

Transitional Market Dynamics in Complex Environments*

C. Lanier Benkard
Stanford University and NBER

Przemyslaw Jeziorski
U.C. Berkeley

Gabriel Y. Weintraub
Stanford University

September, 2024

Abstract

This paper presents a new approach to modeling transitional dynamics in dynamic models of imperfect competition, a crucial yet often neglected aspect of empirical models in industrial organization that seek to understand market responses to policy and environmental changes. We introduce Nonstationary Oblivious Equilibrium (NOE), a computationally efficient equilibrium concept based on a mean-field approximation designed to model short- and medium-run market dynamics. Addressing potential limitations of NOE in more concentrated markets or under aggregate shocks, we propose a variant, NOE with Re-solving (RNOE). RNOE modifies firms' strategies by re-computing NOE as industry states get realized; an iterative process inspired by real-world industry practice that has behavioral appeal. We show the potential of NOE and RNOE by applying them to an empirical setting of technology adoption and to two classic dynamic oligopoly models, demonstrating that, in a wide variety of settings of empirical interest, they generate equilibrium behavior that is close to Markov perfect equilibrium in both the short and long runs.

1 Introduction

Many empirical applications in industrial organization (IO) study markets with imperfect competition and forward looking agents. Modeling and understanding strategic behavior in these markets is crucial to understanding such classic topics in IO as R&D investment, entry/exit, mergers and acquisitions, and learning-by-doing. In recent decades, there has been an explosion in the literature studying dynamic models of imperfect

*This paper grew out of (and substantially extends) our 2008 working paper "Nonstationary Oblivious Equilibrium". We would like to thank Ben Van Roy for numerous fruitful conversations and valuable insights provided at the beginning stages of this project. This paper has also benefitted from conversations with Nikhil Agarwal, Paulo Somaini, Ali Yurukoglu, and seminar participants at various conferences and institutions. Correspondence: lanierb@stanford.edu; przemekj@haas.berkeley.edu; gweintra@stanford.edu.

competition in diverse empirical settings (e.g., see Aguirregabiria et al. (2021) for an extensive survey of the literature). Because of the complexity of even the simplest models of dynamic games, however, researchers are often presented with a choice between dynamic realism (forward-looking firm behavior) and static complexity (firm heterogeneity or allowing for a sufficient number of players or state variables). The canonical solution concept used in this literature, Markov Perfect Equilibrium (MPE), frequently precludes obtaining both goals simultaneously. Thus, many applications are forced to compromise on one dimension or the other, sacrificing some economic realism.

In recent years, the literature has developed methods for analyzing complex dynamic games that can help alleviate this trade-off. These advancements have typically been obtained either through computational approximations to MPE¹ or through the introduction of new, possibly more realistic, equilibrium concepts.² These new models have been shown to work well under various conditions when the object of study is *long-run industry behavior*. However, in many applications central to IO, we are interested in short- and medium-run industry behavior, and a model of long-run behavior may not provide a realistic approximation to these. A natural example is finite horizon problems starting from a particular initial state of interest, for which focusing on the stationary distribution over the recurrent class may not be appropriate (e.g., see Buchholz, 2022; Chen and Jeziorski, 2022). Other examples include applications to policy and regulation in which a policy change alters the equilibrium behavior and long-run distribution of states. If the stationary distribution of the industry changes, it may be equally important to compute the new stationary distribution as to describe the transitional dynamics between the old and new stationary distributions. For example, we may want to know how long it takes for an R&D subsidy to drive a meaningful amount of investment, or how long it takes for sufficient entry to occur after a merger. Transitional dynamics may also be important when studying new markets, where the central questions often surround entry dynamics or product adoption speed.

This paper proposes a framework for studying the transitional dynamics of imperfectly competitive markets with forward-looking firms or agents. With this motivation, we introduce a novel equilibrium concept that we call Nonstationary Oblivious Equilibrium (NOE). The framework has two goals: (i) to obtain a conceptually and computationally light model of transitional dynamics that closely follows MPE in empirically realistic models for which MPE is infeasible to compute; and (ii) to obtain a behavioral model that is a plausible approximation of the decision making process that firms may use.

The canonical form of our model extends oblivious equilibrium (OE) to model transitional dynamics

¹See, for example, Pakes and McGuire (2001), Farias et al. (2012), and Sweeting (2013).

²See Weintraub et al. (2008), Fershtman and Pakes (2012), Benkard et al. (2015) and Ifrach and Weintraub (2016).

for industries with many firms (e.g., tens of firms) that are not too concentrated. OE was introduced in Weintraub et al. (2008) to study the long-run behavior of the industry; our approach naturally extends it to study transitional dynamics. The critical feature of NOE is that firms rely on a mean-field approximation of the short- and medium-run industry state trajectory, in addition to an approximation based on an OE stationary distribution for the long run. Some of the approaches discussed above to study long-run dynamics can potentially be modified to analyze transitional dynamics (e.g., see the discussions in Ifrach and Weintraub, 2016 and Farias et al., 2012). However, we believe that NOE, and the extension discussed below of RNOE, provide a simpler and computationally lighter alternative to those approaches. In addition, since they explicitly model transitional dynamics, we think they are a more natural starting point for studying short-run dynamics.

In NOE, firms optimize against a deterministic trajectory for the industry state. This is a reasonable behavioral assumption when there are enough firms so that idiosyncratic shocks approximately average out at the industry level. However, when there are a small number of firms, and particularly when there are one or two dominant firms, the actual evolution of the industry can deviate quite a bit from its expectation. In NOE, firms also optimize against a deterministic trajectory for the aggregate shock, which may not be a good behavioral assumption if the aggregate shock is highly variable. Both of these issues can cause NOE strategies and transitional dynamics to deviate from those of MPE.

Thus motivated, we introduce an equilibrium concept that extends NOE, that we call NOE with Re-solving (RNOE). We find that RNOE closely follows MPE for oligopolistic markets with a small and medium number of firms in both the short and long runs, even in the presence of a highly variable aggregate shock. In a RNOE, firms use initial period NOE strategies in every period, i.e., in each period they “wake up” and observe the *realized* industry state (as opposed to the expected state), and then reoptimize, solving for a new NOE given that new initial state. They then play the reoptimized initial period NOE strategies for that period. This process is repeated in every period until steady state is attained. Note that this model is behavioral because when firms compute their optimal strategies they do not account for the fact that they will reoptimize in future periods. Not accounting for future reoptimization is what makes RNOE computation so simple: the computational cost of RNOE is equivalent to the cost of solving NOE multiple times (once for each time period until steady state is attained). Meanwhile, RNOE is useful because it allows firms to incorporate new information in every period about the current industry state and the current level of the aggregate shock, both of which may differ significantly from the expected state when the number of firms is small.

Such a heuristic can also mimic real-world firm decision making. Similar approaches that use re-solving,

dubbed as model predictive or receding horizon control, have been widely used for decades in the optimization of industrial processes such as chemical plants, oil refineries, and power systems (Camacho and Bordons, 2004). Re-solving techniques based on large-scale ‘fluid’ approximations applied to single-agent problems have also become popular in business applications, including network revenue management (Phillips, 2021), retailing (Acimovic and Farias, 2019), and in modern online platforms such as bidding in online advertising auctions (Balseiro and Gur, 2019) and online matching (Vera and Banerjee, 2021). Balseiro et al. (2023) provides a survey of such ‘fluid certainty equivalent control heuristics’ in dynamic resource allocation problems arising in Operations Research applications, discussing conditions under which re-solving can achieve provably significantly better performance than using the ‘fluid’ approximation just once. Cai and Judd (2023) provides a recent application of these ideas in single-agent dynamic stochastic problems arising in economics. Because re-solving is widely used in industrial and practical applications, we believe RNOE provides an appealing behavioral model.

After introducing NOE and RNOE, we consider two applications. First, we apply NOE to compute the equilibria of the Ryan and Tucker (2012b) network adoption model. Ryan and Tucker (2012b) consider forward-looking agents choosing whether to adopt teleconferencing technology in the presence of network effects. They introduce a game with many agents and dozens of observable agent types depending on geographical location, department, and seniority. The setting leads to a natural economic question of what is the optimal adoption subsidy. On the one hand, the firm wants to subsidize key leaders who exert positive externalities on other types. On the other hand, the firm wants to subsidize types with relatively more significant adoption costs. Computing the optimal subsidy requires obtaining counterfactual adoption paths by recomputing the adoption equilibrium. However, in this case directly applying MPE is highly impractical due to the high-dimensional state space. Common approaches to obtaining tractability would be to reduce the number of agent types or to assume that the agents are myopic. Applying NOE allows the researcher to retain the full type heterogeneity while also modeling forward-looking behavior. Computing the optimal adoption subsidy highlights the benefits of our approach, since both the stationary distribution (i.e., the total number of adopters) and transitional dynamics (i.e., speed of adoption) are of potential interest to the firm. Overall, the optimal subsidy, calculated using NOE, increases adoption of the technology by more than 20% compared to a subsidy that targets all types uniformly.

Second, we apply NOE and RNOE to two classic dynamic oligopoly models with only a few active firms: differentiated products Bertrand competition with product R&D, and Cournot competition with investment that reduces marginal cost. We consider a wide range of parameterizations that generate symmetric as well as concentrated industry structures. Our most extreme parameterization delivers a “winner-take-all” industry

in which one firm with superior marginal cost takes over the whole market for extended periods of time, and where the identity of the leading firm periodically changes depending on random investment outcomes. Such industries are notoriously hard to analyze using mean-field approximations because they ignore whether the industry leader is present in any given period. We demonstrate that NOE delivers transitional dynamics and welfare outcomes close to those of MPE in all but the most concentrated industry configurations. In the concentrated industry configurations, “re-solving” delivers significant performance improvements with a moderate additional computational burden. RNOE also closely matches MPE in industries with aggregate shocks.

Re-solving techniques have been used in the context of specific game-theoretic settings such as dynamic pricing games of perishable assets (Gallego and Hu, 2014). In addition, model predictive control has been applied in ‘mean-field games’ arising in the control literature (Degond et al., 2014). However, as far as we know, we are the first to apply re-solving techniques based on large-scale approximations to study more generic oligopolies in complex dynamic environments of the type that are of interest in IO.³

Since we introduced NOE in a previous version of this paper, it has found application in many empirical settings (e.g., see Wilson, 2012; Qi, 2013; Pavanini, 2014; Saeedi, 2019; Sweeting, 2015; Igami, 2017; Bian, 2018; Caoui, 2023; Dimitrellos, 2022; George, 2022; Buchholz, 2022; Johnston et al., 2023). It has been applied to study transitional dynamics in markets with many firms such as banking, real estate, ride sharing, and e-commerce. We believe that this version of the paper improves upon NOE and also broadens the set of applications to which it can be applied. Our hope is that this paper will further stimulate researchers to use NOE and RNOE to study transitional dynamics.

The paper is organized as follows. Section 2 introduces a canonical dynamic game with aggregate shocks and defines MPE. In Section 3, we introduce NOE and an algorithm to compute it. In this section, we also introduce the concept of re-solving. Section 4 contains an empirical application calibrated using the estimates from Ryan and Tucker (2012b). Section 5 presents numerical experiments of canonical dynamic oligopoly models. Section 6 concludes.

2 A Dynamic Model of Imperfect Competition

In this section we formulate a model of an industry in which are firms competing in an oligopolistic marketplace. The model encompasses the Ericson and Pakes (1995) model (henceforth, EP) and closely follows

³In the macroeconomics literature, a concept similar to NOE was introduced in a model with a continuum of agents by Costantini and Melitz (2008) – a non-stationary extension of Melitz (2003). In contrast, NOE (and RNOE) are intended to also study models with a finite number of firms such as those that are of interest in IO.

Weintraub et al. (2008).

2.1 Model and Notation

We introduce the main elements of our model.

Time Horizon and Agents. We index time periods with non-negative integers $t \in \{0, 1, 2, \dots\}$. Agents are either active or inactive. Inactive agents can enter the game and become active. In contrast, any active agent can exit the game. The set of active agents (incumbents) at time t is denoted by S_t . Each agent in S_t is assigned a unique positive integer-valued index, denoted by i .

State Space. Firm heterogeneity is reflected through agent states, x_{it} . We assume that the state space is finite. Without loss of generality, we denote the values of the agent state using non-negative integers: $x_{it} \in \mathcal{X} = \{0, 1, 2, \dots, \bar{x}\}$. To fix an interpretation of the model, in the remainder of this section, we will refer to the active set of agents as *an industry*, to an agent as *a firm*, and to an agent's state as its *quality level*. For ease of exposition, we would suggest that readers have in mind an example where single product firms compete in a logit demand model (and this will also be one of our examples below). However, agents might be individuals or groups of individuals. States might more generally reflect productivity, capacity, the size of a firm's consumer network, or any other aspect of the firm that affects its profits. If the agents are individuals, the state may include any utility-relevant characteristics such as the amount of savings, or a status of costly technology adoption (see Section 4).

We define the *industry state*, s_t , to be a vector that lists the number of incumbent firms at each quality level in period t , for each possible quality level $x \in \mathcal{X}$. The state space $\bar{\mathcal{S}} = \left\{ s \in \mathbb{N}^{|\mathcal{X}|} \mid \sum_{x \in \mathcal{X}} s(x) < \infty \right\}$ is the set of all possible industry states. For an incumbent firm i , we define $s_{-i,t}$ to be the state of the *competitors* of firm i ; $s_{-i,t}$ is simply equal to s_t with firm i subtracted out. That is, $s_{-i,t}(x) = s_t(x) - 1$ if $x_{it} = x$, and $s_{-i,t}(x) = s_t(x)$, otherwise.⁴

In addition to the industry state we also allow for a common state (aggregate shock), z_t . We assume that z_t is a non-negative integer bounded by \bar{z} and denote the set of possible aggregate shocks as $\bar{\mathcal{Z}} = \left\{ z \in \mathbb{N}_+ \mid z \leq \bar{z} \right\}$.

Although in principle there are a countable number of industry states, for some of the technical arguments we will also consider an extended state space $\mathcal{S} = \left\{ s \in \mathfrak{R}_+^{|\mathcal{X}|} \mid \sum_{x \in \mathcal{X}} s(x) < \infty \right\}$. This notion will be useful, for example, when considering the expected value of the industry state. Similarly, we introduce

⁴Note that because we focus on symmetric and anonymous equilibrium strategies in the sense of Doraszelski and Pakes (2007), we restrict attention to industry states for which the identities of firms do not matter.

an extended aggregate shock space $\mathcal{Z} = \{z \in \mathfrak{R}_+ \mid z \leq \bar{z}\}$.

Actions. Each period, active firms can take simultaneous actions $\iota_{it} \in \mathcal{I}$. In the interpretation of the model given above, we think of these actions as investments in quality. However, in a broader interpretation actions may include pricing, capacity investment, or technology adoption. The action space \mathcal{I} may be a convex subset of the real line, for example if ι is investment amount or price, but it may also be a finite set such as $\{0, 1\}$, for example if the action is technology adoption.

Single-Period Profit Function. Each period, each incumbent firm earns profits on a spot market. A firm's single period expected profit is denoted by $\pi(x_{it}, s_{-i,t}, \iota_{it}, z_t)$, and depends on its quality level x_{it} , its competitors' state, $s_{-i,t}$, the action taken ι_{it} , and the level of the aggregate shock z_t . If the agents are individuals, profit π reflects agents' utility functions. We assume the single-period profit function is a continuous function of investment. We also assume that the single-period profit function is a continuous function of the industry state and the aggregate shock over the extended state spaces. Note that the current formulation presumes a stationary model since π does not depend on t , but the framework can be easily extended to a non-stationary setup in which profits depend on t , as long as they converge uniformly to a limit profit function as t grows.

To simplify exposition and analysis, we follow the EP model and further assume that the single-period profit function takes the form:⁵

$$\pi_t(x_{it}, s_{-i,t}, \iota_{it}, z_t) = \bar{\pi}(x_{it}, s_{-i,t}, z_t) - c(\iota_{it}, x_{it}, z_t).$$

The cost of investment is given by a non-negative function $c(\iota_{it}, x_{it}, z_t)$ that depends on the firm's individual state x_{it} , the investment level ι_{it} , and potentially the level of aggregate shock z_t . The dependence of the investment cost on the aggregate shock z_t allows for low-investment-cost and high-investment-cost periods.

Exit process. The model also allows for entry and exit. In each period, each incumbent firm observes a positive real-valued sell-off value ϕ_{it} that is private information to the firm. If the sell-off value exceeds the value of continuing in the industry, then the firm may choose to exit, in which case it earns the sell-off value and then ceases operations permanently. We assume that the random variables $\{\phi_{it} \mid t \geq 0, i \geq 1\}$ are i.i.d. and have a well-defined density function with finite moments.

Entry process. The model can also accommodate a variety of different entry processes. Specifically, the number of entrants can be fixed or random, entry costs can be fixed or random, and the entry state can be

⁵Our formulation follows EP in which players are coupled in the the single-period profit function only through their states. It is possible to allow the single-period profits to also depend on the competitors' actions (e.g., in a learning-by-doing model). This requires an extension of NOE that we will not present here to simplify the exposition.

predetermined, randomly determined, or even controlled (meaning the entrant can either choose the entry state or at least influence it). For concreteness, we assume the entry model of Weintraub et al. (2008). In that model, there are a large number of potential entrants who play a symmetric mixed strategy to enter the industry after paying a setup cost κ . In that case, we obtain that the number of firms entering during period t is a Poisson random variable that is conditionally independent of $\{\phi_{it}, |t \geq 0, i \geq 1\}$, conditioned on s_t .⁶ Entrants do not earn profits in the period that they enter. They appear in the following period at state x^e and can earn profits thereafter.

We denote the expected number of firms entering at industry state s_t , by $\lambda(s_t, z_t)$, and assume that entry rates are bounded by a constant $\bar{\lambda} > 0$. This state-dependent entry rate will be endogenously determined in equilibrium, and our solution concept will require that it satisfies a zero expected profit condition. As is common in this literature and to simplify the analysis, we assume potential entrants are short-lived and do not consider the option value of delaying entry. Potential entrants that do not enter the industry disappear and a new generation of potential entrants is created in the next period.

Transition dynamics. Given a firm's strategy μ and state at time t , the firm's transition to a state at time $t + 1$ is described by the following Markov kernel:

$$(2.1) \quad \mathcal{P}_\mu \left[x_{i,t+1} = x' \mid x_{it} = x; \iota_{it} = \iota, s_t = s \right].$$

Uncertainty in state transitions may arise, for example, due to the risk associated with a research and development endeavor or a marketing campaign. We assume that transitions are independent across firms conditional on the industry state and investment levels. We also assume these transitions are independent of all previously defined random quantities. The dependence of the kernel \mathcal{P}_μ on the industry state allows, for example, for the existence of spillover effects across firms. We also allow for deterministic transitions. We assume that the kernel is a continuous function of investment ι .⁷

Aggregate shocks are assumed to evolve in an exogenous independent fashion according to an irreducible Markov process with kernel $\mathcal{P} \left[z_{i,t+1} = z' \mid z_{it} = z \right]$. We assume that this process admits a unique invariant distribution.

Timing of Events. In each period, events occur in the following order: (1) Each incumbent firm observes its sell-off value and then makes exit and investment decisions; (2) The number of entering firms is determined and each entrant pays the entry cost; (3) Incumbent firms compete in the spot market and receive profits; (4)

⁶To avoid pathological entry behavior we also assume that $\kappa > \beta \cdot \bar{\phi}$, where $\bar{\phi}$ is the expected net present value of entering the market, investing zero and earning zero profits each period, and then exiting at an optimal stopping time.

⁷Like the single period-profit function, the transition kernel could also depend on competitors' actions.

Exiting firms exit and receive their sell-off values; (5) Individual states of each firm $x_{i,t+1}$ are determined, new entrants enter, and the industry takes on a new state s_{t+1} . New aggregate shock z_{t+1} is revealed to all firms.

Firms' objective. Firms aim to maximize expected discounted profits. The interest rate is assumed to be positive and constant over time, resulting in a constant discount factor of $\beta \in (0, 1)$ per time period. Finally, we assume that for all competitors' decisions and all continuation values, a firm's one time-step ahead optimization problem to determine its optimal investment has a unique solution.⁸

2.2 Equilibrium

As a model of industry behavior, we first introduce pure strategy Markov Perfect Equilibrium (MPE), in the sense of Maskin and Tirole (1988). We further assume that equilibrium is symmetric, such that all firms use a common investment/exit strategy. In MPE, incumbent firms follow a Markov (investment and exit) strategy that is self-generating in the sense of being optimal when all competitor firms follow the same strategy. At the same time, the equilibrium entry mixed strategy among potential entrants yields an entry rate such that entrants are exactly indifferent between entering and not entering. We now introduce some notation required to state these two properties precisely.

We denote each incumbent firm's investment strategy as $\iota(x_{it}, s_{-i,t}, z_t)$, which is a function of the firm's own state, its competitors' states and the common state. Incumbent firms also have an exit strategy, $\rho(x_{it}, s_{-i,t}, z_t)$, that takes the form of a cutoff rule: firm $i \in S_t$ exits at time t if and only if the sell-off value is greater than ρ , i.e., $\phi_{it} \geq \rho(x_{it}, s_{-i,t}, z_t)$.⁹ Let \mathcal{M} denote the set of exit/investment strategies $\mu = (\iota, \rho)$, and Λ denote the set of entry rate functions. An entry rate function $\lambda \in \Lambda$, takes the value $\lambda(s_t, z_t)$ at time period t .

We define the value function $V(x, s, z | \mu', \mu, \lambda)$ to be the expected net present value for a firm at state x when its competitors' state is s , the aggregate shock takes value z , given that its competitors each follow a common strategy $\mu \in \mathcal{M}$, the entry rate function is $\lambda \in \Lambda$, and the firm itself follows strategy $\mu' \in \mathcal{M}$:

$$V(x, s, z | \mu', \mu, \lambda) = E_{\mu', \mu, \lambda} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \pi(x_{ik}, s_{-i,k}, l_{ik}, z_k) + \beta^{\tau_i-t} \phi_{i, \tau_i} \mid x_{it} = x, s_{-i,t} = s, z_t = z \right],$$

⁸This assumption is similar to the unique investment choice admissibility assumption in Doraszelski and Satterthwaite (2010) that is used to guarantee the existence of an equilibrium to the model in pure strategies. It is satisfied by many of the commonly used specifications in the literature.

⁹Weintraub et al. (2008) shows that there always exists an optimal exit strategy of this form even among very general classes of exit strategies. Furthermore, without loss of generality throughout the paper, we restrict our attention to cut-offs lower than $\sup_{x,s,z} \pi(x, s, z) / (1 - \beta) + \bar{\phi}$.

where i is taken to be the index of a firm at quality level x at time t , τ_i is a random variable representing the time at which firm i exits the industry, and the subscripts of the expectation indicate the strategy followed by firm i , the strategy followed by its competitors, and the entry rate function. In an abuse of notation, we will use the shorthand $V(x, s, z|\mu, \lambda) \equiv V(x, s, z|\mu, \mu, \lambda)$, to refer to the expected discounted value of profits when firm i follows the same strategy μ as its competitors.

A MPE to our model comprises of an investment/exit strategy $\mu = (\iota, \rho) \in \mathcal{M}$, and an entry rate function $\lambda \in \Lambda$ that satisfy the following conditions:

1. Incumbent firm strategies represent a MPE:

$$(2.2) \quad \sup_{\mu' \in \mathcal{M}} V(x, s, z|\mu', \mu, \lambda) = V(x, s, z|\mu, \lambda) \quad \forall x \in \mathcal{X}, \forall s \in \bar{\mathcal{S}}, z \in \bar{\mathcal{Z}}.$$

2. At each state, either entrants have zero expected profits or the entry rate is zero (or both):

$$\begin{aligned} \sum_{s \in \bar{\mathcal{S}}, z \in \bar{\mathcal{Z}}} \lambda(s, z) (\beta E_{\mu, \lambda} [V(x^e, s_{-i, t+1}, z_{t+1}|\mu, \lambda)|s_t = s, z_t = z] - \kappa) &= 0 \\ \beta E_{\mu, \lambda} [V(x^e, s_{-i, t+1}, z_{t+1}|\mu, \lambda)|s_t = s, z_t = z] - \kappa &\leq 0 & \forall s \in \bar{\mathcal{S}}, z \in \bar{\mathcal{Z}} \\ \lambda(s, z) &\geq 0 & \forall s \in \bar{\mathcal{S}}, z \in \bar{\mathcal{Z}}. \end{aligned}$$

Doraszelski and Satterthwaite (2010) establish existence of an equilibrium in pure strategies for a closely related model. We do not provide an existence proof here because it is long and cumbersome and would replicate this previous work. With respect to uniqueness, we presume that our model may have multiple equilibria.¹⁰

There are a wide variety of dynamic programming algorithms available that can be used to solve for MPE to our model. However, these algorithms require compute time and memory that grow at least proportionately with the number of relevant industry states, which in turn grows quickly (high-order polynomial) with the number of individual states and firms in the industry. In the empirical models we have in mind, the number of industry states is typically so large that it would not be possible to store the strategy function on a computer, let alone compute it. This difficulty motivates our alternative approach.

¹⁰Besanko et al. (2010) provide an example of multiple equilibria in their closely related model.

3 Nonstationary Oblivious Equilibrium

Oblivious equilibrium (OE) as introduced in Weintraub et al. (2008) offers a way to approximate *long-run* Markov perfect industry dynamics. Using a similar idea of averaging effects, in this section we introduce a new equilibrium concept, nonstationary oblivious equilibrium (NOE), that can be used to approximate the *short-run* dynamic behavior (transitional dynamics) of an industry starting from an initial state of interest. To simplify notation, we assume that the industry is at the initial state of interest, denoted (s_0, z_0) , at time period $t = 0$. Each initial state generates different transitional dynamics and thus different associated NOE. We provide an important extension of NOE in Section 3.6 by introducing re-solving.

OE is based on the idea that when there are a large number of firms (and no aggregate shocks), simultaneous changes in individual firm quality levels average out such that in the long run the industry state remains roughly constant over time. Hence, in OE, firms assume that the industry state is constant and corresponds to the long-run expected state given the strategies being played. This dramatically simplifies computation because in OE strategies only depend on the firms' own states.

Based on a similar notion, we introduce a method to model the short-run behavior of an industry that starts from a given state of interest. If there are a large number of firms (and no aggregate shocks), the industry state starting from a given initial state roughly follows a deterministic trajectory. In this setting, each firm can make near-optimal decisions based only on its own quality level and by knowing the deterministic trajectory followed by the industry state. With this motivation, we consider restricting firm strategies so that each firm's decisions depend only on the firm's quality level and the time period. We call such restricted strategies *nonstationary oblivious* since they involve decisions made without full knowledge of the circumstances — in particular, the state of the industry. Note that nonstationary oblivious strategies generalize oblivious strategies because they depend on the time period.

It is worth noting that we do not expect nonstationary oblivious strategies to necessarily perform well in the presence of aggregate shocks unless their evolution is deterministic or has small variance around its mean. Weintraub et al. (2010) extends the concept of OE so that strategies are a function of the aggregate shock. Although we could employ a similar method with NOE, we choose not to because it leads to a more complex equilibrium concept and calculation. Instead, we propose re-solving as an alternative avenue to deal with aggregate shocks.

3.1 Nonstationary Oblivious Strategies and Entry Rate Functions

Let us start by denoting $\tilde{\mathcal{M}} \subset \mathcal{M}$ as the set of oblivious strategies, that is, strategies that are only a function of the firm's own state. A nonstationary oblivious strategy is a sequence of oblivious strategies, so let $\tilde{\mathcal{M}}_{ns} = \tilde{\mathcal{M}}^\infty \subset \mathcal{M}^\infty$ denote the set of nonstationary oblivious strategies. Hence, if $\mu \in \tilde{\mathcal{M}}_{ns}$, then $\mu = \{\mu_0, \mu_1, \dots\}$, where for each time period $t \geq 0$, $\mu_t \in \tilde{\mathcal{M}}$ is an oblivious strategy. To illustrate, if firm i uses strategy $\mu \in \tilde{\mathcal{M}}_{ns}$, then at time period t , firm i takes action $\mu_t(x_{it})$, so the action depends both on the time period and the firm's own state. In a NOE firms will make decisions assuming that the industry state, as well as the aggregate shock evolve deterministically. Therefore, under this assumption, the time period determines the industry state and common shock.

Similarly, let $\tilde{\Lambda}_{ns} = \tilde{\Lambda}^\infty \subset \Lambda^\infty$ denote the set of nonstationary oblivious entry rate functions, where $\tilde{\Lambda} \subset \Lambda$ denotes the set of oblivious entry rates. A nonstationary oblivious entry rate function is a sequence of oblivious entry rates. Hence, if $\lambda \in \tilde{\Lambda}_{ns}$, then $\lambda = \{\lambda_0, \lambda_1, \dots\}$, where for every period $t \geq 0$, λ_t is a real nonnegative number that represents the entry rate at time period t .

3.2 Sequence of Expected states and Value Functions

Suppose that all firms use a common strategy $\mu \in \tilde{\mathcal{M}}_{ns}$ and, with some abuse of notation, let an individual firm transition kernel for period t be denoted by $\mathcal{P}_{\mu_t}(x'|x, s)$ (see equation (2.1)). Note that at each period, firms' transitions are independent conditional on the current industry state. If there were an infinite number of firms, though each evolves stochastically, the percentage of firms that transition from any individual state to another would be deterministic. Similarly, the percentage of firms that exit would be deterministic. Motivated by this fact, for $\mu \in \tilde{\mathcal{M}}_{ns}$, $\lambda \in \tilde{\Lambda}_{ns}$, we define the following sequence of (expected) industry states:

$$(3.1) \quad \tilde{s}_{t+1}(x'; \mu, \lambda, s_0) = \begin{cases} \sum_{x \in \mathcal{X}} \mathcal{P}_{\mu_t}(x'|x, \tilde{s}_t) \tilde{s}_t(x; \mu, \lambda, s_0) + \lambda_t & \text{if } x = x^e \\ \sum_{x \in \mathcal{X}} \mathcal{P}_{\mu_t}(x'|x, \tilde{s}_t) \tilde{s}_t(x; \mu, \lambda, s_0) & \text{otherwise.} \end{cases}$$

Note that the sequence of states is defined for every possible starting state $s_0 \in \bar{\mathcal{S}}$ by setting $\tilde{s}_0 = s_0$. In what follows we suppress the dependence of the sequence $\{\tilde{s}_t : t \geq 0\}$ on s_0 to simplify the notation.

We also define the sequence of expected aggregate shock states:

$$(3.2) \quad \vec{z}_{t+1}(z'; z_0) = \sum_{z \in \mathcal{Z}} \mathcal{P}(z'|z) \vec{z}_t(z),$$

$$(3.3) \quad \tilde{z}_{t+1} = \sum_{z' \in \mathcal{Z}} z' \vec{z}_{t+1}(z'; z_0).$$

Equation (3.2) determines the probabilities that the aggregate state is in each individual state $z \in \mathcal{Z}$ at a given time period according to the Markov transition kernel $\mathcal{P}(z'|z)$. Then (3.3) computes the expected value of the aggregate shock in that time period given the probability distribution. Note that the sequence of expected aggregate states $\{\tilde{z}_t : t \geq 0\}$ is defined for every possible starting state $z_0 \in \overline{\mathcal{Z}}$ by setting \vec{z}_0 to be the canonical vector with a one for value z_0 . In what follows we suppress the dependence of the sequence $\{\tilde{z}_t : t \geq 0\}$ on z_0 to simplify the notation.

For nonstationary oblivious strategies $\mu', \mu \in \tilde{\mathcal{M}}_{ns}$, a nonstationary oblivious entry rate function $\lambda \in \tilde{\Lambda}_{ns}$, and an initial industry state (s_0, z_0) , we define a *nonstationary oblivious value function* for period t :

$$(3.4) \quad \tilde{V}_t(x|\mu', \mu, \lambda, s_0, z_0) = E_{\mu'} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \pi(x_{i,k}, \tilde{s}_{-i,k}(\mu, \lambda), l_{i,k}, \tilde{z}_k) + \beta^{\tau_i-t} \phi_{i,\tau_i} \Big| x_{it} = x \right].$$

Note again that expected trajectories \tilde{s} and \tilde{z} depend on the starting state (s_0, z_0) , thus all nonstationary oblivious value functions depend on the starting state as well. These value functions should be interpreted as the expected net present value of a firm that is at quality level x at time t and follows nonstationary oblivious strategy μ' , under the assumption that, for all $t \geq 0$, its competitors' state will be given by $\tilde{s}_{-i,t}(\mu, \lambda)$ and the aggregate shock level will be given by \tilde{z}_t at time t .¹¹

The competitors' state $\tilde{s}_{-i,t}$ is obtained from the average industry state \tilde{s}_t by subtracting company i . A simple way to do this subtraction to proportionally subtract one firm, that is, set

$$\tilde{s}_{-i,t}(x) = \tilde{s}_t(x) - \frac{\tilde{s}_t(x)}{\sum_{x'} \tilde{s}_t(x')}.$$

If there is no entry and exit this method presumes that firm i is a uniform sample from firms in the starting state s_0 .¹²

¹¹The expectation over industry states and aggregate shocks is taken inside the profit function. Instead, Benkard et al. (2015) introduces a "simulated" version of OE, where the expectation over industry states is taken outside the profit function. This version tends to work better in industries with a small number of firms at the cost of additional simulations to compute such expectations. While a similar approach could be taken for NOE, we do not explore it in the current paper to avoid additional computational costs and because we focus on the use of re-solving to study industries with a small number of firms.

¹²Alternatively we could compute $\tilde{s}_{-i,t}$ by conditioning on the actual state of the subtracted firm i at time t , x_{it} . In such a

Although the firm's state trajectory only depends on the firm's own strategy μ' , the nonstationary oblivious value function remains a function of the competitors' strategy μ and the entry rate λ through the expected state trajectory $\tilde{s}_t(\mu, \lambda)$. We abuse notation by using $\tilde{V}_t(x|\mu, \lambda, s_0, z_0) \equiv \tilde{V}_t(x|\mu, \mu, \lambda, s_0, z_0)$ to refer to the nonstationary oblivious value function when firm i follows the same strategy μ as its competitors.

3.3 Equilibrium

Equipped with the above machinery, we now define our new solution concept. To avoid pathological behavior in which an entry rate grows large and is followed by massive exit, we restrict all entry rates to be less than a predetermined upper bound $\bar{\lambda}$.¹³ For a given starting state (s_0, z_0) , a *nonstationary oblivious equilibrium (NOE)* consists of a strategy $\mu \in \tilde{\mathcal{M}}_{ns}$ and an entry rate function $\lambda \in \tilde{\Lambda}_{ns}$ that satisfy the following conditions:

1. Firm strategies optimize a nonstationary oblivious value function:

$$(3.5) \quad \sup_{\mu' \in \tilde{\mathcal{M}}_{ns}} \tilde{V}_0(x|\mu', \mu, \lambda, s_0, z_0) = \tilde{V}_0(x|\mu, \lambda, s_0, z_0), \quad \forall x \in \mathcal{X}.$$

2. At every period of time, the nonstationary oblivious expected value of entry is zero or boundary conditions are satisfied. For all $t \geq 0$,

$$\begin{aligned} \lambda_t \in (0, \bar{\lambda}) \text{ implies } \beta \tilde{V}_{t+1}(x^e|\mu, \lambda, s_0, z_0) - \kappa &= 0 \\ \beta \tilde{V}_{t+1}(x^e|\mu, \lambda, s_0, z_0) - \kappa < 0 \text{ implies } \lambda_t &= 0 \\ \beta \tilde{V}_{t+1}(x^e|\mu, \lambda, s_0, z_0) - \kappa > 0 \text{ implies } \lambda_t &= \bar{\lambda}. \end{aligned}$$

Note that the optimization of \tilde{V}_0 implies, by dynamic programming principles, that firms optimize \tilde{V}_t for all $t \geq 0$.

We note that if we ignore aggregate shocks or if they follow a deterministic trajectory, a similar argument to the one provided in Weintraub et al. (2008) shows that NOE approximates MPE as the industry becomes large. We do not provide a proof to avoid replicating the argument.

case one would set $\tilde{s}_{-i,t} = E[s_{-i,t}|x_{it} = x]$. This formulation is computationally impractical since $\tilde{s}_{-i,t}$ would need to be recomputed for every x_{it} , which requires many additional computations. However, we have experimented with this method and found little difference in comparison with proportional subtraction. Consequently, we use proportional subtraction throughout the paper because of its computational simplicity.

¹³We assume $\bar{\lambda} \geq (\sup_{x,s,t,z} \pi(x, s, t, z)/(1 - \beta) + \bar{\phi})/\kappa$. It is simple to show that under mild conditions an OE entry rate must be less than $\bar{\lambda}$. Moreover, in our computational experiments we never observed NOE entry rates growing beyond $\bar{\lambda}$, so in practice the restriction was not binding. For this reason and to simplify the explanation, $\bar{\lambda}$ is omitted in the description of the algorithm in Section 3.5. We make use of $\bar{\lambda}$, however, in the existence proof that follows.

3.4 Existence of NOE that Become Stationary

In this paper, we focus on NOE that become stationary as time progresses. That is, we focus on NOE $(\mu, \lambda) \in \tilde{\mathcal{M}}_{ns} \times \tilde{\Lambda}_{ns}$ that converge to an OE $(\tilde{\mu}, \tilde{\lambda}) \in \tilde{\mathcal{M}} \times \tilde{\Lambda}$ in the following sense: for all $x \in \mathcal{X}$, $\lim_{t \rightarrow \infty} \mu_t(x) = \tilde{\mu}(x)$, and $\lim_{t \rightarrow \infty} \lambda_t = \tilde{\lambda}$. In this subsection, we show the existence of such NOE.

First, we define the set of converging nonstationary strategies and entry rate functions:

$$\begin{aligned} \hat{\mathcal{M}}_{ns} &= \{ \mu \in \tilde{\mathcal{M}}_{ns} : \text{for which there exists } \tilde{\mu} \in \tilde{\mathcal{M}}, \text{ such that, for all } x \in \mathcal{X}, \lim_{t \rightarrow \infty} \mu_t(x) = \tilde{\mu}(x) \}, \\ \hat{\Lambda}_{ns} &= \{ \lambda \in \tilde{\Lambda}_{ns} : \text{for which there exists } \tilde{\lambda} \in \tilde{\Lambda}, \text{ such that, } \lim_{t \rightarrow \infty} \lambda_t = \tilde{\lambda} \}. \end{aligned}$$

Note that $\tilde{\mu}$ and $\tilde{\lambda}$ in the definitions above are oblivious and entry rate strategies (not necessarily equilibrium). We make the following additional assumptions that we keep throughout this section. First, we assume that the investment space \mathcal{I} is a convex and compact subset of \mathfrak{R}_+^n , for some $n \geq 1$. Second, we assume that the random variables $\{\phi_{it} | t \geq 0, i \geq 1\}$ have support \mathfrak{R}_+ .

We prove the following result under these assumptions, the model assumptions introduced in Section 2.1, and the assumption that entry rates are bounded. The primary challenge in establishing existence lies in the fact that we are working within infinite-dimensional spaces. To address this, we employ the Brouwer-Schauder-Tychonoff fixed point theorem. A detailed proof can be found in the appendix.

Theorem 3.1. *For any given initial state (s_0, z_0) , there exists a NOE $(\mu, \lambda) \in \hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$, such that for all $x \in \mathcal{X}$, $\lim_{t \rightarrow \infty} \mu_t(x) = \tilde{\mu}(x)$ and $\lim_{t \rightarrow \infty} \lambda_t = \tilde{\lambda}$, where $(\tilde{\mu}, \tilde{\lambda}) \in \tilde{\mathcal{M}} \times \tilde{\Lambda}$ is an OE.*

We do not rule out the existence of multiple NOE; we focus on the study of NOE that are selected by the algorithm described below.

3.5 Algorithm to Compute NOE

The theorem above is important because it justifies an algorithm for computing NOE that become stationary and converge to OE. We impose this form of convergence in the algorithm, and then solve backwards for NOE. In this way, the problem of finding a NOE is reduced to a finite-horizon problem.

Suppose that we are mostly interested in the behavior of the industry in the interval between time periods $t = 0$ and $t = \underline{T}$. Let $\tilde{V}, \tilde{\mu}, \tilde{\lambda}, \tilde{s}, \tilde{z}$ be a (stationary) OE value function, strategy, entry rate, expected industry state, and expected common state, respectively. Let $\bar{T} := \min\{t | \beta^{t-\underline{T}} \max_{x \in \mathcal{X}} \tilde{V}(x) \leq \delta\}$, where $\delta > 0$ is a predetermined precision. We assume that there is a finite time horizon of length \bar{T} after which NOE

coincides with OE. More specifically, for all $t > \bar{T}$, $\mu_t = \tilde{\mu}$, $\lambda_t = \tilde{\lambda}$, $\tilde{s}_t = \tilde{s}$ and $\tilde{z}_t = \tilde{z}$. In addition, for $t > \bar{T}$, firms earn profits according to the OE value function. This simplification should not have a significant impact on the behavior of the industry for the time periods of interest between $t = 0$ and $t = \underline{T}$. After this reduction computing a NOE is simple; it requires solving finite-horizon one-dimensional dynamic programming problems.

At each iteration of the algorithm, we (1) compute the strategies that maximize the nonstationary oblivious value functions (step 12) and (2) we compute new entry rates depending on the extent of the violation of the zero-profit conditions (step 18) while also checking the boundary condition for the entry rate (steps 15 and 16). Strategies and entry rates are updated “smoothly” (steps 22 and 23). The parameters N_1 , N_2 , γ_1 , and γ_2 are set after some experimentation to speed up convergence. One step of the algorithm resembles a best-response operator against the current trajectory of expected industry states.

If $\epsilon_0 = 0$ and $\delta = 0$, and the termination condition of the outer loop is satisfied with $\epsilon_1 = \epsilon_2 = 0$, we have found an NOE. Small values of ϵ_0 , ϵ_1 , and ϵ_2 allow for small errors associated with limitations of numerical precision.

3.6 Nonstationary Oblivious Equilibrium with Re-solving (RNOE)

As discussed previously, in industries with a large number of firms and absent an aggregate shock (or if the aggregate shock follows a deterministic trajectory or one close to its mean), the expected state remains close to its expectation and, therefore, typically NOE strategies are close to those of MPE. However, if these conditions fail, the industry state may significantly diverge from its expectation, and NOE strategies may be far from optimal.

To mitigate potential approximation errors, we draw on established concepts from Operations Research (OR) that are frequently applied in practice. In OR, it is common practice to approximate large dynamic programs by employing their deterministic equivalents. Once a deterministic policy is derived, the algorithm then “grows” the solution forward, allowing firms to re-optimize when such optimization is advantageous, particularly when the observed state diverges from its expected value and new information becomes available. This method enables firms to refine their policies to more closely align with a full information best response.

Re-solving is computationally feasible because each time the firm is permitted to re-optimize, it tackles an approximate problem that is significantly simpler than the full dynamic programming solution. Re-solving allows the firm to minimize losses when the industry state diverges from its expectation.

Algorithm 1 Nonstationary Oblivious Equilibrium Solver with Starting State (s_0, z_0)

- 1: Compute OE $(\tilde{\mu}, \tilde{\lambda})$
 - 2: $\lambda_t := \tilde{\lambda}$, for all t .
 - 3: $\mu_t := \tilde{\mu}$, for all t .
 - 4: Define $\tilde{V}_{\bar{T}+1}(x|\mu^*, \mu, \lambda, s_0, z_0) := \tilde{V}(x)$, for all x, μ^*, μ , and λ .
 - 5: Compute \tilde{z}_t for $t \in \{0, \dots, \bar{T}\}$
 - 6: $n := 1$.
 - 7: **repeat**
 - 8: Compute $\tilde{s}_t(\mu, \lambda)$ for $t \in \{0, \dots, \bar{T}\}$.
 - 9: $\Delta_0 := 0; \Delta_1 := 0$.
 - 10: $t := \bar{T}$.
 - 11: **repeat**
 - 12: Choose $\mu_t^* \in \tilde{\mathcal{M}}$ to maximize $\tilde{V}_t(x|\mu^*, \mu, \lambda, s_0, z_0)$ simultaneously for all x .
 - 13: $\psi_t = \beta \tilde{V}_{t+1}(x^e|\mu^*, \mu, \lambda, s_0, z_0) - \kappa$
 - 14: $\Delta_0 = \max(\Delta_0, \psi_t)$.
 - 15: **if** $\lambda_t > \epsilon_0$ **then**
 - 16: $\Delta_1 = \max(\Delta_1, -\psi_t)$.
 - 17: **end if**
 - 18: $\lambda_t^* := \lambda_t(\beta \tilde{V}_{t+1}(x^e|\mu^*, \mu, \lambda, s_0, z_0))/\kappa$.
 - 19: Let $t := t - 1$.
 - 20: **until** $t = 0$.
 - 21: $\Delta_2 := \|\mu - \mu^*\|_\infty$.
 - 22: $\mu := \mu + (\mu^* - \mu)/(n^{\gamma_1} + N_1)$.
 - 23: $\lambda := \lambda + (\lambda^* - \lambda)/(n^{\gamma_2} + N_2)$.
 - 24: $n := n + 1$.
 - 25: **until** $\Delta_0 \leq \epsilon_1$ and $\Delta_1 \leq \epsilon_1$ and $\Delta_2 \leq \epsilon_2$.
-

The computational tractability comes with some concessions. Similar to other re-optimization methods, RNOE prevents firms from considering future re-solving opportunities when making present decisions. This limitation might introduce some dynamic inconsistency in firm behavior. However, such an approach may more closely reflect real-world decision-making, as it mimics the optimization techniques frequently employed in practice.

Formally, in the basic NOE, the firms use an exact state s_0 at $t = 0$ and an approximation \tilde{s}_t for each $t > 0$. In RNOE firms incorporate information about the actual realized state s_t for $t > 0$, and a basic NOE is *re-solved* in every time period.

More specifically, suppose that the initial state is (s_0, z_0) . The main idea is that at time $t = 0$ firms implement (μ_0, λ_0) , the NOE strategy and the entry rate at $t = 0$ given the initial state (s_0, z_0) . Given this strategy and the initial state, at time $t = 1$, the *actual realized* industry state is s_1 . Similarly, the aggregate shock evolves to z_1 . Then, firms re-solve for NOE with starting state (s_1, z_1) and firms implement the first-period strategy and entry rate associated to the new NOE. This procedure is repeated again and again.

One advantage of RNOE with respect to NOE is that the former can self-correct if the industry state and aggregate shock depart from their expectation. This correction becomes even more important under the presence of aggregate shocks because, even with many firms the aggregate shock is not subject to averaging. Hence, typically, the aggregate shock trajectory may experience significant random fluctuations around its expectation. Re-solving alleviates this concern because the actual value of the aggregate shock is updated every period.

Now, we formally define RNOE. A RNOE to our model is comprised of an investment/exit strategy $\mu = (\iota, \rho) \in \mathcal{M}$, and an entry rate function $\lambda \in \Lambda$ that satisfy the following:

$$(\mu(s, z), \lambda(s, z)) = (\mu_0^{s,z}, \lambda_0^{s,z}), \forall s \in \bar{\mathcal{S}}, z \in \bar{\mathcal{Z}},$$

where for every $s \in \mathcal{S}$ and $z \in \mathcal{Z}$, $(\mu_0^{s,z}, \lambda_0^{s,z})$ is the first-period strategy of an NOE with initial starting state (s, z) . Note that RNOE strategies are stationary, Markovian, and are a function of both the industry state and the aggregate shock level. However, RNOE is computationally practical because the computation of strategies at different states and shocks can be completely decoupled. This becomes apparent when simulating the industry according to RNOE strategies, in which NOE only gets computed (or re-solved) along the path of visited states and aggregate shocks, as formalized in the next algorithm.

Algorithm 2 Simulating industry according to RNOE strategies from initial state (s, z)

- 1: $t := 0, s_0 := s, z_0 := z$
 - 2: **repeat**
 - 3: Compute NOE from initial state (s_t, z_t) using Algorithm 1. Let $(\tilde{\mu}_0, \tilde{\mu}_1, \dots)$ and $(\tilde{\lambda}_0, \tilde{\lambda}_1, \dots)$ be the NOE.
 - 4: $(\mu_t, \lambda_t) := (\tilde{\mu}_0, \tilde{\lambda}_0)$
 - 5: Simulate one period evolution of the industry from (s_t, z_t) when firms use strategy (μ_t, λ_t) and let s_{t+1} be the realized state and z_{t+1} be the realized shock.
 - 6: $t := t + 1$
 - 7: **until** $t = \bar{T}$
-

In the next two sections, we provide two applications of (R)NOE. First, an empirical application to compute an optimal subsidy for a technology adoption with network effects. The model and parameters are set to the estimates of Ryan and Tucker (2012a). This exercise aims to showcase the usefulness of NOE in a real-world application. The empirical setting was chosen to involve a large state-space (so that MPE cannot be easily computed), and to emphasize the importance of transitional dynamics as opposed to steady-state dynamics (so OE may not provide the best answer). Also, the setting intentionally involves a large number of agents to provide a best case application for pure NOE without the need of re-solving.

Second, we apply NOE to model transitional dynamics in markets with small numbers of firms. We compare the transitional dynamics of NOE to those of MPE via numerical experiments, including a few canonical oligopoly models based on Ericson and Pakes (1995). Since the number of agents we consider is small, we also compute the NOE solution with Re-solving. We demonstrate that while in general NOE provides good average approximations of MPE transitional dynamics, RNOE improves the approximation point-wise, i.e., for states that are far from the industry average. Furthermore, RNOE improves the approximation to MPE in highly concentrated industries and in industries with aggregate shocks.

4 Application to Technology diffusion

We apply NOE to an empirical application due to Ryan and Tucker (2012a) (henceforth, RT), who study the diffusion of technology in a market with network externalities and a large number of agents. In particular, RT examines the adoption of video conferencing technology in a large multinational firm with more than 2,000 heterogeneous workers. The primary focus of their exercise is to estimate the size and heterogeneity of the network effect and adoption cost.

A typical characteristic of network adoption models is that the agents' optimization problem is extremely complex. Each agent needs to compute the expected sum of discounted utility of joining the network, which in turn depends on expectations about the future installed base. An additional complication here is that the agents do not value the adoption of coworkers equally, but care more about adoption of some type of coworkers than others. Thus, to correctly compute the current value of adoption, the agents need to know the adoption states of the different coworker types and make predictions about their adoption times. With this heterogeneity the state space becomes enormous, making exact computation of MPE impossible even for a relatively small number of agents.

Ryan and Tucker (2012a) circumvent this difficulty by estimating the model using the two-stage estimation method of Bajari et al. (2007) that does not require solving for MPE. The paper then computes a set of seeding exercises that also do not require recomputing equilibria. Because MPE computation was infeasible, the paper did not consider counterfactuals that perturb model primitives, such as the optimal allocation of adoption cost subsidies.

More generally, the literature has employed several alternative methods of analyzing the network good adoption problem (see Bramoullé et al., 2020, for the survey of the literature). The most common approach is to use a static model with fulfilled expectations (e.g., see Katz and Shapiro, 1985). In such a model, similarly to OE, the agents optimize against the stationary distribution of adopters. Another approach is to

accommodate nonstationary dynamics via the assumption of myopic agents (e.g., see Mele, 2017). Finally, Björkegren (2019) models network adoption with forward-looking agents and circumvents computational issues by computing a full information equilibrium in which agents have perfect information about all future periods and commit at time zero to a publicly announced future adoption time. The NOE framework, on the other hand, allows researchers to compute equilibria with forward-looking agents and full heterogeneity, under substantially weaker assumptions.

The RT paper is an excellent testing bed for NOE, for three reasons: (i) computing an MPE is difficult without abstracting away much of the economically relevant information about the network topology; (ii) the set of agents is large, and therefore, NOE behavior and dynamics should be close to those of MPE; and (iii) computation of NOE is very simple — so simple, in fact, that the code easily fits in the published appendix. We proceed to describe the formal RT model.

4.1 Model

Consider a continuum of agents with mass N living over an infinite, discrete-time horizon.¹⁴ Each period, each agent decides whether to adopt a teleconferencing technology. Let $\mu_{it} = 1$, if agent i adopts the technology at time t , and $\mu_{it} = 0$ otherwise.

Agents are partitioned into $M = 64$ discrete types. At $t = 0$ each agent is assigned one of persistent types m_i and a persistent one-time adoption cost $c_i \in \mathcal{C}$ drawn from a distribution F_{m_i} . Type m_i is public information and cost c_i is agent’s private information. The agent’s state is given by the tuple $x_{it} = (m_i, c_i, a_{it})$, where a_{it} is the adoption status. The adoption of the technology is perpetual, that is, $a_{it} = 1$, if $a_{i,t-1} = 1$ or $\mu_{it} = 1$. There is no entry of new agents nor exit of incumbents.

We consider cut-off strategies. Under any profile of such strategies, for a *given users’ type*, users within that type adopt in a sequence, such that lowest cost users adopt first and the highest cost users adopt last. In such case, the fraction of adopters is a sufficient statistic for the cost distribution of non-adopters. We can abstract from explicitly tracking the beliefs of agents about the distribution of non-adopters.

As in our base model, we focus on symmetric and anonymous equilibrium, so we restrict attention to industry states that do not depend on agents’ identities. Thus motivated, with some abuse of notation we denote the industry state as the mass of adopters of type m as $s_t(m)$, for all possible types m . The industry state is an element of the set $\mathcal{S} = [0, N_1] \times [0, N_2] \times \cdots \times [0, N_M]$, where N_m is the total number of agents

¹⁴A model with continuum of agents is a modification of our earlier model with finite number of agents. Assuming a continuum of agents greatly simplifies the structure of beliefs. Specifically, with a continuum of agents the current number of adopters is a sufficient statistic for the cost distribution of nonadopters.

of type m .

In addition, there is an aggregate shock, z_t , common to all agents that follows a deterministic trajectory, reflecting common time trends and seasonality of the value of video conferencing. The tuple (s_t, z_t) contains the payoff relevant information of the game and fits the framework laid out in Section 3.

4.1.1 Agents' Objective and Adoption

The agent derives the utility from making phone calls to other agents who also adopted the technology. We denote the per-period value of using the technology as $U_m(s_{-i,t}, z_t)$. In particular, the value depends on the mass of other adopters segmented by types, $s_{-i,t}$ and the aggregate shock z_t .

The utility, U_m , reflects that the value of adoption technology depends on the number of other adopters, as well as on the topology of the network. For instance, a member of the Marketing department may have larger interconnection value with other members of the same department; while a member of the Finance team may have a large value to contact employees from Operations. The exact micro model of U_m is in Appendix B. The parameters of this model have been calibrated to the results of Ryan and Tucker (2012b) and reflect an interdependence graph identified using revealed preferences of the agents from the empirical data.

Agents maximize expected discounted utility given by

$$E \sum_{t=s}^{\infty} \beta^{t-s} [a_{i,t} U_m(s_{-i,t}, z_t) - \mu_{i,t} c_i]$$

We consider only equilibria in cut-off adoption strategies. Hence, an agent adopts the technology if the continuation value of doing so is greater than the adoption cost, and does not adopt otherwise. That is, $\mu_{it} = 1$, if and only if $c_i \leq \bar{c}(x_{it}, s_{-i,t}, z_t)$, where $\bar{c}(x_{it}, s_{-i,t}, z_t)$ is the continuation value of adopting for firm i at time t .

4.2 Nonstationary Oblivious Equilibrium

Consider a nonstationary strategy $\mu \in \tilde{\mathcal{M}}_{ns}$. The strategy is of the form $\mu_t(x_{it})$. A nonstationary strategy generates an expected path of network adoption, \tilde{s}_t . For every agent type m , The sequence $\tilde{s}_t(m)$ is bounded and increasing, thus it has a limit. Denote this limit by $\tilde{s}(m)$.

Consider a NOE $\tilde{\mu}$. It is easy to show that, since both $\{\tilde{s}_t : t \geq 0\}$ and $\{z_t : t \geq 0\}$ have a limit, $\{\tilde{\mu}_t : t \geq 0\}$ must also have a limit and that this limit must be OE. Thus, every NOE of the RT model

converges to a OE. OE of the RT model is relatively straightforward to compute and corresponds to a fulfilled expectation static adoption problem (similar to Katz and Shapiro (1985)), in which adoption fraction for each type m solves the following Bellman equation:

$$\tilde{s}(m) = N_m F_m(U_m(\tilde{s})/(1 - \beta)).$$

This equation is typically solved by fixed-point iteration. We compute NOE by simulated best response iteration, which is an adaption of our general algorithm to account for the heterogeneity in the adoption cost. Using this algorithm we can obtain counterfactual adoption paths for any set of primitives. We are particularly interested in computing optimal adoption subsidies that generate economically desirable adoption paths from an initial state in which no one has adopted. In the Appendix we describe the algorithm to compute NOE for this model, which is around 20 lines of code.

4.3 Calibration of the Model and Subsidies

We start by calibrating our model to the estimates reported in Ryan and Tucker (2012a). The RT paper does not report two statistics: exact population sizes for 16 of the 64 types in the data, and values of the estimated aggregate utility shifters z_t . Since we know the total number of workers, we allocate the unaccounted agents uniformly across the 16 omitted types. To obtain z_t , we use a full solution nested-fixed point calibration procedure. We choose z_t that minimizes the mean square error between the number of adopters for each t predicted by our model and that reported in Ryan and Tucker (2012a).

We consider a practically relevant case in which the firm can commit to a type-dependent cash adoption subsidy $\gamma = (\gamma_1, \dots, \gamma_{64})$. The firm's objective is to maximize its payoff function $G(\tilde{s}_0, \tilde{s}_1, \dots, \tilde{s}_T)$ within a subsidy budget B . There are two practical problems for the firm when choosing the amount of the subsidy γ : (i) Since the number of adopters is endogenous, forecasting the total cost of the subsidy for a given per capita subsidy is nontrivial, (ii) Because of the network effects, the optimal allocation of money across types is hard to determine. The former is a cost management issue that could be mitigated with some loss of efficiency by stopping the subsidy once its cost reaches the budget. The latter issue, however, is more substantive and harder to eliminate. The firm does not want to subsidize types that would adopt without the subsidy, and wants to target influencers (types that generate large network externalities). Adoption decisions of all types are interdependent and are determined by a series of complicated feedback loops resulting from a rich specification of the link synergy matrix. For this reason, it is practically impossible to determine which types to subsidize by examining the primitives of the game (link synergy matrix and adoption cost structure)

without computing the equilibrium. Thus, in order to determine which types to subsidize, we propose that the firm solves the equilibrium adoption game using NOE, and solves the following constrained optimization problem:

$$(4.1) \quad \begin{aligned} \max_{\gamma} \quad & G(\tilde{s}_1, \dots, \tilde{s}_T) \\ \text{s.t.} \quad & \sum_m \gamma_m \tilde{s}_{mT} \leq B \quad \text{and} \quad D(\gamma) = 0, \end{aligned}$$

where $D(\gamma)$ are other practical constraints on the subsidy structure. Henceforth, we compute solutions to the above problem that maximize the discounted number of adopters. Formally we set

$$G(\tilde{s}_0, \tilde{s}_1, \dots, \tilde{s}_T) = \sum_{t=0}^T \delta^t \sum_m \tilde{s}_t(m)$$

for large enough T . The parameter δ captures the trade-off between short term adoption and long-term adoption. If δ is close to 1 the company places more weight on long term adoption rates; and if δ is small short-term adoption rates receive more weight. We set $\delta = 0.9$.

Since a first-best subsidy requires us to optimize over a 64-dimensional vector γ , which is impractical, we consider two types of second-best subsidies that we numerically optimize. First, we consider a benchmark cost management case in which $\gamma_m = \gamma_{m'}$ for all m and m' . Next, we consider a case where the subsidies are optimally allocated across 4 different hierarchical functions within the company, which implies that $\gamma_m = \gamma_{m'}$ if m and m' correspond to the same hierarchical function. Even in case of a simple uniform subsidy we need to solve the model repeatedly to obtain the counterfactual adoption paths under various subsidies. Thus, this exercise is only possible if the model can be solved relatively quickly. To obtain workable computation speed, we utilize NOE.

Figure 1 contains the comparison of a simple subsidy and a more complicated subsidy optimally targeting by hierarchical function. The simple subsidy allocates a 0.95 util¹⁵ uniform transfer for any adopter of the network. There are four hierarchical roles within the firm: Associate (AS), Vice President (VP), Director (DR), and Managing Director (MD). The optimal subsidy subsidizes the two middle roles: VPs by 1.96 utils, and DRs by 1.18 utils. The subsidy does not subsidize ASs and MDs.

In total, there are 2,169 potential adopters. A uniform subsidy increases adoption by 436 agents in the short run (month 2). The targeted subsidy increases adoption by 557 clients, or 28% more than the uniform

¹⁵The RT paper does not have a price or income measure that can be used to convert utils into meaningful units such as dollars. For reference, mean adoption costs range between -0.7 to 2.2 utils across individuals, depending on employee subtype.

subsidy. In the long run, the uniform subsidy results in 240 more adopters, as compared to 294 more adopters resulting from the targeted subsidy. In other words, the targeted subsidy has 23% higher effectiveness in increasing the overall adoption than the uniform subsidy. Overall, computing targeted subsidies allows for better targeting that delivers substantially higher short- and long-term adoption rates for the same overall cost of the subsidy. We believe the application shows the potential for using NOE in economically relevant settings with a large number of agents.

5 Application to Dynamic Oligopoly

This subsection contains computational experiments that examine the transitional dynamics generated by (R)NOE in markets characterized by a small number of firms and varying levels of concentration. The experiments have two main goals. First, we aim to compare the transitional industry paths generated by NOE with the transitional paths generated by MPE and OE. Since OE is designed to approximate long-run dynamics, its short-run dynamics can be far from those of MPE even if its long-run dynamics match closely. We demonstrate this issue in several different models, and then show that the transitional dynamics generated by NOE are more realistic and closer to those of MPE.

Our second goal is to investigate the value of re-solving in improving the modeling (relative to both NOE and OE) of both long- and short-run industry dynamics. For example, as demonstrated in Benkard et al. (2015), OE can fail to approximate the stationary distribution of MPE in highly concentrated industries with an unusually high turnover of industry leaders. Also, the OE approximation may fail in the presence of serially correlated aggregate shocks that cannot be averaged away.

5.1 Models Analyzed

We numerically analyze two structurally distinct dynamic oligopoly models from the past literature that represent a range of potential applications: (i) a Quality-Ladder model with a differentiated products logit demand model and Bertrand competition with investments in product quality; (ii) a homogeneous products Cournot model with investments that reduce marginal cost. Although it would be straightforward to implement entry and exit in these models, to simplify the exposition and computation of MPE, we omit them from our specifications. We describe the models in detail next.

QUALITY-LADDER MODEL. This price-quality competition model most closely matches the differentiated products oligopoly models that are commonly used in empirical applications in IO and related fields.¹⁶ In

¹⁶This model closely follows Weintraub et al. (2010) that in turn is an extension of the model in Pakes and McGuire (1994).

this model, single product firms produce products that have an endogenous quality state that can be improved through investment. Consumer demand is modeled through a standard logit demand model. There are m consumers in the market. In each period t , consumer j receives utility u_{ijt} from consuming the good produced by firm i given by:

$$u_{ijt} = \theta_1 \ln(x_{it} + 1) + \theta_2 \ln(Y - p_{it}) + \nu_{ijt}, \quad i \in S_t, \quad j = 1, \dots, m,$$

where Y is the consumer's income, and p_{it} is the price of the good produced by firm i . ν_{ijt} are i.i.d. random variables distributed Gumbel that represent unobserved characteristics for each consumer-good pair. There is also an outside good that provides consumers zero utility. We assume consumers buy at most one product each period and choose the product that maximizes utility. Under these assumptions, our demand system is a classic logit model.

Let $N(x_{it}, p_{it}) = \exp(\theta_1 \ln(x_{it} + 1) + \theta_2 \ln(Y - p_{it}))$. Then, the expected market share of each firm is given by:

$$\sigma(x_{it}, s_{-i,t}, p_t) = \frac{N(x_{it}, p_{it})}{1 + \sum_{j \in S_t} N(x_{jt}, p_{jt})}, \quad \forall i \in S_t.$$

We assume that firms set prices in the spot market. If there is a constant marginal cost c , the Nash equilibrium of the pricing game satisfies the first-order conditions,

$$(5.1) \quad Y - p_{it} + \theta_2(p_{it} - c)(\sigma(x_{it}, s_{-i,t}, p_t) - 1) = 0, \quad \forall i \in S_t.$$

There is a unique Nash equilibrium in pure strategies, denoted p_t^* Caplin and Nalebuff (1991). Expected profits are given by:

$$\pi_m(x_{it}, s_{-i,t}) = m\sigma(x_{it}, s_{-i,t}, p_t^*)(p_{it}^* - c), \quad \forall i \in S_t$$

where m is the market size.

Firms can also invest $\iota \geq 0$ in order to improve their product quality over time. A firm's investment is successful with probability $\frac{a\iota}{1+a\iota}$, in which case the quality of its product increases by one level. The firm's product also depreciates one quality level at random with probability δ , independently each period. Combining the investment and depreciation processes, it follows that the transition probabilities for a firm

in state x that does not exit and invests ι are given by:

$$\mathcal{P} \left[x_{i,t+1} \mid x_{it} = x, \iota \right] = \begin{cases} \frac{(1-\delta)a\iota}{1+a\iota} & \text{if } x_{i,t+1} = x + 1 \\ \frac{(1-\delta)+\delta a\iota}{1+a\iota} & \text{if } x_{i,t+1} = x \\ \frac{\delta}{1+a\iota} & \text{if } x_{i,t+1} = x - 1 . \end{cases}$$

QUANTITY-COST COMPETITION. As we show below, the Quality-Ladder model generates dynamics that are fairly straightforwardly modeled using just OE and NOE. Thus, we also consider an additional model that better highlights the weaknesses of OE, and allows us to explore both the value and the limits of NOE and RNOE. As shown in Benkard et al. (2015), the quantity-cost model generates highly concentrated, near monopoly markets that also have a fringe of smaller firms, and a significant turnover of leading firms, features that are difficult for OE to replicate.

Consider an industry with an homogeneous product and quantity setting, for which the state of each firm represents its marginal cost of production. The industry has a linear inverse demand function

$$P(q_t) = m_1 - \sigma \sum_i q_{it} ,$$

where q_{it} is the quantity produced by firm i at time period t . The marginal cost for firm i in state x_{it} is

$$\text{MC}(x_{it}) = \gamma_0 \exp(-(\gamma_1 x_{it} - \gamma_2)).$$

With this specification, we can generate highly concentrated markets. However, in periods of high concentration, lagging firms stop producing altogether, leading to a monopoly. Thus, it is difficult to generate highly concentrated markets in which small firms continue to produce. Monopoly markets are close to single agent problems, thus they are well modeled by OE, and less useful in our context. To generate concentrated oligopoly markets, we take a shortcut and assume that there is also a second market that always yields $m_2 + \alpha x_{it}$ of profits. One can think of the first market as representing the “national” market. All firms compete in quantities in this market. The second market may represent a “local” market. Each firm captures its local market entirely regardless of the national market structure. The second market provides small firms with some market share that incentivizes investment even when there is a very efficient large firm. Period

profits for firm i at time period t in state x_{it} are given by

$$\Pi(q_{-i,t}, x_{it} | s_{-i,t}) = \max_{q_{it}} \{P(q_t)q_{it} - \text{MC}(x_{it})q_{it} + m_2 + \alpha \log(x_{it})\}$$

Firms set quantities simultaneously conditioning on the observed state, and the spot market is assumed to be in static Nash equilibrium. In this model, investment improves the state by reducing marginal cost. Transitions are modeled as in the Quality-Ladder model above.

5.2 Comparison of OE, (R)NOE and MPE

This section compares the long-run and transitional dynamics of OE, (R)NOE, and MPE in the classic dynamic oligopoly models described above. Both models are capable of generating a wide range of empirically relevant industry structures. The critical parameter in both models is the cost of investment. As this cost increases, the industry moves between symmetric oligopoly, in which most firms occupy the upper portion of the state space, to highly concentrated near-monopoly markets, in which all but one firm are in the lowest industry state. Because of the nature of the tails of the extreme value distribution, the logit Quality-Ladder model tends to generate less extreme industry structures than the Quantity-Cost model, with milder differences in profits across states. The Quantity-Cost model can generate industries with arbitrary levels of concentration. This model enables us to study some extreme near-monopoly industry configurations that are arguably not very empirically relevant, and for which the dynamic oligopoly framework may not always be the most suitable modeling choice. Nevertheless, it is informative to analyze such extreme cases in order to explore the limits of the proposed modeling techniques.

For both models, we present transitional equilibrium dynamics towards the equilibrium stationary distribution starting from an industry state in which all firms are in the lowest state. We think of this as representing the time path of a new industry starting from its inception, far from its long run distribution of states. Because we compute NOE that converges to OE, its long-run dynamics are the same as OE. However, its short run dynamics are often very different to those of OE. Because RNOE re-solves every period, its *long-run dynamics* need not be the same as OE and NOE. We first investigate the short-run dynamics of NOE and RNOE in cases where OE provides an excellent long-run approximation to MPE. We then move on to the cases where OE's long-run dynamics are less realistic.

For both models, we assume that there are 8 firms and individual states are given by $\{1, \dots, 8\}$. We compare the simulated trajectories of consumer surplus, producer surplus, and expected total investment across the proposed equilibrium concepts. For each simulated trajectory r , we set the initial state, $s_{r0} =$

$\{8, 0, 0, 0, 0, 0, 0, 0\}$. We use similar parameters parameter values as in Benkard et al. (2015), see Table 1. The parameters range from values that generate typical industry dynamics, to the values that generate atypical dynamics i.e., abnormally high rate of industry turnover for which OE fails; see extensive discussion in Benkard et al. 2015.

To obtain R trajectory draws,¹⁷ denoted by $\{s_{rt}\}_{t=0}^T$, we start at s_{t0} and simulate the industry for $T = 100$ periods using the corresponding equilibrium strategy. NOE is computed using backward induction from OE with $\bar{T} = 100$. For example, in the case of MPE, we apply full Markov strategies at each current state s_{rt}^{MPE} to obtain the draw of the next industry state $s_{r,t+1}^{MPE}$. Similarly, for OE, we apply the OE strategy to obtain the next industry state. For the case of NOE, we compute a NOE strategy starting from the initial state s_0 , $\{\tilde{\mu}_t\}_{t=0}^{\bar{T}}$ and apply this NOE strategy to simulate the transitions from s_{rt}^{NOE} to $s_{r,t+1}^{NOE}$. Finally, for RNOE, we re-solve for a new NOE starting at s_{rt} and use $\tilde{\mu}_0$ from the new NOE to obtain the next state. In this case, we never use $\tilde{\mu}_t$ for $t > 0$; instead, we re-solve for new NOE by changing the initial state as t increases. To increase the numerical efficiency, we store NOE(s) for various initial states in the memory and pull from memory/re-solve as necessary.

5.2.1 Quality-Ladder Model

Figure 2 compares average industry statistics between OE and (R)NOE for an industry with low investment cost. In the stationary distribution under MPE, the industry has an average of 5.2 firms in the upper portion of the state space (states 5-8) and the industry is relatively competitive with a 1,463 Herfindahl-Hirschman Index (HHI). All of the equilibrium concepts closely match MPE in the long run. In the short run, we see that OE produces lower equilibrium investment than MPE. This lower investment is unsurprising since OE responds to the average long run industry state rather than the true realized short-run industry state. In MPE, in early periods firms know that there are no large firms, so they are eager to invest a lot to capture market share. In OE, even in the initial periods firms optimize against the long-run distribution that prescribes 5.2 large firms, and therefore they underinvest in the short run. NOE mostly corrects this bias. In NOE, firms optimize against a mean field approximation of the industry evolution over time, and invest almost the same as in MPE. RNOE, NOE, and MPE are almost indistinguishable.

Figure 3 examines an industry with a larger investment cost. In this industry, there are 3.1 firms in the upper portion of the state space, and the HHI is 1,743. In this more concentrated industry, OE and NOE now slightly differ from MPE in the long run. Resolving corrects this issue and RNOE is virtually identical to

¹⁷ R , with and without re-solving, is chosen to obtain a given precision of the industry statistics, at most 10%-wide 95% confidence interval, but not less than 2,000 draws.

MPE in the long run. OE does not perform well in the short run for the same reasons as above. NOE is still massively better than OE in the short run, but even in the short run it is slightly hampered by the fact that it converges to OE in the long run, and OE is beginning to fail slightly. Resolving solves both of these issues and RNOE is numerically indistinguishable from MPE in both the short and long runs.

We explore this point further by considering an industry with an even larger investment cost, depicted in Figure 4. This industry has 1.5 firms in the upper portion of the state space, and an HHI of approximately 1,900. Note that even with extreme product quality differences the logit demand model cannot endogenously generate high HHI's in this dynamic model. Below we use the other model to explore markets with extreme HHI's. For this industry, OE, and thus NOE, produce approximately 10% long-run error in average consumer surplus and between 3% and 5% error in producer surplus and average investment. Again, the short-run approximation error of NOE is similar to the long-run error. RNOE reduces the short- and long-run approximation errors in consumer and producer surplus by over 30%, though some small differences remain. RNOE delivers numerically negligible approximation error in average investment in all but the initial periods.

The Quality-Ladder model is generally incapable of producing very concentrated and numerically stable industries because of the fat tails in the logit model. To investigate the limits of (R)NOE in more highly concentrated industries, we now consider the Quantity-Cost model.

5.2.2 Quantity-Cost Model

In this subsection, we examine the Quantity-Cost model. The comparison of short-run dynamics for the industry with low investment cost is presented in Figure 5. This industry has an average of 4.5 firms in the upper part of the state space and an HHI of 2,230 in the MPE long-run. Importantly, this industry has a higher concentration index than any of the parameterizations of the Quality-Ladder model, despite having on average as many as 4.5 firms in the upper portion of the state space. A high HHI is more natural in this model because firms in the lower portion of the state space nearly shut down when more efficient competitors can price below their marginal cost. Despite the industry being more concentrated, OE is essentially identical to MPE in the long run here. As in the Quality-Ladder model, since OE is close to MPE in the long run, NOE is also close to MPE in the short run.

We omit the intermediate investment cost case (see the online appendix, Figure 1) as it is similar to the low investment cost case but with a slight difference for OE and NOE in the long run, and instead move to an extremely concentrated industry, shown in Figure 6. This industry represents a near monopoly, with

0.93 firms in the upper portion of the state and an HHI of 9,033. Importantly, this industry has very unique dynamics because the identity of the monopoly firm periodically changes, and it is important for firms to identify these key moments when they may be able to take over the market. It is hard to imagine a more challenging industry for OE and NOE. It also seems unlikely that this type of model would be used to study such an industry in an empirical setting. Nevertheless it is useful to look at how the various models behave in such an extreme environment. For this industry, OE and NOE are around 10% different from MPE in the long run. Similarly to before, the short run performance of NOE is hampered when OE does not do well in the long run. In contrast, RNOE mostly solves these issues, and is within about 2% of MPE in both the short and long runs in all three metrics.

5.2.3 Aggregate shocks

We now extend the Quality-Ladder model to accommodate aggregate shocks. Specifically, we allow for three demand states that change the market size: low, medium, and high. The medium state is a baseline case that has the same parameter levels as the baseline Quality-Ladder model. The low state results in 50% lower than baseline demand, and the high state results in 50% higher than baseline demand. The transition matrix between states is given by:

$$\begin{pmatrix} 0.8 & 0.2 & 0.0 \\ 0.1 & 0.8 & 0.1 \\ 0.0 & 0.2 & 0.8 \end{pmatrix}.$$

The stationary distribution assigns 0.5 probability to the medium state and 0.25 to the low and the high state. Expected profits in the stationary distribution are equal to the baseline profits at the medium state.¹⁸

We consider two extreme cases of low and high investment costs. The comparison of different equilibria for different values of starting aggregate shock and low investment cost is presented in Figures 7. In this case, OE performs well in the long run, and both NOE and RNOE perform well in the short and long run. Consumer surplus and producer surplus show similar patterns (see Section 2 in the Online Appendix).

To investigate a more complicated case where OE fails to provide accurate long-run approximations, we consider a case of a more concentrated industry with a higher investment cost. The results are presented in Figure 8. We find that the long-run gap between OE and MPE is large. NOE works well in the first few periods but is eventually hampered by the poor long run performance of OE. In comparison, re-solving

¹⁸In this exercise, firms track only the average value of the aggregate shock in OE and NOE. It is possible to compute versions of OE and NOE in which firms track the exact value of the aggregate in addition to their state. These implementations, while computationally more expensive, would better approximate MPE.

allows firms to use an updated policy every time the value of the shock changes, which delivers an excellent approximation to MPE. Intuitively, RNOE is piecing together the very short-run dynamics of NOE, where NOE does perform well. We obtain similar results with respect to the performance of RNOE in markets with aggregate shocks for both Cournot models.

5.2.4 Equilibrium distribution of investment paths

We have seen that RNOE matches MPE well when averaging over firms and states. It is reasonable to ask if this is the case because errors across different states at a given period average out or because RNOE approximates investment levels for individual states well. With this motivation, this subsection considers an error metric that captures point-wise sample-path differences, that is, a mean absolute difference (MAD) over states.

Formally, RNOE and MPE investment levels for each t are random variables due to the variability of the states. MAD_t is defined as $E[|\sum_i \iota_{it} - \sum_i \tilde{\iota}_{it}|]$. The expectation is taken with respect to the MPE distribution of states (s_t, z_t) conditional on (s_0, z_0) .

We apply MAD to the Quality-Ladder model to see if this metric reveals differences in distributions that are potentially hidden by analyzing only the average investment across firms. The results are presented in Figure 9. Recall that the differences in means between NOE and MPE were negligible for the Quality-Ladder model. In contrast, we find that the MAD difference between NOE and MPE investment increases over time and asymptotes to approximately 15%. In other words, NOE (and OE) investment is on average 15% different from MPE investment point-wise, across firms, but the differences cancel out in expectation. RNOE delivers significant point-wise improvements. MAD with re-solving starts at 5% at $t = 0$ and declines to less than 2% of MPE investment over time.

5.3 Summary of findings

We now summarize the main findings from the numerical experiments. The OE model, while often aligning closely with the MPE in the long run, often diverges significantly from MPE in the short-run. NOE provides a much better tool for analyzing short run dynamics.

In addition, we find that when OE is close to MPE in the long run, NOE is typically also close to MPE in the short run. This holds true especially in industries characterized by low- to medium-high concentration levels. In such industries, NOE emerges as a viable substitute for MPE, delivering comparable short-run dynamics but allowing for richer and more realistic models. NOE is also simple to program and compute

even in complex models.

In industries with particularly high concentration levels or in the presence of aggregate shocks, OE often deviates from MPE in important ways in the long run, and NOE similarly tends to deviate from MPE in the short run. In such cases, the researchers should consider re-solving. We found that RNOE performed well in essentially all the cases considered, and is also closer to MPE point-wise. Thus, RNOE would be particularly useful when the research objective includes not just average short-term outcomes, but also the distribution of outcomes.

6 Conclusion

Recent decades have witnessed a surge in data availability and computational power, enabling the modeling and estimation of complex game-theoretical systems that closely approximate individual and firm decision-making. The parameters estimated from these models are often used to generate counterfactual predictions intended to inform managers and policymakers. However, the solution concepts currently applied to these complex models are not very tractable. Consequently, even if a realistic version of the model can be formulated and estimated using observed data, the model often needs to be further simplified to obtain counterfactual results. These simplifications limit the application of counterfactuals to smaller markets or unrealistic synthetic versions of larger markets, raising two critical questions: How much insight is lost when the models are simplified? More fundamentally, are the current solution concepts, which are intractable for academic researchers operating at the frontier, even well-suited to reflect real world firm behavior?

Dynamic games are an example in which complexity increases quickly with the richness of the model. In order to capture the economic realities of empirical applications, it may be necessary to include features like firm heterogeneity, multi-product firms, nonstationarity, and aggregate shocks. However, the dimension of the model's state space grows quickly with the number of players and the number of states per player, and this causes the canonical model of Markov Perfect Equilibrium (MPE) to become intractable. One possible solution is to apply a long-run mean-field approximation, which performs well in approximating long-run behavior of games with enough players, but may be inaccurate in the short run, or if the number of players is small.

This paper proposes a new solution concept called Nonstationary Oblivious Equilibrium (NOE), which has two main objectives. First, it is a computationally attractive solution that delivers short-run paths close to or even numerically indistinguishable from MPE, even if the number of players is small. We show that NOE typically works well for short run dynamics whenever a mean field approximation (such as OE) works well

for long run dynamics. NOE allows researchers to use a model for counterfactuals that does not compromise the complexity of the state space.

Second, NOE, and especially its re-solving variant, offers a behaviorally attractive model of firm conduct. Although theoretically attractive, MPE requires an extreme level of rationality that seems likely to be beyond the computational and reasoning capacity of most firms. In contrast, NOE and RNOE are based on concepts from Operations Research that have been frequently implemented in real firms and markets.

RNOE is grounded in the simple idea that, when firms are unable to compute complete solutions to their optimization problems, they apply approximations such as mean field approximations. However, they do not stick to these approximations forever. They update their approximations when the realized state of the world significantly deviates from the one assumed when computing the initial solution. This approach breaks down a large computational problem into a series of simple ones, and is orders of magnitude less computationally intensive than computing MPE.

We conclude that (R)NOE appears to offer an attractive behavioral model that mimics real-world decision-making of firms while also closely matching MPE dynamics. Notably, RNOE provides a computationally feasible method for modeling dynamics in concentrated oligopolistic industries of particular interest in IO.

References

- Acimovic, J. and V. Farias (2019). The fulfillment-optimization problem. In *Operations Research & Management Science in the Age of Analytics*. INFORMS.
- Aguirregabiria, V., A. Collard-Wexler, and S. P. Ryan (2021). Chapter 4 - dynamic games in empirical industrial organization. In *Handbook of Industrial Organization, Volume 4*, Volume 4 of *Handbook of Industrial Organization*, pp. 225–343. Elsevier.
- Aliprantis, C. D. and K. C. Border (2006). *Infinite Dimensional Analysis* (Third ed.). Springer.
- Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica* 75(5), 1331 – 1370.
- Balseiro, S., O. Besbes, and D. Pizarro (2023). Survey of dynamic resource constrained reward collection problems: Unified model and analysis. *Operations Research*.
- Balseiro, S. and Y. Gur (2019). Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science* 65(9), 3952 – 3968.
- Benkard, C. L., P. Jeziorski, and G. Y. Weintraub (2015). Oblivious equilibrium for concentrated industries. *The RAND Journal of Economic* 46(4), 671–708.
- Bertsekas, D. P. (2001). *Dynamic Programming and Optimal Control, Vol. 2* (Second ed.). Athena Scientific.
- Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica* 78(2), 453–508.
- Bian, B. (2018). *Essays on Empirical Industrial Organization*. The Pennsylvania State University.

- Björkegren, D. (2019). The adoption of network goods: Evidence from the spread of mobile phones in rwanda. *The Review of Economic Studies* 86(3), 1033–1060.
- Bramoullé, Y., H. Djebbari, and B. Fortin (2020). Peer effects in networks: A survey. *Annual Review of Economics* 12, 603–629.
- Buchholz, N. (2022). Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *The Review of Economic Studies* 89(2), 556–591.
- Cai, Y. and K. Judd (2023). A simple but powerful simulated certainty equivalent approximation method for dynamic stochastic problems. *Quantitative Economics* 12(2), 651 – 687.
- Camacho, E. and C. Bordons (2004). *Model Predictive Control*. Springer London.
- Caoui, E. H. (2023). Estimating the costs of standardization: Evidence from the movie industry. *The Review of Economic Studies* 90(2), 597–633.
- Caplin, A. and B. Nalebuff (1991). Aggregation and imperfect competition - on the existence of equilibrium. *Econometrica* 59(1), 25 – 59.
- Chen, N. and P. Jeziorski (2022). Consequences of dynamic pricing in competitive airline markets. Technical report.
- Costantini, J. A. and M. J. Melitz (2008). *The dynamics of firm-level adjustment to trade liberalization*. Harvard University Press.
- Degond, P., M. Herty, and J.-G. Liu (2014). Mean-field games and model predictive control. Working Paper.
- Dimitrellos, P. (2022). *Essays on Sponsored Search Auctions and Online Platforms*. Ph. D. thesis, University of Maryland, College Park.
- Doraszelski, U. and A. Pakes (2007). A framework for applied dynamic analysis in IO. In *Handbook of Industrial Organization, Volume 3*. North-Holland, Amsterdam.
- Doraszelski, U. and M. Satterthwaite (2010). Computable markov-perfect industry dynamics. *The RAND Journal of Economics* 41(2), 215–243.
- Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies* 62(1), 53 – 82.
- Farias, V., D. Saure, and G. Weintraub (2012). An approximate dynamic programming approach to solving dynamic oligopoly models. *The RAND Journal of Economics* 43(2), 253–282.
- Fershtman, C. and A. Pakes (2012). Dynamic games with asymmetric information: A framework for empirical work. *The Quarterly Journal of Economics* 127(4), 1611–1661.
- Gallego, G. and M. Hu (2014). Dynamic pricing of perishable assets under competition. *Management Science* 60(5), 1241–1259.
- George, S. (2022). *Essays on the Market for IPv4 Addresses*. Ph. D. thesis, University of Pennsylvania.
- Ifrach, B. and G. Y. Weintraub (2016, 09). A Framework for Dynamic Oligopoly in Concentrated Industries. *The Review of Economic Studies* 84(3), 1106–1150.
- Igami, M. (2017). Estimating the innovator’s dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998. *Journal of Political Economy* 125(3), 798–847.
- Johnston, S., Y. Liu, and C. Yang (2023). An empirical analysis of the interconnection queue. Technical report, National Bureau of Economic Research.

- Katz, M. L. and C. Shapiro (1985). Network externalities, competition, and compatibility. *The American economic review* 75(3), 424–440.
- Marsden, J. E. and M. J. Hoffman (1993). *Elementary Classical Analysis* (Second ed.). W. H. Freeman and Company.
- Maskin, E. and J. Tirole (1988). A theory of dynamic oligopoly, I and II. *Econometrica* 56(3), 549 – 570.
- Mele, A. (2017). A structural model of dense network formation. *Econometrica* 85(3), 825–850.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695 – 1725.
- Pakes, A. and P. McGuire (1994). Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. *RAND Journal of Economics* 25(4), 555 – 589.
- Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric Markov perfect equilibrium, and the ‘curse’ of dimensionality. *Econometrica* 69(5), 1261 – 1281.
- Pavanini, N. (2014). *Imperfect competition and market structure with asymmetric information: the Italian banking sector*. Ph. D. thesis, University of Warwick.
- Phillips, R. (2021). *Pricing and Revenue Optimization*. Stanford Business Books.
- Qi, S. (2013). The impact of advertising regulation on industry: The cigarette advertising ban of 1971. *The RAND Journal of Economics* 44(2), 215–248.
- Royden, H. (1988). *Real Analysis* (Third ed.). Pearson Education.
- Ryan, S. and C. Tucker (2012a, March). Heterogeneity and the dynamics of technology adoption. *Quantitative Marketing and Economics (QME)* 10(1), 63–109.
- Ryan, S. P. and C. Tucker (2012b). Heterogeneity and the dynamics of technology adoption. *Quantitative Marketing and Economics* 10(1), 63–109.
- Saeedi, M. (2019). Reputation and adverse selection: theory and evidence from ebay. *The RAND Journal of Economics* 50(4), 822–853.
- Sweeting, A. (2013). Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica* 81(5), 1763–1803.
- Sweeting, A. (2015). A model of non-stationary dynamic price competition with an application to platform design. Technical report.
- Vera, A. and S. Banerjee (2021). The bayesian prophet: A low-regret framework for online decision making. *Management Science* 67(3), 1368 – 1391.
- Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008). Markov perfect industry dynamics with many firms. *Econometrica* 76(6), 1375–1411.
- Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2010). Computational methods for oblivious equilibrium. *Operations Research (Special Issue in Computational Economics)* 58(4), 1247–1265.
- Wilson, N. E. (2012). Uncertain regulatory timing and market dynamics. *International Journal of Industrial Organization* 30(1), 102–115.

A Proof of Theorem 3.1

A.1 Preliminaries

To simplify notation we embed the aggregate shock state into the industry state, that is, in an abuse of notation, $s \in \mathcal{S} \times \mathcal{Z}$. We also abuse notation to consider a restricted state space $\mathcal{S} = \left\{ s \in \mathbb{R}_+^{|\mathcal{X}|+1} \mid \sum s(x) < N \right\}$. Under our assumptions, this restriction is done with out loss of generality for N large enough. We endow the set \mathcal{S} with a norm. We define a set $\mathcal{S}^\infty = \{(s_0, s_1, s_2, \dots) : s_t \in \mathcal{S}; \text{ for which there exists } s \in \mathcal{S}, \lim_{t \rightarrow \infty} s_t = s\}$, endowed with a metric compatible with the product topology. The elements of \mathcal{S}^∞ are denoted by $S = \{s_0, s_1, \dots\}$. We also endow $\tilde{\mathcal{M}}$ and Λ with norms and endow $\hat{\mathcal{M}}_{ns}$ and $\hat{\Lambda}_{ns}$ with a metric compatible with the product topology.

We define a new set of strategies $\mathcal{M}_{\mathcal{S}} : \mathcal{X} \times \mathcal{S}^\infty \rightarrow \mathcal{I} \times [0, \sup_{x,s,\ell} \pi(x, s, \ell)/(1 - \beta) + \bar{\phi}]$. For all strategies $\mu \in \mathcal{M}_{\mathcal{S}}$, individual states $x \in \mathcal{X}$, and sequence of industry states $S \in \mathcal{S}^\infty$, we define a new value function:

$$\bar{V}(x, S | \mu) = E_\mu \left[\sum_{k=0}^{\tau_i} \beta^k (\pi(x_{ik}, s_k, \ell_{ik}) + \beta^{\tau_i - k} \phi_{i, \tau_i} | x_{i0} = x) \right].$$

Let $\bar{V}(x, S) = \sup_{\mu \in \mathcal{M}_{\mathcal{S}}} \bar{V}(x, S | \mu)$, $\forall x \in \mathcal{X}, S \in \mathcal{S}^\infty$. Note that the state space of this dynamic programming problem is uncountable. However, because single-period profits and expected sell-off value are bounded, the supremum can always be attained simultaneously for all x and S by a common strategy μ (Bertsekas, 2001).

We define a translation operator $G : \mathcal{S}^\infty \rightarrow \mathcal{S}^\infty$, such that, $G(S) = \{s_1, s_2, \dots\}$ and define the following Bellman operator for a value function V defined over $\mathcal{X} \times \mathcal{S}^\infty$:

(A.1)

$$TV(x, S) = \bar{\pi}(x, s_0) + E \left[\max_{\ell \in \mathcal{I}} \left\{ \phi_{i\ell}, \sup_{\ell \in \mathcal{I}} (-c(\ell, x, s_0) + \beta E [V(x_{i,t+1}, G(S)) | x_{it} = x, \ell_{it} = \ell]) \right\} \right],$$

for all $x \in \mathcal{X}$ and $s \in \mathcal{S}^\infty$. Note that the dependence of the cost function on s_0 encodes its potential dependence on z_0 . The Bellman operator greedily optimizes investment and the exit rule with respect to the value function V . Accordingly, the greedy investment and cut-off exit strategies with respect to \bar{V} , $\bar{\iota} : \mathcal{X} \times \mathcal{S}^\infty \rightarrow \mathcal{I}$ and $\bar{\rho} : \mathcal{X} \times \mathcal{S}^\infty \rightarrow [0, \sup_{x,s,\ell} \pi(x, s, \ell)/(1 - \beta) + \bar{\phi}]$, are respectively given by:

$$(A.2) \quad \bar{\iota}(x, S) = \underset{\ell \in \mathcal{I}}{\operatorname{argmax}} \left(-c(\ell, x, s_0) + \beta E [\bar{V}(x_{i,t+1}, G(S)) | x_{it} = x, \ell_{it} = \ell] \right),$$

$$(A.3) \quad \bar{\rho}(x, S) = \max_{\ell \in \mathcal{I}} \left(-c(\ell, x, s_0) + \beta E [\bar{V}(x_{i,t+1}, G(S)) | x_{it} = x, \ell_{it} = \ell] \right).$$

Under our assumptions, these strategies exist and are unique. We denote $\bar{\mu}(x, S) = (\bar{i}(x, S), \bar{p}(x, S))$.

Finally, we define the following two operators. First, for all $\mu \in \hat{\mathcal{M}}_{ns}$ and $\lambda \in \hat{\Lambda}_{ns}$ we define the operator H_1 as:

$$H_1(\mu, \lambda) = (\{s_t\}_{t=0}^{\infty}, \lambda),$$

where $\{s_t\}_{t=1}^{\infty}$ is determined by equations (3.1), (3.2), and (3.3), given (μ, λ) and the initial state s_0 . The first component of the operator H_1 maps an initial state, and sequence of strategies and entry rates into a sequence of expected states. The second component applies the identity to the sequence of entry rates.

Second, for all $S \in \mathcal{S}^{\infty}$ and $\lambda \in \hat{\Lambda}_{ns}$, we define the following operator:

$$H_2(S, \lambda) = \left\{ \{\bar{\mu}(x, G^t(S))\}_{x \in \mathcal{X}}, \max\{0, \min\{\lambda_t + \beta \bar{V}(x^e, G^{t+1}(S)) - \kappa, \bar{\lambda}\}\} \right\}_{t=0}^{\infty}.$$

The operator H_2 maps a sequence of states S and entry rates λ into a sequence of optimal strategies and updated entry rates.

A.2 Outline of Proof

We prove Theorem 3.1 at the end of the section by leveraging the lemmas we prove in Section A.3. We provide an outline here. For $\mu \in \hat{\mathcal{M}}_{ns}$ and $\lambda \in \hat{\Lambda}_{ns}$, we define the operator $H(\mu, \lambda) = H_2 \circ H_1(\mu, \lambda)$. It is simple to verify that a fixed point of H is a NOE. The proof uses the Brouwer-Schauder-Tychonoff's theorem (Aliprantis and Border, 2006) to show that a fixed point of H in fact exists in $\hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$. The main steps of the proof are the following:

1. Prove that H_1 is a continuous operator (Lemma A.1).
2. Prove that H_2 is a continuous operator (Lemma A.3).
3. Prove that $H_2 \circ H_1$ maps elements from $\hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$ into itself (Lemmas A.1 and A.3 together).
4. Show that $\hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$ is compact.
5. Show that if NOE strategies and entry rates converge as time progresses, they converge to OE strategies and entry rates.

A.3 Lemmas

Lemma A.1. *The operator H_1 maps elements from $\hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$ into $\mathcal{S}^{\infty} \times \hat{\Lambda}_{ns}$ and is continuous.*

Proof. Consider $(\mu, \lambda) \in \hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$ with limit $(\tilde{\mu}, \tilde{\lambda})$. Let s_t be the expected state at time t given by the recursions (3.1), (3.2), and (3.3). Because for each t the transition kernels are continuous functions of strategies $\{\mu_i\}_{i=0}^t$, s_t is continuous in μ for all t .

Now let us show that $\{s_t\}_{t=0}^{\infty}$ is a converging sequence. Note that because profits are bounded and ϕ_{it} has support in \mathfrak{R}_+ , there is a probability uniformly bounded away from zero over all strategies $\mu \in \hat{\mathcal{M}}_{ns}$, states $x \in \mathcal{X}$, and time periods $t \in \mathbb{N}$, of exiting the industry at each time. Therefore, for all $\mu \in \hat{\mathcal{M}}_{ns}$, $\sup_{t \geq 0} \sigma(P_{\mu_t}) < 1$, where $\sigma(P)$ is the spectral radius of the matrix P .

Let \tilde{s}_t be the expected state of the industry at time period t under oblivious strategy and entry rate $\tilde{\mu}$ and $\tilde{\lambda}$, starting from an empty industry and the aggregate shock invariant distribution. Because $\sup_{t \geq 0} \sigma(P_{\tilde{\mu}}) < 1$, \tilde{s}_t converges.

Again, because $\sup_{t \geq 0} \sigma(P_{\mu_t}) < 1$ and transition kernels are continuous, it is simple to show that $\lim_{t \rightarrow \infty} \|s_t - \tilde{s}_t\| = 0$, because transient behavior vanishes. These two limits together imply that $\{s_t\}_{t=0}^{\infty}$ converges, which completes the proof. □

Now, we show that the operator H_2 is continuous. First, we show a preliminary lemma.

Lemma A.2. *The value function \bar{V} is the unique solution of Bellman's equation $V = TV$ within the class of bounded functions. Moreover, \bar{V} and $\bar{\mu}$ are continuous over $\mathcal{X} \times \mathcal{S}^{\infty}$.*

Proof. Because single-period profits, investments, and expected sell-off value are bounded, \bar{V} is the unique solution of Bellman's equation $V = TV$ within the class of bounded functions (Bertsekas, 2001). We use the contraction mapping theorem to prove that, additionally, \bar{V} is continuous. In particular, we show that T has a fixed point within the class of *bounded and continuous* functions. Because the fixed point is \bar{V} , \bar{V} is continuous.

Let $C_b(\mathcal{X} \times \mathcal{S}^{\infty}, \mathfrak{R})$ be the space of continuous and bounded real-valued functions with domain $\mathcal{X} \times \mathcal{S}^{\infty}$. Recall that $C_b(\mathcal{X} \times \mathcal{S}^{\infty}, \mathfrak{R})$ is a complete metric space with the metric defined by the supremum norm (Marsden and Hoffman, 1993). Also, recall that T is a contraction in the supremum norm because $\beta < 1$ (Bertsekas, 2001). Now, we show that T maps elements from $C_b(\mathcal{X} \times \mathcal{S}^{\infty}, \mathfrak{R})$ into $C_b(\mathcal{X} \times \mathcal{S}^{\infty}, \mathfrak{R})$.

Take a function $V \in C_b(\mathcal{X} \times \mathcal{S}^{\infty}, \mathfrak{R})$. By definition (A.1)

$$TV(x, S) = \bar{\pi}(x, s_0) + E \left[\max \left\{ \phi_{it}, \sup_{\iota \in \mathcal{I}} (-c(\iota, x, s_0) + \beta E[V(x_{i,t+1}, G(S)) | x_{it} = x, \iota_{it} = \iota]) \right\} \right]$$

The operator G is trivially continuous. The cost function, the value function V , and the transition kernel are continuous by assumption. The random variable ϕ_{it} is absolutely continuous. Moreover, ι is optimized over a compact space. By Berge's Maximum Theorem, $TV(x, S)$ is continuous. Furthermore, TV is bounded because profits, investment costs, and expected sell-off value are uniformly bounded over all states. Therefore, T maps elements from $C_b(\mathcal{X} \times \mathcal{S}^\infty, \mathfrak{R})$ into $C_b(\mathcal{X} \times \mathcal{S}^\infty, \mathfrak{R})$.

Using the contraction mapping theorem (Marsden and Hoffman, 1993), we conclude that T has a fixed point among the class of continuous and bounded functions, therefore, \bar{V} is continuous. Using Berge's Maximum Theorem again, we conclude that $\bar{\iota}$ and $\bar{\rho}$ are also continuous. \square

Lemma A.3. *The operator H_2 maps elements from $\mathcal{S}^\infty \times \hat{\Lambda}_{ns}$ into $\hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$, and is continuous.*

Proof. By Lemma A.2 and because the operator G is continuous, for each $t \geq 0$,

$\{\{\bar{\mu}(x, G^t(S))\}_{x \in \mathcal{X}}, \max\{0, \min\{\lambda_t + \beta \bar{V}(x^e, G^{t+1}(S)) - \kappa, \bar{\lambda}\}\}\}$ is continuous in (S, λ) . Hence, H_2 is continuous.

Now, $\lim_{t \rightarrow \infty} G^t(S) = \{s, s, s, \dots\}$, for some $s \in \mathcal{S}$, because $S \in \mathcal{S}^\infty$. Hence, using the continuity of $\bar{\mu}$ (Lemma A.2), $\lim_{t \rightarrow \infty} \{\bar{\mu}(x, G^t(S))\}_{x \in \mathcal{X}} = \tilde{\mu}$, for some $\tilde{\mu} \in \tilde{\mathcal{M}}$. By a similar argument, the second component of H_2 also converges. Hence, H_2 maps elements from $\mathcal{S}^\infty \times \hat{\Lambda}_{ns}$ into $\hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$. \square

A.4 Proof Theorem 3.1

Proof. Because the range of $\mu \in \tilde{\mathcal{M}}$ and $\lambda \in \tilde{\Lambda}$ are bounded, by Tychonoff Theorem (Royden, 1988), $\hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$ is a compact set with the product topology. It is also a convex set and a subset of a locally convex Hausdorff space. By Lemmas A.1 and A.3, H is a continuous operator that maps elements from $\hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$ into itself. Hence, by Brouwer-Schauder-Tychonoff' theorem, there exists a fixed point (μ, λ) of H in the set $\hat{\mathcal{M}}_{ns} \times \hat{\Lambda}_{ns}$. The fixed point is a NOE, s.t. for all x , $\lim_{t \rightarrow \infty} \mu_t(x) = \tilde{\mu}(x)$ and $\lim_{t \rightarrow \infty} \lambda_t = \tilde{\lambda}$.

To finish the proof we show that $(\tilde{\mu}, \tilde{\lambda}) \in \mathcal{M} \times \Lambda$ is an OE. By the argument in the proof of Lemma A.1, the sequence S in the first component of $H_1(\mu, \lambda)$ converges to \tilde{s} , the long-run expected state under oblivious strategy and entry rate $(\tilde{\mu}, \tilde{\lambda})$. Note that A.3, $\mu_t(\cdot) = \bar{\mu}(\cdot, G^t(S))$. Taking $\lim_{t \rightarrow \infty}$, using the fact that $\lim_{t \rightarrow \infty} G^t(S) = \{\tilde{s}, \tilde{s}, \dots\}$, and that $\bar{\mu}$ is continuous (Lemma A.2), we conclude that $\tilde{\mu}(\cdot)$ is an OE strategy. Because \bar{V} is continuous (Lemma A.2), the associated OE value function is $\bar{V}(x, \{\tilde{s}, \tilde{s}, \dots\})$. Because λ is a fixed point of H and taking $\lim_{t \rightarrow \infty}$, if $\tilde{\lambda} \in (0, \bar{\lambda})$, then $\beta \bar{V}(x^e, \{\tilde{s}, \tilde{s}, \dots\}) = \kappa$. If $\bar{V}(x^e, \{\tilde{s}, \tilde{s}, \dots\}) - \kappa < 0$, then $\tilde{\lambda} = 0$. Similarly, if $\bar{V}(x^e, \{\tilde{s}, \tilde{s}, \dots\}) - \kappa > 0$, then $\tilde{\lambda} = \bar{\lambda}$. Hence, $\tilde{\lambda}$ is an OE entry rate. The result follows. \square

B Technology adoption: Single-period payoffs

Suppose that agent i adopts the technology. Denote by $s_{-i,t}$ the number of other adopters, which is a vector of natural numbers. The agent derives the utility from making phone calls to other agents that also adopted, denoted by a set $A(s_{-i,t})$. After making $k - 1$ calls, the users decide if to make an additional k -th call and who to call. The utility of the next phone call to each potential recipient $j \in A(s_{-i,t})$ is given by a vector of values

$$u_{i,j,k,t} = \gamma_{ij} - \theta'_3 \eta_{ikt} + \theta_4(k - 1) + z_t + \epsilon_{ijkt}.$$

The utility is composed of three parts: the static connection utility γ_{ij} , a component depending on previous phone calls in a sequence $-\theta_3 \eta_{ikt} + \theta_4(k - 1)$, and the utility shock $z_t + \epsilon_{ijkt}$.

The γ_{ij} measures the link synergy between users i and j that depends on the types m_i and m_j . The values γ_{ij} are given by a 64×64 matrix, summarizing the network topology. The term η_{ikt} denotes the vector denoting number of users of the same type called in the previous $k - 1$ calls. The expression $-\theta'_3 \eta_{ikt}$ represents the notion that users may get diminishing returns by repeatedly calling the same group of employees. The term $\theta_4(k - 1)$ shifts the utility down linearly in the number of calls, capturing overall diminishing returns of using the video-conferencing technology. The term ϵ_{ijkt} is an idiosyncratic shock to the value of calling user j as k -th call in the sequence. The outside option of not calling is set to zero.

For each period, the users start with an empty call sequence and add calls as long as and incremental utility of the additional call is greater than zero. The length of the sequence is determined as:

$$(B.1) \quad K = \min\{k \in \mathbb{N} : \max_{j \in A(s_{-i,t})} u_{i,j,k,t} < 0\} - 1$$

Users are myopic when considering calls and myopically add the call of the most current utility, until the value of the best additional call is negative (see RT for the discussion of this assumption). The addition of marginal adopters affects the utility of the calling sequence in two ways: (i) the extra adopter can be of a type that previously did not exist in the network, which allows new connection synergies γ_{ij} to be realized; (ii) each extra adopter generates additional draw of ϵ_{ijkt} , which increases the length of the calling sequence.

Denote a chosen set of phone calls by user i and time t as Ω_{it} . The expected per-period utility of adoption is given by

$$U(s_{-i,t}, z_t) = E_{\Omega_{it}} \sum_{k=1}^K u_{i,j_k,k,t},$$

where K is the total number of calls in the sequence Ω_{it} . The expectation is taken over possible random

calling sequences.

The above discussion follows the original RT model for which the number of agents is finite. Our model is a variation with the continuum of agents, so we extend the utility function to the set \mathcal{S} . We apply a simple linear interpolation by setting $U(s_{-i,t}, z_t) = EU_{it}(s'_{-i,t}, z_t)$. Each element $s'_{j,t}$ of $s'_{-i,t}$ is a IID Bernoulli random variable that obtains $\lceil s_{j,t} \rceil$ with probability $s_{j,t} - \lfloor s_{j,t} \rfloor$, and $\lfloor s_{j,t} \rfloor$ otherwise. The operators $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ and ceiling and floor functions, respectively.

We assume that $\lim_{t \rightarrow \infty} z_t = z < \infty$, and denote the limit of the utility function by

$$\bar{U}_i = \lim_{t \rightarrow \infty} U(s_{-i,t}, z_t).$$

The limit exists because $s_{-i,t}$ is increasing in t , U is increasing in $s_{-i,t}$, and U is bounded from above (by the full adoption value).

C Code for RT model

C.1 Oblivious Equilibrium

```
n = 1
while (tol > 1e-6)
    [u K] = profit(stilde);
    Vtilde = u / (1 - beta);
    for j = 1:64
        stilde_new(j) =
            N(j) * normcdf(Vtilde(j), mu(j), sigma(j));
    end
    tol = norm(stilde_new - stilde);
    stilde = stilde + (stilde_new - stilde) / (n^0.5 + 1);
end
```

C.2 Nonstationary OE

```
while (norm(stilde_noe - stilde_noe_new) > 1e-9)
    % Value of adoption
    Va(T, :) = U(T, :) / (1 - beta);
    for t = T-1:-1:1
        Va(t, :) = profit(stilde_noe(t, :)) + beta * Va(t+1, :);
    end

    stilde_noe_new = zeros(T+1, M);
    % Reseed the pseudo random numbers
    randn('seed', 9077984732);
    % Simulate adoption shares for each type
    for m = 1:M
        for r = 1:R
```

```

% Value of waiting
Vwait=zeros(T+1,1);
% Draw adoption cost
F=mu(m)+sigma(m)*randn;
for t=T:-1:1
    % Adopt at t if not adopted before t?
    [Vwait(t) a(t)]=max([Va(t,m)-F beta*Vwait(t+1)]);
end
% Determine earliest adoption date
at=find(a==1,1);
% Adoption is implemented next period
if (~isempty(at))
    stilde_noe_new(at+1:T,m)=stilde_noe_new(at+1:T,m)+1;
end
end
end
stilde_noe_new=(stilde_noe_new/R);
end

```

D Tables and Figures




Figure 1: Comparison of uniform feasible subsidy and optimal subsidy.

	Quality Ladder	Cost Competition
Market size (M)	100.0	200.0
Demand slope (σ)	-	10.0
Marginal cost (γ_0)	-	1.0
Marginal cost (γ_1)	-	0.50
Marginal cost (γ_2)	-	5.0
Fixed cost (F)	-	0.0
Second demand (m_2)	-	5.0
Second demand quality slope (α)	-	75.0
Income (Y)	1.0	-
Quality sensitivity (θ_1)	0.8	-
Price sensitivity (θ_2)	0.5	-
State multiplier (ψ)	1.0	-
Marginal cost	0.5	-
Investment effectiveness	3	1
Depreciation probability	0.2	0.4
Investment cost	[8-15]	[50-155]
Discount factor	0.9	0.9
States per firm	8	8

Table 1: Parameters for numerical simulations




Figure 2: Comparison of the dynamics in predicted consumer surplus, producer surplus and average investment: Quality-Ladder model (low investment cost).




Figure 3: Comparison of the dynamics in predicted consumer surplus, producer surplus and average investment: Quality-Ladder model (moderate investment cost).




Figure 4: Comparison of the dynamics in predicted consumer surplus, producer surplus and average investment: Quality-Ladder model (high investment cost).




Figure 5: Comparison of the dynamics in predicted consumer surplus, producer surplus and average investment: Quantity-Cost competition (low investment cost).




Figure 6: Comparison of the dynamics in predicted consumer surplus, producer surplus and average investment: Quantity-Cost competition (high investment cost).




Figure 7: Comparison of the dynamics in average investment: Quality-Ladder model with aggregate shocks (low investment cost). The panels show low, medium and high value of the shock at $t = 0$, respectively.




Figure 8: Comparison of the dynamics in average investment: Quality-Ladder model with aggregate shocks (high investment cost). The panels show low, medium and high value of the shock at $t = 0$, respectively.




Figure 9: Average absolute deviation error: Quality-Ladder model (low investment cost)