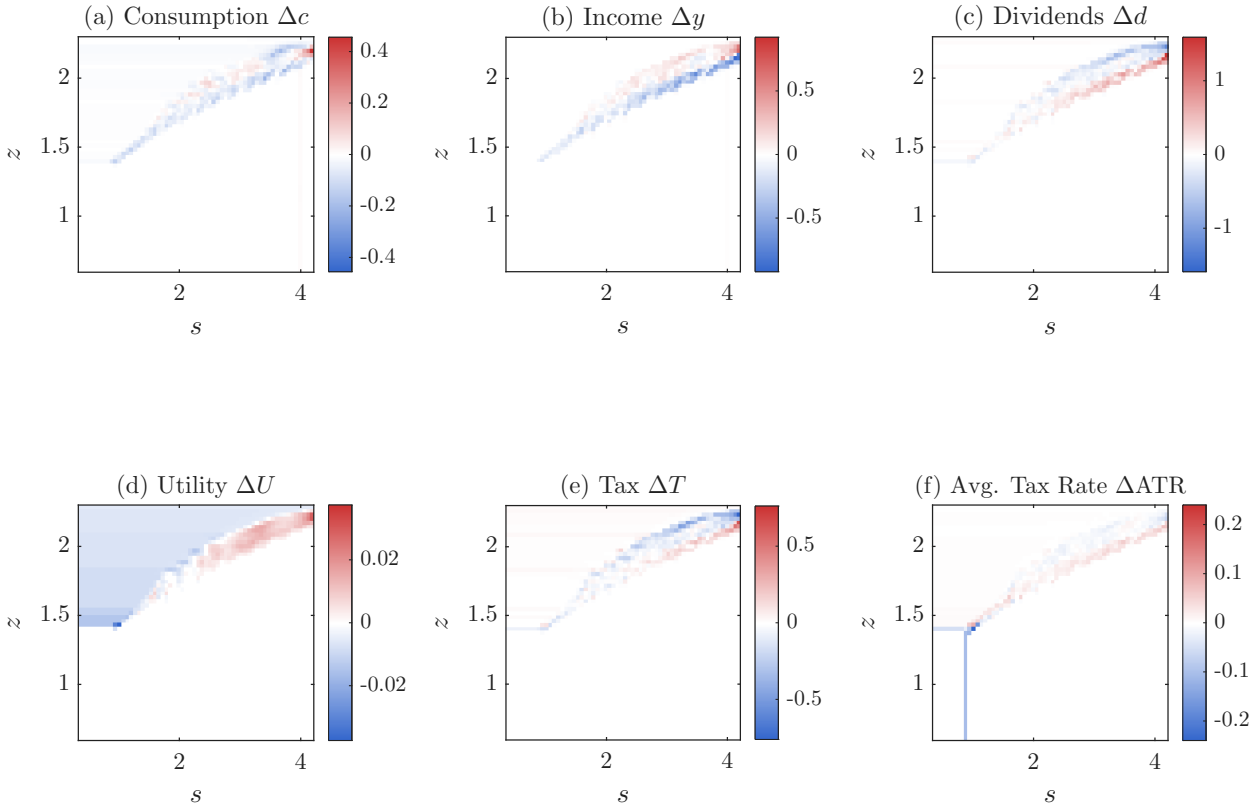
Figure D.1: Occupational assignment and randomization over the type space



Notes: Occupational lottery in the benchmark solution ($G > 0$, $R = 2$). For each type and shock realization the linear program chooses a lottery $\mu_{i,a,b,r}$ over admissible income–dividend bundles; since occupations are exclusive, summing the weights on bundles with $d > 0$ gives the probability that the type is assigned to entrepreneurship. The figure reports that probability, averaged over shock realizations. It equals zero for 66.2% of types and one for 24.3%; the remaining 9.5%, concentrated in a narrow band along the occupational frontier, are assigned a genuinely random occupation.

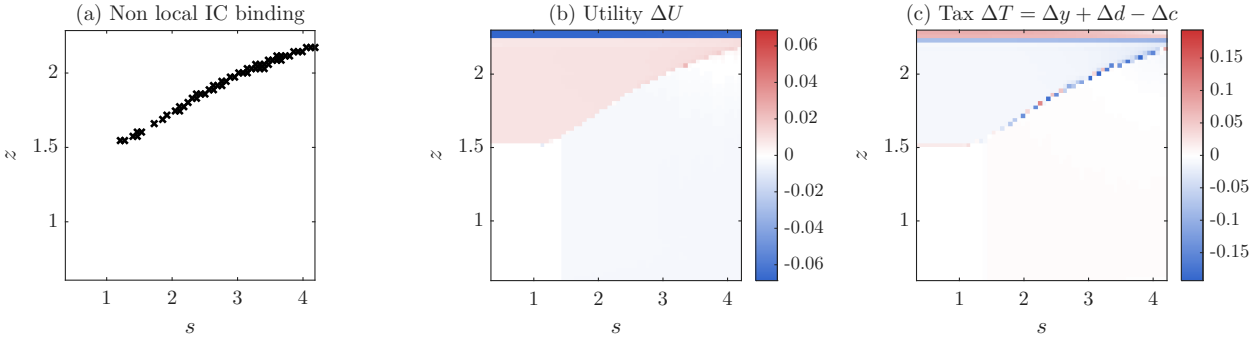
E Additional Figures: Incentive Constraints

Figure E.1: Effects of binding non-local IC constraints: additional panels



Notes: The figure compares the solution in which all global IC constraints are imposed with the solution imposing only local IC constraints. Panel (a) reports the difference in consumption. Panel (b) reports the difference in labor income, Δy, between the global and local-IC-only solutions, computed as global minus local. Panel (c) reports the corresponding difference in dividends, Δd, again computed as global minus local. Similarly, the remaining three panels report the differences in utility, tax, and average tax rate.

Figure E.2: Effects of binding non-local IC constraints with linear consumption utility



Notes: The figure compares the solution in which all global IC constraints are imposed with the solution imposing only local IC constraints, in the stochastic economy with linear consumption utility, $U(c) = c$. Panel (a) identifies types for which a non-local IC constraint is violated by more than 10^{-2} in the local-IC-only solution. Panel (b) reports the difference in indirect utility, ΔU, between the global and local-IC-only solutions, computed as global minus local. Panel (c) reports the corresponding difference in net tax liabilities, ΔT, again computed as global minus local.

F Appendix: Own-Account Extension

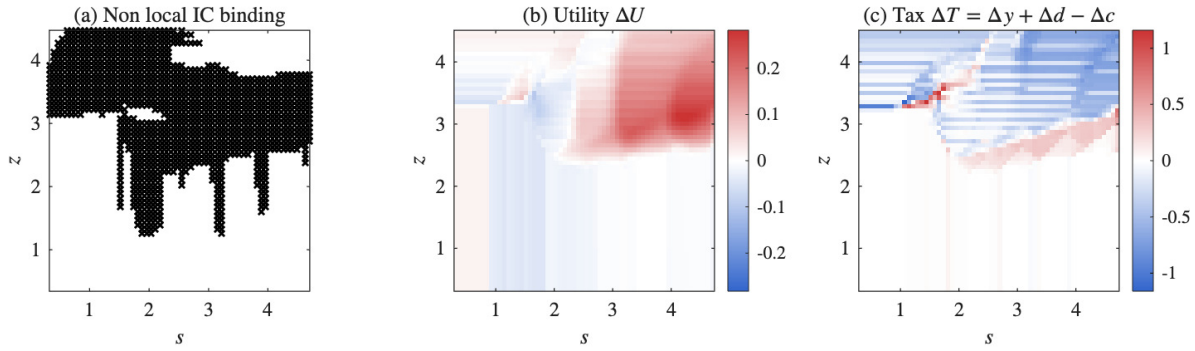
This appendix collects the calibration and the additional quantitative results for the extension of Section 6, in which self-employed agents choose ex post between operating own-account and hiring workers.

Figure F.1 repeats for the extended model the comparison between the solution that imposes all global incentive constraints and the solution that imposes only local ones. The exercise is the same as in the benchmark, but the answer is different in degree. Under the local-only solution some non-local constraint is violated for 41.1% of types, against 2.9% in the benchmark. The reason is the own-account branch: a type who can produce alone as well as by hiring has a second way to deliver any given dividend, and that second route opens deviations toward types that are far away in (s, z). Local constraints, which only compare a type with its immediate neighbours, do not see them.

The violated region is nonetheless a region and not a haze. It is connected, sits at high entrepreneurial ability, and coincides with the part of the type space in which panels (b) and (c) show the allocation actually moving when global constraints are imposed — precisely where the choice between producing alone and hiring is close to indifferent. Away from that area the local-only solution is globally incentive compatible. The welfare cost of ignoring the global constraints stays below 0.1%, so the failure is one of characterization rather than of magnitude: the first-order approach misdescribes the mechanism over a sizeable part of the type space while costing little in welfare terms. This is what motivates solving the problem as a linear program with constraint generation rather than through local conditions.

That the own-account branch is what drives this can be shown directly. Figure F.2 repeats the exercise while varying the solo-production penalty δ, holding every other parameter — including

Figure F.1: Global versus local incentive constraints in the extended model

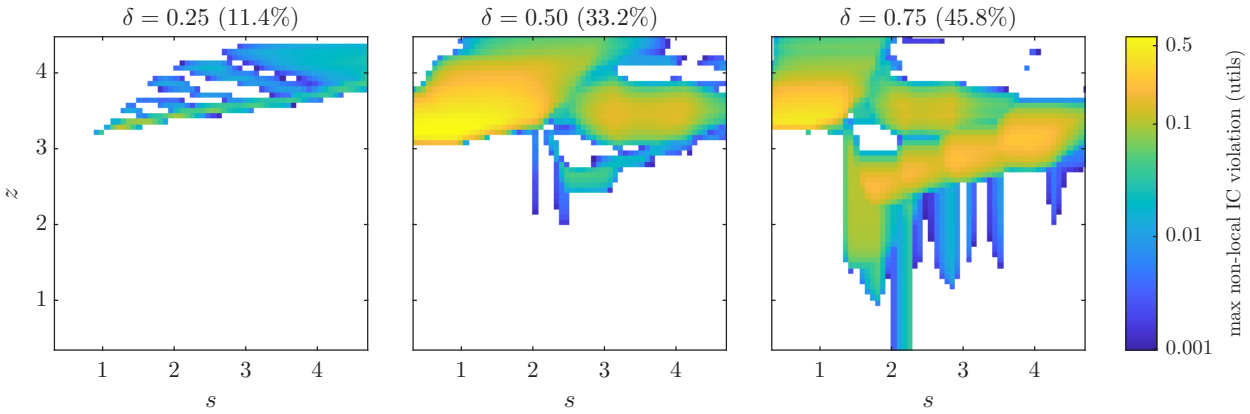


Notes: Difference between the solution with all global incentive constraints imposed and the solution imposing only local constraints, in the extended model. Panels report the binding non-local constraints, the difference in indirect utility, and the difference in net tax liabilities. Panel (a) marks types for which some non-local constraint is violated by more than 10^{-2} under the local-only solution. This occurs for 41.1% of types, against 2.9% in the benchmark. The affected types are not spread uniformly: they form a connected region at high entrepreneurial ability, coinciding with the area in which panels (b) and (c) show the allocation responding to global incentive compatibility. The own-account branch is what generates them, by giving a type a second way to deliver a given dividend and hence a profitable deviation toward a distant type. The associated welfare cost remains small: imposing all global constraints changes welfare by less than 0.1%.

the revenue requirement G — at its calibrated value. The parameter governs how much output a self-employed agent forgoes by producing alone: at $\delta = 0$ own-account production yields nothing, so the technology collapses to the one-branch economy of the benchmark, while at $\delta = 1$ producing alone is as productive as hiring. These are comparative statics, not re-calibrated economies; only $\delta = 0.668$ matches the observed employer share.

The share of types at which some non-local constraint is violated rises and then falls in δ, from 11.6% at $\delta = 0$ to 45.8% at $\delta = 0.75$ and 29.7% at $\delta = 1$, and so does the size of the violations. The two ends of the range are informative. At $\delta = 0$ the violated set is the thin band along the occupational frontier familiar from the benchmark, and it is mostly dim: few types have a substantially profitable distant deviation. This is the expected result and a check on the measurement, since with the second branch switched off the extension reproduces the benchmark's incentive structure. As δ rises the band thickens, develops interior regions in which the best available deviation is worth 20 to 100% of consumption, and sends tendrils down into the interior of the type space, reaching types with substantially lower entrepreneurial ability. The violations decay smoothly toward the edge of the region, so the boundary in the figure is a level set of a continuous object rather than a sharp frontier. The non-local constraints are therefore not a generic feature of two-dimensional screening but a specific consequence of giving agents two technologies for producing the same dividend.

Figure F.2: Size of the non-local incentive violations as the solo-production penalty varies



Notes: Colour reports, at each type, the largest violation of a non-local incentive constraint under the local-only solution, for a different value of the solo-production penalty δ. The scale is logarithmic and measured in utils; since $U(c) = \log c$ and utility is additively separable, a violation of v is exactly equivalent to scaling the type's consumption by e^v, so 0.1 corresponds to a gain of about 10.5% of consumption. Types below 10^{-3} (roughly 0.1% of consumption) are left white; the numerical floor of the solver lies an order of magnitude below that threshold. All other parameters, including G, are held at the values of Table 6; the calibrated value is $\delta = 0.668$. Panel titles report the share of types above the threshold. Both the global and the local problem are re-solved at each δ.

F.1 Calibration of the extended model

Tables 6 and 7 report the eleven calibrated parameters and the eleven targeted moments. Relative to the benchmark, the extension adds the solo-production penalty δ, identified by the employer share of the self-employed, and the effort-disutility scale η_0, identified by average hours. The occupational moments are taken from the administrative tax data of Bhandari et al. (2026). Panel B of Table 7 reports untargeted moments.

F.2 Additional results for the extended model

Figures F.3–F.5 report, for the extended model, the comparative statics with respect to the dispersion of entrepreneurial ability, the affiliation between labor skill and entrepreneurial ability, and the quasi-linear benchmark. In each case the conclusions of the benchmark model carry over: subsidies extend deeper into the entrepreneurial region when high-ability types are scarce, concentrate on superstars under positive affiliation, and give way to larger firms when the planner's redistributive motive is switched off.

Table 6: Calibrated parameters, own-account extension

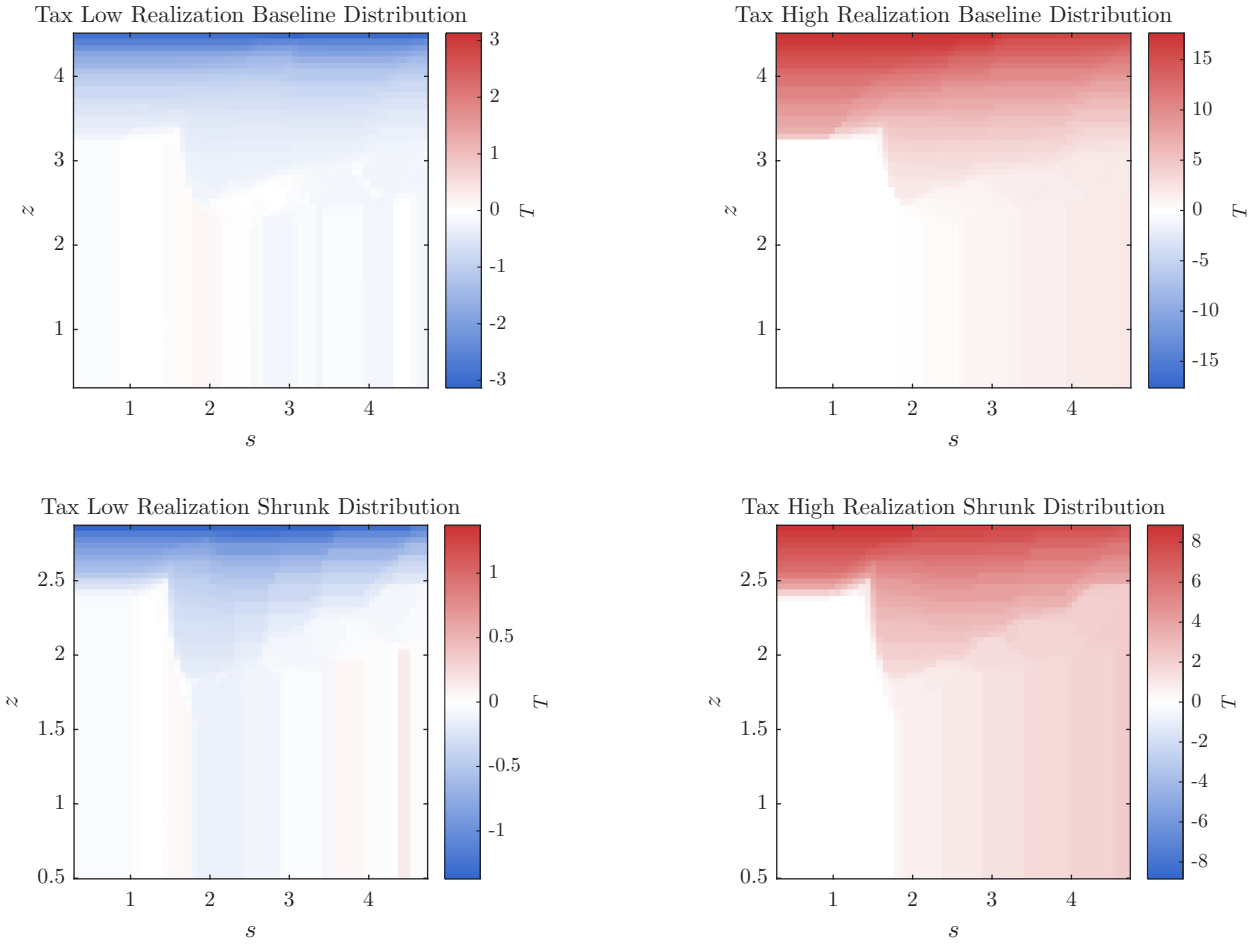
Parameter	**Symbol**	**Value**	**Target moment**	**Source**
Set externally				
Risk aversion	γ	1.00	Log utility	Standard
Inverse Frisch	ψ	2.00	Frisch = 0.50	Standard
Ability correlation	ρ	0.00	Independence	Standard
Calibrated jointly				
Worker fixed cost	κ_1	0.214	Labor participation	Bureau of Labor Statistics (2024b)
Self-empl. entry cost	κ_E	0.422	Self-employment share	Bhandari et al. (2026)
Span of control	α	0.343	Avg. employees per employer	U.S. Census Bureau (2021)
Solo-production penalty	δ	0.668	Employer share of self-employed	Bhandari et al. (2026)
Effort-disutility scale	η_0	36.36	Average hours	Heathcote et al. (2014)
Shock volatility	σ_ε	0.755	Top 10% business-income share	Bhandari et al. (2026)
Skill tail (Fréchet)	ξ_s	2.30	$\log(p_{90}/p_{50})$ wage	Piketty et al. (2022)
Ability tail (Fréchet)	ξ_z	2.39	Gini of business income	Piketty et al. (2022)
Tax progressivity	τ_L	0.056	ETR bottom 50%	Piketty et al. (2022)
Tax level	λ_L	0.728	ETR middle 40%	Piketty et al. (2022)
Home production	c_0	0.009	Participation elasticity	Chetty (2012)
Statutory business tax (not calibrated)				
Bracket 1	T_B	0%	$d \leq$ \$50k	IRC §1(h)
Bracket 2	T_B	15%	\$50k–\$550k	IRC §1(h)
Bracket 3	T_B	20%	$d >$ \$550k	IRC §1(h)

Table 7: Targeted moments: model vs. data, own-account extension

Moment	**Model**	**Data**	**Identifies**
Panel A: Targeted moments			
Labor participation	0.641	0.620	κ_1
Self-employment share (active)	0.109	0.103	κ_E
Employer share of self-employed	0.362	0.329	δ
Avg. employees per employer	9.754	10.000	α
Average hours	0.429	0.333	η_0
Top 10% business-income share	0.512	0.530	σ_ε
$\log(p_{90}/p_{50})$ wage	0.893	0.997	ξ_s
Gini of business income	0.671	0.650	ξ_z
ETR bottom 50%	0.158	0.148	τ_L
ETR middle 40%	0.198	0.214	λ_L
Participation elasticity	0.282	0.250	c_0
Panel B: Untargeted moments (validation)			
Inactivity rate	0.359	0.380	—
CV of business income	1.590	1.200	—
Employer share (active)	0.039	0.034	—
Self-employed mass (uncond.)	0.070	0.064	—
Wage Gini	0.313	0.520	—
ETR top 10%	0.178	0.249	—
Avg. ETR labor	0.191	0.234	—
Top 10% wage share	0.236	0.375	—
Business income share (employers)	0.454	0.280	—
Self-employment income share	0.624	—	—

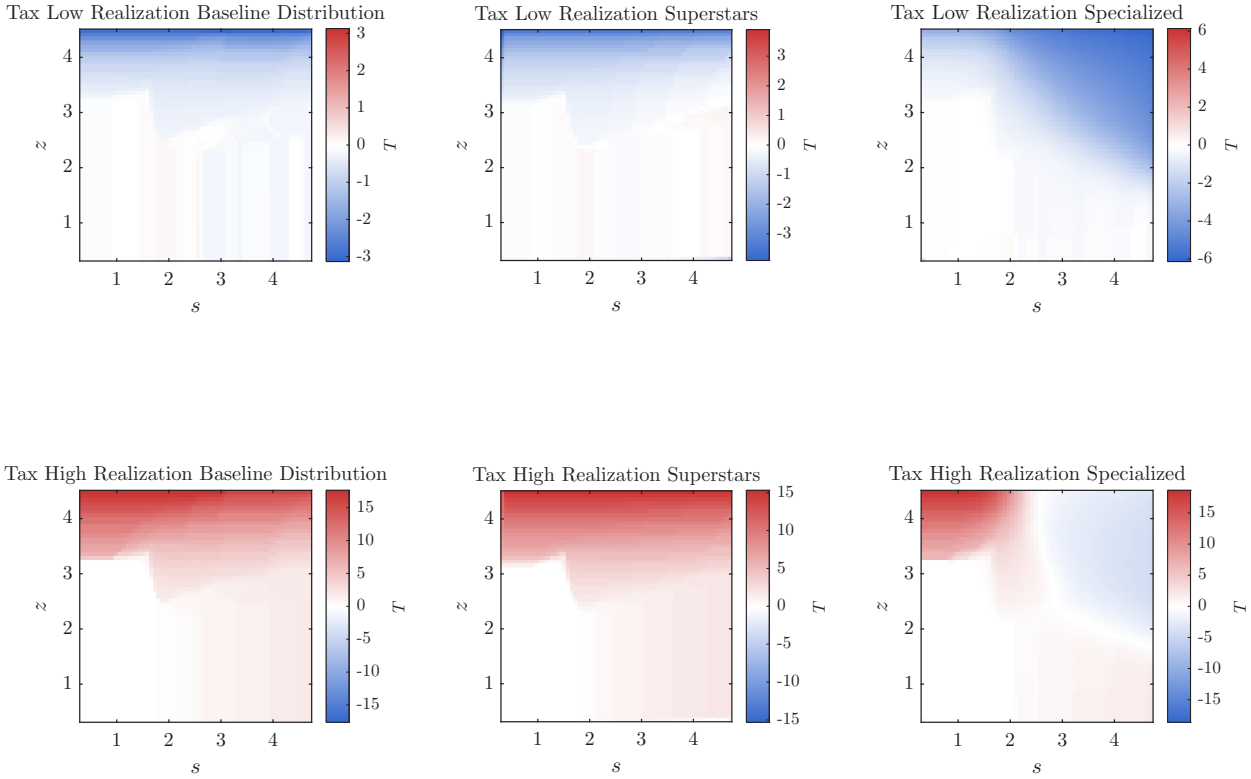
Notes: Calibration evaluated on a 150×150 type grid with $R = 2$ shocks. Tax moments are model-relevant effective tax rates (individual income tax + payroll + corporate) from the DINA distributional accounts of Piketty et al. (2022), 2018 data. Participation elasticity: Hicksian extensive-margin elasticity with respect to the net-of-tax rate (Chetty, 2012). SMM loss = 0.1371.

Figure F.3: Extended model: scarcity of high-ability entrepreneurs



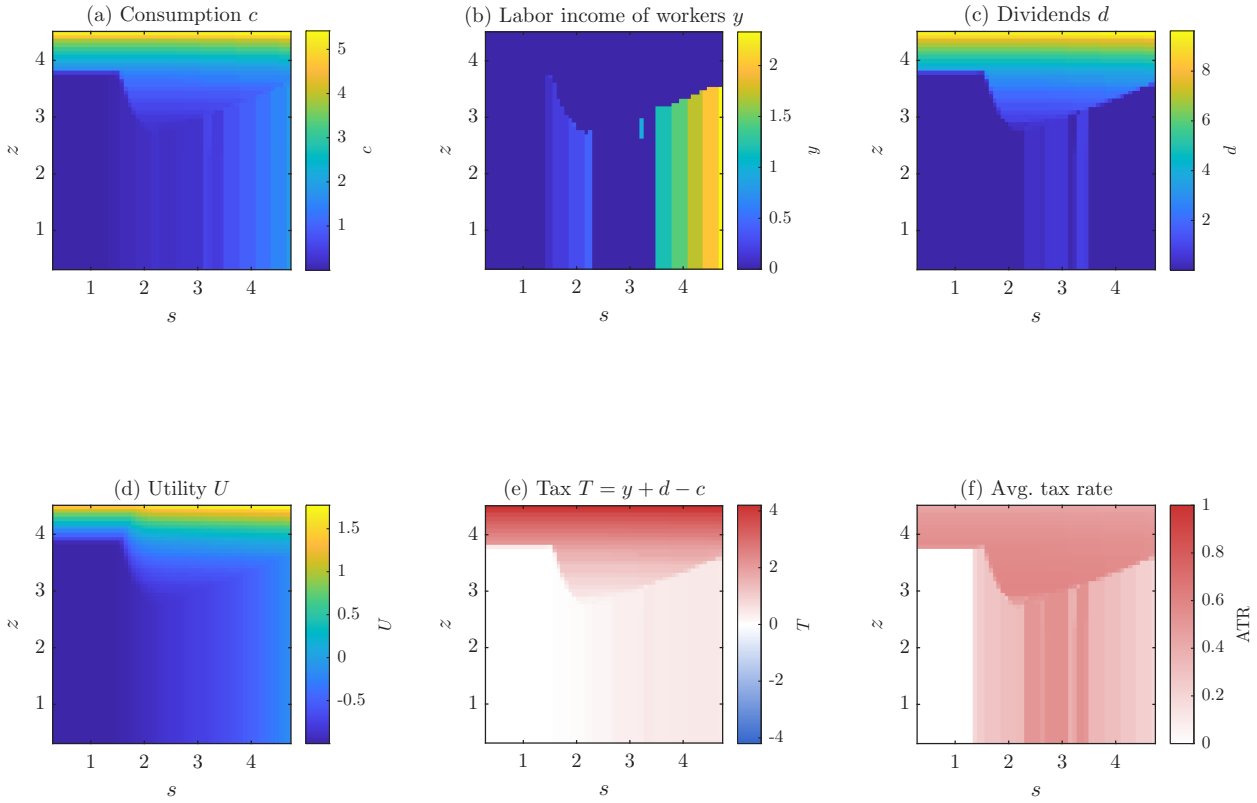
Notes: Net tax liabilities by shock realization in the extended model, baseline dispersion of entrepreneurial ability (top) versus compressed dispersion (bottom).

Figure F.4: Extended model: affiliation between skill and ability



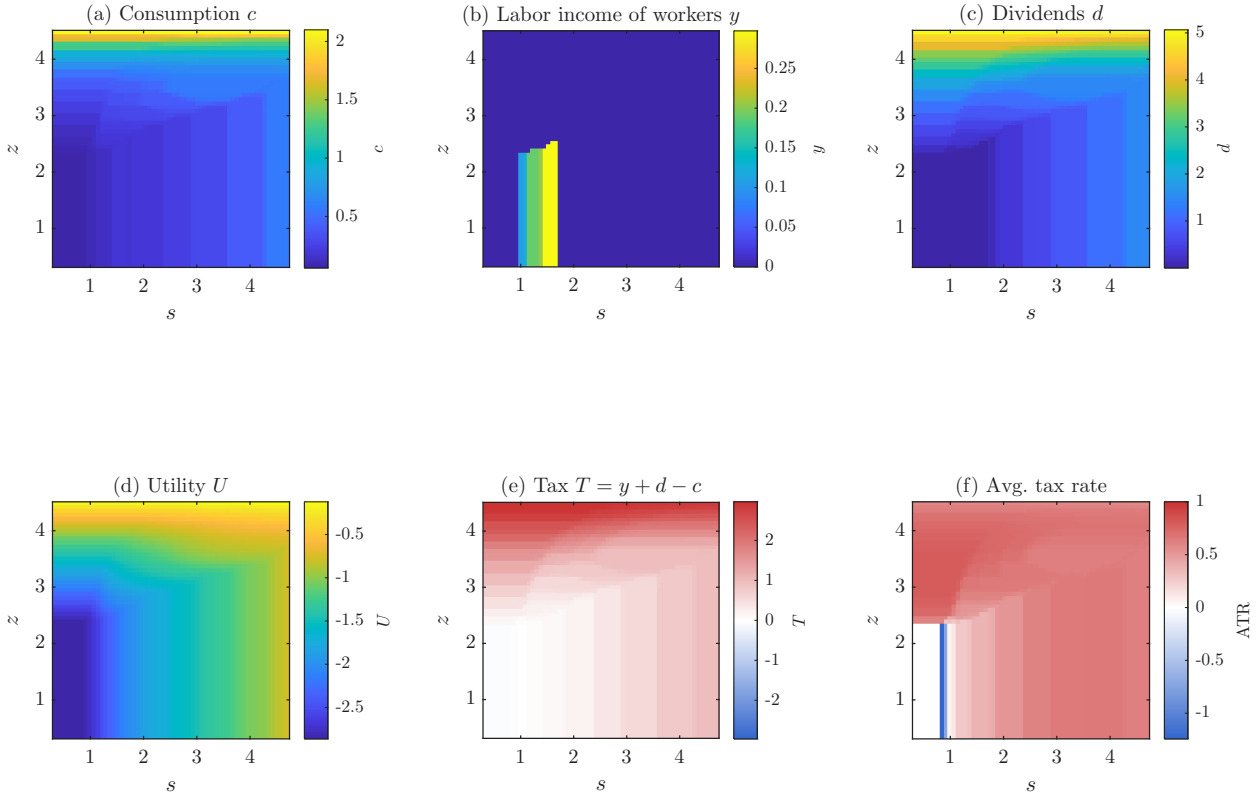
Notes: Net tax liabilities by shock realization in the extended model, under independence (left), positive affiliation (middle), and negative affiliation (right).

Figure F.5: Extended model: quasi-linear consumption utility



Notes: Optimal allocation in the extended model when consumption utility is linear ($U(c) = c$), which switches off the planner's direct redistributive motive. As in Figure 10, labor income is shown for types assigned to paid work with probability at least 0.9 and set to zero for the others.

Figure F.6: Extended model without entrepreneurial risk



Notes: As Figure 10, for the deterministic economy ($R = 1$).