Approximating Transition Dynamics with Discrete Choice

Anmol Bhandari* David Evans† Ellen McGrattan‡

February 9, 2025

Abstract

We develop a method to analyze policy reforms in environments with discrete choice, such as occupational choice, default, entry, and exit decisions. Computing transition paths in these settings is computationally challenging, particularly in models with substantial heterogeneity and many endogenous states. We extend perturbation methods to handle discrete choice by appropriately tracking both "small" changes conditional on discrete choices and "big" changes due to switches in discrete choices. The method is fast, scalable, and efficient, providing good initial estimates for global solution methods. We demonstrate our method by analyzing optimal business taxation in a model with occupational choice between paid- and self-employment.

1 Introduction

The analysis of policy reforms in dynamic economies has long been a central focus of macroeconomics. Although substantial progress has been made in analyzing policies in representative agent frameworks, the incorporation of

*University of Minnesota
†University of Oregon
‡University of Minnesota

rich heterogeneity and discrete choice remains computationally challenging. This challenge is particularly acute when studying important policy questions involving occupational choice, firm entry and exit, default decisions, or migration choices. In these settings, policy changes affect both continuous choices (e.g., consumption, investment) and discrete choices (e.g., whether to become an entrepreneur), with the latter often generating nonconvexities that complicate traditional solution methods.

The computational difficulties arise from three key features that frequently arise in modern macroeconomic models. First, discrete choices create discontinuities in policy functions, making standard perturbation methods inapplicable. Second, substantial heterogeneity in assets, productivity, and other characteristics requires tracking high-dimensional distributions. Third, the interaction between discrete and continuous choices generates rich dynamics that is difficult to characterize, particularly along transition paths. These challenges have led researchers to either abstract from important model features or rely on computationally intensive global solution methods that become intractable in richer settings.

This paper develops a novel approach to analyzing policy reforms in environments with discrete choice by extending perturbation methods to handle non-convexities. Our key insight is that changes in policy affect economic outcomes through two distinct channels in models with discrete choice: (1) "small" changes in continuous choices conditional on discrete decisions, and (2) "big" changes resulting from switches in discrete choices. By appropriately tracking both types of changes, we can characterize transition dynamics in rich heterogeneous agent economies while maintaining computational tractability.

The method introduces preference shocks that smooth out discrete choices, ensuring that derivatives used in perturbation analysis are well-defined. This smoothing allows us to leverage the computational advantages of perturbation methods while capturing the essential features of the underlying discrete choice problem. Importantly, we show how to track both the direct effect of policy changes on individual choices and the indirect effects working through changes in equilibrium prices and distributions.

Our approach offers several advantages over existing methods. First, it is computationally efficient and scalable to high-dimensional state spaces, making it feasible to analyze rich models with substantial heterogeneity. Second, it provides a natural decomposition of policy effects into intensive and extensive margin responses, facilitating economic interpretation. Third, even when global solution methods are feasible, our method generates good initial guesses that can accelerate their convergence. Fourth, the approach readily handles general equilibrium effects and transition dynamics, which are often abstracted from in discrete choice settings.

We demonstrate the utility of our method by analyzing optimal business taxation in a model with occupational choice between paid employment and entrepreneurship. This application illustrates several key features that make standard solution methods challenging: (1) a discrete occupational choice, (2) rich heterogeneity in assets and productivity, (3) financial frictions that generate nonconvexities, and (4) general equilibrium price effects. Despite these complexities, our method efficiently characterizes transition dynamics and optimal policy.

The analysis yields several important insights on business taxation. First, we find that optimal business tax rates are generally negative when financial frictions are severe, reflecting the value of subsidizing entrepreneurship to overcome credit constraints. Second, the optimal subsidy is larger when credit constraints are tighter, highlighting the interaction between tax policy and financial frictions. Third, accounting for transition dynamics substantially affects welfare calculations, underscoring the importance of analyzing policies in a dynamic rather than steady-state framework.

1.1 Related Literature

Our work contributes to two important strands of literature in macroeconomics: the analysis of discrete choice in dynamic economies and computational methods to solve heterogeneous agent models. We bridge these literatures by developing techniques that can handle both rich heterogeneity and discrete choices while maintaining computational tractability.

The incorporation of discrete choice in macroeconomic models has a parallel history, beginning with Rust (1987)'s seminal work on dynamic discrete choice. In the context of occupational choice, Quadrini, Vincenzo (2000) and Cagetti and De Nardi (2006) developed influential models that examined the interaction between financial constraints and entrepreneurship. However, these studies primarily focus on steady-state analysis rather than transition dynamics. More recent work by Kaplan et al. (2018) emphasizes the importance of discrete choices in understanding household behavior and the transmission of monetary policy, while highlighting the computational challenges that arise when combining discrete choice with rich heterogeneity.

Our methodological approach builds on two recent advances in computational methods. First, the projection-perturbation framework of Reiter (2009), which opened the door to analyzing aggregate shocks in heterogeneous agent economies. Second, the sequence-space approaches developed by Boppart et al. (2018) and refined by Auclert et al. (2021), which showed how to efficiently compute first-order impulse responses to an unanticipated shock. We particularly build on Bhandari et al. (2023), who show how to unify and extend these insights by characterizing equilibrium dynamics through sequences of linear systems for directional derivatives. Our contribution here is to extend their framework to accommodate discrete choices through a novel application of generalized functions based on techniques developed by Kanwal (1998). This allows us to properly account for both the intensive margin adjustments in continuous choices and the extensive margin responses arising from discrete choices.

The role of financial frictions in shaping discrete choices has been extensively studied following Evans and Jovanovic (1989). Recent work by Catherine (2019) emphasizes the importance of transition dynamics in evaluating entrepreneurial decisions. Our framework allows for analysis of how policy reforms interact with financial frictions in determining occupational choice and capital allocation, while properly accounting for transition dynamics and general equilibrium effects.

The remainder of the paper is organized as follows. Section 2 presents the model environment. Section 3 develops our methodology for approximating transition dynamics with discrete choice. Section 4 discusses the numerical implementation. Section 5 applies the method to analyze optimal business taxation. Section 6 concludes.

2 Model Environment

We consider an incomplete-markets economy with occupational choice, building on Aiyagari (1994). Time is discrete and infinite. The economy is populated by a continuum of infinitely-lived agents indexed by $i \in [0, 1]$ who face idiosyncratic risk and make occupational choices between paid employment and entrepreneurship. A representative firm operates in the corporate sector. The government levies differential tax rates on wage and business income and provides lump sum transfers.

2.1 Individual State Variables and Preferences

Each agent i at time t is characterized by their asset holdings $a_{i,t-1}$, labor productivity $\theta^w_{i,t}$, and entrepreneurial productivity $\theta^b_{i,t}$. The productivity processes follow a joint Markov chain:

$$\theta_{i,t} = [\log \theta^w_{i,t}, \log \theta^b_{i,t}]' = \rho_\theta \theta_{i,t-1} + \varepsilon^\theta_{i,t} \tag{1}$$

where $\varepsilon^{\theta}_{i,t}$ is a vector of innovations with variance-covariance matrix Σ_θ. Agents have identical preferences over consumption streams represented by:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(c_{i,t}) \tag{2}$$

where $\beta \in (0,1)$ is the discount factor and $U(\cdot)$ is a strictly increasing, strictly concave period utility function.

2.2 Occupational Choice

In each period, agents choose between paid employment ($d_{i,t} = 1$) and entrepreneurship ($d_{i,t} = 0$). This choice is affected by a preference shock $\eta_{i,t}$ that follows an i.i.d. Type 1 Extreme Value (Gumbel) distribution Γ with scale parameter σ_η. The occupational choice problem can be written as:

$$\begin{aligned} v_t(\theta^w_{i,t}, \theta^b_{i,t}, a_{i,t-1}, \eta_{i,t}) = \max_{d_{i,t}\in\{0,1\}} \; & d_{i,t}(v^w_t(\theta^w_{i,t}, \theta^b_{i,t}, a_{i,t-1}) + \eta_{i,t}) \\ & + (1 - d_{i,t}) v^b_t(\theta^w_{i,t}, \theta^b_{i,t}, a_{i,t-1}) \end{aligned} \tag{3}$$

where v^w_t and v^b_t are the value functions associated with paid employment and entrepreneurship, respectively.

2.3 Workers

If an agent chooses paid employment, they supply their labor productivity unit $\theta^w_{i,t}$ inelastically to the labor market at wage rate W_t. Their problem is:

$$v^w_t(\theta^w_{i,t}, \theta^b_{i,t}, a_{i,t-1}) = \max_{c_{i,t}, a_{i,t} \geq 0} U(c_{i,t}) + \beta \mathbb{E}_t[v_{t+1}(\theta^w_{i,t+1}, \theta^b_{i,t+1}, a_{i,t}, \eta_{i,t+1})] \tag{4}$$

subject to the budget constraint:

$$c_{i,t} + a_{i,t} = (1 + R_t - \delta) a_{i,t-1} + (1 - \tau^w) W_t \theta^w_{i,t} + T_t \tag{5}$$

where R_t is the interest rate, δ is the depreciation rate, τ^w is the labor income tax rate, and T_t is a lump-sum transfer.

2.4 Entrepreneurs

Entrepreneurs operate a decreasing returns to scale technology using capital $k_{i,t}$ and labor $n_{i,t}$ as inputs. Their production function is:

$$y_{i,t} = \theta^b_{i,t} k^\alpha_{i,t} n^\nu_{i,t} \tag{6}$$

where $\alpha + \nu < 1$, capturing span-of-control limitations as in Lucas (1978). These entrepreneurs face a collateral constraint that limits capital investment to a multiple of their assets:

$$k_{i,t} \leq \chi a_{i,t-1} \tag{7}$$

where $\chi \geq 1$ determines the maximum leverage ratio. The dynamic program to be solved is given by:

$$v^b_t(\theta^w_{i,t}, \theta^b_{i,t}, a_{i,t-1}) = \max_{c_{i,t}, a_{i,t} \geq 0} U(c_{i,t}) + \beta \mathbb{E}_t[v_{t+1}(\theta^w_{i,t+1}, \theta^b_{i,t+1}, a_{i,t}, \eta_{i,t+1})] \tag{8}$$

and is subject to the following budget constraint:

$$c_{i,t} + a_{i,t} = (1 + R_t - \delta) a_{i,t-1} + (1 - \tau^b) \pi_{i,t} + T_t \tag{9}$$

where τ^b is the business income tax rate and $\pi_{i,t}$ are profits. In this case, profits are defined as follows:

$$\pi_{i,t} = \max_{n_{i,t}, k_{i,t} \leq \chi a_{i,t-1}} \{\theta^b_{i,t} k^\alpha_{i,t} n^\nu_{i,t} - R_t k_{i,t} - W_t n_{i,t}\}. \tag{10}$$

2.5 Corporate Sector

A representative firm operates in the corporate sector with a constant returns to scale technology:

$$Y_t^c = \Theta (K_t^c)^\alpha (N_t^c)^{1-\alpha}, \tag{11}$$

where Θ is aggregate productivity, K_t^c is corporate capital, and N_t^c is corporate labor. The firm's profit maximization yields factor prices:

$$R_t = \alpha \Theta (K_t^c / N_t^c)^{\alpha - 1} \tag{12}$$

$$W_t = (1-\alpha)\Theta (K_t^c / N_t^c)^{\alpha}. \tag{13}$$

2.6 Government

The government levies differential tax rates on wage income (τ^w) and business income (τ^b). Tax revenues are rebated as lump-sum transfers:

$$T_t = \int \tau^b \pi_{i,t}(1-d_{i,t})di + \int \tau^w W_t d_{i,t} \theta_{i,t}^w di. \tag{14}$$

2.7 Market Clearing and Equilibrium

Finally, we impose market clearing conditions in capital and labor markets. The conditions are given by:

$$K_t^c + \int k_{i,t} di = \int a_{i,t-1} di \tag{15}$$

$$N_t^c + \int (1-d_{i,t}) n_{i,t} di = \int d_{i,t} \theta_{i,t}^w di. \tag{16}$$

Given initial conditions $\{a_{i,-1}, \theta_{i,0}\}_{i \in [0,1]}$, a competitive equilibrium consists of prices $\{R_t, W_t\}_{t \geq 0}$, individual choices $\{c_{i,t}, a_{i,t}, d_{i,t}, k_{i,t}, n_{i,t}\}_{t \geq 0, i \in [0,1]}$, corporate choices $\{K_t^c, N_t^c\}_{t \geq 0}$, and transfers $\{T_t\}_{t \geq 0}$ such that: (i) given prices and transfers, individual choices solve the worker and entrepreneur problems; (ii) corporate choices maximize profits; (iii) the government budget constraint holds; and (iv) markets clear.

3 Methodology

In this section, we develop the approach to approximate equilibrium transition paths in economies with discrete choice. We first introduce notation to represent the equilibrium conditions compactly and then describe our perturbation-based approximation method.

3.1 Notation and Equilibrium Representation

To analyze transition dynamics, we need notation that captures both aggregate and individual states and choices. Let $Z_t = [A_{t-1}, \Omega_t]'$ denote the aggregate state vector, where A_{t-1} represents predetermined aggregate variables and Ω_t is the distribution over individual states. Individual states consist of $(a_{i,t-1}, \theta_{i,t})$, where $a_{i,t-1}$ is the predetermined asset position and $\theta_{i,t}$ captures idiosyncratic productivity. For each occupation $d \in \{w, b\}$, let $x_{i,t}^d$ denote the vector of individual choices conditional on occupation, with $a_{i,t}^d$ and $v_{i,t}^d$ being elements representing next period assets and value functions, respectively. The occupational choice is determined by the indicator function:

$$d_{i,t} = \iota(v_{i,t}^b - v_{i,t}^w \leq \eta_{i,t}), \tag{17}$$

where $\iota(\cdot)$ is the indicator function. Given this choice, realized individual choices are:

$$x_{i,t} = d_{i,t} x_{i,t}^w + (1 - d_{i,t}) x_{i,t}^b. \tag{18}$$

Let X_t denote the vector of aggregate variables that are not distributions, with $A_t \in X_t$ being the subset of predetermined variables. We define $Y_t = [A_{t-1}, X_t, X_{t+1}]'$ as the vector of aggregate variables relevant for period t decisions. Then the equilibrium conditions can be represented compactly as:

$$F^d(\theta_{i,t}, a_{i,t-1}, x^d_{i,t}, \mathbb{E}_{i,t}x_{i,t+1}, Y_t) = 0 \quad \text{for all } i, t, d \in \{w, b\} \tag{19}$$

$$G(\int x_{i,t}di, Y_t) = 0 \quad \text{for all } t \tag{20}$$

where F^d captures individual optimality conditions for each occupation and G represents all other equilibrium conditions.

3.2 Approximating Transition Paths

Consider a policy reform that moves the economy from an initial steady state characterized by $Z_{-1} = [A_{-1}, \Omega_{-1}]'$ to a new steady state $Z^* = [A^*, \Omega^*]'$. We approximate the transition path using perturbation methods around the post-reform steady state. The key idea is to scale deviations of the initial state from the new steady state using a parameter σ:

$$Z_0 = Z^* + \sigma(Z_{-1} - Z^*) \tag{21}$$

This allows us to characterize the equilibrium path $X_t(\sigma)$ as a function of σ, where: $\sigma = 0$ corresponds to starting at the post-reform steady state and $\sigma = 1$ corresponds to the actual transition from the steady state prior to a policy reform.

We use Taylor expansions with respect to σ to approximate the transition path. The first-order approximation takes the form:

$$X_t = \overline{X} + \hat{X}_t + O(|Z_{-1} - Z^*|^2) \tag{22}$$

where $\overline{X}$ is the post-reform steady state value and $\hat{X}_t$ represents first-order adjustments along the transition path. The first-order adjustments $\hat{X}_t$ are determined by a sequence of directions $\hat{Z}_t$ that satisfy:

$$\hat{Z}_0 := Z_{-1} - Z^* \tag{23}$$

$$\hat{Z}_t := \overline{Z}_Z \cdot \hat{Z}_{t-1} \tag{24}$$

$$\hat{X}_t := \overline{X}_Z \cdot \hat{Z}_t \tag{25}$$

where subscript Z denotes the Frechet derivative with respect to aggregate states.

3.3 Handling Discrete Choice

A key challenge in problems with discrete choices is non-differentiabilities in policy functions. We handle this by explicitly tracking two types of changes, namely, "small" changes in choices conditional on occupation and "big" changes due to switches in occupation. More specifically, consider decomposing the change in aggregate variables as:

$$\left(\int \overline{x} d\Gamma d\Omega^* \right)_Z \cdot \hat{Z}_t = \int \hat{x}_t d\Omega^* + \int (\overline{x}^w - \overline{x}^b)\Gamma'(\overline{\kappa})\hat{\kappa}_t d\Omega^* \tag{26}$$

where $\hat{x}_t = \overline{\Gamma}^w \overline{x}^w_Z \cdot \hat{Z}_t + \overline{\Gamma}^b \overline{x}^b_Z \cdot \hat{Z}_t$ captures "small" changes; $\hat{\kappa}_t = \overline{\kappa}_Z \cdot \hat{Z}_t$ represents shifts in the discrete choice threshold; and $\Gamma'(\overline{\kappa})$ is the density at the threshold, determining the mass of agents near the margin.

We turn next to characterizing the individual policy functions and then to the distribution dynamics.

3.4 Computing Individual Policy Function Derivatives

Individual policy function derivatives $\{\hat{x}^d_t\}$ for each occupation $d \in \{w, b\}$ are characterized by:

$$\hat{x}^d_t(a, \theta) = \sum_{s=0}^{\infty} \text{x}^d_s(a, \theta)\hat{Y}_{t+s} \tag{27}$$

where the coefficients $\{\mathsf{x}_s^d\}$ satisfy the recursive system:

$$\mathsf{x}_0^d(a,\theta) = -(\mathsf{F}_x^d(a,\theta) + \mathsf{F}_{x^e}^d(a,\theta)\mathbb{E}[\overline{x}_a|a,\theta]\mathsf{p})^{-1}\mathsf{F}_Y^d(a,\theta) \tag{28}$$

$$\mathsf{x}_{s+1}^d(a,\theta) = -(\mathsf{F}_x^d(a,\theta) + \mathsf{F}_{x^e}^d(a,\theta)\mathbb{E}[\overline{x}_a|a,\theta]\mathsf{p})^{-1}\mathsf{F}_{x^e}^d(a,\theta)\mathbb{E}[\mathsf{x}_s|a,\theta]. \tag{29}$$

These coefficients are computed recursively using linear algebra operations with pre-computed basis matrices. This approach avoids numerical differentiation and is computationally efficient.

3.5 Evolution of the Distribution

The evolution of the distribution must account for both marginal adjustments within occupations and discrete switches between occupations. We can express the change in measure as:

$$\frac{d}{da}\frac{d}{d\theta}\hat{\Omega}_t = -\frac{d}{da}\hat{\omega}_t^{(a)} + \hat{\omega}_t \tag{30}$$

where $\hat{\omega}_t^{(a)}$ captures "small" changes in state variables and $\hat{\omega}_t$ captures "big" changes from occupation switching. These components evolve according to:

$$\hat{\omega}_{t+1}^{(a)} = \mathcal{L}^{(a)} \cdot \hat{\omega}_t^{(a)} + \mathcal{M}^{(a),w} \cdot \hat{a}_t^w + \mathcal{M}^{(a),b} \cdot \hat{a}_t^b \tag{31}$$

$$\hat{\omega}_{t+1} = \Lambda \cdot \hat{\omega}_t + \mathcal{L} \cdot \hat{\omega}_t^{(a)} + \mathcal{M} \cdot \hat{\kappa}_t. \tag{32}$$

3.5.1 Time-Shift Operators

The time-shift operators map distributions from the current period to the next period. These operators come in two varieties, each handling different types of changes in the distribution. The first operator, $\mathcal{L}^{(a)}$, handles small changes and is defined as:

$$\mathcal{L}^{(a)} \cdot \hat{\omega}^{(a)}\langle a', \theta' \rangle = \iint \sum_d \bar{\Lambda}^d(a',\theta',a,\theta)\bar{\Gamma}^d(a,\theta)\overline{a}_a^d(a,\theta)\hat{\omega}^{(a)}(a,\theta)da d\theta. \tag{33}$$

The operator $\mathcal{L}^{(a)}$ incorporates three essential components that determine how the distribution evolves. The steady-state transition density $\bar{\Lambda}^d(a', \theta', a, \theta)$ governs the evolution of states, while $\bar{\Gamma}^d(a, \theta)$ represents the probability of choosing occupation d at state (a, θ). The term $\overline{a}^d_a(a, \theta)$ captures how changes in the current state affect the next period's state.

The second time-shift operator, Λ, handles changes in the distribution directly and is given by:

$$\Lambda \cdot \hat{\omega}\langle a', \theta' \rangle = \iint \sum_d \bar{\Lambda}^d(a', \theta', a, \theta)\Gamma^d(a, \theta)\hat{\omega}(a, \theta) da d\theta. \tag{34}$$

This operator differs from $\mathcal{L}^{(a)}$ in that it operates directly on distribution changes rather than state variable changes, providing a mechanism for tracking how the overall distribution evolves over time.

3.5.2 Policy-Shift Operators

The policy-shift operators translate changes in individual policies into changes in the distribution. The first type, $\mathcal{M}^{(a),d}$, handles policy changes within occupations and is defined as:

$$\mathcal{M}^{(a),d} \cdot \hat{a}^d\langle a', \theta' \rangle = \iint \bar{\Lambda}^d(a', \theta', a, \theta)\bar{\Gamma}^d(a, \theta)\omega^*(a, \theta)\hat{a}^d(a, \theta) da d\theta. \tag{35}$$

This operator combines several key elements to capture policy changes. The transition density $\bar{\Lambda}^d$ specific to occupation d interacts with the probability $\bar{\Gamma}^d$ of being in that occupation. These are then weighted by the steady-state distribution ω^* and the changes in policy functions $\hat{a}^d$ for the given occupation.

The second policy-shift operator, $\mathcal{M}$, specifically handles occupation switching and is given by:

$$\mathcal{M} \cdot \hat{\kappa}\langle a', \theta' \rangle = \iint (\bar{\Lambda}^w - \bar{\Lambda}^b)\Gamma'(\bar{\kappa})\omega^* \hat{\kappa} da d\theta. \tag{36}$$

This operator plays a crucial role in capturing the effects of changes in the occupation switching threshold κ. The term $\Gamma'(\bar{\kappa})$ represents the density of agents at the switching threshold, while the difference $(\bar{\Lambda}^w - \bar{\Lambda}^b)$ captures how transition dynamics differ between workers and business owners.

3.5.3 Aggregation Operators

The aggregation operators serve to map from the distribution to aggregate variables, with the overall aggregation expressed as:

$$\int \overline{x} d\hat{\Omega}_t = \mathcal{I}^{(a)} \cdot \hat{\omega}_t^{(a)} + \mathcal{I} \cdot \hat{\omega}_t. \tag{37}$$

The operator $\mathcal{I}^{(a)}$ handles small changes and consists of two terms:

$$\mathcal{I}^{(a)} \cdot \hat{\omega}^{(a)} = \iint \sum_d \overline{x}_a^d \Gamma^d \hat{\omega}^{(a)} da d\theta + \iint (\overline{x}^w - \overline{x}^b) \Gamma'(\bar{\kappa}) \bar{\kappa}_a \hat{\omega}^{(a)} da d\theta. \tag{38}$$

The first term in this operator captures the direct effect of state changes on aggregates within each occupation, while the second term accounts for how state changes affect occupation switching decisions. Together, these terms provide a complete picture of how small changes affect aggregate outcomes.

The operator $\mathcal{I}$ handles distribution changes directly and is defined as:

$$\mathcal{I} \cdot \hat{\omega} = \iint \sum_d \overline{x}^d \bar{\Gamma}^d \hat{\omega} da d\theta. \tag{39}$$

This operator weights individual choices by occupation probabilities to compute aggregate quantities, providing a direct link between distribution changes and aggregate outcomes.

The aggregation operators possess several important properties that facilitate computation. Their linearity allows for the use of standard linear algebra techniques, while their inherent sparsity—as transitions are only possible between nearby points in the state space—enables efficient computation. The operators are constructed using steady-state objects ($\bar{\Lambda}^d$, $\bar{\Gamma}^d$, ω^*), which contributes to their computational efficiency. Furthermore, they provide a clean separation between intensive (within-occupation) and extensive (switching) margins of adjustment.

In numerical implementation, these operators are constructed as sparse matrices using the steady-state transition matrix. The sparsity pattern reflects the structure of possible transitions in the state space, with non-zero elements corresponding to feasible transitions. The time-shift operators exhibit particular sparsity, as transitions typically occur only between nearby points in the state space. While the policy-shift operators inherit this sparsity pattern, they may have additional non-zero elements due to occupation switching. The aggregation operators are less sparse but are applied less frequently in the computation.

This operator-based approach offers several distinct advantages. It provides a clear separation of different economic forces, allowing us to distinguish between intensive and extensive margin adjustments. The implementation through sparse matrix operations ensures computational efficiency, while each component maintains a natural interpretation in terms of economic behavior. The framework also scales effectively to problems with larger state spaces or additional occupational choices, making it a powerful tool for analyzing complex economic environments.

3.6 Computing the Aggregate Path

The sequence $\{\hat{X}_t\}_t$ is determined by a system of linear equations that combines all the previously described operators with equilibrium conditions. For all periods t, the system takes the form:

$$\mathsf{G}_Y \hat{Y}_t + \mathsf{G}_x \sum_{s=0}^{\infty} \mathsf{J}_{t,s} \hat{Y}_s + \mathsf{G}_x \mathsf{J}_t^{TD} = 0, \tag{40}$$

where the vector $\hat{Y}_t$ collects current and adjacent period responses:

$$\hat{Y}_t = \left[\mathsf{P}\hat{X}_{t-1}, \hat{X}_t, \hat{X}_{t+1}\right]^{\mathrm{T}}. \tag{41}$$

The Jacobian terms $\mathsf{J}_{t,s}$ evolve recursively according to:

$$\mathsf{J}_{t,s} = \mathsf{J}_{t-1,s-1} + \sum_d \mathcal{IL}_t \cdot \mathcal{M}^{(a),d} \cdot \mathsf{px}_s^d + \mathcal{I} \cdot \left(\bar{\Lambda}\right)^{t-1} \cdot \mathcal{M} \cdot \left(\mathsf{v}_s^b - \mathsf{v}_s^w\right), \tag{42}$$

with

$$\mathcal{IL}_t = \mathcal{IL}_{t-1} \cdot \mathcal{L}^{(a)} + \mathcal{I} \cdot \left(\bar{\Lambda}\right)^{t-2} \cdot \mathcal{L}. \tag{43}$$

The initial conditions for these recursions are given by:

$$\mathsf{J}_{0s} = \int \mathsf{x}_s d\Omega^* + \int (\bar{x}^w - \bar{x}^b)\Gamma'(\bar{\kappa})\kappa_s d\Omega^* \tag{44}$$

and

$$\mathsf{P}\hat{X}_{-1} = \mathsf{P}\left(\overline{X}_{-1} - \overline{X}\right). \tag{45}$$

The term J_t^{TD} captures the pure distributional effects:

$$\mathsf{J}_t^{TD} = \mathcal{I} \cdot \Lambda^t \cdot \hat{\omega}_0. \tag{46}$$

Finally, there is a transversality condition that requires that $\lim_{t\to\infty} \hat{X}_t = 0$.

3.6.1 Numerical Implementation

The practical implementation of the method requires several key components. The user provides steady-state policy functions (typically as splines) and equations describing the competitive equilibrium. The first-order approximation requires computing G_x, G_Y, J_t^{TD}, and the sequence $\{\mathsf{J}_{t,s}\}_{t,s}$.

The matrices G_x and G_Y are obtained through automatic differentiation of the equilibrium conditions G evaluated at the steady state. For $\mathsf{J}_{t,s}$ and J_t^{TD}, we construct the required operators $\mathcal{I}^{(a)}$, $\mathcal{I}$, $\mathcal{L}^{(a)}$, $\mathcal{L}$, $\mathcal{M}^{(a)}$, and $\mathcal{M}$ as sparse matrices using the steady-state transition matrix. This sparsity is essential for computational efficiency.

The coefficients $\{\mathsf{x}_t\}_t$ are computed recursively using linear algebra with pre-computed basis matrices. This approach avoids the need for numerical differentiation, making the computation both faster and more accurate. Specifically, for any occupation d, these coefficients satisfy:

$$\begin{aligned}\mathsf{x}_0^d(a,\theta) &= -\left(\mathsf{F}_x^d(a,\theta) + \mathsf{F}_{x^e}^d(a,\theta)\,\mathbb{E}\left[\overline{x}_a|a,\theta\right]\mathsf{p}\right)^{-1}\mathsf{F}_Y^d(a,\theta)\\ \mathsf{x}_{s+1}^d(a,\theta) &= -\left(\mathsf{F}_x^d(a,\theta) + \mathsf{F}_{x^e}^d(a,\theta)\,\mathbb{E}\left[\overline{x}_a|a,\theta\right]\mathsf{p}\right)^{-1}\mathsf{F}_{x^e}^d(a,\theta)\,\mathbb{E}\left[\mathsf{x}_s|a,\theta\right].\end{aligned}$$

The final step involves solving the linear system for $\{\hat{X}_t\}_{t=0}^T$ by truncation at some sufficiently large horizon T. This truncation is justified by the transversality condition, which ensures that the approximation becomes arbitrarily accurate as T increases.

4 Empirical Application

5 Application: Optimal Business Taxation

5.1 Overview

To demonstrate our method, we analyze optimal business taxation in an economy with occupational choice between paid employment and entrepreneurship. This application builds on the literature studying fiscal policy in incomplete market settings pioneered by Aiyagari (1994) and extended to entrepreneurial economies by Quadrini, Vincenzo (2000), Cagetti and De Nardi (2006), and Meh (2005). Recent work by Guvenen et al. (2019) has emphasized the importance of firm heterogeneity for optimal tax policy, while Kitao (2008) specifically examines entrepreneurial taxation with financial frictions. Our analysis extends these insights by fully incorporating transition dynamics and occupational choice.

5.2 Model Environment

The model combines elements from the incomplete markets framework of Aiyagari (1994) and the entrepreneurial choice setting of Evans and Jovanovic (1989). Workers and entrepreneurs face uninsurable productivity risk and make forward-looking savings decisions. Financial frictions, modeled as in Buera et al. (2011), limit entrepreneurs' ability to leverage their businesses. This creates potentially important interactions between tax policy and credit constraints, as highlighted by Scheuer (2014).

5.3 Calibration

The period utility function takes the CRRA form:

$$U(c) = \frac{c^{1-\mathrm{ra}}}{1-\mathrm{ra}} \tag{47}$$

where we set risk aversion $\mathrm{ra} = 2$.

The productivity processes for both paid employment and entrepreneurship follow finite-state Markov chains calibrated to match empirical earnings distributions. Labor income risk parameters are set following Storesletten et al. (2004), while entrepreneurial income processes are calibrated to match the right tail of the income distribution as documented by Bhandari et al. (2024).

A key parameter governing financial frictions is the collateral constraint parameter χ, which determines the maximum leverage ratio entrepreneurs can utilize, following the approach of Evans and Jovanovic (1989) and Buera et al. (2011). Our baseline calibration sets $\chi = 1.5$, implying a maximum leverage ratio of $\frac{1}{3}$. For robustness, we also consider $\chi = 1.0$ (no leverage) and $\chi = 2.0$ (high leverage).

The preference shocks $\eta_{i,t}$ follow a Gumbel distribution with scale parameter σ_ε close to zero, building on the discrete choice literature (Rust, 1987). This

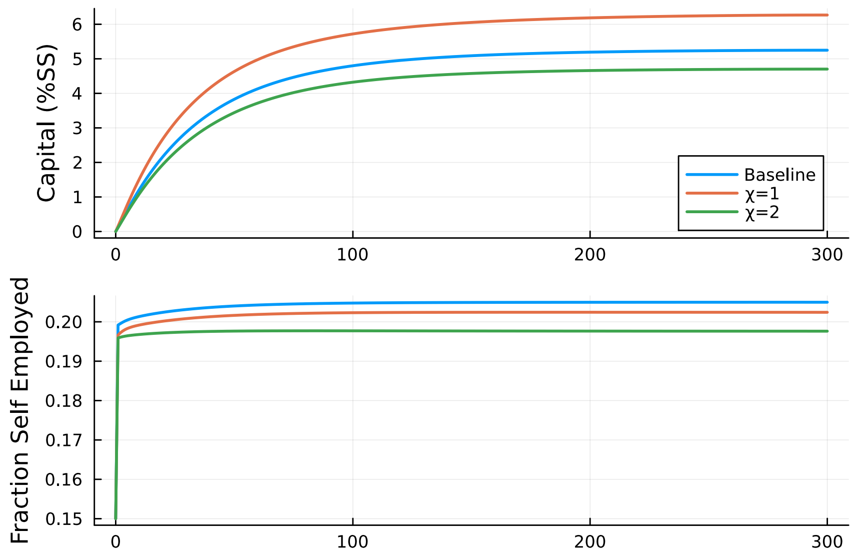

Figure 1: Transition path for capital and fraction of self-employed after a reduction in business tax

approximates a pure discrete choice model while maintaining smoothness needed for our perturbation approach.

5.4 Transition Dynamics

Figure 1 illustrates the transition dynamics following a reduction in the business tax rate from 30 to 10 percent for different levels of financial constraints. The analysis of such transitions builds on methods developed by Huggett (1997) and extended by Dyrda and Pugsley (2019) to study tax reforms. Panel A shows the path of aggregate capital (relative to initial steady state), while Panel B displays the evolution of the entrepreneurship rate.

Several key patterns emerge. First, we find that the reduction in business taxes leads to substantial capital accumulation, with the aggregate capital stock increasing by 4.5-6 percent in the long run depending on financial constraints. The transition is gradual, taking approximately 100-150 periods to reach the new steady state. Second, the long-run capital response is monotonically decreasing in the collateral constraint parameter χ. When $\chi = 1$ (no leverage), the capital stock increases by about 6 percent. With $\chi = 2$ (high leverage), the increase is only about 4.5 percent. This pattern aligns with findings of Midrigan and Xu (2014) about how financial frictions amplify policy effects. Third, the entrepreneurship rate jumps immediately by about 4-5 percentage points and then continues to rise gradually. The initial jump reflects the direct effect of lower taxes on business returns, while the subsequent increase captures the dynamic effects of capital accumulation, similar to the channels highlighted by Kitao (2008). Finally, tighter financial constraints (lower χ) are associated with larger increases in entrepreneurship rates. This reflects how constraints create a pool of "marginal entrepreneurs" who switch occupation when taxes are reduced.

5.5 Optimal Tax Policy

We now characterize optimal business tax policy accounting for transition dynamics. The welfare criterion is the population-weighted average of lifetime utility:

$$W_0 \equiv \int v_{0,i,t} di. \tag{48}$$

Several economic forces shape the optimal tax rate, building on insights from Saez and Zucman (2019) and Scheuer (2014). Positive business taxes redistribute from entrepreneurs to workers. Capital taxes reduce inefficient precautionary savings, as emphasized by Aiyagari (1994). Business taxes distort extensive-margin occupational choice and the allocation of capital across business owners, as discussed by Buera et al. (2011) and Guvenen et al. (2019).

Table 1 presents optimal tax rates assuming differences in the severity of financial frictions. The key finding is that optimal business tax rates are generally negative, especially when financial frictions are severe. This extends results from Meh (2005) and Kitao (2008) by showing how transi-

Maximum Leverage	χ	Optimal τ	Welfare Gain (%C)
0	1.0	-0.21	2.5
33	1.5	-0.08	1.8
100	2.0	0.00	1.6

Table 1: Optimal Business Tax Rates and Welfare Gains

tion dynamics and financial frictions interact to shape optimal policy. The welfare gains from tax optimization are moderate, ranging from 1.6 to 2.5 percent of consumption, similar in magnitude to gains found by Moll (2014) in related contexts.

The key finding is that optimal business tax rates are generally negative, especially when financial frictions are severe. This extends results from Meh (2005) and Kitao (2008) by showing how transition dynamics and financial frictions interact to shape optimal policy. The welfare gains from tax optimization are moderate, ranging from 1.6 to 2.5 percent of consumption, similar in magnitude to gains found by Moll (2014) in related contexts.

6 Conclusion

This paper develops a new approach for analyzing policy reforms in environments with discrete choice. By extending perturbation methods to handle discrete choices, we provide a tractable way to compute transition dynamics in rich heterogeneous agent models. The method's key innovation is properly accounting for both "small" changes conditional on discrete choices and "big" changes from switches in discrete choices.

Our application to business taxation demonstrates the method's utility. We find that optimal business tax rates are negative when financial frictions are severe, reflecting the value of subsidizing entrepreneurship to overcome credit constraints. The analysis also highlights the importance of accounting for transition dynamics when evaluating policy reforms.

Future work could extend the method to handle multiple discrete choices or apply it to other settings like default, housing, or labor market policies. The approach may also prove useful for estimation of structural models with discrete choice.

References

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3), 659-684.

Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2021). Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5), 2375-2408.

Bhandari, A., Bourany, T., Evans, D., & Golosov, M. (2023). A perturbational approach for approximating heterogeneous agent models. Working Paper.

Bhandari, A., Kass, T., May, T. J., McGrattan, E. R., & Schulz, E. (2024). On the nature of entrepreneurship. *Internal Revenue Service Working Paper.*

Boppart, T., Krusell, P., & Mitman, K. (2018). Exploiting MIT shocks in heterogeneous-agent economies: The impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, 89, 68-92.

Buera, F. J., Kaboski, J. P., & Shin, Y. (2011). Finance and development: A tale of two sectors. *American Economic Review*, 101(5), 1964-2002.

Cagetti, M., & De Nardi, M. (2006). Entrepreneurship, frictions, and wealth. *Journal of Political Economy*, 114(5), 835-870.

Catherine, S. (2019). Keeping options open: What motivates entrepreneurs? *Working Paper.*

Dyrda, S., & Pugsley, B. W. (2019). Taxes, private equity, and evolution of income inequality in the US. *Working Paper*.

Evans, D. S., & Jovanovic, B. (1989). An estimated model of entrepreneurial choice under liquidity constraints. *Journal of Political Economy*, 97(4), 808-827.

Guvenen, F., Kambourov, G., Kuruscu, B., Ocampo-Diaz, S., & Chen, D. (2019). Use it or lose it: Efficiency gains from wealth taxation. *Working Paper*.

Huggett, M. (1997). The one-sector growth model with idiosyncratic shocks: Steady states and dynamics. *Journal of Monetary Economics*, 39(3), 385-403.

Meh, C. A. (2005). Entrepreneurship, wealth inequality, and taxation. *Review of Economic Dynamics*, 8(3), 688-719.

Midrigan, V., & Xu, D. Y. (2014). Finance and misallocation: Evidence from plant-level data. *American Economic Review*, 104(2), 422-458.

Moll, B. (2014). Productivity losses from financial frictions: Can self-financing undo capital misallocation? *American Economic Review*, 104(10), 3186-3221.

Kanwal, R. P. (1998). Generalized functions: Theory and technique. *Springer Science & Business Media*.

Kaplan, G., Moll, B., & Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3), 697-743.

Kitao, S. (2008). Entrepreneurship, taxation and capital investment. *Review of Economic Dynamics*, 11(1), 44-69.

Quadrini, V. (2000). Entrepreneurship, Saving, and Social Mobility. *Review of Economic Dynamics*, 3(1), 1-40.

Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3), 649-665.

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, 55(5), 999-1033.

Saez, E., & Zucman, G. (2019). Progressive wealth taxation. *Brookings Papers on Economic Activity*, 2019(2), 437-511.

Scheuer, F. (2014). Entrepreneurial taxation with endogenous entry. *American Economic Journal: Economic Policy*, 6(2), 126-163.

Storesletten, K., Telmer, C. I., & Yaron, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3), 695-717.