Why Are Underperforming Firms Rarely Acquired?

Wenyu Wang and Yufeng Wu*

November 11, 2016

Abstract

Only 4.5% US firms in the bottom performance quintile are taken over annually, and the association between a firm's performance and its subsequent takeover exposure is close to zero. These observations cast doubts on the effectiveness of the takeover market to reallocate resources towards more efficient users and better management. In this paper, we revisit this problem by estimating a dynamic model in which takeovers are pursued either to enhance firm performance or to create control benefit for managers . Our estimates suggest that the takeover market is overall efficient with most value-enhancing mergers consummated quickly. Managers' entrenchment blocks less than 10% of the profitable deals. Meanwhile, an efficient takeover market triggers an ongoing selection effect so that underperforming firms with more entrenched managers survive longer. Even if this selection effect is small each period, it accumulates over time and is amplified as the economy evolves. As a result, underperforming firms become overrepresented by managers with high control benefit, leading to the weak takeover-performance sensitivity in data.

Key words: mergers and acquisition, asset reallocation, private benefit of control, structural estimation

*Wenyu Wang (wenywang@indiana.edu) are with Kelley School of Business, Indiana University, and Yufeng Wu (yufengwu@illinois.edu) is with College of Business, University of Illinois Urbana-Champaign. We thank Andrey Golubov, Dimitris Petmezas and Nickolaos Travlos for providing us with the data on adjusting revelation effect in equity and mixed payment deals. We thank Ron Kaniel, Janet Gao, Neil Pearson and Toni Whited for their helpful comments.

Introduction

The literature on mergers theorizes that they play an important role in shifting economic resources toward more efficient firms and management. According to the Q-theory of mergers, high-valuation firms acquire low-valuation firms to unlock efficiency gains. Underperforming firms are most likely to be targeted and acquired because the unexploited opportunity to cut costs and increase earnings are likely to be the highest among such firms.

Despite the focus on the theoretical work, only a small fraction of underperforming firms are acquired annually in the data.[1] Examining the full spectrum of firms also reveals a weak association between a firm's performance and its subsequent takeover exposure: a firm's performance, measured using different metrics, is only weakly related to its probability of being acquired. Agrawal and Jaffe (2003), adjusting for firm size, industry, and past performance, conclude that "there is little evidence of pre-acquisition underperformance for the entire sample of targets". This weak association is often interpreted as the takeover market being a relatively inefficient channel for asset reallocation. The evidence is more consistent with another influential view in the literature which states that takeover decisions are mainly influenced by the conflicts of incentives between the principles and the managers of the firms.

In this paper, we revisit these conflicting views by estimating a structural model. Our goal is two-old. First, we examine why underperforming firms are rarely acquired, or in other words, what leads to a weak association between a firm's performance and its takeover exposure (henceforth referred to as *takeover-performance sensitivity*). Second, we ask whether the weak takeover-performance sensitivity implies that the takeover market is inefficient in reallocating assets and restoring firm performance.

The model we consider is an industry partial equilibrium model with entry and exits. In the model, a firm exits the market through either natural death or acquisition. Mergers and acquisitions happen for two reasons. First, there is heterogeneity in firms' efficiency in utilizing

[1] On average, only about 4.5% of firms in the bottom performance quintile, as measured by return-on-assets (ROA) or operating cash flow scaled by total book assets (OCF), are acquired each year.

assets. Acquirers can transfer their technology and organizational structure to the target after paying a stochastic integration cost. This type of merger allows good firms to create value through purchasing the underperformers. Second, there is heterogeneity in managers' empire-building incentives. High-entrenchment managers derive larger utility from each unit of asset under their management. Therefore, these managers have incentives to acquire assets from low-entrenchment managers who do not value their control rights as much. Mergers pursued purely on this basis are likely to be value-destroying due to positive integration cost.

In a static setting, there may be two reasons why underperforming firms are rarely acquired. On the one hand, it could be that the benefit from reallocating resources across firms is limited, due to low dispersion in efficiency or a high cost of integration. On the other hand, it is also possible that there is wide dispersion in managers' empire-building incentives, in which case assets are reallocated towards managers with high private control benefit, instead of those who utilize assets more efficiently. This effect could render the efficiency gains from mergers a second-order consideration, therefore weakening the tie between takeover exposure and target firm performance.

Examining the model in a dynamic setting points to a third possibility for why underperforming firms are rarely acquired–the ongoing selection effect. The takeover market induces selection of firms based on two dimensions–firms' operating efficiency and managers' private benefit of control. Therefore, underperforming firms with low-entrenchment managers are likely to exit the market quickly while the ones with high-entrenchment managers can survive longer. This selection process naturally creates a negative correlation between firms' performance and managers' control benefit among the existing firms in the economy. This selection effect may not be particularly strong for any particular period. However, it accumulates over time and the negative correlation is amplified as the economy evolves. Hence, in the steady state, underperforming firms are overrepresented by those with high private benefit of control. Such control benefit serves as a barrier to takeover attempts and significantly attenuates the takeover-performance sensitivity.

The model imposes no priors on the relative importance of the takeover motives, and we let the data tell us how they affect firms' decisions regarding mergers and acquisitions. We do so by estimating the model's parameters using the simulated method of moments (SMM). Our dataset includes all firm-year observations covered by the CRSP-Compustat Merged database (CCM) and the mergers and acquisitions involving U.S. public firms during the years 1980 to 2014. The key parameters of interest are the dispersion in managers' private benefit of control, as well as the mean and standard deviation of the potential efficiency gains from mergers. We choose the parameter values by minimizing the distance between the model-implied moments with the same set of moments calculated on the actual data.

In our estimation, the distribution of operating efficiency is identified directly off the cross-sectional distribution of firms' ROA. The implicit assumption is that both firms engaged in mergers and standalone firms share, ex ante, the same law of motion for operating profit. Takeovers are influenced by the ex post realizations of an individual firm's efficiency. The dispersion in managers' control benefit is identified off the fraction of value-destroying mergers in the economy. Value-destroying mergers may happen in the model when the acquirer manager enjoys much higher control benefit than the target manager and decides to overpay the target to consummate a deal even if this deal destroys the combined firm's value. Therefore, a larger fraction of value-destroying mergers indicates a higher dispersion of control benefit across firm managers. The combined firm's announcement return helps identify the stochastic integration costs. Overall, our result suggests that the model can match targeted data moments related to firm operating performance and takeover activities. The model also fits a wide range of moments that are not explicitly targeted in our SMM estimation, including the duration of underperformance and the probability that a firm exits the bottom performance quintile through self-recovery, death, or acquisition.

Using the estimated model as a laboratory, we examine how the selection effect influences takeover-performance sensitivity in the simulated data. We find that once the selection based on managers' private benefit is properly controlled for, the model-predicted takeover-performance sensitivity becomes five times more significant; this effect is mainly a result of firms with well below-average operating performance. Intuitively, this is because most underperforming firms with low private control benefit are acquired quickly, leaving mostly those with high private benefit to survive in the steady state. Therefore, the resistance stemming from managers' private control benefit is strongest among these underperforming firms. Controlling for this omitted variable allows us to restore a much stronger takeover-performance relation.

We present more empirical evidence supporting the existence of the selection effect in actual data. Previous studies use pooled regressions to explore both the cross-sectional and time-series correlation between firm performance and takeover exposure. Our model suggests that the selection effect mainly attenuates the cross-sectional correlation but has a negligible effect on time-series correlation if managers' private benefit of control is highly persistent.[2] Consistent with this model prediction, we show that the takeover-performance sensitivity becomes much larger in magnitude (almost ten times larger) and highly significant (with a t-stat of -9.1) after we control for the firm fixed effects in the regression, limiting the identification to time-series variations only. On the contrary, when we control for the year fixed effects, the estimated sensitivity becomes even weaker than that in the pooled regression.

Next, we turn to the question of whether the market characterized by both synergistic takeovers and managers' conflict of interest is efficient. To do so, we perform a counterfactual analysis using the estimated model, where we assume all managers have zero control benefit. After eliminating managers' private benefit of control, all mergers happen for value-enhancing reasons, and some profitable deals that were blocked by entrenched target managers in the baseline model can now be consummated. If managers' private benefit of control creates a large inefficiency in the takeover market, eliminating it should significantly increase firm value. We find the average firm valuation ratio (i.e., the market-to-book equity ratio) rises by 1% in

[2]The selection effect attenuates the cross-sectional correlation between firm performance and takeover exposure, because due to this selection effect, underperforming firms surviving in the economy are overrepresented by those with high control benefit, and thus they appear much more difficult to acquire than what the Q-theory of M&A would suggest. The selection effect, however, has no effect on the time-series correlation between firm performance and takeover exposure if managers' private benefit of control is constant. This is because given a firm manager's private benefit of control, the firm always becomes more vulnerable to takeovers when its performance deteriorates.

this counterfactual exercise. Is the 1% improvement in firm value large or small? We compare it with a benchmark in which the whole M&A market is shut down. When takeovers are banned in the model economy, the average firm value drops by 6.4%. We therefore conclude that managers' private benefit of control results in a relatively small inefficiency, amounting to about 13% of the total value created in the takeover market.

It is worth emphasizing the only key friction in our model is the managers' private control benefit. This single friction, nevertheless, generates two seemingly contradictory predictions – a very weak performance-takeover sensitivity and a relatively efficient market for asset reallocation. How can these findings be reconciled? Our model provides a plausible explanation: in each period, only a small fraction of the firms that receive negative shocks and become underperforming have excessively entrenched managers who block the value-enhancing deals. The efficiency losses due to managers' control benefit are therefore small in the economy. However, the small fraction of underperforming firms with high control benefit accumulate over time and become overrepresented in the pooled cross-section of firms, creating a perception that underperforming firms are not acquired more frequently than other firms. This perception is misleading because most underperforming firms are actually acquired quickly once they become underperforming and exit the economy. The fact that the underperforming firms left in the economy are hard to acquire does not necessarily imply that the takeover market is essentially inefficient.

Our paper contributes to three streams of literature. First, it adds to the large body of literature on takeover motives. Eisfeldt and Rampini (2006), Jovanovic and Rousseau (2008), and Eberly and Wang (2009) , among others, model the takeover market as a channel to efficiently reallocate assets towards more advanced technology and better management. On the other hand, Jain (1985), Morck, Shleifer, and Vishny (1990), and Hartzell, Ofek, and Yermack (2004) emphasize the conflicts of interest between owners and managers. They argue that many acquisitions are undertaken following managerial empire building incentives, and that they serve little economic purpose. Maksimovic and Phillips (2001, 2002) reexamine these conflicting views using plant-level data. They conclude that the timing and pattern of observed asset sales are consistent with efficient allocation of resources through the takeover market, and are likely to be profit maximizing for firms. Our paper differs from Maksimovic and Phillips (2001, 2002) in two important aspects. First, while their paper includes partial-firm sales, our paper focuses on full-firm mergers and acquisitions, where the relation between takeover and target prior performance is the weakest , and hence there are greater doubts regarding asset reallocation efficiency. Second, by relying on a structural model, we can talk about not only the completed deals, but also the ones blocked by managers' private control benefit. This approach allows us to analyze the welfare gains/losses from the takeover market more comprehensively.

Our work is also related to the literature on firm entry and exits and industry dynamics. In his seminal work, Hopenhayn (1992) characterizes the long-run equilibrium in an industry with heterogeneous firms. Dimopoulos and Sacchetto (2015) extend Hopenhayn's model by incorporating mergers and show that sizable synergistic benefits are created through firms' takeover decisions. Our setup deviates from theirs by modeling takeover as a selection mechanism, based not only on firm performance, but also on managers' empire-building incentives. This dynamic selection effect implies that underperforming firms with high managerial control benefit are more likely to stay in the economy. The high control benefit then acts as barriers for subsequent takeover bids and attenuates the relation between target performance and its takeover exposure.

Methodologically, our paper belongs to the growing literature that employs structural model calibration or estimation to answer standard corporate finance questions in capital investment, leverage choice, CEO turnover, and market competition, as summarized in Strebulaev and Whited (2011). We extend the application of structural estimation to studying firms' takeover decisions, with a special focus on the weak takeover-performance sensitivity. We also relate our findings to the overall efficiency of the takeover market and quantify the gains created through asset reallocation.

The paper is organized as follows: Section 1 presents the model setup and model solution;

Section 2 introduces the data and discusses how the model parameters are estimated; Section 3 presents the estimation results and uses the estimated model to investigate the takeover-performance sensitivity and takeover market efficiency; and Section 4 concludes.

1 Model

The model is in discrete time and infinite horizon. There are two types of agents in the model: firms' shareholders and managers. We assume that both the shareholders and the managers are risk-neutral, the firms are fully equity-financed, and there is no friction in raising capital or distributing residual earnings.

1.1 Basic Setup

The model features a cross section of heterogeneous firms. Each firm is characterized by the book value of assets, B, profitability, α, and its manager's private benefit of control on each unit of capital, π. The firm has constant return to scale technology; hence the total profit generated in each period is αB; the total private benefit of control enjoyed by its manager is πB. We assume that the book value of assets, B, remains constant over time when a firm stands alone.[3]

A firm's profitability α evolves following a mean-reverting process:

$$\alpha^{'} = \mu_\alpha + \phi_\alpha(\alpha - \mu_\alpha) + \upsilon_\alpha \epsilon, \tag{1}$$

where μ_α is the long-run population mean of firms' profitability, ϕ_α measures the persistence, and υ_α captures the standard deviation of the profitability innovations.

A firm's incumbent manager derives a private benefit of control from running the firm. We assume that the private benefits among the cross section of managers are normally distributed

[3] This specification follows Taylor (2010) and effectively assumes that a firm distributes all residual earnings to its shareholders each period.

with mean μ_π and variance σ_π^2. When a firm first hires a manager, it takes a random draw from the distribution $N(\mu_\pi, \sigma_\pi^2)$, and the realization of the manager's private benefit stays constant as long as she remains in her position. As in Albuquerque and Schroth (2015), we do not model a manager's private benefit as coming directly from the firm's cash flows. Instead, we interpret it as a psychological benefit that the manager derives from overseeing the firm. This benefit will be forgone if the manager loses her control right, through either a firm closure or takeover.

If a takeover takes place, the target assets are transferred to the acquirer and inherit the acquirer's profitability. Therefore, in our model, mergers can create value by transferring assets to the user who operates them more efficiently, consistent with the neoclassic Q-theory of M&A. Meanwhile, mergers can also create additional gains or losses beyond the transfer of profitability. We model such additional gains or losses generated in mergers as sB. We specify s based on each unit of target assets and assume it follows the distribution below:

$$s = \mu_s + \delta|\alpha_T - \alpha_A| + \sigma_s \varepsilon. \tag{2}$$

This specification assumes that s is pair-specific. δ captures the correlation between s and the performance differentials between the target and the acquirer.

Figure 1 describes the timeline of the model. At the beginning of each period, a given fraction of new firms enters into the market. The current period profitability,α, is realized for both the new and existing firms, and it is observed by the incumbent manager and the shareholders. The shareholders first decide whether to shut down the firm or not. If the shareholders decide to shut down the firm, they receive their reservation value $\underline{U}B$. If the shareholders decide to keep the firm running for this period, then the firm produces, and its manager learns the opportunities in the takeover market–the firm can be matched with either a potential acquirer or a potential target in each period, with probability ρ_A and ρ_T, respectively. Since firms may play different roles in the model, we use the subscript T to denote the target firm and A to denote the acquirer firm.

Let $U(B, \alpha, \pi)$ be the value function of the firm at the beginning of each period *before* the opportunities in the takeover market are discovered, which should satisfy the following Bellman equation:

$$U(B,\alpha,\pi) \quad = \quad \left\{\alpha B+\beta E\left[U(B,\alpha^{'},\pi)\right]+\beta\left(\rho_{A}E\left[\Delta^{U}_{Acq}\right]+\rho_{T}E\left[\Delta^{U}_{Tar}\right]\right)\right\}\cdot(1-1_{Cls})+\underline{U}B\cdot 1_{Cls}, \quad (3)$$

where$E\left[\Delta^{U}_{Acq}\right]$ and $E\left[\Delta^{U}_{Tar}\right]$ represent the shareholders' expected gains from the takeover market when the firm acts as an acquirer or a target, respectively. We defer the discussion of $E\left[\Delta^{U}_{Acq}\right]$ and $E\left[\Delta^{U}_{Tar}\right]$ to the next subsection.

The intuition of equation 3 is straightforward. If shareholders decide to keep the firm running this period (i.e., $1_{Cls}=0$), they get the current period profit flow plus the discounted continuation value of the firm when it operates as an independent entity. The shareholders also capture the expected gains from the merger market. If the shareholders decide to shut down the firm (i.e., $1_{Cls}=1$), they get their reservation value $\underline{U}B$. Shareholders' close-down decision, 1_{Cls}, is defined as:

$$1_{Cls}=\begin{cases}1 & if\, U(B,\alpha,\pi)<\underline{U}B\\ 0 & otherwise.\end{cases} \quad (4)$$

That is, the shareholders decide to shut down the firm if the value they expect to get from the firm drops below their reservation value, $\underline{U}B$.

1.2 Takeover Decisions

A firm's merger and acquisition decisions are made by its incumbent manager, whose objective is to maximize his own utility. We use $V(B,\alpha,\pi)$ to denote the manager's utility. The Bellman equation for his problem is given by:

$$V(B,\alpha,\pi) \quad = \quad \left\{\kappa\alpha B+\pi B+\beta E\left[V(B,\alpha^{'},\pi)\right]+\beta\left(\rho_{A}E\left[\Delta^{V}_{Acq}\right]+\rho_{T}E\left[\Delta^{V}_{Tar}\right]\right)\right\}\cdot\left(1-1_{Cls}\right), \quad (5)$$

where $\kappa\alpha B+\pi B$ is the flow utility that the manager derives in each period. The manager's flow utility has two components: the cash flow to the shareholders and the private benefit of control that herself derives. κ measures how the incumbent manager weighs the shareholders' gains relative to her private control benefit: A larger value of κ implies a smaller agency conflict and a better alignment of interests between the incumbent manager and the firm's shareholders. In general, κ is determined by two factors. The first is the manager's inherent preference. Some managers highly value taking control of a large empire relative to other things; the second factor is the discipline that the manager faces. For example, a manager with a high empire-building incentive might face close monitoring by the board. Therefore, she lacks the ability to implement her decision if it is not in the interest of the shareholders. In this case, we should expect high κ. On the contrary, if it is costly for the shareholders to launch a control challenge against a possibly entrenched manager, then κ will be, ceteris paribus, small. Since utility functions can be scaled without affecting model solutions, we normalize κ to one, following Taylor (2010). With this normalization, κ should be interpreted as the incumbent manager's utility (disutility) of seeing shareholders gain (lose) one dollar.

Suppose an acquirer indexed by (B_A, α_A, π_A) meets a potential target indexed by (B_T, α_T, π_T). Equation 7 below define the merger gains accrued to the incumbent managers of the acquirer and target firms, respectively:

$$\Delta^V_{Acq} = \left\{E\left[V(B_A + B_T, \alpha_A^{'}, \pi_A)\right] + B_T s - E\left[V(B_A, \alpha_A^{'}, \pi_A)\right] - P\right\} \cdot 1_{Buy} \tag{6}$$

$$\Delta^V_{Tar} = \left\{P - E\left[V(B_T, \alpha_T^{'}, \pi_T)\right]\right\} \cdot 1_{Sell}. \tag{7}$$

Similarly, we can define the merger gains that accrue to the acquirer and target shareholders:

$$\Delta^U_{Acq} = \left\{E\left[U(B_A + B_T, \alpha_A^{'}, \pi_A)\right] + B_T s - P - E\left[U(B_A, \alpha_A^{'}, \pi_A)\right]\right\} \cdot 1_{Buy} \tag{8}$$

$$\Delta^U_{Tar} = \left\{P - E\left[U(B_T, \alpha_T^{'}, \pi_T)\right]\right\} \cdot 1_{Sell}. \tag{9}$$

To close the loop, we introduce how the offer price, P, is determined in a takeover. We assume that the target and acquirer managers bargain via a standard Nash game to determine the transaction price. Let θ be the bargain power allocated to the target manager. Then the transaction P should be set such that:

$$\theta \times \mathit{\Delta}^V_{Acq} \quad = \quad (1-\theta) \times \mathit{\Delta}^V_{Tar}. \tag{10}$$

The acquirer manager decides to purchase another firm if the deal enhances her own utility:

$$1_{Buy} = \begin{cases} 1 & if\ \mathit{\Delta}^V_{Acq} > 0 \\ 0 & otherwise. \end{cases} \tag{11}$$

Similarly, for a target manager, we have:

$$1_{Sell} = \begin{cases} 1 & if\ \mathit{\Delta}^V_{Tar} > 0 \\ 0 & otherwise. \end{cases} \tag{12}$$

A merger deal is consummated if and only if it is approved by both the acquirer and target managers. The rent-splitting rule specified by equation 10 ensures that the acquirer and target managers always agree on whether a deal should be pursued or not.

1.3 Steady-state Equilibrium

We characterize the model economy as a joint distribution of firms' state variables $\Gamma(\alpha, \pi)$.[4] In general, $\Gamma(\alpha, \pi)$ evolves over time, driven by three forces. First, the law of motion, described in equation 1, governs the dynamics of firm profitability α. Second, mergers and acquisitions endogenously change firms' α and π. Third, we assume that the new firms entering the economy each period follow the entrant distribution $\mathit{\Upsilon}(\alpha, \pi)$. We assume $\mathit{\Upsilon}(\alpha, \pi)$ follows a joint normal distribution:

$$\mathit{\Upsilon}(\alpha, \pi) \sim N\left(\begin{bmatrix} \mu_\alpha \\ \mu_\pi \end{bmatrix}, \begin{bmatrix} \sigma^2_\alpha & \rho_{\alpha\pi}\sigma_\alpha\sigma_\pi \\ \rho_{\alpha\pi}\sigma_\alpha\sigma_\pi & \sigma^2_\pi \end{bmatrix} \right).$$

[4] We do not track the distribution of B in model equilibrium, because the dynamics of firm size is not the main interest of this paper. Instead, we use the empirical distribution of B in our estimation. This approach significantly speeds up the model solution and allows us to empirically capture the effect of firm size in the estimation. Extending the model to incorporate the distribution of B as part of the steady-state equilibrium is feasible, but our results will remain quantitatively similar as long as the extended model can fit the empirical distribution of B in steady state.

Denote $Q(\alpha', \pi'|\alpha, \pi)$ as the transition matrix driven by the law of motion of α and firms' decisions to merge and closedown. The dynamics of $\Gamma(\alpha, \pi)$ can be derived as:

$$\Gamma(\alpha', \pi') = \int_{\alpha,\pi} Q(\alpha', \pi'|\alpha, \pi)\Gamma(\alpha, \pi) d\alpha d\pi + \Upsilon(\alpha, \pi). \tag{13}$$

We define a steady-state equilibrium $(U, V, 1_{Cls}, 1_{Sell}, 1_{Buy}, \Gamma)$ such that:

1. The shareholders' value function U and the decision to close down the firm 1_{Cls} solve equations 3 and 4, given the shareholders' expectation of the incumbent manager's decision rule 1_{Sell} and 1_{Buy} and the joint distribution $\Gamma(\alpha, \pi)$;

2. The incumbent manager's value function V and her decision regarding M&A, 1_{Sell} and 1_{Buy}, solve equations 5, 11, and 12, given her expectation of the shareholders' decision rule 1_{Cls} and the joint distribution $\Gamma(\alpha, \pi)$;

3. The transition matrix $Q(\alpha', \pi'|\alpha, \pi)$ is constructed from the law of motion defined in equation 1 and the dynamics of α and π caused by M&A;

4. The joint distribution $\Gamma(\alpha, \pi)$ evolves following equation 13 and remains constant; that is, $\Gamma(\alpha', \pi') = \Gamma(\alpha, \pi)$;

5. All agents' expectations are fulfilled in equilibrium (i.e., rational expectation).

To solve the steady-state equilibrium, we start with an initial distribution of firms, $\Gamma_0(\alpha, \pi)$. Given $\Gamma_0(\alpha, \pi)$ as the perceived distribution, we solve the manager's and shareholders' value functions and policy functions. We then compute the new distribution $\Gamma^{'}(\alpha, \pi)$ using equation 13. We use the updated distribution as the new perceived distribution $\Gamma(\alpha, \pi)$ and repeat the steps above until the distribution converges; that is, $\Gamma^{'}(\alpha, \pi) = \Gamma(\alpha, \pi)$.

1.4 Model Solutions

To illustrate the model implications, we solve the model with the parameters set to their estimated values.[5] We simulate a panel of firms according to the model solution and show their movements in an industry equilibrium via a heat map. Panel A corresponds to a three-dimensional heat map with the firms' beginning-of-period profitability on the x-axis and their managers' private benefit of control on the y-axis. On the z-axis, the probability that a firm gets acquired in any specific period is measured by the depth of the color– the darker color means that a firm is less likely to get acquired and thus is more likely to survive in the economy while the lighter color suggests the opposite. The takeover market induces a selection based on the firms' operating efficiency and the managers' private control benefit. Managers with high private control benefit and good firm performance seize limited gains from being taken over, and thus their firms are more likely to survive the selection. Therefore, we see downward sloping level curves in the probability of survival, especially in the lower left region of the plot where the takeover threats are the most relevant.

Panel B is constructed in a way similar to Panel A, except that the z-axis now measures a firm's cumulative probability of exiting the economy over a 10-year horizon. Similar to Panel A, the heat map features downward sloping indifference curves, and the gradients become even steeper as the selection effect amplifies overtime (refer to the legends for the z-axis located at the right of the plots). For most underperforming firms in the economy, their takeover exposure is high, and they are likely to exit the market over time via value-enhancing mergers; these firms' takeover likelihood decreases monotonically with their managers' entrenchment level, which serves as takeover deterrence and allows the firms to survive longer. Therefore, the selection effect produces a negative correlation between a firm's performance and its manager's private benefit of control among the existing firms in the economy.

Next, we take a snapshot of the industry equilibrium at a given point in time and calculate the profitability, private benefit of control, and other characteristics for different groups of firms:

[5] We defer the discussion of the parameter estimates to the next section.

The full sample covers all firms from the simulation, the target (acquirer) sample covers firms that choose to be targets (acquirers) in completed deals, and the bankrupt sample covers the firms that are shut down. Table 4 provides the summary statistics for the four groups of firms. The results show that the firms that outperform the average and whose managers have high private control benefits tend to become acquirers, while the firms with the opposite characteristics (slightly below-average performance and low managerial control benefit) tend to accept the bids and act as targets. There is a special group of firms who underperform substantially but are unlikely to get acquired because of their managers' high private control benefits. These firms tend to linger longer in the economy. They will eventually be shut down if, over time, they fail to recover by themselves or through profitable takeover deals.

2 Estimation

This section describes the data and sample construction, the simulated method of moments (SMM) estimator, and the intuition behind the estimation method.

2.1 Data and Sample Construction

We obtain firms' accounting data from CRSP-Compustat Merged Database (CCM), and our sample covers the period from 1980 to 2014. We measure a firm's performance using its operating income before depreciation scaled by the total book value of assets. We also collect information on other firm characteristics, such as firm size, market-to-book equity ratio, market leverage, and cash holdings. All variables are winsorized at the 1% level. Detailed variable definitions are summarized in Table 2.

Data on M&A come from Thomson Reuters SDC Platinum. We examine all bids announced between 1980 and 2014. To be included in the final sample, a bid has to satisfy the following criteria:

1. The announcement date falls between 1980 and 2014;
2. The targets are U.S. publicly traded firms;
3. The deal can be clearly classified as successfully completed or a failure, and the date of the bid completion or the bid withdrawal is available;
4. The acquirer seeks to acquire more than 50% of the target's shares in order to gain control of the firm and holds less than 50% of the target's shares beforehand;
5. The deal value exceeds one million dollars;
6. The deal is classified as a merger, not a tender offer or a block trade.[6]

Following Betton et al. (2008), we construct the sample of control contests based on the raw M&A data. Specifically, we say that a control contest begins with the first bid for a given target and continues until 126 trading days have passed without any additional offer. Each time an additional offer for the target is identified, the 126 trading day search window rolls forward. We take the winner of each control contest as the acquirer in the data, and for the control contests that have no winner, we take the bidder with the highest offer price as the acquirer. Over 95% of the control contests contain only one announcing bidder in our sample.

We define our main variables associated with M&A as follows. We measure the offer premium as the offer price per share divided by the target stock price 22 trading days before the bid announcement, minus one. The offer premium data provided by SDC include some large outliers. Following Bates and Lemmon (2003) and Officer (2003), we drop observations with offer premium lower than zero and cap the offer premium at two. We measure the acquirer, the target, and the combined firm announcement-period abnormal returns using the market model. To capture the price run-up caused by possible information leakage before the takeover

[6]We follow Betton et al. (2008) in classifying the deal type: If the tender flag is "no," and the deal form is a merger, then the deal is a merger. If the tender flag is "no," the deal form is "acquisition of majority interest," and the effective date of the deal equals the announcement date, then the deal is classified as a control block trade. If the tender flag is "yes," or if the tender flag is "no" and it is not a block trade, then the deal is a tender offer.

announcement, we also include the abnormal return earned within 22 trading days prior to the acquisition.

Our final sample includes a panel of 232,381 firm-year observations. Table 3 presents the summary statistics of the sample. Panel A compares the characteristics of all firms covered by the CCM database with those of the acquirers and targets. Consistent with the findings in previous studies, the acquirers are on average much larger than the targets in size; the acquirers also have higher market-to-book equity ratios and better operating performance. Consistent with Rhodes-Kropf and Robinson (2008), we also find in our sample that the targets are not bad firms in general: They are larger than an average Compustat firm, and their operating performance is comparable to that of the average firm. Panel B shows the deal characteristics of all mergers and acquisitions in our sample. The target pre-acquisition value averages 24% of the acquirer pre-acquisition value, with significant cross-sectional variation. Offer premiums are high and average 44% of the targets' pre-acquisition market value. A typical bid is composed of 34% cash and 66% equity, and a large fraction of the bids in the data are paid with full cash or full equity. The acquirer announcement-period abnormal returns, computed based on the market model, are on average -1%. The target firms earn much higher announcement-period abnormal returns, which average 29%. To account for possible information leakage before bid announcements, we include the 22-day price run-up in computing the announcement-period abnormal returns. The combined firm's announcement return is calculated as the value-weighted average of the acquirer's and the target's announcement return. The method of payment may reveal acquirers' misvaluation and lead the market to reprice the acquirers' standalone value on takeover announcements (i.e., the revelation effect). The revelation effect confounds the acquirer announcement return as a measure of the merger gains. To correct for the revelation effect, we follow Golubov et al. (2015) and make adjustment to the acquirer announcement returns based on the payment method and the deal and firm characteristics.

Firms enter or exit the CCM Database for various reasons during their life cycles. To track firm dynamics in the CCM database, we use the variable *Research Company Reason for Deletion* (i.e., *dlrsn*) provided by CCM. *dlrsn* is a categorical variable that indicates the reason why a firm loses CCM coverage. Common reasons for a firm to lose CCM coverage include being acquired (mergers and acquisitions or leveraged buyout), going bankrupt or being liquidated, going private, and failing to fit the original reporting format. We then merge *dlrsn* with the panel of firms in our sample; thus each firm in our sample is either still operating by the end of the sample period or exited the CCM at some point in time during the sample period with its *dlrsn* flag indicating the reason for the exit. In some rare cases, firms that drop out of the CCM have no associated *dlrsn* flag, and we label their exit reason "unknown."

With the *dlrsn* flag, we can track how a firm evolves over time, including its performance dynamics and possible exit. Specifically, we first sort each firm-year observation into performance quintiles, with the bottom quintile containing the most underperforming firm-year observations. We then compute the transition matrix of firm performance across different quintiles. Panel A of Table 4 presents the matrix, where several important findings emerge. First, we note that firm performance is persistent, and the probability that a firm stays in the same performance quintile in the next period ranges from 66% to 81%. Second, the probability of a firm being acquired does not correlate much with the firm's performance. Regardless of the firm's current period performance, there is about a 4.5% probability that the firm will be acquired. Underperforming firms are much more likely to go bankrupt or exit due to other reasons.

We then explore more details regarding underperforming firms. We analyze once a firm receives a negative shock and enters the bottom performance quintile, how long it stays there (i.e., the duration) and how it moves out of the bottom quintile. Panel B of Table 4 summarizes the distribution of the duration. Most firms do not stay long in the bottom performance quintile. The average duration is only 2.54 years, with a standard deviation of 2.60 years. About 25% of firms stay for less than two years and then either recover by themselves or exit the economy through acquisition or bankruptcy. The distribution, however, is highly skewed; 10% of the firms remain alive but keep underperforming for more than 10 years.

It is also interesting to examine the duration of underperformance broken down by the different reasons of exits. Panel C of Table 4 summarizes the results. The first row reports the probability that an underperforming firm eventually exits the bottom quintile through a specific channel. For example, there is about a 53% chance an underperforming firm eventually recovers by itself and transits into the upper quintile after an average of 2.27 years; there is about a 13% chance the firm is eventually acquired and exits the economy, and the average duration it stays in the bottom quintile before being acquired is 2.67 years. Overall, self-recovery and M&A appear to be the most common channels for a firm to exit the bottom performance quintile.

2.2 Estimator

We estimate the model using the simulated method of moments (SMM), which chooses parameter values that minimize the distance between the moments generated by the model and their analogs in the data. In this section, we present the data moments used in the estimation and explain how they help identify the model parameters. There are 12 parameters in the model: β is the subjective discount factor; μ_a and ϕ_α are the long-run mean and persistence of a representative firm's profitability, and υ_α captures the standard deviation of profitability innovations; μ_π and σ_π are the mean and standard deviation of the private control benefits among the population of managers, and $\rho_{\alpha\pi}$ captures the correlation coefficient between firm performance and managerial private control benefit among new entrants; μ_s and σ_s are the mean and standard deviation of unobservable synergy, s, and δ captures the relation between s and the acquirer-target performance differential; the probability that a firm meets with a target or an acquirer is ρ, the target manager's bargain power is θ, and the shareholders' reservation utility is U.

2.3 Identification and Selection of Moments

We start with the parameters that are standard in the literature. We set the annual subjective discount factor β to 0.9, a value commonly used in the literature. We choose the parameter θ=0.5 so that the target and the acquirer have intrinsically equal bargain power. Notice that, equal bargain power, however, does not translate into equal gains between the target and the acquirer because their managers' private control benefits can differ, which also influence how profits are shared between the two parties. We also assume that an entrant's profitability is uncorrelated with its manager's private benefit of control by setting $\rho_{\alpha\pi} = 0$ in our baseline estimation. We will relax these assumptions later in the robustness test. We have 10 parameters left to estimate.

Since we conduct a structural estimation, identification requires choosing moments whose predicted values are sensitive to the model's underlying parameters. We use 12 moments to identify the remaining 11 model parameters. Our identification strategy ensures that there is a unique parameter vector that makes the model fit the data as closely as possible.

Since we estimate all parameters in one big SMM system, we essentially allow each moment to respond to all model parameters, but we explain below which moments are the most important for identifying each parameter. When we map the model to the data, we compare the model implied moments, generated in the steady-state equilibrium, to the empirical data moments, implicitly assuming that the economy we observe is in steady state.[7] It is worth noting that the model does not directly specify the joint distribution of α and π in the steady state. Instead, we specify the distribution of α and π among the potential entrants. The joint distribution in the steady state is then determined jointly by the distribution of the entrants that we specify exogenously and the endogenous dynamics in the industry equilibrium through firm closure and M&A decisions. With a given set of model parameters, the joint distribution of α and π in the steady state can be calculated from the model solution and used to construct corresponding moments.

[7] This assumption is standard in previous studies that use models on heterogeneous agents.

The intuition of the parameter identification is as follows. The first three parameters, μ_a, ϕ_α, and υ_α, control the dynamics of firms' profitability process. We assume a firm's profitability follows an AR(1) process; thus, we identify ϕ_α from the auto-regressive coefficient, obtained from a panel regression that regresses a firm's profitability on its lagged value in steady state, υ_α, is estimated from the residual of the panel regression, and μ_a is identified from the average firm profit,

The private benefit of control, π, and the additional synergistic benefits or integration costs, s, are unobservable and can affect merger outcomes. Therefore, we infer their underlying distributions from various observed deal characteristics. First, notice that a higher value of μ_π implies higher average control benefits for both the acquirer and the target managers. In which case, the target manager demands a higher offer premium to sell the firm, and the acquirer manager is more willing to pay for the control rights, thus boosting the transaction price. Although a high value of μ_π does not change the real gains from the merger, it alters how such gains are split between the target and the acquirer. Following Ahern (2012), we measure the fraction of the merger gains captured by the target shareholders using the target's dollar value gain minus the acquirer's dollar value gain, scaled by the total pre-acquisition market value of the acquirer and the target.

The parameter σ_π controls the dispersion of private benefits across firm managers. If σ_π is zero, all managers have the same level of private benefit of control, and the sole driver for mergers will be the differences in operating efficiency. Therefore, all mergers will be value-enhancing. When σ_π increases, more mergers might be driven by the acquirer manager's high empire-building motives instead of the value-related reasons, making value-destroying mergers more likely. So we use the fraction of value-destroying mergers to identify off σ_π. We define a merger as value-destroying if the combined firm's announcement return, after adjusting for the revelation effect, is negative.

μ_s and σ_s determine the average level and dispersion of s, which captures the present value of additional gains and losses from mergers beyond the transfer of operating efficiency. A positive

s can be interpreted as an additional synergistic benefit coming from asset complementarity, the scale of economy, or enhanced market power, while a negative s should be understood as capturing the integration costs. Obviously, a large value of μ_s pushes up the total value created in mergers and acquisitions, and a larger σ_s suggests that the total return would be more volatile. Therefore, we use the mean and standard deviation of the combined target and acquirer returns to identify these two moments. Note that σ_s does not affect the fraction of value-destroying mergers, which is the main data moment we use to identify σ_π. This property allows us to pin down σ_π and σ_s separately.

Last, the conditional probability of low profitability firms acquiring high profitability firms helps to identify δ, the incidence of mergers pins down the matching probability, ρ, and the fraction of firms exiting the economy each period determines the shareholders' reservation utility $\underline{\mathrm{U}}$. We also include the takeover-performance sensitivity directly into our identification, as it is the key moment that our model mechanisms aim to explain.

3 Empirical Results

We begin by assessing how the model fits the data, and then we present the parameter estimates. We use the estimated model to analyze why underperforming firms are rarely acquired and quantify the welfare consequences associated with various features of the takeover market.

3.1 Model Fit

Table 5 compares empirical and model-implied moments. Panel A presents moments we target to match in the SMM estimation. The model is able to match most moments fairly closely—both M&A and firm closure (bankruptcy) are rare events in the model. The model predicts that only 4.1% of firms are acquired each year and that 1% of firms are closed down. Target firms are estimated to gain more from M&A than acquirers, even in dollar value. The model implies that target firms, on average, gain 4.9% more than acquirers in terms of the combined firm's value. Since targets are usually much smaller than acquirers and the merged entities, this number translates into a significant proportional gain to target shareholders. Some M&A are driven by managers' empire building incentives and, thus, might hurt shareholders' value. In the model, we find that about 16.8% of deals are value-destroying in the sense that the combined firm value decreases post-merger. The empirical analog is 19.6%, which is slightly higher but still economically close to what the model predicts. The model also fits very well with the average combined firm announcement return. This result, combined with the fit of the overall merger rate, ensures that our model captures the total value created by M&A in the economy.

Panel B of Table 5 shows how the model matches additional moments that are not targeted in estimation. The model-implied target announcement return averages 22.5%, and the model-implied acquirer announcement return is slightly negative, consistent with the patterns observed in the data. The model also comes close to matching the average market-to-book equity ratio and profitability for acquirers, targets, and all firms in the economy.

The next nine untargeted moments describe firm dynamics in the industry equilibrium. Specifically, they measure how long a firm stays underperforming (i.e., underperformance duration) and how an underperforming firm evolves over time. We define a firm as underperforming if its profitability falls into the bottom performance quintile and as exiting the underperforming group if it recovers by itself and moves to upper performance quintiles, is acquired by or merged with another entity, or goes bankrupt (chapter 11 or chapter 7). These moments are of particular importance in testing the model's predictions because they are driven by all model factors and are particularly sensitive to the relative importance of the takeover motives that the model is designed to capture. We do not explicitly target these moments in our estimation but leave them as the critical moments for validating the model post-estimation.

The model matches all nine untargeted moments well. Among all firms in the model simulation, approximately 68% of them exit the bottom performance quintile within three years, 22% of them stay underperforming for three to five years before exiting, and the remaining 10% underperform for more than five years. Therefore, the findings suggest that underperformance is usually not persistent because most underperforming firms either recover by themselves in a couple of years or are acquired and eliminated soon from the economy. However, there is also a non-trivial fraction of firms that have remained underperforming in the economy for a significant period of time.

The model also tracks how firms evolve over time once they become underperforming: over 79% of underperforming firms eventually recover by themselves, 15% are acquired, and the remaining 6% are closed down.[8] Hence, most underperformance is temporary, and firms are able to rescue themselves when they are hit by bad shocks. Meanwhile, the takeover market plays an important role in reallocating assets from the persistent underperformers to more efficient users in the economy.

3.2 Parameter Estimates

Table 6 reports the parameter estimates. The long run mean of firm profitability, μ_α, is 0.021, which is lower than the observed average firm profit in the economy. The takeover market contributes largely to this difference: most underperforming firms (except those with very high managerial control benefit) are acquired. As a result, firms with good performance are more likely to survive, while firms with bad performance are more likely to exit. This effect trims the lower tail of firms' profitability distribution, leading to a higher average profit in the steady state economy. Firm performance is estimated to be persistent, with an annual autoregressive coefficient of 0.809. Shocks to firm profitability have a large standard deviation

[8] They represent the cumulative probabilities for an underperforming firm to recover by itself, be acquired, or close down during its underperformance period. Therefore, they are different from the probabilities reported in Table 5, which represent the annual probabilities for an average firm to be acquired or shut down. The empirical values of the probabilities for an underperforming firm to exit through different channels, reported in Panel B of Table 5, are also slightly different from those reported in Panel C of Table 4. The main reason for this difference is that the model does not feature exits classified as "Other Reasons" in the data. To make a fair comparison between the model and data, we remove the category of "Other Reasons" and redistribute its probability mass proportionally to other categories of exits.

of about 0.1. These parameter estimates suggest that firm profitability features a persistent component and a quite volatile transitory component, which is consistent with prior studies (Cooper and Haltinwanger, 2006; Whited and Warusawitharana, 2016).

In the model, we specify the private benefit of control as a flow variable which an incumbent manager enjoys every period when she runs the firm. For an average entrant, its manager's private benefit of control is estimated to be 0.0033 per unit of capital, which is equal to $8\% = \frac{0.0033}{0.041}$ of the profit generated from each unit of capital in the steady state. In other words, our estimate implies that the disutility an average manager suffers from losing her control of the firm is as high as having her see the firm's profitability drop by 8% permanently. The estimated private benefit of control varies dramatically across firm managers, with a standard deviation of 0.032. It is worth noting that the estimates of private benefit, reported above, are for entrants, which can be quite different from that for firm in the steady state. As we show in the next section, managers' private benefit of control in the steady state is correlated with firm performance due to M&A and firm closure, even though the correlation is zero among entrants, as we assume in the estimation.

The last three parameters characterize s (i.e., the additional gains or integration costs generated in M&A beyond the transfer of efficiency). The estimated average s is -11%, which suggests that merging two firms randomly is expected to destroy 11% of the target firm value. The value of s also depends on the acquirer-target performance differential, as captured by the parameter δ. δ is estimated to be negative, suggesting that it is more costly to integrate two firms with very different performance. This finding seems reasonable: On the one hand, it takes more efforts for the acquirer managers to integrate target operation and improve target asset productivity if the target underperforms substantially pre-merger. This cost can therefore be interpreted as an integration cost in combining the two firms. On the other hand, as Rhodes-Kropf and Robinson (2008) suggest, pairing up firms with similar performance or valuation may create more value through asset complementarity. A negative δ in the model can generate the "like-buys-like" effect in Rhodes-Kropf and Robinson (2008).

3.3 Takeover-performance Sensitivity

Using the estimated model as a laboratory, we first investigate how a weak takeover-performance sensitivity emerges in the economy. We then explore the economic implications of a weak takeover-performance sensitivity and link it to the takeover market efficiency.

Consistent with previous studies, we find in the data that a firm's operating performance has very weak, if any, prediction power on its subsequent takeover exposures: underperforming firms are not acquired more frequently than other firms. Our estimated model is able to replicate this empirical pattern in the simulated data: Even though the Q-theory of M&A is a pivotal motive in the model, the model-implied takeover-performance sensitivity is almost zero. This result is striking, but dissecting the model dynamics reveals the underlying mechanism.

First note that previous studies calculate the empirical measure of takeover-performance sensitivity by pooling all firm-year observations in the economy. As its analog, we compute the model-implied takeover-performance sensitivity based on the cross-section of firms in the steady state. The characteristics of firms in the steady state, therefore, are critical in understanding the takeover-performance sensitivity. In the model, firm characteristics in the steady state are largely driven by the M&A market: Outperforming firms are less likely to be targeted regardless of their managers' private benefit of control; whereas, underperforming firms with less entrenched managers are more likely to be acquired than those with highly entrenched managers. The M&A market, therefore, induces a selection effect on firms' performance and their managers' private benefit of control. This selection effect is particularly strong for underperforming firms and generates a negative correlation between firms' performance and managerial control benefit among the surviving firms.

The selection effect, discussed above, exists in each period of a model economy, and as the economy evolves to reach the steady state, this effect is accumulated and amplified. Hence, even if the selection effect is moderate within each period, the accumulated effect on firms can, nevertheless, be substantial in the steady state. Consider a thought experiment in which 100 firms receive negative shocks and become underperforming each period, and only ten of them have very entrenched managers who are highly resistant to takeovers. If we temporarily leave firm closure and self-recovery out of this thought experiment, then 90 of the 100 firms will be acquired sooner or later, while the remaining 10 firms, influenced by their managers' private control benefit, will continue to underperfom in the economy. As the economy evolves, the firms with high-entrenchment managers accumulate and eventually represent a very large fraction of the underperforming group—after N periods, then there will be in total $100+10\times N$ underperforming firms in the economy: 100 new underperformers where 10% of them have entrenched managers, plus $10 \times N$ underperformers inherited from previous periods, all of which have managers with high control benefits. Such high control benefits act as a barrier for future takeover bids and make these underperforming firms hard to acquire. Adding back firm closure and self-recovery allows the model economy to reach a steady state but does not change the intuition of this thought experiment. Overall, underperforming firms in the steady state are overrepresented by firms with excessively high private benefit of control, which in turn makes them more difficult to be acquired. As a result, this model delivers the striking implication that, even though a relatively small fraction of merger decisions are either induced or blocked by managers' control benefit, the effect can significantly weaken the takeover-performance sensitivity in the steady state.

Next, we explore the above implication quantitatively using the estimated model. We carry out an experiment in which we simulate a cross-section of firms and add them into the steady state model economy.[9] We assume that this group of experimental firms accounts for a negligible fraction of the model economy and, therefore, do not affect the steady state. These experimental firms are randomly drawn, with their operating efficiency and private control benefit from the steady state marginal distribution of α and π, respectively. In this way, we ensure that these firms have the same first and second moments for operating efficiency and private control benefit as the other firms in the steady state equilibrium. At the same

[9]Note that, different from the thought experiment we present above, this quantitative exercise strictly follows the estimated model and takes into account firm closure and self-recovery in the simulated path.

time, because α and π are drawn independently from their marginal distributions, we ensure that they are uncorrelated among the experimental firms. We then start the simulation and track the evolution of the experimental firms. We are particularly interested in the following three questions. First, how does the average private benefit of control evolve for the subset of underperforming firms? Second, how does the probability of being acquired among these underperformers change over the simulation path? Third, how does the takeover-performance sensitivity evolve for the full spectrum of experimental firms? We measure takeover-performance sensitivity as the correlation between a firm's performance in a given year and a dummy variable of it being acquired the next year.

Figure 3 presents the results. Panel A plots the evolution of the three variables we are interested in. In order to plot them in the same figure, we normalize their values in the steady state to be one. The dotted blue line depicts the evolution of the average private benefit of control for underperforming experimental firms. Due to the selection effect, underperforming firms with low private benefit are acquired over time, leaving in the economy those with high private benefit. In the steady state, the average private benefit of control for underperforming firms is six times as large as that at the beginning of this simulation when the selection effect is absent.

Responding to the significant increase in the private benefit of control among underperforming firms, the probability of them being acquired drops quickly. The solid red line in Figure 3 shows that the M&A rate for underperforming firms is almost halved five years down the simulation path. The large drop in the M&A rate for underperforming firms leads to a sharp decline in the takeover-performance sensitivity calculated for the full spectrum of experimental firms, as shown by the dashed black line in Panel A of Figure 3.

Panel B of Figure 3 compares the distribution of underperforming firms' private benefit of control in different stages of the simulation path. The dashed blue line presents the distribution when simulation starts, and the solid red line depicts the distribution in the steady state. We find that the distribution of underperforming firms' private benefit of control migrates towards the right as the group of experimental firms evolves toward the steady state. Notice that the steady state firms have the same average private control benefit as the experimental firms at the beginning of the simulation. Therefore, the results presented in Panel B are driven purely by the negative correlation between α and π that arises endogenously in the steady state due to selection.

According to the Q-theory of M&A, takeover-performance sensitivity is driven by both cross-sectional and time-series correlations between firm performance and takeover exposure. Cross-sectionally, bad firms are more likely to be targeted; over time, firms are more vulnerable to takeover threats when their performance deteriorates. A pooling linear or logit regression does not distinguish the sources of the correlations. The selection effect analyzed in our model mainly attenuates the cross-sectional correlation, but it has little effect on the time-series correlation. Therefore, the model predicts that takeover-performance sensitivity should be more pronounced if we focus on identifying the time-series results.

To test this prediction, we compare results from three regression specifications with different fixed effects. We use the linear probability model because results with fixed effects are more difficult to interpret in logit models. In each regression specification, we pool together all firm-year observations and regress a firm's dummy variable of receiving a bid during the current year on its prior year's performance. In each regression, we control a set of common variables that are documented to affect firm takeovers. The baseline specification follows most previous studies and controls for no fixed effect; the time-series specification controls for firm fixed effect and, thus, identifies the correlation mainly from the time series comovement between firm performance and takeover exposure; and the cross-sectional specification controls for year fixed effect and identifies the correlation based on the cross-sectional variation.

The loading on firm performance is -0.003 in the baseline specification, with a t-stat of -1.4. This result is consistent with previous findings that the takeover-performance sensitivity is slightly negative but insignificant in the pooled regression without any fixed effect. When we control for firm fixed effect in the time-series specification, the loading on firm performance becomes -0.029, with a t-stat of -9.1. This large and highly significant coefficient lends strong support to the model's prediction that takeover-performance sensitivity is more pronounced when we focus on identifying the time-series correlation. On the contrary, when we control for the year fixed effect in the cross-sectional specification, the loading is reduced to -0.001, with a t-stat of -0.66, which is even weaker than that in the baseline specification. These results are consistent with our prediction that the dynamic selection effect mainly attenuates the cross-sectional correlation between firm performance and takeover exposure.

Though controlling the firm fixed effect partially recovers the takeover-performance sensitivity concealed by the selection effect, it is not a substitute to the structural approach we pursue in the paper. First, it only captures the time-series correlation, which contributes to a relatively small fraction of the true takeover-performance sensitivity.[10] Second, it is silent on the underlying mechanism and does not show how the takeover-performance sensitivity arises endogenously in the industry equilibrium though continuous selection. Third, it cannot be used for analyzing the takeover market efficiency because no counterfactual could be done to investigate how firms' M&A activities would change if certain characteristics of the economy were altered or completely eliminated.

3.4 Efficiency of the Takeover Market

So far, through the model, we have demonstrated that the takeover market induces a selection effect on managers' private benefit of control, which in turn, gives birth to a weak takeover-performance sensitivity. The model suggests that even a moderate selection effect within each period can accumulate over time and become amplified in the steady state. Next, we quantify the magnitude of the selection effect based on managers' private benefit, and more importantly, we use this result to assess the efficiency of the takeover market.

The efficiency that we are trying to capture in this paper is only regarding firms' asset reallocation. According to the q-theory of mergers, the takeovers market can generate efficiency

[10] Because the panel contains a much larger number of firms cross-sectionally than the number of time periods.

gains by reallocating assets from sub-optimal users (the persistently underperforming firms) towards users with more advanced technologies and better management. Our model is built on the q-theory of mergers but with an added important friction— managers' private control benefit. Inefficiency can arise in our model when a value-enhancing takeover is blocked due to the target manager's excessive private benefit of control or a value-destroying takeover is consummated to serve the acquirer manager's empire-building motive. Note that the takeover market can perform other important functions than just reallocating assets—for example, the takeover market can serve as an implicit threat to managers so that they always behave and run their corporations in the way that maximizes shareholders' value. Our measure for takeover market efficiency does not take into account such an ex-ante disciplinary role. If the takeover market only functions through implicit threat but not explicit transactions, then our measure will conclude that the takeover market is not an efficient channel for reallocating assets and control rights across firms. Therefore, our results should always be interpreted as setting a lower boundary for the gains of the takeover market.

To assess the (in)efficiency of the takeover market, we first perform a counterfactual analysis in which we eliminate the private benefit of control for all managers in the estimated model. Specifically, we take the estimated model and set the mean and standard deviation of π to be zero. Doing so counterfactually assumes that managers of all entrants enjoy no private benefit of control and that their interests are perfectly aligned with shareholders. We then compute firms' average valuation ratio (i.e., market-to-book equity ratio) in this counterfactual economy and compare it with the one in our baseline model. We find that completely eliminating the private benefit of control from the economy increases firms' average valuation ratio by 1%. To better assess the magnitude of this result, we compare it with another counterfactual benchmark in which we completely shut down the takeover market. When the takeover market is eliminated from the economy, the average valuation ratio declines by 6.4%, indicating that, for an average firm, 6.4% of its market value can be attributed to the existence of an active takeover market. Combining findings from the two counterfactual analyses, we conclude that managers' private benefit of control leads to small inefficiencies in the takeover market and destroys about $13\% = \frac{1\%}{1\%+6.4\%}$ of the total value created in M&A.

Why does the private benefit of control significantly attenuate the takeover-performance sensitivity but contribute little to the takeover market inefficiency? Our model helps reconcile these seemingly contradictory findings: Overall, managers' private benefit of control does not play a significant role in driving M&A. Each period, among the firms that receive negative shocks and become underperforming, only a small fraction of them possess high private benefit of control and manage to block value-creating takeovers. Hence, the private benefit of control does not hurt takeover market efficiency to a large extent, and the majority of value-enhancing deals still consummate each period on a *dynamic* basis. These underperforming firms with highly entrenched managers, however, accumulate over time and represent a disproportionally large fraction of underperforming firms in the steady state. They significantly attenuate the takeover-performance sensitivity, which takes a snapshot of the steady state and measure the relation on a *static* basis. A weak takeover-performance sensitivity, therefore, does not necessarily suggest that the takeover market is inefficient.

4 Conclusions

Despite its prevalence in theoretical studies, the q-theory of mergers receives weak empirical support. Only a very small fraction of underperforming firms is acquired every year, and the association between a firm's performance and its subsequent takeover exposure (i.e., takeover-performance sensitivity) is, at best, weakly negative. In this paper, we revisit this problem by estimating a dynamic model in which takeovers are pursued either to boost firm performance or create control benefit for managers. Our estimation suggests that the takeover market is overall efficient. It promptly reallocates assets from most underperforming firms towards more efficient users and adds significant value to the economy. Meanwhile, an efficient takeover market triggers an ongoing selection effect so that underperforming firms with more entrenched managers survive longer. Even if this selection effect is small in each period, it accumulates over time and is amplified as the economy evolves. As a result, underperforming firms become overrepresented by managers with high control benefit, which deters future takeover bids and leads to the weak takeover-performance sensitivity in the data.

References

Agrawal, A. and J. F. Jaffe: 2003, 'Do takeover targets underperform? Evidence from operating and stock returns'. *Journal of Financial and Quantitative Analysis* **38**(04), 721–746.

Ahern, K. R.: 2012, 'Bargaining power and industry dependence in mergers'. *Journal of Financial Economics* **103**(3), 530 – 550.

Albuquerque, R. and E. Schroth: 2015, 'The value of control and the costs of illiquidity'. *The Journal of Finance* **70**(4), 1405–1455.

Bates, T. W. and M. L. Lemmon: 2003, 'Breaking up is hard to do? An analysis of termination fee provisions and merger outcomes'. *Journal of Financial Economics* **69**(3), 469–504.

Betton, S., B. E. Eckbo, and K. S. Thorburn: 2008, 'Corporate takeovers'. *Elsevier/North-Holland Handbook of Finance Series*.

Golubov, A., D. Petmezas, and N. G. Travlos: 2015, 'Do Stock-Financed Acquisitions Destroy Value? New Methods and Evidence'. *Review of Finance* p. rfv009.

Officer, M. S.: 2003, 'Termination fees in mergers and acquisitions'. *Journal of Financial economics* **69**(3), 431–467.

Rhodes-Kropf, M. and D. T. Robinson: 2008, 'The market for mergers and the boundaries of the firm'. *The Journal of Finance* **63**(3), 1169–1211.

Strebulaev, I. A. and T. M. Whited: 2011, 'Dynamic models and structural estimation in corporate finance'. *Foundations and Trends in Finance* **6**(1-2).

Taylor, L. A.: 2010, 'Why are CEOs rarely fired? Evidence from structural estimation'. *The Journal of Finance* **65**(6), 2051–2087.

Figure 1: Timeline of the Model

This figure illustrates the timeline of the model. At the beginning of period t, firm I's profitability is realized and observed by agents in the model. The firm's manager then observes the takeover opportunities: the firm either meets a potential acquirer or a potential target. The managers of the two matched firms then decide whether to consummate the deal and if so, the transaction price the acquirer should pay the target. If the deal is consummated, target shareholders obtain a terminal payoff equal to the transaction price, target assets are transferred to the acquirer's control, and the target manager loses her position. If the deal fails, both firms remain standalone. After the takeover decisions are made, shareholders decide on whether to close down the firm by the end of period t.

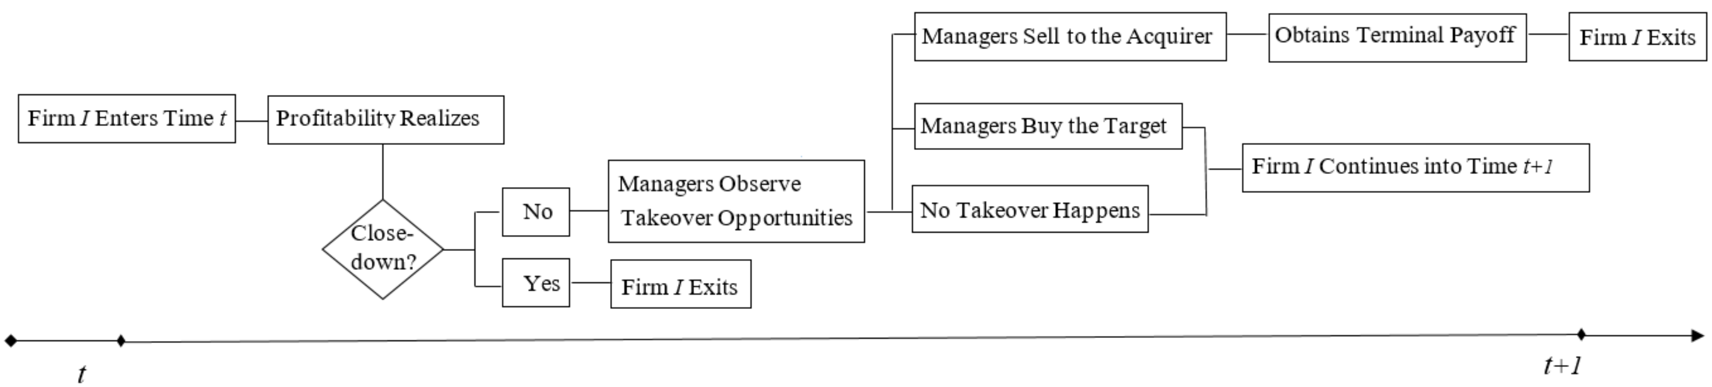

Figure 2: Takeover Likelihood and Firm Characteristics

Panel A corresponds to a 3 dimensional heat map with firms' begin-of-period profitability on the x-axis and their managers' private benefit of control on the y-axis. The legend for the z-axis heat map is located at right of the plot and represents the probability that a firm gets acquired. Panel B is constructed in a way similar to Panel A, except that the z-axis now measures the cumulative probability that a firm exits an economy over a ten year horizon.

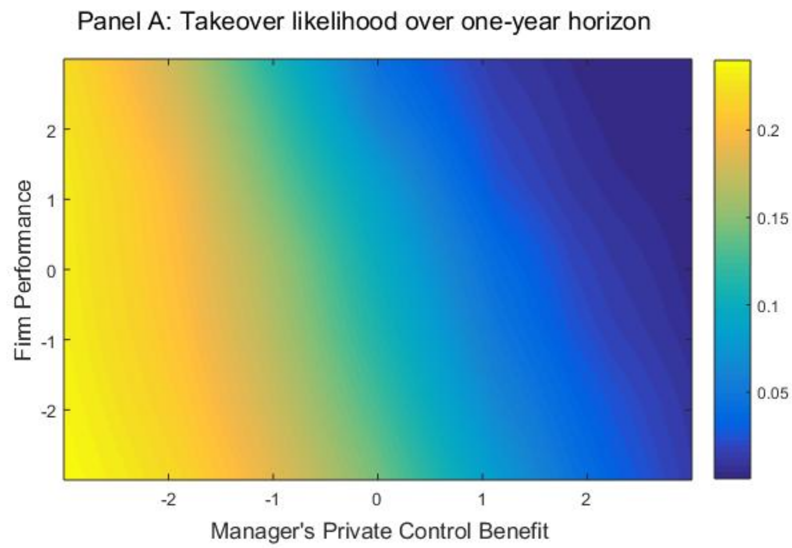

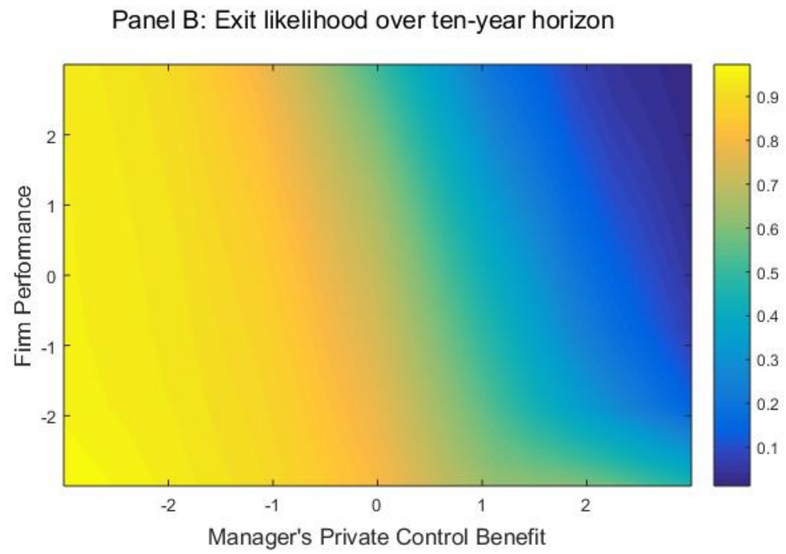

Figure 3: Evolution of Firm Characteristics

This figure plots the evolution of firm characteristics in model simulation. Panel A presents the the M&A rate and the average private benefit of control for the underperforming firms as well as the takeover-performance sensitivity computed as a correlation between a firm's performance and its subsequent takeover exposure in the full sample. Panel B compares the distribution of private benefit of control among underperforming firms at the first round of simulation and at the steady state.

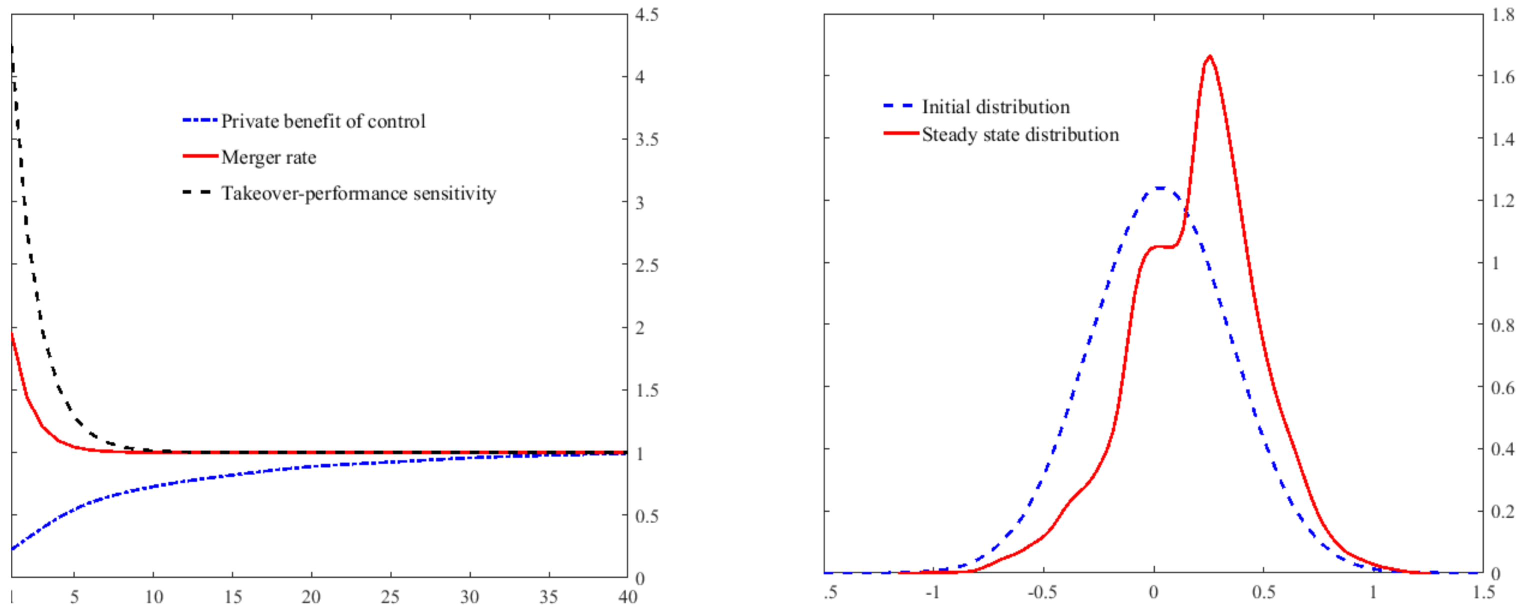

Table 1: Model Implied Firm Characteristics

This table presents the model implications on characteristics of different firms. The model is solved with parameters set to their estimated values. We simulate a panel of firms and their takeover decisions from the estimated model and compute the profitability, private benefit of control, and the unobservable synergy net of integration cost for different groups of firms. Full sample covers all firms in our simulation, target (acquirer) sample covers firms that choose to be targets (acquirers) in the takeover market, and bankrupt sample covers firms that are closed down.

	Full Sample	Target	Acquirer	Bankrupt
Average market to book ratio	1.948	1.810	2.047	0.766
Std. of market to book ratio	0.532	0.469	0.529	0.122
Average profitability	0.041	0.019	0.098	-0.175
Std. of profitability	0.162	0.137	0.145	0.066
Average private benefit of control	0.029	0.007	0.059	0.055
Std. of private benefit of control	0.048	0.042	0.036	0.093
Average additional synergy/integration cost	-0.110	-0.053	-0.053	-0.123
Std. of additional synergy/integration cost	0.092	0.090	0.090	0.087

Table 2: Variable Definition

This table defines the construction of variables. All variables in lower case letters are obtained directly from CRSP-Compustat Merged Database with the same variable name and all variables in upper case letters are variables we define.$PSLV$ takes the value of $pstkl$, $pstkrv$, or $pstk$ in sequence if the corresponding variable is not missing and $PSLV$ is set to zero if all three variables are missing. ME is defined as $prcc\_c * cshpri$; BE is defined as $seq + txditc - PSLV$ and we set BE as missing if it is negative. CAR is the cumulative abnormal return computed using the market model; $TranVal$ is the transaction value the acquirer proposes to pay to the target shareholders; $CashVal$ is the total value of cash in the bid; $ME_{Tar,-22}$ is the target firm's market equity value measured 22 tradings before the date of bid announcement. Payment method may reveal acquirer misvaluation and leads the market to reprice the acquirer's standalone value (i.e., the revelation effect). This revelation effect confounds the acquirer announcement period return as a proper measure of merger gains. We follow Golubov et al. (2015) and make adjustment to the acquirer's announcement-period abnormal return to remove the revelation effect. The adjusted acquirer announcement return is $AcqRet\_adj$. The combined firm announcement return is computed from the acquirer announcement return and hence is adjusted accordingly, which is denoted as $CombRet\_adj$.

Notation	Variable	Definition
α	operating performance	$\frac{oibdp}{at}$
$Size$	logarithm of market equity	$ln\,(ME)$
$MktLv$	market leverage	$\frac{dltt+dlc}{dltt+dlc+ME+PSLV-txditc}$
MB	market-to-book equity ratio	$\frac{ME}{BE}$
$Cash$	cash holdings	$\frac{che}{at}$
$TarRltSz$	target market equity to acquirer market equity	$\frac{TarME}{AcqME}$
$OfferPrem$	offer premium	$\frac{TranVal}{ME_{Acq,-22}} - 1$
$CashFrac$	fraction of cash payment in the deal	$\frac{CashVal}{TranVal}$
$AcqRet$	acquirer abnormal return	acquirer $3-day\,CAR$ plus 22-day run-up, adjusted
$TarRet$	target abnormal return	target $3-day\,CAR$ plus 22-day runup, adjusted
$CombRet$	the combined firm abnormal return	the combined firm $3-day\,CAR$ plus 22-day runup

Table 3: Summary Statistics

This table presents the summary statistics of firm and deal characteristics. Panel A presents firm characteristics for all firms in the sample (with subscript *All*), acquirers (with subscript *Acq*), and targets (with subscript Tar). Panel B presents deal characteristics in our sample. Variable definitions are summarized in Table 2.

Panel A. Firm Characteristics

Variable	Mean	Stdev.	p5	p25	p50	p75	p95
$Size_{All}$	4.96	2.27	1.35	3.35	4.85	6.48	8.98
$MktLv_{All}$	0.28	0.26	0.00	0.03	0.21	0.46	0.80
MB_{All}	3.30	5.97	0.44	0.98	1.64	3.05	10.71
$Cash_{All}$	0.17	0.21	0.00	0.02	0.08	0.23	0.67
α_{All}	0.04	0.23	-0.39	0.02	0.09	0.15	0.27
$Size_{Acq}$	7.37	1.90	4.19	6.00	7.39	8.76	10.65
$MktLv_{Acq}$	0.27	0.23	0.00	0.07	0.22	0.46	0.70
MB_{Acq}	3.37	4.29	0.85	1.47	2.22	3.71	8.88
$Cash_{Acq}$	0.14	0.17	0.01	0.03	0.06	0.17	05.4
α_{Acq}	0.10	0.10	0.01	0.03	0.09	0.17	0.28
$Size_{Tar}$	5.09	1.76	2.50	3.85	4.95	6.28	8.22
$MktLv_{Tar}$	0.28	0.26	0.00	0.02	0.23	0.49	0.75
MB_{Tar}	2.72	4.29	0.61	1.12	1.72	2.69	7.56
$Cash_{Tar}$	0.16	0.21	0.00	0.02	0.06	0.23	0.68
α_{Tar}	0.06	0.16	-0.17	0.02	0.05	0.13	0.24

Panel B. Deal Characteristics

Variable	Mean	Stdev.	p5	p25	p50	p75	p95
$TarRltSz$	0.24	0.66	0.01	0.04	0.12	0.30	0.77
$OfferPrem$	0.44	0.33	0.06	0.23	0.36	0.57	1.13
$CashFrac$	0.34	0.41	0.00	0.00	0.10	0.73	1.00
$AcqRet$	-0.01	0.13	-0.22	-0.07	-0.01	0.05	0.20
$TarRet$	0.29	0.26	-0.03	0.13	0.25	0.41	0.78
$CombRet$	0.03	0.12	-0.17	-0.03	0.02	0.09	0.24

Table 4: Duration of Under performance and Transition Matrix

This table presents how firms transit among different performance groups and how long underperforming firms survive before they become self-corrected, go bankrupt, or are acquired. Panel A presents the transition matrix of firm performance, and Panel B presents the distribution of survival duration.

Panel A. Transition Matrix

	Performance Transition					Exit		
	1 (Low)	2	3	4	5 (High)	Acquired	Bankruptcy	Other Reasons
1 (Low)	0.759	0.077	0.015	0.007	0.003	0.043	0.018	0.078
2	0.093	0.706	0.115	0.015	0.003	0.045	0.005	0.019
3	0.022	0.136	0.660	0.114	0.007	0.043	0.004	0.013
4	0.009	0.023	0.140	0.679	0.090	0.045	0.002	0.012
5 (High)	0.003	0.003	0.010	0.118	0.811	0.046	0.002	0.008

Panel B. Distribution of Underperformance Duration (in years)

Mean	SD	Percentile						
		1	10	25	50	75	90	99
2.54	2.6	1.0	1.0	2.0	3.0	4.0	5.0	13.0

Panel C. Average Duration of Underperformance, Breakdown by Exit Reasons

	Performance Transition					Exit		
	1 (Low)	2	3	4	5 (High)	Acquired	Bankruptcy	Other Reasons
Prob. (%)	8.08	47.62	3.97	1.38	0.02	12.81	5.33	21.49
Duration (in years)	2.55	2.29	2.80	2.03	2.01	2.67	2.64	2.74

Table 5: Model Fit

This table assesses model fit. Panel A shows how well the model fits 12 targeted moments (i.e., moments used in SMM): Mean of firm performance is the average cross-sectional firm profitability; persistence of performance and variance of performance residual are estimated coefficients of equation 1; probability of being acquired is the total number of firms acquired divided by the total number of firms in the economy; takeover-performance sensitivity is the coefficient obtained by regressing a firm's takeover exposure dummy on its previous year's performance; probability of firm closure is the probability of bankruptcy (chapter 11 and chapter 7); mean and variance of performance for entrants are the average and variance of entrants' profitability; target's fraction of merger gains are the target's dollar value gain/loss minus the acquirer's dollar value gain/loss from the M&A deal scaled by the combined firm's market value pre-merger; probability of value-destroying merger is the fraction of M&As that have negative combined firm announcement returns after adjusting for the revelation effect; mean of the combined firm ann. ret. is the average announcement return for the combined firms; prob. of low-profitability firms acquiring high-profitability firms is the fraction of M&As in which the acquirer's profitability is lower than the target's profitability. Panel B shows how the model fits 18 untargeted moments (i.e., moments not used in SMM).

Panel A. Targeted Moments

Moment	Data		Model
	Empirical value	Standard error	Simulated value
Mean of firm performance	0.041	0.003	0.041
Persistence of firm performance	0.818	0.006	0.818
Variance of performance residual	0.016	0.001	0.016
Probability of being acquired	0.044	0.002	0.043
Takeover-performance sensitivity	-0.010	0.029	-0.016
Probability of firm closure	0.006	0.001	0.006
Target's fraction of merger gains	0.042	0.003	0.049
Prob. of value-destroying mergers	0.196	0.017	0.168
Mean of the combined firm ann. ret.	0.028	0.004	0.025
Std. of the combined firm ann. ret.	0.013	0.004	0.013
Prob. of low-profitability firms acquiring high-profitability firms	0.345	0.014	0.330

Panel B shows how the model fits 18 untargeted moments (i.e., moments not used in SMM). Mean of target (acquirer) ann. ret. is the average announcement returns for targets (acquirers); mean M/B for all firms (acquirers or targets) is the average market-to-book equity ratio for all firms (acquirers or targets); mean profitability for all firms (acquirers or targets) is the average profitability for all firms (acquirers or targets); distribution of underperformance duration describes how long a firm remains underperforming (i.e., stay in the bottom performance quintile) before it exits underperformance; Channels of exiting underperformance describe the fraction of underperforming firms that exit the bottom performance quintile through different channels including self-recovery, being acquired, or going bankruptcy; mean underperformance duration breakdown by exit channels summarizes the average length a firm remains underperforming before it exits through a specific channel.

Panel B. Untargeted Moments		
Moment	Data	Model
	Empirical value	Simulated value
Mean of target ann. ret.	0.293	0.225
Mean of acquirer ann. ret.	-0.009	-0.014
Mean M/B for all firms	2.156	1.942
Mean M/B for acquirers	2.608	2.160
Mean M/B for targets	2.097	1.750
Mean profitability for all firms	0.041	0.041
Mean profitability for acquirers	0.115	0.118
Mean profitability for targets	0.068	0.025
Distribution of underperformance duration		
Less than 3 years	0.625	0.677
3 to 5 years	0.275	0.212
More than 5 years	0.100	0.101
Channels of exiting underperformance		
Self-correction	0.775	0.792
Being acquired	0.159	0.144
Closure (bankruptcy)	0.066	0.064
Mean underperformance duration, breakdown by exit channels		
Self-correction	2.27 year	2.66 year
Being acquired	2.67 year	2.71 year
Closure (bankruptcy)	2.64 year	2.74 year

Table 6: Parameter Estimates

This table reports parameter estimates for the baseline model using the simulated method of moments. ϕ_α is the mean-reverting speed of profitability, $\bar{\alpha}$ is the average long-run mean of profitability, υ_α is the standard deviation of shocks to profitability, μ_α,μ_π,σ_α, and σ_π determine the joint distribution of α and π for entrants, μ_s and σ_s are the average and standard deviation of the unobservable synergy or integration cost s, and δ measures the relation between s and the acquirer-target performance differential.

	ϕ_α	$\bar{\alpha}$	υ_α	μ_α	σ_α	μ_π	σ_π	μ_s	σ_s	δ
Estimate	0.809	0.008	0.099	0.021	0.247	0.003	0.032	-0.110	0.092	-0.992
S.E.	0.025	0.002	0.027	0.007	0.081	0.001	0.007	0.043	0.023	0.246