

Agency, Firm Growth, and Managerial Turnover*

Ronald W. Anderson[†], M. Cecilia Bustamante[‡] and Stéphane Guibaud[§]

October 8, 2012

Abstract

We study managerial incentive provision under moral hazard in a firm subject to stochastic growth opportunities. In our model, managers are dismissed after poor performance, but also when an alternative manager is more capable of growing the firm. The optimal contract may involve managerial entrenchment, such that growth opportunities are foregone after good performance. Firms with better growth prospects have higher managerial turnover and more front-loaded compensation. Firms may pay severance to incentivize their managers to report truthfully the arrival of growth opportunities. By ignoring the externality of the dismissal policy onto future managers, the optimal contract implies excessive retention.

1 Introduction

Firms derive value not only from operating their assets in place, but also from their ability to exploit growth opportunities. The latter source of value creation often involves major transformations of the firm— such as adopting innovative production techniques, changing the organization of labor, developing new products, or venturing into new markets. However incumbent managers may lack the vision or skills necessary to implement these transformations and lead their firm through a new, more profitable phase. In such circumstances, a change of management is needed for the firm to successfully pursue its growth opportunities.

This paper introduces *growth-induced* managerial turnover in a dynamic moral hazard model of the firm. As in previous studies on optimal long-term managerial contracts with limited liability, firms use early termination as a disciplinary device— and managers are replaced following poor performance. But in contrast with previous studies, firms may also fire their managers despite good performance when a change of management is the best or only option to seize a growth

*We are grateful to Bruno Biais, Catherine Casamatta, Peter DeMarzo, Mike Fishman, Denis Gromb, Zhiguo He, Andrey Malenko, our discussants David Dicks and Oliver Spalt, seminar participants at LSE, Toulouse and UT Austin, and conference participants at the meetings of the CEPR European Summer Symposium in Financial Markets (Gerzensee), European Finance Association, European Finance Summit, and Western Finance Association for helpful comments. All responsibility for errors and views expressed is our own.

[†]London School of Economics, r.w.anderson@lse.ac.uk

[‡]London School of Economics, m.bustamante@lse.ac.uk

[§]London School of Economics, s.guibaud@lse.ac.uk

opportunity. We characterize how growth-induced turnover affects the provision of managerial incentives and analyze the determinants of realized firm growth.

In our model, a risk-neutral manager is hired by a risk neutral, long-lived firm. Cash flows are only observable by the manager, who can divert them for his own private benefit. The firm can fire its current manager at any point in time, and replace him at a cost. Take-it-or-leave-it growth opportunities may arrive stochastically in any period. In our baseline model, the arrival of a growth opportunity is contractible and the firm needs to hire a new manager in order to pursue such an opportunity. If the firm decides to take up a growth opportunity, it pays the costs associated with replacing the manager and transforming the firm— possibly but not necessarily including capital investment— and its profitability increases.

We solve for the optimal long-term contracts signed between the firm and its successive managers at the time they are hired. A manager’s expected discounted payoff under the optimal contract (often referred to as the manager’s *promise*) evolves over time in response both to cashflow and growth opportunity realizations. Replacement and compensation decisions are determined by endogenous performance thresholds. A key feature of our analysis is that the continuation value of the firm upon replacing a manager is endogenous (equal to the value of the firm under the newly hired manager), and contingent on the availability of a growth opportunity. This contrasts with most of the existing dynamic contracting models where, upon firing the manager, the firm obtains an exogenously given liquidation value.

Our results in the baseline model are as follows. First, we characterize turnover. The probability of replacing an incumbent manager depends not only on past and current performance, as summarized by the manager’s promise, but also on the availability of a growth opportunity. The conditional probability of managerial turnover is always higher in states of the world where a growth opportunity is available. However a firm may in some circumstances decide to retain its incumbent manager despite the arrival of a growth opportunity. Firms facing only modest growth opportunities or those plagued with severe agency problems optimally choose to do so when a growth opportunity arises after a period of good performance. The optimal contract can therefore involve a particular form of managerial entrenchment. When it occurs, forsaken growth adds to the usual inefficiency that, for *ex ante* incentive provision, managers can be fired upon poor past performance in the absence of a growth opportunity.

Second, we characterize the optimal compensation scheme, and determine how growth opportunities affect managerial compensation. We find that the optimal managerial contract is readily implementable by a system of deferred compensation credit and bonuses. Deferred compensation is used, along with the threat of inefficient replacement, in order to provide incentives in the best possible way. We show that the degree to which firms rely on back-loading of compensation is affected by their growth prospects. Namely, the extent of back-loading decreases with the quality of firms’ growth opportunities. We also find that severance is not required under the optimal contract when the arrival of growth opportunities is verifiable. It is more efficient for the firm, in order to mitigate agency costs, to give zero severance to a departing manager and instead increase the manager’s promise contingent on him being retained.

Third, realized firm growth depends both on the features of the exogenous growth process and on the severity of moral hazard. We distinguish between two endogenous types of firms: *high growth* firms that undertake all growth opportunities, and *low growth* firms that may forego some growth opportunities depending upon the incumbent manager’s past performance. In our model, two firms with similar growth opportunities may end up having very different realized

growth profiles just because they differ in the severity of the agency problem they face. Thus we find for example that improvements in a firm’s corporate governance can transform a firm with given technology and growth opportunities from a *low growth* firm into a *high growth* firm.

Lastly, we identify a new component of agency costs that arises in our framework, which is due to a form of contractual externality. We decompose the firm value into the value of present and future payments to the principal, to the current manager and to all the future managers. When a firm negotiates a contract with an incoming manager, it maximizes the sum of the first two components. It fails to take into account the spillover effect upon the expected amount of time before hiring future managers and thus the present value of compensation received by all future managers. This may result in excessive managerial retention relative to the constrained optimal contract. The agency cost induced by this externality is larger for *low growth* firms, where the arrival of a growth opportunity does not always result in managerial turnover. This externality of the current binding contracts of the firm on its future binding contracts does not arise in earlier papers in the literature, in which firms are liquidated at an exogenous value upon termination of the incumbent, and only, manager of the firm.

We consider two extensions of the baseline model. First, we consider a setting where the arrival of growth opportunities is only observable by the incumbent manager. Under the maintained assumption that growth entails a change of management, the incumbent must be incentivized to truthfully reveal to the firm the realization of a growth opportunity. When the quality of growth opportunities is good enough, the manager is systematically dismissed when he announces that such opportunity becomes available, and he receives a severance pay contingent on the firm’s performance history under his tenure— in stark contrast with the no-severance result that arises when growth opportunities are contractible. Our framework therefore suggests a possible explanation for the widespread use of severance, as a way to assure the goodwill of the incumbent and incentivize him to not bury the news that value-enhancing transformations of the firm have become available when such news may lead to his dismissal.

In another extension, we consider an environment where firms can grow with their incumbent managers, possibly at a different cost than when they grow with a new manager. When growing with the incumbent manager is very costly, e.g., because realizing a growth opportunity would require paying an army of external consultants to help the firm reinvent itself, all the results of the baseline model survive. However, when the costs of growing with the incumbent are sufficiently low, the incumbent may grow the firm, but only if his past performance has been sufficiently good. If instead past performance has been poor, the incumbent will be dismissed, even though this is ex post inefficient.

Our paper contributes to a recent body of work that applies the tools of optimal dynamic contracting to the study of the firm in the presence of agency conflicts. Quadrini (2004), Clementi and Hopenhayn (2006), DeMarzo and Fishman (2007a), Biais et al. (2010), DeMarzo et al. (2011), and Philippon and Sannikov (2011) explore, as we do, the link between dynamic moral hazard and firm growth when growth is contractible.¹ Our framework differs in several dimensions. A key difference is that in these papers a firm always grows with its incumbent manager

¹The extension of our model with privately observed growth opportunities, analyzed in Section 6, is related to He (2008) and Malenko (2012). He (2008) considers an environment where cashflow growth is directly affected by non-observable effort from the manager. Malenko (2012) solves for the optimal investment and auditing policies when the arrival and quality of investment projects are only observed by a division manager who needs the headquarters to finance the working capital.

whereas we put ‘growth-induced’ turnover, i.e., the idea that a change of management may be a pre-requisite for growth, at center stage. Furthermore, we consider growth opportunities which arrive stochastically, we endogenize the state-contingent ‘liquidation’ value of the firm, and focus on managerial turnover rather than firm survival. Finally, our interpretation of growth is slightly different— growth in our framework may take place simply through strategic or organizational change without changes in fixed or working capital. Spear and Wang (2005) and Garrett and Pavan (2012) study optimal termination policies in settings where the firm can dismiss their manager and hire a new one from an external labor market, but they abstract from growth and the economic determinants of turnover they emphasize are different from the ones in our paper.

Our notion that the growth of a firm may require replacing the incumbent manager is found in many early contributions to the managerial theory of the firm. Penrose (1959) discusses why firms may operate successfully under competent managers but may still fail to take full advantage of their opportunities of expansion. Williamson (1966) elaborates on how management constraints affect the realized growth of firms. More recently, Roberts (2004) echoes Penrose by emphasizing the need for different organizational capabilities in the *exploration* and *exploitation* of firms’ investment projects. He discusses a number of business cases where this effect is prominent. In a repeated moral hazard framework but without optimal contracting, Anderson and Nyborg (2011) study the link between managerial replacement and firm growth and show how it is affected by the firm’s choice of debt or equity financing. The idea that firms may need different managers at different times is also present in recent theoretical work on managerial turnover by Eisfeldt and Kuhnen (2012), although not in the specific context of firm growth. They consider a competitive matching model without agency conflicts to explore the role of industry conditions in determining managerial turnover, managerial compensation and the type of CEOs being hired.

The implications of our model are related to the empirical literature on managerial turnover and compensation. In the context of venture capital, Kaplan, Sensoy and Stromberg (2009) provide evidence that the management teams of firms in their early stages of growth (pre-IPO) undergo high turnover. Martin and McConnell (1991), Mikkelsen and Partch (1997), as well as a recent paper by Jenter and Lewellen (2011), study the links between CEO turnover and acquisitions— which can be seen as a source of value creation, and often involve target CEOs being either fired or forced to retire early. Murphy (1985, 2001) and Yermack (2006) document the use of bonuses and severance in U.S. firms, respectively, and Kaplan and Minton (2008) document the joint evolution of CEO compensation and managerial turnover in the U.S. over the recent decades.

As we discuss throughout the paper, the existing empirical findings on managerial turnover and CEO compensation in the literature are compatible with our model. However, the specific empirical predictions of our model of growth-induced turnover and its effect on managerial compensation have not been subjected to direct testing— largely because the existing empirical literature on managerial compensation, severance and turnover has not conditioned upon the growth profiles of firms, either potential or realized, the severity of agency problems, and the contractibility of growth opportunities. Our framework provides guidance on how these factors are likely to affect compensation, severance, and turnover and thus may help in designing more powerful empirical tests. Finding ways to control for these firm characteristics may be challenging, but in our view this should be feasible. For example, one clear prediction of our model is that industries with better growth prospects should experience higher CEO turnover and use more front-loaded compensation schemes. Then sorting industries into those that have and those that

have not undergone a persistent technological shock that creates opportunities for growth should reveal significant differences in turnover experience and compensation policies. In a similar vein, our predictions about severance pay may be tested by comparing the average severance paid to departing managers in fast growing industries as opposed to slow growing industries.

The rest of the paper proceeds as follows. Section 2 describes our baseline model. Section 3 derives the optimal long-term contract, and provides an informal discussion of its main features. Section 4 provides an illustration in the stationary limit of the model. Section 5 employs numerical simulations to further analyze the empirical implications of our model. Section 6 considers an extension where incumbent managers have private information about the arrival of growth opportunities. Section 7 considers an environment where the firm can grow with its current manager. Section 8 concludes. A mathematical appendix includes the proofs of some key results.

2 The baseline model

2.1 Setup

We consider a project that generates a stream of risky cashflows $\{Y_1, Y_2, \dots, Y_T\}$ over T periods (we later consider the stationary limit as T goes to infinity). The project is run by an agent (the *manager*) who is hired by a principal (the *firm*). The agent can underreport cashflows and divert them for his own private benefit. For each unit of diverted cash, he gets $\lambda \leq 1$, so that λ captures the severity of moral hazard. In any period, an incumbent agent can be fired and replaced by a new agent. For simplicity, we normalize the value of an agent's best outside option upon being fired to zero. Agents and principal are risk-neutral with discount rates ρ and $r < \rho$, respectively.

The cashflow generated in period t is $Y_t = \Phi_t y_t$, where y_t is independently and identically distributed with support \mathcal{Y} , $\mathbb{E}(y_t) = \mu$ and $\min(\mathcal{Y}) = 0$, and Φ_t denotes the size of the firm at the beginning of period t . Growth opportunities arrive stochastically over time. The state variable $\theta_t \in \{G, N\}$ describes whether a growth opportunity is available ($\theta_t = G$) or not ($\theta_t = N$) in period t . We assume the process followed by θ_t is independently and identically distributed, with $q \in (0, 1)$ denoting the probability of arrival of a growth opportunity, independent from cashflow realizations. Taking up a growth opportunity involves costs from hiring a new manager and from implementing value-enhancing transformations of the firm. Specifically, if a growth opportunity realizes in period t , given an initial size Φ_t , the firm can grow to a size $(1 + \gamma)\Phi_t$ in period $t + 1$ at a cost of $(\kappa + \chi)\Phi_t$, where $\kappa > 0$ and $\chi \geq 0$ denote the proportional costs of replacing the manager and increasing the scale of the firm, respectively.² If there is no growth opportunity or if an available growth opportunity is not taken up, the size of the firm remains constant; however, the manager may still be replaced at the same scale-adjusted cost, κ . Figure 1 summarizes the timing within each period.

The assumption (relaxed later in Section 7) that growth necessarily entails replacing the incumbent manager is quite natural in circumstances where firm growth requires a new skill set and/or a change in corporate culture. The incumbent manager, whose human capital has to some degree become specific to the firm in its current form during his tenure, will have lost the flexibility to adapt his skills to new requirements. While we have in mind drastic changes of the

²When considering the stationary limit of the model as $T \rightarrow \infty$, we impose that $q\gamma < e^r - 1$ to ensure finite valuation.

firm, as a modeling convenience we capture this as a discrete change in firm size, which scales up the distribution of cashflows. However growth in our model does not necessarily involve an increase in physical capital. Instead, growth could simply be the result of finding better management able to implement a permanent increase in firm productivity.

We focus our analysis on situations where it is first-best efficient to replace management to take up an available growth opportunity, which in the infinite horizon limit of the model amounts to the following parameter restriction

$$\frac{\gamma\mu}{e^r - 1} > \kappa + \chi. \quad (1)$$

Absent a growth opportunity, a manager would never be fired under first best when $\kappa > 0$. As a benchmark, let $V_t(\Phi)$ denote the first-best value of the firm in period t given size Φ , ex-cashflow and before the growth opportunity realization. The sequence of first-best value functions is given recursively by

$$V_t(\Phi) = q [-(\kappa + \chi)\Phi + e^{-r} \{(1 + \gamma)\Phi\mu + V_{t+1}[(1 + \gamma)\Phi]\}] + (1 - q)e^{-r} \{\Phi\mu + V_{t+1}(\Phi)\}.$$

The recursion starts at $V_T(\Phi) = 0$, for all Φ , since at the end of period T there are no further cashflows and the firm expires. In the infinite horizon stationary limit, the homogenous nature of the model allows us to write $V(\Phi) = v^*\Phi$, where

$$v^* = \frac{-q(\kappa + \chi) + e^{-r}(1 + q\gamma)\mu}{1 - e^{-r}(1 + q\gamma)}. \quad (2)$$

2.2 Contracting

We consider optimal second-best contracting when cashflows are non-verifiable and the arrival of growth opportunities is contractible.³ A contract is established between the firm and the manager at the outset of his tenure. When the latter is replaced, the contract is terminated and a new contract is established with a new manager. A contract specifies as a function of history (i.e., the sequence of reported cashflows, and the history of growth opportunity realizations), circumstances under which an agent is fired (i.e., history-contingent firing probabilities), growth is undertaken, and non-negative cash compensation is paid by the principal to agents. Agents have limited liability, and the principal has deep pockets implying that he will not pass up growth opportunities because he is cash constrained. For simplicity, we assume a contractual environment with full-commitment (no renegotiation) and we rule out private savings by the agent.⁴ The amount of diversion is the only decision over which the agent has control. In searching for an optimal contract, we restrict our attention to contracts that induce truthful reporting (since $\lambda \leq 1$ diversion is at least weakly inefficient). An optimal contract is one that gives maximum value to the principal subject to providing a certain expected discounted payoff to the agent, while satisfying incentive compatibility and limited liability constraints. We assume that the contract is designed so as to give an expected discounted payoff of Φw_0 to a manager hired to run the firm at size Φ .

³Section 6 considers the case with non-verifiable cashflows and private information about the arrival of growth opportunities.

⁴DeMarzo and Fishman (2007b), Section 2.1 and Corollary 1, show that if the rate of return available to the agent is less than or equal to r (i.e., private saving is weakly inefficient), even if allowed to do so, the agent would have no incentive to use private savings under the derived optimal contract.

3 The optimal contract

In this section, we characterize managerial turnover, compensation, and realized firm growth under the optimal contract. Our derivation of the optimal contract in Section 3.1 follows and extends the approach of DeMarzo and Fishman (2007b).⁵ An informal discussion of the main features of the optimal contract and its implementation is provided in Section 3.2.

The history of cashflow and growth opportunity realizations can be summarized by two state variables: the current size of the firm Φ , and the agent's size-adjusted expected payoff w . Given this simplified state space, the optimal contracting problem is solved by dynamic programming. To this end, it is useful to introduce a number of value functions (as depicted in Figure 1) to keep track of the principal's expected payoff at different points of time within a period, as a function of the state variables. We let $B_t^y(\Phi, w)$ denote the principal's value under the optimal contract at the beginning of period t , before cashflow realization, given current size Φ and scaled expected payoff w to be delivered to the agent; $B_t^q(\Phi, w)$ denotes the principal's value in period t , after cashflow realization, but before the growth opportunity is realized; $B_{t,\theta}^\ell(\Phi, w)$ denotes the principal's value conditional on a growth opportunity being available or not, before replacement and growth decisions; $B_t^c(\Phi, w)$ denotes the principal's value after the growth/severance decision has been taken but before compensation to the retained agent, conditional on the firm entering period $t + 1$ with size Φ ; and $B_t^e(\Phi, w)$ denotes the principal's value at the end of period t , conditional on the firm entering period $t + 1$ with size Φ and with scaled expected payoff $e^\rho w$ to be delivered to the manager as of the beginning of period $t + 1$. Our assumptions that firm cashflows and costs are all proportional to size guarantee that these value functions are all homogenous in current firm size.

Lemma 1. *All value functions satisfy the following homogeneity property*

$$B_t^i(\Phi, w) = \Phi B_t^i(1, w) \equiv \Phi b_t^i(w), \quad i \in \{y, q, \ell, c, e\}. \quad (3)$$

The analysis is therefore simplified by applying dynamic programming directly onto the size-adjusted value functions. In the end, an optimal contract is entirely characterized by a set of rules specifying the evolution of the state variable w , and a set of policy functions specifying the agent's compensation and the optimal replacement and growth policies as a function of the current value of w and of whether a growth opportunity is currently available or not.

3.1 Properties of the optimal contract

In this subsection, we derive the size-adjusted value functions, the law of motion for the agent's expected payoff w (which following the literature we refer to as the agent's *promise*), and the optimal compensation, replacement and growth policies, proceeding by backward induction.

The recursion starts in the final period with $b_{T,\theta}^\ell(w) = -w$ for $\theta = G, N$. Then for $t \leq T - 1$, taking $b_{t+1,G}^\ell$ and $b_{t+1,N}^\ell$ as given, the value function b_{t+1}^q is obtained from the maximization problem

$$b_{t+1}^q(w) = \max_{w_G, w_N} q b_{t+1,G}^\ell(w_G) + (1 - q) b_{t+1,N}^\ell(w_N), \quad (4)$$

⁵See Green (1987) and Spear and Srivastava (1987) for early applications of recursive techniques in the context of dynamic moral hazard, and DeMarzo and Fishman (2007a) and Biais et al. (2010) for applications involving time-varying firm size.

subject to the promise-keeping condition $qw_G + (1 - q)w_N = w$ and the limited liability constraints, $w_\theta \geq 0$ for $\theta = G, N$. The determination of contingent continuation promises w_G and w_N in (4) is an important feature of the optimal contract in the presence of stochastic growth opportunities. We describe the solution to this problem later in Proposition 2 after having characterized the continuation value functions $b_{t+1,\theta}^{\ell}$ in earlier periods (Proposition 1).

Taking b_{t+1}^q as given, the beginning-of-period value function is obtained as

$$b_{t+1}^y(w) = \max_{\{w^q(y)\}_{y \in \mathcal{Y}}} \mu + \mathbb{E}\{b_{t+1}^q[w^q(y)]\}, \quad (5)$$

where the expectation is taken over the distribution of y , subject to the promise-keeping condition $\mathbb{E}[w^q(y)] = w$, limited liability $w^q(y) \geq 0$, and incentive compatibility

$$w^q(y) \geq w^q(\tilde{y}) + \lambda(y - \tilde{y}), \quad \forall y \in \mathcal{Y}, \forall \tilde{y} \in [0, y]. \quad (6)$$

The following lemma further characterizes the beginning-of-period value function, as well as the impact of the firm's performance on the agent's expected payoff.

Lemma 2. *The value function b_t^y is only defined for $w \geq \lambda\mu$. Moreover, given a beginning-of-period promise w and cashflow realization y , the agent's promise is adjusted as follows*

$$w^q(y, w) = w + \lambda(y - \mu), \quad w \geq \lambda\mu. \quad (7)$$

The agent's continuation promise w^q is increasing in the cashflow realization with sensitivity λ . This guarantees that the agent has no incentive to underreport. Hence the incentive-compatibility condition gives the slope of w^q with respect to y , while the promise-keeping condition gives the level of the schedule. The fact that b_t^y is only defined for $w \geq \lambda\mu$ comes from the interplay between incentive compatibility and limited liability. The beginning-of-period promise w needs to be high enough to guarantee that even for the lowest possible cashflow realization, the continuation promise w^q consistent with incentive-compatibility and promise-keeping constraints remains non-negative.⁶ Given b_{t+1}^y , the end-of-period value function in period t is simply given by

$$b_t^e(w) = e^{-r} b_{t+1}^y(e^{\rho} w), \quad w \geq e^{-\rho} \lambda\mu, \quad (8)$$

where the domain of b_t^e follows directly from that of b_{t+1}^y .

Lemma 3. *For $t < T - 1$, b_t^e is concave in w .*

In a Modigliani-Miller world, increasing the agent's expected payoff would merely amount to redistributing total firm value, and the principal's value would simply be linearly decreasing in the agent's promise with a slope of -1 . In the presence of moral hazard with limited liability, a change in w also affects the principal's value via its impact on the likelihood of inefficient turnover. Under the contract, the principal is committed to firing the agent following a string of bad cashflow realizations even though this may be costly (Proposition 1). When the agent's current promise is low, this *ex post* bad outcome for the principal is relatively likely. Increasing the agent's promise by one dollar actually costs less than one dollar to the principal as this significantly reduces the prospect of a costly turnover. When instead the agent's current promise is relatively high, the prospect of inefficient turnover is already slight and the benefit derived from increasing the agent's promise is also small.⁷

⁶Recall that $\min(\mathcal{Y}) = 0$. More generally, the lower bound of the domain of b^y is $\lambda(\mu - \min(\mathcal{Y}))$.

⁷In the mathematical appendix, we provide a proof to Lemma 3 which takes into account the impact of a change

3.1.1 Cash compensation

The value function b_t^c gives the principal's value contingent on the incumbent manager being retained. The problem at this stage is to determine the optimal amount of cash compensation to the retained manager. Given a current promise $w \geq e^{-\rho}\lambda\mu$, the optimal mix of present versus deferred compensation (captured by c and w^e , respectively) satisfies

$$b_t^c(w) = \max_{c, w^e} -c + b_t^e(w^e) \quad (9)$$

subject to the promise keeping condition $c + w^e = w$, the limited liability condition $c \geq 0$ and $w^e \geq e^{-\rho}\lambda\mu$.

Lemma 4. *Let \bar{w}_t such that $b_t^{e'}(\bar{w}_t) = -1$. The optimal compensation policy is*

$$c_t(w) = \begin{cases} 0, & w \leq \bar{w}_t, \\ w - \bar{w}_t, & w > \bar{w}_t. \end{cases} \quad (10)$$

Therefore, $b_t^c(w) = b_t^e(w)$ for $w \leq \bar{w}_t$ and $b_t^c(w) = b_t^e(\bar{w}_t) - (w - \bar{w}_t)$ for $w > \bar{w}_t$.

The optimal cash compensation to a continuing agent is determined by a basic tradeoff: deferred compensation is preferable because it keeps the agent's promise high and makes inefficient termination less likely, while early compensation is preferable because the agent is more impatient than the principal. Lemma 4 states that it is optimal to defer an agent's compensation until his promise has reached the endogenous threshold \bar{w}_t . Naturally, the compensation threshold \bar{w}_t is determined by comparing the marginal cost for the principal of present versus deferred compensation. By compensating the agent with Δc in period t , the principal's value is $-\Delta c + b_t^e(w - \Delta c)$. For a small Δc , this can be approximated by $b_t^e(w) + \Delta c(-1 - b_t^{e'}(w))$, which implies that non-zero compensation is optimal if and only if $b_t^{e'}(w) < -1$, i.e., if and only if $w > \bar{w}_t$.

3.1.2 Replacement and growth

We now proceed with the construction of $b_{t,\theta}^\ell$ for $\theta = G, N$. At this stage, given the realization of θ and the manager's promise w , the contract specifies the dismissal probability $p_{t,\theta}(w)$, the updated continuation value $w_{t,\theta}^c(w)$ that the incumbent manager gets upon being retained, and a possible severance pay $s_{t,\theta}(w)$ awarded if he is not.

The replacement decision takes into account the continuation value that the principal would get upon replacing the incumbent manager. When there is no growth opportunity available ($\theta = N$), the principal's scaled continuation value is

$$\ell_{t,N} = e^{-r} b_{t+1}^y(w_0) - \kappa. \quad (11)$$

in the agent's promise w on firm growth, which was ignored in the basic intuition above. The key observation in the proof of concavity is that at the next but last period before the end of the firm (period $T - 1$), in order to be able to properly discipline the agent in the last period, there will be circumstances that lead to inefficient replacement. This implies concavity of the value function $b_{T-1,N}^\ell$. One can then show recursively that if the principal's value function is concave at one stage of the firm, the construction of the optimal contract guarantees that the principal's value function is also concave at earlier stages.

If instead a growth opportunity is available ($\theta = G$), the principal's continuation value upon hiring a new manager depends on whether the opportunity is taken up or not. We restrict our attention to situations where the cost of growth (captured by χ) is sufficiently small relative to the benefit of growth (captured by γ), so as to rule out the uninteresting case where the firm would never grow under second best. Hence the principal's continuation value scaled by current size is

$$\ell_{t,G} = e^{-r}(1 + \gamma)b_{t+1}^y(w_0) - (\kappa + \chi) > \ell_{t,N}, \quad (12)$$

and $p_{t,G}(w)$ can also be interpreted as the probability of growing conditional on a growth opportunity being available.

The optimal replacement/growth and severance policies are obtained by considering the following constrained maximization problem, separately for $\theta = G$ and $\theta = N$:

$$b_{t,\theta}^\ell(w) = \max_{p,s,w^c} p(\ell_{t,\theta} - s) + (1 - p)b_t^c(w^c) \quad (13)$$

subject to the promise keeping condition $ps + (1 - p)w^c = w$, the limited liability condition $s \geq 0$, $w^c \geq e^{-\rho}\lambda\mu$, and $p \in [0, 1]$. To analyze this problem, it is useful to introduce for $\theta \in \{G, N\}$,

$$\delta_{t,\theta} = \sup_{w \geq e^{-\rho}\lambda\mu} \frac{b_t^c(w) - \ell_{t,\theta}}{w}, \quad (14)$$

and

$$\underline{w}_{t,\theta} = \begin{cases} \inf\{w \geq e^{-\rho}\lambda\mu \text{ s.t. } b_t^c(w) \leq \ell_{t,\theta}\}, & \text{if } \delta_{t,\theta} > -1, \\ \infty, & \text{otherwise.} \end{cases} \quad (15)$$

Graphically, $\delta_{t,\theta}$ and $\underline{w}_{t,\theta}$ are determined by finding the line of maximum slope relating the termination point $(0, \ell_{t,\theta})$ to the curve representing the value function b_t^c .⁸ The slope of this line gives $\delta_{t,\theta}$, while $\underline{w}_{t,\theta}$ is defined as the value of w at the intersection/tangency point if $\delta_{t,\theta} > -1$ and $\underline{w}_{t,\theta} = \infty$ otherwise. Note that $\ell_{t,G} > \ell_{t,N}$ implies $\delta_{t,G} < \delta_{t,N}$. Furthermore, since $\kappa > 0$, the slope $\delta_{t,N}$ is typically positive.

Proposition 1. *For any realization of $\theta \in \{G, N\}$, the optimal replacement policy can be described as follows:*

- (i) *if $\delta_{t,\theta} > -1$, the probability of the incumbent agent being replaced is*

$$p_{t,\theta}(w) = \begin{cases} 1 - w/\underline{w}_{t,\theta}, & 0 \leq w < \underline{w}_{t,\theta}, \\ 0, & w \geq \underline{w}_{t,\theta}. \end{cases} \quad (16)$$

The agent receives no severance pay upon being fired, $s_{t,\theta}(w) = 0$, $\forall w < \underline{w}_{t,\theta}$, and his continuation promise upon being retained is

$$w_{t,\theta}^c(w) = \begin{cases} \underline{w}_{t,\theta}, & 0 < w < \underline{w}_{t,\theta}, \\ w, & w \geq \underline{w}_{t,\theta}, \end{cases} \quad (17)$$

Hence

$$b_{t,\theta}^\ell(w) = \begin{cases} \ell_{t,\theta} + \delta_{t,\theta}w, & 0 \leq w \leq \underline{w}_{t,\theta}, \\ b_t^c(w), & w \geq \underline{w}_{t,\theta}. \end{cases} \quad (18)$$

⁸See Figures 3 and 4.

- (ii) if $\delta_{t,\theta} \leq -1$, the incumbent manager is replaced with probability one independently of the agent's promise, $p_{t,\theta}(w) = 1$ for all $w \geq 0$. Upon being replaced, the manager receives $s_{t,\theta}(w) = w$, and

$$b_{t,\theta}^{\ell}(w) = \ell_{t,\theta} - w, \quad \forall w \geq 0. \quad (19)$$

Proposition 1 shows that the two values $\underline{w}_{t,N}$ and $\underline{w}_{t,G}$ defined by Equation 15 are key determinants of the optimal replacement policy. The value of $\underline{w}_{t,N} < \infty$ serves as the dismissal threshold when no growth opportunity is available. If the manager's current promise lies above this threshold, $w \geq \underline{w}_{t,N}$, then he is retained for sure, as indicated by Equation 16. If rather the operating performance has been so poor that the manager's promise is below the threshold, $w < \underline{w}_{t,N}$, then he is at risk of being fired. In effect, he is given a lottery whereby with some probability he is dismissed (although this is *ex post* inefficient) and receives no further payment from the firm, or alternatively he stays with the firm and is awarded a continuation promise increased to the threshold amount, $\underline{w}_{t,N}$. The probability of dismissal in that case is chosen so that the lottery is fair (i.e., its expected value equals the agent's promise w), and is therefore decreasing in w .

When a growth opportunity is available, the dismissal policy essentially depends on how good the growth opportunity is— which determines whether $\delta_{t,G}$ is greater or lower than -1 . In the case of modest growth opportunities, the logic of the dismissal decision is similar as above, except for the fact that it is made by comparing the agent's promise to the threshold $\underline{w}_{t,G}$, which by construction satisfies $\underline{w}_{t,G} > \underline{w}_{t,N}$. In particular, an incumbent manager might be retained despite the availability of a growth opportunity, and the probability of this happening is increasing in the agent's promise. When instead the benefit of growing is large enough (so that $\delta_{t,G} \leq 1$), the incumbent manager is systematically replaced when a growth opportunity is available, independently of past performance.

3.1.3 Contractual response to the arrival of a growth opportunity

We now close the derivation of the optimal contract by characterizing how the realization of a growth opportunity affects the agent's expected payoff. This involves solving the optimization problem entering in the definition of b_t^q , as stated in (4).

Proposition 2. *For a given post-cashflow promise w , the contingent continuation promises (w_G, w_N) in period t are characterized as follows:*

- (a) If $\delta_{t,G} > -1$
- (i) if $w < (1 - q)\underline{w}_{t,G}$, $w_G = 0$ and $w_N = w/(1 - q)$;
 - (ii) if $(1 - q)\underline{w}_{t,G} \leq w < \underline{w}_{t,G}$, $w_G = (w - (1 - q)\underline{w}_{t,G})/q$ and $w_N = \underline{w}_{t,G}$;
 - (iii) if $\underline{w}_{t,G} \leq w \leq \bar{w}_t$, $w_G = w_N = w$;
 - (iv) if $w > \bar{w}_t$, any combination of w_G and w_N such that $w_G \geq \bar{w}_t$, $w_N \geq \bar{w}_t$, and $qw_G + (1 - q)w_N = w$ can be chosen.
- (b) If $\delta_{t,G} \leq -1$
- (i) if $w \leq (1 - q)\bar{w}_t$, $w_G = 0$ and $w_N = w/(1 - q)$;

(ii) if $w > (1 - q)\bar{w}_t$, any combination of $w_G \geq 0$ and w_N such that $w_N \geq \bar{w}_t$ and $qw_G + (1 - q)w_N = w$ can be chosen.

Proposition 2 describes how contingent promises (w_G, w_N) are optimally set, subject to the constraint of delivering in expectation to the agent his post-cashflow promise w . Part (a) of the proposition, illustrated in Figure 2, applies in situations where the agent is not systematically replaced when a growth opportunity is available. For low levels of the agent's post-cashflow promise, the promise is allocated entirely to the state of the world where no growth opportunity is available. This is clearly optimal since a higher promise in the no-growth state reduces the likelihood of inefficient turnover, while a lower promise in the growth state increases the probability that growth be pursued if a growth opportunity becomes available.⁹ When w reaches $(1 - q)\underline{w}_G$, keeping the agent's promise in the no-growth state pegged at $\underline{w}_{t,G}$ while allocating any marginal increase in the agent's promise to the growth state is optimal as the marginal cost of increasing the promise in the no-growth state is equal to the marginal cost of increasing the promise in the growth state— both being equal to $\delta_{t,G}$.¹⁰ Thus for low levels of the post-cashflow promise w , the arrival of a growth opportunity is bad news for the incumbent manager. For higher levels w , the continuation promise w_θ is independent of whether growth is available or not. The agent will be retained for sure, and by keeping the agent's promise at w the principal equalizes the marginal cost of the promise across the growth and no-growth states.¹¹ Part (b) of Proposition 2 applies when the agent is systematically replaced upon realization of a growth opportunity, in which case the marginal cost of an increase in w_G is constant and equal to -1 . Then setting $w_G = 0$ is always optimal, and strictly so as long as $b_N^{\ell'}(w/(1 - q)) \geq -1$, i.e., $w \leq (1 - q)\bar{w}_t$.

We conclude with a couple of observations which follow directly from our analysis, by combining Propositions 1 and 2.

Corollary 1. *For a given post-cashflow promise $w > 0$, $p_{t,G}[w_G(w)] \geq p_{t,N}[w_N(w)]$.*

This first result states that, conditional on a given performance history, the probability of turnover is weakly higher in the presence of a growth opportunity than in its absence.¹² A strict inequality holds unless $w \geq \underline{w}_{t,G}$ and $\delta_{t,G} > -1$. In addition, we have the following result:

Corollary 2. *There always exists an optimal contract under which the agent receives no severance pay upon being replaced.*

Corollary 2 establishes that severance pay plays no material role in the optimal dynamic contract.¹³ Positive severance pay can never arise in the absence of a growth opportunity, or even

⁹In particular for $(1 - q)\underline{w}_{t,N} < w < (1 - q)\underline{w}_{t,G}$, allocating the promise entirely to the no-growth state remains optimal so as to ensure that any available growth opportunity is taken up with probability one.

¹⁰Part (a-ii) of Proposition 2 implies that, conditional on surviving, the agent carries a promise of $\underline{w}_{t,G}$ into the compensation phase, independently of whether a growth opportunity realizes or not.

¹¹Part (a-iv) of Proposition 2 shows that for very high values of $w > \bar{w}_t$, any (w_G, w_N) such that $b_G^{\ell'}(w_G) = b_N^{\ell'}(w_N) = -1$ is optimal. Figure 2 assumes that $w_G = w_N = w$ in that region.

¹²In particular, in the case $\delta_{t,G} > -1$, part (a-i) of Proposition 2 along with Proposition 1 imply that when the agent's post-cashflow promise w is below $(1 - q)\underline{w}_{t,G}$, the probability of replacing the agent to take up an available growth opportunity is equal to one.

¹³The analysis of Section 6 shows that this result hinges crucially on the contractibility of growth opportunity realizations (Corollary 3).

upon realization of such an opportunity as long as $\delta_{t,G} > -1$. Indeed in both circumstances, part (i) of Proposition 1 applies. The only circumstance, although somewhat artificial, where severance pay could arise under an optimal contract is if $\delta_{t,G} \leq -1$, and the firm has had good recent performance so that the agent's post-cashflow promise is above $(1 - q)\bar{w}_t$. In that case, combining part (ii) of Proposition 1 and case (b-ii) of Proposition 2, it appears that the principal is indifferent between giving a non-zero severance pay to the agent contingent on $\theta_t = G$, or zero severance and a higher continuation payoff contingent on $\theta_t = N$.

3.2 Discussion of the optimal contract

Having formally derived the optimal contract in our setting, it is useful to summarize it informally and to discuss how it can be implemented in practice. The optimal contract between the firm and its manager sets out the conditions under which the manager will be compensated during his tenure at the firm and also those which will lead to his leaving the firm. These terms and conditions are chosen to maximize the value of payoffs to the firm's owners subject to incentivizing the manager to truthfully report realized cashflows. Compensation and replacement/retention decisions are made over time as a function of the value of payments promised to the manager (scaled by current firm size), w_t , which evolves under the influence of the firm's operating performance (a manager's promise is increasing in the firm's performance under his tenure) and growth opportunity realizations. The contractual features in force in period t are summarized by three threshold values— the *dismissal thresholds*, $\underline{w}_{t,N}$ and $\underline{w}_{t,G}$, and the *bonus threshold*, \bar{w}_t .

Since the replacement decision is made after the availability of a growth opportunity (or lack thereof) has been observed, dismissal thresholds are conditioned on such opportunity being available or not. The dismissal threshold $\underline{w}_{t,N}$ determines the replacement decision when there is no growth opportunity available. If the manager's current promise w_t lies above this threshold, he knows that he will be retained. If instead the firm's performance under his tenure has been so poor that his current promise is below $\underline{w}_{t,N}$, then he is at risk of being fired with no further payments from the firm. If he survives this, he stays with the firm and is awarded a continuing promise that is increased to the dismissal threshold amount. The intuition for why there is zero severance pay in this case is that by reducing the payment upon dismissal to zero the principal is able to increase the promise to the agent if he survives the dismissal threat, thus reducing the agency problem faced by the firm subsequently. The probability of dismissal is chosen so that the lottery is fair and is decreasing in the agent's promise.

When a rather modest growth opportunity is available, the dismissal decision follows a similar logic, although it is made by comparing the agent's promise to the dismissal threshold $\underline{w}_{t,G}$, which is higher than that without growth. Again, if the manager's promise is below the threshold $\underline{w}_{t,G}$ he is given a fair lottery in which, if he is dismissed, he leaves the firm with no further compensation, and if he survives, his expected discounted payoff is increased to $\underline{w}_{t,G}$. If the manager's promise is above the threshold $\underline{w}_{t,G}$ he knows he is safe. Notice however that retaining the incumbent in the face of a growth opportunity is inefficient, i.e., the firm passes up a positive NPV project. A form of agency-induced managerial *entrenchment* can therefore arise in our setting following periods of sustained good performance.¹⁴ When instead the benefit of growing

¹⁴This possibility (which only arises in our model when $\delta_{t,G} > -1$) is in contrast with the result of Casamatta

is large, the incumbent manager is always replaced for the sake of growth, independently of past performance ($w_{t,G} = \infty$). Conditional on performance history, the risk of dismissal is always weakly higher upon realization of a growth opportunity than if no such opportunity materializes.

In any period, if the agent has survived the replacement phase, he may be entitled to cash compensation, as determined by the bonus threshold \bar{w}_t . If the adjusted promise w_t of a surviving agent lies above \bar{w}_t , a bonus is awarded in that period equal to the excess $w_t - \bar{w}_t$, and the agent's continuing promise is reduced to the threshold amount \bar{w}_t . Otherwise, if $w_t \leq \bar{w}_t$, the agent receives no compensation in that period and continues with his promise w_t , which is adjusted to $e^\rho w_t$ at the beginning of the next period as a fair compensation to the agent for his payoff being delayed.

The promise that the agent takes into a period undergoes two adjustments prior to the replacement and compensation phases. First, upon the report of the cashflow for the period, the agent's promise is adjusted linearly as described in Equation 7, the cashflow sensitivity being set so as to provide the right incentives for the agent not to divert. Then depending on whether a growth opportunity realizes or not, the promise is further adjusted as described in Proposition 2. By allocating a given post-cashflow promise w between the 'no-growth' state ($\theta = N$) and the 'growth' state ($\theta = G$), the principal is effectively determining the probabilities of inefficient and efficient turnover. In circumstances where the firm faces moderate growth opportunities, when the cashflow performance has been poor and w is low, it is optimal to allocate the promise entirely to the no-growth state, so as to reduce the chances of a costly turnover while increasing those of efficient turnover. Past the point where a costly turnover is avoided for sure, it is when the marginal cost of increasing the promise in the no-growth state becomes equal to that of increasing the promise in the growth state (due to a reduction in the probability of taking up an available growth opportunity) that the principal starts allocating some of the promise to the times where a growth opportunity realizes. For higher values of the post-cashflow promise, the continuation promises are independent of the growth opportunity realization, as the agent will be retained for sure and the firm will not grow in any state of the world. In circumstances where the benefit of growing is so large that the incumbent manager is systematically dismissed upon realization of a growth opportunity, it is always (at least weakly) optimal for the principal to allocate the promise entirely to the no-growth state, so as to reduce the likelihood of inefficient turnover.

The optimal contract calls for zero severance pay to a dismissed manager under most circumstances (in particular if a manager is not dismissed upon growth), and positive severance is always at least weakly dominated by no severance (Corollary 2). Severance is suboptimal because agents are risk neutral and severance payments have no agency cost-reducing benefit once the agent leaves the firm. By reducing the severance and increasing the promise to the agent in the case he is retained, the agent can be made as well off but the principal can be made better off because the prospect of a subsequent costly liquidation is made more remote. Our zero-severance result relies crucially on the assumption that growth opportunities are both exogenous and contractible. We show in Section 6 that growth-induced turnover can result in positive severance

and Guembel (2010) who find that moral hazard, combined with reputational concerns and managerial legacies, leads to managerial entrenchment after *poor* performance. In a setting where managers privately observe shocks to their productivity, Garrett and Pavan (2012) also find that the optimal contract allows some form of entrenchment in the sense that retention decisions become more permissive with time.

with risk neutral agents if the principal has to incentivize the agent to truthfully report the arrival of a growth opportunity.

The optimal contract we have just described can be implemented fairly directly using standard employment contracts, and there is some evidence that features of our optimal contracts are used in practice. Our contract specifies an indefinite term with both the manager and the firm having the right to terminate at will.¹⁵ Actual employment contracts are often written in this way.¹⁶ In practice, it is not unheard of that following a period of poor performance when the manager was thought to be under threat of dismissal, the firm instead retains the manager and gives him an improved compensation package as a vote of confidence. This is analogous to the award of deferred compensation of $\underline{w}_\theta - w$ when the manager survives a dismissal threat. The bonus calculation in the optimal contract is very much like the typical contract that was found by Murphy (2001) in his study of the bonus contracts of large U.S. firms. The key parameters he identifies are the performance target, the pay-performance-sensitivity (pps), and the bonus threshold. In our contracts, these are μ , λ , and \bar{w}_t respectively.

Finally, our analysis implies that it is useful to distinguish two categories of firms depending upon the quality of their growth prospects. The tenure of an incumbent manager will be heavily dependent upon the type of firm he is running. A *high growth* firm is one that will undertake growth any time it has an opportunity to do so, thus generating a lot of growth-induced turnover. Other firms, which for simplicity we call *low growth* firms even though in practice they may grow quite fast, do not always take up an available growth opportunity.¹⁷ While *high growth* firms fully realize their growth potential, *low growth* firms optimally waste part of theirs. In the face of a growth opportunity, a firm of the latter type will retain its incumbent manager and forsake growth when the firm's performance history under the manager's tenure has been good.¹⁸ The reason is that when the manager has accumulated a high promised compensation, the cost of taking up the growth opportunity (which would include an expensive severance package) is prohibitive. For lower values of w , the probability of taking a growth opportunity is decreasing in w . That is, the better the recent operating performance, the less likely that the firm will take up a growth opportunity.

4 Optimal stationary contract

We now consider our model in the stationary limit where $T \rightarrow \infty$. This is a useful simplification because the key features of the optimal contract, adjusting for changes of scale as the firm grows, will be constant over the life of the firm. This allows us to better understand the relationship between these contract features and the deep underlying characteristics of the firm, in particular,

¹⁵Our setup could easily be extended to incorporate a positive reservation value for the agent. With zero reservation value and limited liability, inducing the agent to remain in the contract is never an issue.

¹⁶Of course, some employment laws may constrain this, e.g., by imposing a mandatory notice period which may vary with the tenure.

¹⁷Proposition 1 shows that the distinction between the two types of firms depends crucially on whether the value of $\delta_{t,G}$ in Equation 14 is higher (*low growth* firms) or lower (*high growth* firms) than -1 . In Section 4.3, we provide a mapping of *high growth* vs. *low growth* firms in the parameter space in the stationary limit of the model.

¹⁸This result contrasts with DeMarzo and Fishman (2007a) who find that investment is increasing in the agent's promise because the return on investment is high then.

the severity of managerial moral hazard and the frequency of growth opportunities.

To do this, we solve numerically for the value functions and associated replacement, growth, and compensation policies by iterating backward until convergence for a large value of T . We assume size-adjusted cashflows are independently, identically and uniformly distributed on $\{0, 1, 2, \dots, 20\}$, with mean $\mu = 10$. The moral hazard parameter is $\lambda = 0.9$. Discount rates for the principal and the agent are such that $e^r - 1 = 6.5\%$ and $e^\rho - 1 = 7\%$. The cost of firing and replacing a manager is equal to 2% of annual mean cashflow ($\kappa = 0.2$), while the cost of scaling-up is set to 20% of annual mean cashflow ($\chi = 2$). We set the scale adjusted reservation compensation for a new manager at $w_0 = 14$. Other parameter values to be specified are q and γ , capturing the likelihood and the magnitude of growth opportunities, respectively.

4.1 High growth and low growth firms

Our analysis in Section 3.1 shows that the optimal stationary contract is entirely summarized by three time-independent threshold values \underline{w}_N , \underline{w}_G and \bar{w} . Consider first the case where $q = 0.1$ and $\gamma = 0.25$. In this case, the optimal stationary thresholds are $\underline{w}_N = 8.42$, $\underline{w}_G = \infty$ and $\bar{w} = 26.06$. The fact that $\underline{w}_G = \infty$ indicates that it is optimal to grow and replace the agent with probability 1 whenever a growth opportunity is available. That is, this is a *high growth* firm. Figure 3 represents the corresponding stationary value functions. The fact that the value function b_G^ℓ decreases linearly with slope -1 and lies above b^c for all values of w indicates graphically that this is a case of high growth. The agent's compensation threshold $\bar{w} = 26.06$ means that an agent who enters the job with an expected discounted payoff of $w_0 = 14$ must experience a sustained run of good cashflow realizations before receiving any cash compensation.

Suppose instead $\gamma = 0.1$, while all other parameters are kept the same. The optimal stationary thresholds become $\underline{w}_N = 8.42$, $\underline{w}_G = 18.06$ and $\bar{w} = 33.29$. Having reduced the rate at which the firm can grow upon arrival of a growth opportunity, we now have a firm which does not systematically take up efficient growth opportunities when available, and never does so if w is above the threshold $\underline{w}_G = 18.06$. This is a *low growth* firm. Figure 4 shows the stationary value functions in this case. Note that b_G^ℓ initially decreases linearly with slope greater than -1 and is tangent to b^c at $\underline{w}_G = 18.06$. The bonus threshold in the low growth benchmark firm is higher than in the high growth benchmark (33.29 versus 26.06). Later when we simulate the model we will see that on average compensation will arrive much later for the agent in this lower growth case.

4.2 Sensitivity of contract terms

The realized earnings and growth performance of firms are the result of managers' and owners' responses to cashflow shocks and to the arrival of growth opportunities, and these reactions will be shaped by the terms of the contract as set out in the pay-performance sensitivity and in the thresholds, \underline{w}_N , \underline{w}_G and \bar{w} . Thus understanding how these thresholds are affected by changes in the deep parameters of the model is an important step toward understanding how the earnings and growth experience of firms is determined.

Figure 5 depicts the three thresholds as functions of the severity of moral hazard, λ , and the frequency of growth opportunities, q , for a firm with a finite \underline{w}_G , that is, for a low growth firm. The dismissal threshold in the absence of growth opportunities is given analytically by $\underline{w}_N =$

$e^{-\rho}\lambda\mu$. That is, the non-growth dismissal threshold is linearly increasing in λ and independent of q . Intuitively, in the face of increased moral hazard, the principal will increase the dismissal threshold, thereby increasing the risk of disciplinary dismissal.

Next consider the impact of λ on the bonus threshold, \bar{w} . It is increasing in λ reflecting an increased benefit of deferred compensation. This is because the inefficient termination threshold is higher and the pay-performance sensitivity increases, implying that it takes a shorter run of poor performance for the no-growth dismissal threat to be active.

An increase in λ also results in a higher dismissal threshold \underline{w}_G . This is beneficial to the principal as it implies that the probability that the firm will take an available growth opportunity increases, and in case an available growth opportunity is not taken up, the surviving agent's promise is reset to a higher level, $w = \underline{w}_G$, which makes subsequent inefficient liquidation less likely. This also contributes to shift the distribution of promises to the right, making it more likely that an agent's promise will reach the (higher) compensation threshold, a necessary condition for the initial promise w_0 to be fulfilled.

We turn next to the impact of q on \bar{w} and \underline{w}_G , again for low growth firm. A higher q causes a fall in the bonus threshold, \bar{w} , implying that cash payouts will be made following a shorter run of good performance. This follows because, a higher q implies higher unconditional probability of early termination, with no severance pay. Thus in order to deliver the reservation value, w_0 , ex ante, the cash compensation needs to be paid earlier. Furthermore, for the same reason, in order to increase the probability of getting to the bonus threshold the growth dismissal threshold, \underline{w}_G , decreases because this decreases the probability of dismissal, conditional on $\theta = G$.

Finally, for high-growth firms, by definition $\underline{w}_G = \infty$. The sensitivities of \underline{w}_N and \bar{w} are similar to those in the low-growth case and for similar reasons. Again, in our framework, $\underline{w}_N = e^{-\rho}\lambda\mu$. The bonus threshold \bar{w} is increasing in λ and decreasing in q , as is the case for low-growth firms. \bar{w} falls with an increase in q because the marginal benefit of deferred compensation decreases as q increases. This is because as q increases it is more likely that a growth opportunity will arrive soon, in which case it will be taken up for sure—and the likelihood of inefficient replacement is therefore reduced.

4.3 What makes a firm grow fast?

Our baseline examples in Section 4.1 show that two firms that differ only in the size of the growth opportunity will have very different contracts for top management. These differences translate into very different policies toward growth opportunities with high-growth firms undertaking all opportunities that present themselves and low-growth firms undertaking opportunities only if incumbent management is not performing well.

It is also the case that differences only in agency costs may result in very different growth experiences. To see this, consider an example of two firms that have the same size of their growth opportunities ($\gamma = .125$), the same probability of having a stochastic growth opportunity $q = 10\%$, and only differ in the degree of moral hazard λ . All other parameters are as in our baseline cases. In this example, our model predicts that the firm with $\lambda = 0.5$ grows at an average rate of 1.25%. This is because it is a high-growth firm that undertakes all the growth opportunities that arise. Meanwhile, the firm with $\lambda = 1.0$ grows at an average rate of around

0.41%.¹⁹ Stated otherwise, suppose the two firms start out life with identical scale of operations. Fifty years on (i.e., at $t = 50$), the expectation is that the firm with low agency problems will have a scale (measured by the mean cashflow rate) that is 52% larger than the firm with more severe agency problems.²⁰

This holds for other parameters as well. That is, we may have two firms that differ only slightly in their deep parameters, with one a high-growth firm and the other a low-growth firm. Figure 6 depicts regions of the parameter space corresponding to high-growth firms and low-growth firms. All parameters are set as in the second baseline case (low-growth firm) of Section 4.1 except for the two parameters depicted in the diagram.

To summarize, small differences in parameters can result in dramatically different growth and turnover behavior. Growing firms need a sustained flow (high q) of good opportunities (high γ) for improving their technology and expanding markets. They need to manage transitions well (low κ , low χ). And they need to keep agency problems under control, for example, through increased monitoring (low λ).

5 Turnover, compensation timing, agency costs

5.1 Simulating the model

We now simulate the model to understand its implications for management turnover and the relative importance of deferred compensation. Simulations also allow us to assess the importance of the agency costs due to the contracting imperfections present in this framework.

Specifically we draw repeatedly a sequence of cashflows and growth opportunity realizations, keeping track of compensation, growth and termination decisions commanded by the optimal contract. We then characterize these histories using a variety of summary statistics. We focus on three statistics that are of particular interest. First we calculate the average longevity or ‘tenure’ of managers, which is inversely related to the replacement frequency. Second we calculate the unconditional probabilities of efficient termination (i.e. fire the agent to undertake growth) and inefficient termination (i.e. fire the agent without growing) as the corresponding realized sample frequencies. Third, to measure the extent to which the optimal contract relies on deferred compensation, we calculate the average duration of the agent’s compensation conditional on the agent receiving non-zero compensation during his tenure in the firm. This is calculated as the weighted average of the tenure years at which the agent receives positive payments, with weights calculated as the ratio of discounted payment to the sum of discounted payments.

For example, consider the results for the benchmark cases given in Section 4.1. For the *high growth* firm with $\gamma = 0.25$, average tenure of an agent is 8.6 years. The average probability of efficient termination is 10% per year, reflecting the fact that for a *high growth* firm any available growth opportunity is undertaken. The probability of inefficient termination is about 1.57% per year. And the average duration of compensation is 7.1 years.

¹⁹The latter statement is based on simulations.

²⁰Note that an improvement in corporate governance, if it induces a fall in λ , can potentially eliminate agency-induced entrenchment.

In contrast for the *low growth* firm with $\gamma = 0.1$, the average tenure is 109 years.²¹ The probability of inefficient termination is 0.25% which is lower than the probability of efficient termination (0.66%). The average duration of compensation is 20.4 years. Comparing results for the two cases, we see that *high growth* firms receive compensation earlier on average than do agents in *low growth* firms.

5.2 Comparative statics

In this section, we further explore predictions from our model in terms of its comparative statics with respect to some key parameters. Specifically, we solve our model for alternative values of these parameters and then simulate the model assuming the same realizations for underlying cashflow shocks and growth opportunities. We record the histories of management turnover, whether turnover takes place for growth or for disciplinary reasons, and the compensation histories for each of the firm’s managers. The parameters we vary are q , the probability of having a stochastic growth opportunity, and λ , the severity of agency problems. The default values of these parameters take on when the other parameter is varied are $q = 0.1$ and $\lambda = 0.9$. Other parameters are as in Section 4.1.

5.2.1 Management turnover

In our model managers are replaced either to facilitate growth or because a history of poor operating results leads to dismissal. The exact conditions under which managers are replaced are sensitive to both the growth prospects of the firm and to the severity of agency problems faced by the firm.

Representing the quality of the growth prospects by the frequency of arrival of growth opportunities, q , we show the sensitivity to this parameter of average manager tenure. This is depicted in the left panel of Figure 7 for a high growth firm with $\gamma = 0.25$. From the figure we see that as the probability of growth opportunity in a year rises from 5% to 25% the average tenure of the agent declines from 15 years to about 4 years. A similar negative sensitivity to increases in q holds for low growth firms (e.g., with $\gamma < 0.1$), with the difference that, for a given q , the average tenure is much higher.

Thus tenure falls and turnover rises for firms with better growth prospects. To our knowledge this hypothesis has not been submitted to direct empirical testing. However, there is some indirect evidence which is supportive of the hypothesis. Specifically, Mikkelsen and Partch (1997) compare top management turnover intensity in two successive five-year periods with very different mergers and acquisitions activity. They find that in the active take-over period of 1984-1988, 33% of firms in the sample underwent complete management changes (i.e., replaced all of the president, CEO and Chairman); whereas this intensity was only 17% in the subsequent period 1989-1993 when take-over activity was low. Interestingly their notion of complete management corresponds better to our model which associates turnover and major changes of direction than does most of the literature which has focused exclusively on CEO turnover. While they do not specifically make a link of management turnover and firm growth, the two periods they cover coincide with very different experiences of firm growth. On the aggregate level, average annual

²¹The model could be easily extended to obtain a more realistic turnover rate, e.g., by introducing stochastic death, legal retirement age, or time-variation in managers’ reservation values.

total factor productivity growth in the U.S. was 1.35% in the 1984-88 period, whereas between 1989 and 1993 it was only 0.61%.²²

In the right panel of Figure 7 we see the consequences of increasing the severity of managerial moral hazard. As the rent extraction efficiency (λ) of the agent rises the average longevity declines. This is a reflection of the fact that the optimal contract relies more heavily on the threat of termination in the face of more severe moral hazard. A similar pattern is found for low growth firms as well.

5.2.2 Efficient and inefficient replacement probabilities

As already noted, turnover may occur for growth or for discipline. These two kinds of managerial turnover are affected differently by changes in the firm's underlying characteristics. To distinguish these effects, we calculate the average frequency of these two types of turnover in the simulated histories and plot these as functions of q and λ in Figure 8. The top row pertains to the high growth case, with $\gamma = 0.25$ as above. In high growth firms the unconditional probability of replacement for reasons of growth are higher than the probability of disciplinary replacement. Since all growth opportunities are taken up in these firms, this frequency increases linearly in q .

The effect of more severe agency problems on dismissal frequencies in high growth firms is given in the upper right panel of Figure 8. Since all growth opportunities are taken up, changes in λ have no effect on the efficient dismissal probability. The probability of inefficient dismissal is slightly increasing in λ . This reflects an increased reliance on the termination threat when moral hazard is more severe.

The sensitivities of dismissal probabilities for low growth firms are given in the bottom row of Figure 8. As for high growth firms, the probability of efficient dismissal increases when growth arrives more frequently, i.e., as q increases. In the right panel, the probability of inefficient replacement increases with λ , reflecting greater reliance on the dismissal threat (increased \underline{w}_N). The decline in the unconditional probability of efficient dismissal, seen in the figure, comes from a shift to the right of the distribution of promises under the optimal contract.

5.2.3 Compensation duration

To assess the consequence of changing parameters for the reliance on deferred compensation, we calculate the realized duration of compensation from bonuses during agents' tenure. These sensitivities are given in Figure 9. From the top row we see that for both high and low growth firms an increase in q reduces the duration of compensation. That is, when growth opportunities arrive more frequently, firms optimally rely on more front-loading of compensation. The effect works through the lower bonus threshold for both high-growth and low-growth firms.

The second row of Figure 9 shows the effect of increasing λ . For both high growth firms and low growth firms the average duration of compensation rises as λ rises. The reason for this is that a higher λ increases bonus threshold, \bar{w} . Managers receive compensation only after a sustained run of good performance.

Again, to our knowledge, there are no empirical studies that directly test whether these effects on the timing of compensation hold. However, Kaplan and Minton (2008) find evidence

²²This is also reflected in the evolution of U.S. annual non-residential investment spending, which increased by 28% and 12.5% over the two sample periods, respectively (based on annual U.S. National Income Statistics).

of an acceleration of CEO turnover since 1990, especially after 2000, a period that also saw very rapid increases in the amount of top management compensation. They argue that the observed increases in CEO pay are compensation for shorter tenure. This is consistent with our theory in which high growth will be associated with shorter tenure and more front-loading of compensation.

5.3 Second-best value and agency costs

The second-best value of the firm in our framework can be computed as the expected present value of all cashflows that accrue to the principal and to all managers who successively run the firm under optimal contracts as set out in Section 3. Two subtleties should be noted in calculating this second best value. First, cashflows to agents are discounted at the agents' discount rate, ρ ; whereas, the firm discounts cashflows at rate r . Since $\rho > r$, the promise to an agent is worth less to the agent than it costs the firm. Second, the calculation of agent cashflows includes payments to *all* agents, both current and future. Thus in the stationary case we can write the size-adjusted, beginning-of-period second-best value of the firm as

$$v(w) = b^y(w) + w + f(w), \quad (20)$$

where $f(w)$ denotes the expected present value of payoffs to future agents as a function of the current agent's promised value, w .²³ To assess the extent of agency costs, the total value of the firm under the optimal contract $v(w)$ can be compared to the beginning-of-period, first-best value of the firm, $\mu + v^*$, for v^* defined in (2).

Figure 10 depicts firm value and its components under the second best optimal contact for the *high growth* ($\gamma = 0.25$ in the top panel) and *low growth* firms ($\gamma = 0.1$ in the bottom panel) as set out in Section 4.1. The left panel gives the value for the principal and the incumbent agent, $b(w) + w$. The middle panel gives the present value of compensation to future agents who are not party to the current contract but who are affected by the current contract and the current promise to the incumbent agent, $f(w)$. The right panel gives the sum of all these components, that is, the second best value of the firm, $v(w)$. In the left panels of Figure 10 we see that for both high and low growth firms the combined value to the principal and the incumbent manager is increasing and concave in the promise to this manager. This reflects the relaxation of agency problems affecting the two parties to the current contract, and this is an effect already seen in previous dynamic agency models. Interestingly, the total second-best firm value, taking into account the effect on future managers, is not increasing and concave in w . This is seen in the right panel of Figure 10 where, for both high-growth and low-growth firms, $v(w)$ becomes decreasing beyond a certain point. This can happen because, as depicted in the central panel of Figure 10, the present value of payoffs to future agents, $f(w)$, is decreasing in the current promise. In the case of low growth firms there are two separate effects. A higher promise w tends to decrease the probability that the incumbent will be replaced for disciplinary reasons. And it also reduces the probability of replacing the agent in order to undertake growth. Both effects contribute to delaying the hiring of the next and subsequent managers, thereby reducing the present value of their future compensation. In the case of high growth firms, by definition, growth opportunities are undertaken whenever they appear, independently of w . Thus only the

²³The last term, $f(w)$, does not appear in the earlier contributions to the literature on optimal long-term contracts where there is a single agent and the 'liquidation' value of the firm is exogenous.

first effect is present. This is the reason that the value $f(w)$ is less sensitive to changes in w in the high growth case than in the low growth case. Note that as w increases from 10 to 30, $f(w)$ declines by about 5 for the high-growth firm and by about 9 for the low-growth firm.

The second-best values in Figure 10 can be compared to the corresponding first best values $(\mu + v^*)$ of 260.4 and 189.4, respectively. Agency costs amount to roughly 5% of first-best value for the high growth case and about 13% in the low growth case. That is, agency costs represent about fifteen months of expected cashflows for the high-growth firm and about thirty months of expected cashflows for the low-growth firm. In our framework, the total loss of value induced by the non-contractibility of cashflows has four distinct causes. First, as in previous studies of agency in a dynamic setting, under the optimal contract the firm will dismiss managers for disciplinary reasons following a series of poor cashflow realizations even though this is ex post inefficient. Second, there is an inefficiency due to the reliance on deferred compensation when managers are more impatient than firms, $\rho > r$. Third, under the optimal contract the firm will sometimes retain an incumbent manager and pass-up growth opportunities even though growth is efficient under first best. This is the component of agency costs due to managerial entrenchment, or ‘under-investment’. Finally, there is a more subtle form of agency costs, which is due to the fact that at the time of agreeing a contract with an incoming manager the firm does not take into account the spill-over effect on the timing of future managers’ hiring.

Specifically, if the current agent will be succeeded by future agents at stochastic stopping times τ_i , $i = 1, 2, 3, \dots$, the expected present values of the amounts they will receive, $\mathbb{E}[e^{-\rho\tau_i}\Phi_{\tau_i}w_0]$, is affected by the retention policy currently in place since it affects the distribution of stopping times. In low growth firms especially, by not taking into account the value of payoffs to future agents, the optimal second-best contract results in excessive retention relative to the optimal constrained planner solution. As noted in the Introduction, this effect is absent in the previous literature.

To get a sense of the possible significance of omitting this spill-over effect, we have solved for the constrained planner’s problem, where the objective includes the expected discounted payoffs to future managers. This can be done along the lines of Section 3.1 with the continuation values $\ell_{t,N}$ and $\ell_{t,G}$ augmented by $e^{-\rho}w_0$ and $e^{-\rho}(1 + \gamma)w_0$, respectively. When we do so for the low-growth benchmark case above, we obtain a total firm value of about 188, i.e., just 1% less than the first-best value and significantly greater than the second-best value. The main reason for this large increase in value is that for the parameters chosen under the constrained optimal planner’s solution all investment opportunities are taken up, i.e., with this change in policy the firm becomes a high-growth firm. However, even when applying the same comparison to the firm that is high-growth under the second-best policy, the improvement in firm value is noticeable. It goes from about 247.3 under the second best policy to 258.4 under the constrained planners’ policy which is again less than 1% less than first-best.

6 When growth opportunities are non-verifiable

In this section, we consider an extension of our baseline model where the incumbent manager is privately informed about the arrival of a growth opportunity, i.e., θ is only observable by the

manager.²⁴ This corresponds to situations where the incumbent manager knows what transformations could improve the firm's future prospects, but is also aware that he would be unable to implement these transformations himself. Analyzing the optimal second-best contract in this extended environment clarifies the extent to which our no-severance result in Section 3 relies on the contractibility of growth opportunities.

We are looking for the optimal contract that implements truth telling, i.e., under which the incumbent truthfully reports not only cashflows but also the arrival of a growth opportunity. The additional truth telling constraint enters in the definition of the value function $b_t^q(w)$ given $b_{t,G}^\ell$ and $b_{t,N}^\ell$. Namely, we now have

$$b_t^q(w) = \max_{w_G, w_N} qb_{t,G}^\ell(w_G) + (1 - q)b_{t,N}^\ell(w_N),$$

subject to promise keeping $qw_G + (1 - q)w_N = w$, limited liability $w_\theta \geq 0$, and

$$w_G \geq w_N. \tag{21}$$

The inequality constraint (21) guarantees that the agent has no incentive to conceal the arrival of a growth opportunity lest it should result in his dismissal. It should be noted that the contingent continuation promises (w_G, w_N) as described in Proposition 2 typically violate incentive compatibility. Instead, we have the following result:

Proposition 3. *When the realization of θ_t is only observable by the manager, the continuation promises (w_G, w_N) under the optimal contract satisfy*

$$w_G = w_N = w \tag{22}$$

for any given post-cashflow promise w , and $b_t^q(w) = \mathbb{E}_\theta[b_{t,\theta}^\ell(w)]$. All other aspects of the recursive representation of the optimal contract are obtained along the lines of Section 3.

The principal would rather have set $w_G \leq w_N$ if the truth telling constraint (21) was removed, hence that constraint holds as an equality under the optimal contract. An immediate implication of Proposition 3 is the following:

Corollary 3. *When only managers can observe the realization of a growth opportunity, the managers of high growth firms (i.e., when $\delta_G < -1$) are replaced with probability one upon realization of a growth opportunity and receive a severance pay equal to their post-cashflow promise.*

Corollary 3 shows that positive severance can become an essential part of the optimal contract when growth opportunities are non-contractible. This is in contrast with the no-severance result obtained in the baseline model (Corollary 2). Note however that positive severance pay only arises in high growth firms. In low growth firms, dismissed managers leave the firm with zero severance under any circumstance. It should also be noted that the optimal contract does not set some fixed severance amount ex ante. Instead, severance is contingent on performance history and increases with past performance.

²⁴Effectively, the incumbent has private information about how the firm's expected productivity under his tenure compares relative to what it could be under new management. This feature is also present in Inderst and Mueller (2010), and Garrett and Pavan (2012).

We believe this result is particularly significant since in our model both the principal and the manager have risk neutral preferences. As discussed in Spear and Wang (2005), it is reasonable to expect severance pay to arise when managers are risk averse. Thus risk aversion and the non-contractibility of cash flows would represent two distinct and potentially self-reinforcing reasons for severance pay. The existing empirical literature on severance has not attempted to explore these issues. The study by Yermack (2006) provides evidence that the payment of severance packages to departing CEOs is a widespread practice. He interprets his findings as providing support for a “damage control” view of severance pay (as a means of avoiding dismissed managers making trouble for the firm under the management of their successors) as well as for “bonding theories” of severance pay (as a means of providing insurance to the managers for their human capital). One version of the latter theory is that of Inderst and Mueller (2010) who show that severance may be part of the optimal contract to discourage managers from concealing adverse information about their productivity. Our analysis suggests a related but alternative explanation to why departing managers are paid severance in the context of growth. When the incumbent CEO has privileged information about the creation of growth opportunities he might be tempted to hide the news about a value-improving transformation of the firm that would require a new manager. Severance serves to give him a stake in welcoming such growth.

To disentangle these alternative explanations, future empirical work will need to find adequate controls for risk aversion, the contractibility or not of growth opportunities, the significance of potential growth opportunities, and whether or not a growth opportunity was present when severance was paid. At the firm level, sorting observed turnover cases into those occurring in a growing firm and those absent growth should reveal clear differences in the frequency and the size of severance awards. At the industry level, and all else equal, the corresponding testable implication of the model is that fast growing industries should provide higher severance awards to their departing managers than slow growing industries.

7 When the incumbent can grow the firm

Our analysis so far proceeded under the maintained assumption that in order to pursue an opportunity to grow, the incumbent manager had to be replaced. In this section, we consider an environment where upon the arrival of a growth opportunity, the firm can grow either with a new manager or with the incumbent manager. Growth opportunities are contractible. The analysis clarifies under which circumstances our conclusions from the baseline model hold, and how they need to be modified in other cases.

We let χ^i denote the (size-adjusted) cost of taking the growth opportunity with the incumbent manager, and χ^n the cost of growing with a new manager.²⁵ The derivation of the optimal contract follows the same logic as in Section 3.1, except for the construction of b_G^ℓ .²⁶ The key novel feature of the optimal contract in the extended environment is that, whenever the firm retains an incumbent manager at a time a growth opportunity is available, it now needs to choose

²⁵We assume that $\gamma\mu/(e^r - 1) > \min(\chi^i, \chi^n + \kappa)$, so that the first-best policy in steady state involves taking all growth opportunities. Under first best, the firm grows with new managers if and only if $\chi^n + \kappa < \chi^i$.

²⁶For notational convenience, we drop all time subscripts in this section.

optimally whether to grow or not. Formally, we define

$$\bar{b}_G^\ell(w) = \max_{p,s,w^c} p(\ell_G - s) + (1-p)b^c(w^c) \quad (23)$$

subject to the promise keeping condition $ps + (1-p)w^c = w$ the limited liability condition $s \geq 0$, $w^c \geq e^{-\rho}\lambda\mu$, and $p \in [0, 1]$. We also define

$$\hat{b}_G^\ell(w) = \max_{p,s,w^c} p(\ell_G - s) + (1-p)[(1+\gamma)b^c(w^c) - \chi^i] \quad (24)$$

subject to the alternative promise keeping condition $ps + (1-p)(1+\gamma)w^c = w$. The value function \bar{b}_G^ℓ corresponds to the case where upon retaining its incumbent manager the firm does not take up the growth opportunity. The value function \hat{b}_G^ℓ corresponds to the alternative case where, if retained, the incumbent manager does implement the growth opportunity.²⁷ Note that ℓ_G , the continuation value upon replacement contingent on $\theta = G$, is generally defined as²⁸

$$\ell_G = \max\{e^{-r}(1+\gamma)b^y(w_0) - \kappa - \chi^n; e^{-r}b^y(w_0) - \kappa\}. \quad (25)$$

Whenever the cost of growing with a new manager χ^n is sufficiently small (relative to γ), if a new manager is hired at a time a growth opportunity is available, growth is implemented ($\ell_G > \ell_N$). For high values of χ^n , the firm never grows with a new manager ($\ell_G = \ell_N$).

7.1 When the incumbent never grows the firm

We start our analysis of the extended model by noting that under some circumstances the firm will never grow with an incumbent manager and that in this case the results of Sections 3-5 go through. Indeed if it is prohibitively costly to grow with an incumbent manager (χ^i very large), a firm would never choose to do so and would only ever grow with new managers — as long as the costs of doing so (captured by χ^n) are reasonably low. Our analysis of the baseline model directly applies to such configurations.

Proposition 4. *When χ^i is large, the firm never grows with an incumbent manager ($b_G^\ell = \bar{b}_G^\ell$). If moreover χ^n is relatively small, all the results of Section 3 apply.*

7.2 When the incumbent sometimes grows the firm

In the remainder of this section, we turn our attention to situations where the cost of growing with the incumbent χ^i is sufficiently low relative to the gains from growth, so that it can be optimal for the firm to sometimes grow with an incumbent manager.²⁹ Our next proposition

²⁷Note that in that case, the probability of managerial replacement $p_G(w)$, which appears as p in (24), no longer coincides with the probability of growing conditional on $\theta = G$.

²⁸In levels, we have $L_G(\Phi) = \max\{e^{-r}B^y[(1+\gamma)\Phi, w_0] - \Phi(\kappa + \chi^n); e^{-r}B^y(\Phi, w_0) - \Phi\kappa\}$ and $\hat{B}_G^\ell(\Phi, w)$ is obtained by maximizing $p\{L_G(\Phi) - \Phi s\} + (1-p)\{B^c[(1+\gamma)\Phi, w^c] - \chi^i\Phi\}$ over severance $s \geq 0$ adjusted for current size Φ , dismissal probability $p \in [0, 1]$, and continuation promise $w^c \geq 0$ adjusted for expanded size $(1+\gamma)\Phi$. The promise keeping condition is $p\Phi s + (1-p)(1+\gamma)\Phi w^c = \Phi w$. The definitions of the continuation value ℓ_G and value function \hat{b}_G^ℓ , both adjusted for current firm size, follow from homogeneity.

²⁹We focus on situations where $\hat{b}_G^\ell > \bar{b}_G^\ell$ everywhere, ignoring situations that could potentially arise where $\hat{b}_G^\ell(w) > \bar{b}_G^\ell(w)$ if and only if w is above some threshold.

describes the construction of the value function b_G^ℓ and the associated replacement and severance policies conditional on $\theta = G$ in such configurations. Note that the value function b_N^ℓ along with the policy functions $p_N(w)$, $s_N(w)$ and $w_N^c(w)$ are obtained along the lines of Proposition 1, as before.

Proposition 5. *When χ^i is low, the firm sometimes grows with incumbent managers. Taking the continuation value function b^c as given, let $\underline{w}_G \equiv (1 + \gamma)e^{-\rho}\lambda\mu$ and*

$$\hat{b}^c(w) \equiv (1 + \gamma)b^c\left(\frac{w}{1 + \gamma}\right) - \chi^i, \quad w \geq \underline{w}_G. \quad (26)$$

The probability of managerial turnover conditional on $\theta = G$ is

$$p_G(w) = \begin{cases} 1 - (w/\underline{w}_G), & 0 \leq w < \underline{w}_G, \\ 0, & w \geq \underline{w}_G. \end{cases} \quad (27)$$

Severance pay conditional on $\theta = G$ is $s_G(w) = 0, \forall w$, and the continuation value to the retained manager (scaled by next period size) is

$$w_G^c(w) = \begin{cases} \underline{w}_G/(1 + \gamma), & 0 < w < \underline{w}_G, \\ w/(1 + \gamma), & w \geq \underline{w}_G. \end{cases} \quad (28)$$

Finally

$$b_G^\ell(w) = \hat{b}_G^\ell(w) = \begin{cases} \ell_G + \delta_G w, & 0 \leq w < \underline{w}_G, \\ \hat{b}^c(w), & w \geq \underline{w}_G, \end{cases} \quad (29)$$

where the slope of \hat{b}_G^ℓ for $w < \underline{w}_G$ is given by $\delta_G \equiv \frac{\hat{b}^c(\underline{w}_G) - \ell_G}{\underline{w}_G}$.

Proposition 5 shows that when χ^i is sufficiently low, if a growth opportunity arises after a period of sustained good performance, the incumbent manager is retained and grows the firm for sure (Eq. 27). If instead the recent performance of the firm has been relatively poor, the incumbent manager is at risk of being dismissed. If he survives this threat, he is allowed to grow the firm. If not, he leaves the firm with zero severance. Whether or not new management grows the firm depends on their ability to implement the transformations that are required. If χ^n is high, the firm installs the new manager but passes up the available growth opportunity. When χ^n is relatively low, growth is implemented for sure, either with or without the incumbent. Note however, that growing with a new manager is inefficient relative to first best when $\chi^n + \kappa > \chi^i$.

Figure 11 depicts stationary value functions b_G^ℓ and b_N^ℓ for parameter values such that Proposition 5 applies, i.e., the firm grows with the incumbent manager if a growth opportunity arises after sustained good performance.³⁰ Threshold values are $\underline{w}_N = 8.41$, $\underline{w}_G = 9.25$, and $\bar{w} = 44.93$. In that example $\ell_G > \ell_N$, i.e., if turnover occurs at times a growth opportunity is available, the firm will grow with its new manager. Moreover, $\delta_N = 0.77 > \delta_G = 0.69$, which captures the fact that replacement is slightly more inefficient ex-post in the absence of growth.

To close the analysis, we characterize the adjustment of an agent's expected payoff to the arrival of a growth opportunity — which determines whether managers benefit or not from the

³⁰Here we assume $\chi^i = \chi^n = 2$, $\kappa = 7.5$, $\gamma = 0.1$, $q = 0.2$, and the other parameters are as in the benchmark case of Section 4.1, i.e., $\lambda = 0.9$, $e^r - 1 = 6.5\%$, $e^\rho - 1 = 7\%$, $\mu = 10$ and $w_0 = 14$.

arrival of a growth opportunity. The next proposition is analogous to Proposition 2, accounting for the fact that the marginal benefit of an increased promise to the manager conditional on the arrival of a growth opportunity is typically increased when the incumbent sometimes ends up implementing growth.

Proposition 6. *For a given post-cashflow promise w , the contingent continuation promises (w_G, w_N) in period t are characterized as follows.*

- (a) *for low χ^n so that $\delta_N > \delta_G$,*
- (i) *if $w < (1 - q)\underline{w}_N$, $w_G = 0$ and $w_N = \frac{w}{1-q}$;*
 - (ii) *if $(1 - q)\underline{w}_N \leq w < q\underline{w}_G + (1 - q)\underline{w}_N$, $w_G = \frac{w - (1-q)\underline{w}_N}{q}$ and $w_N = \underline{w}_N$;*
 - (iii) *if $q\underline{w}_G + (1 - q)\underline{w}_N \leq w \leq (1 + \gamma q)\bar{w}$, $w_G = \frac{1+\gamma}{1+\gamma q}w$ and $w_N = \frac{1}{1+\gamma q}w$;*
 - (iv) *if $w > (1 + \gamma q)\bar{w}$, any combination of w_G and w_N such that $w_G \geq (1 + \gamma)\bar{w}$, $w_N \geq \bar{w}$, and $qw_G + (1 - q)w_N = w$ is optimal.*
- (b) *for high χ^n so that $\delta_N < \delta_G$,*
- (i) *if $w < q\underline{w}_G$, $w_G = w/q$ and $w_N = 0$;*
 - (ii) *if $q\underline{w}_G \leq w \leq q\underline{w}_G + (1 - q)\underline{w}_N$, $w_G = \underline{w}_G$ and $w_N = \frac{w - q\underline{w}_G}{1-q}$;*
- and (iii) and (iv) of case (a) apply for higher values of w .*

The general logic that runs through Proposition 6 is the same as in Proposition 2, namely that in the optimal contract the principal puts his promise to the agent where it counts most. Figure 12 illustrates this logic under the assumption that the cost of growing with the new manager χ^n is low, so that $\delta_N > \delta_G$ and part (a) of Proposition 6 applies. For the purpose of this example, we take $\underline{w}_N = 10$, $\underline{w}_G = 14$, $q = .25$, and $\gamma = .4$. For low values of w , the entire promise is allocated to the no-growth state so as to minimize the chances of a relatively inefficient turnover in the absence of growth. However, once the post-cashflow promise w has risen sufficiently so that $w_N(w) = \underline{w}_N$, the agent will be retained for sure in the no-growth state. The concern then becomes to avoid inefficient turnover upon growth, and therefore any marginal increase in w is allocated to the growth state while keeping the no-growth promise pegged at \underline{w}_N . Finally when the post-cashflow promise has risen to $q\underline{w}_G + (1 - q)\underline{w}_N$, the incumbent will be retained for sure. Above that threshold, any promise to the agent is delivered in the form of continuation promises (w_G, w_N) that equalize the marginal cost to the principal across the growth and no-growth states, $b_G^{\ell'}(w_G) = b_N^{\ell'}(w_N)$.

Figure 13 illustrates Proposition 6 under the assumption of high χ^n so that $\delta_G > \delta_N$ and part (b) applies. In this configuration, turnover when $\theta = G$ is most inefficient, and for low values of w , the post-cashflow promise is allocated entirely to the growth state until it is sure that the incumbent will be retained for sure if a growth opportunity realizes. Then for higher w the incumbent's no-growth promise takes on positive values, and so forth. Notice that in case (b) where new managers are relatively bad at growing the firm, the arrival of a growth opportunity is good news for the incumbent, in contrast with what we found in the benchmark model.

Together with the dismissal thresholds \underline{w}_G and \underline{w}_N , the updating of the agent's promise conditional on θ outlined in Proposition 6 determines how the probability of managerial turnover is affected by the realization of a growth opportunity, for a given cashflow history. Figure 14

exhibits dismissal probabilities conditional on post-cashflow promise and on the realization or not of a growth opportunity. Note that in situations where the incumbent is at a comparative advantage at growing the firm (bottom panel of Figure 14), for low values of the post-cashflow promise, the arrival of a growth opportunity reduces the probability of dismissal. This is in contrast with the baseline model (Corollary 1).

7.3 Numerical example

To conclude this section, we illustrate by way of a simple numerical example the turnover, growth and compensation policies under the optimal contract, for the same parameter values as in Figure 11.³¹ The example illustrates under which circumstances the firm finds it optimal to grow with the incumbent manager, and how managerial turnover is affected by past and current cashflow realizations and the availability of a growth opportunity.

Table 1 presents the evolution of the contractual promise to the incumbent manager for a particular path of scale adjusted cashflows and growth opportunity realizations. At the beginning of the episode we consider, at $t = 6$, the manager is still running the firm at its initial size (normalized to one), and has accumulated a high promise as a consequence of sustained good performance. His promise w_N^ℓ is much higher than the dismissal threshold ($34.16 > 8.41$), but not high enough to warrant a bonus ($34.16 < 44.93$). Thus he continues into period $t = 7$ carrying a promise that has been augmented from previous period to take into account the manager's rate of time preference, ρ . A good cashflow realization leads to an upward adjustment of the agent's promise, and when a growth opportunity then presents itself, the promise is increased still further. Given the high promise level, the manager is retained and is allowed to grow the firm. Notice that his scale-adjusted promise is reduced (from $w_G^\ell = 42.75$ to $w^c = 38.86$) to reflect that in the future he will be running a larger firm and therefore will be facing a high expected cashflow implying higher compensation.³² Subsequently, the firm is operated at a scale of 1.1 and following another good cashflow in period $t = 8$ the agent has accumulated a sufficiently high promise to be awarded a bonus.

After period $t = 8$, the firm goes through several periods of sustained poor performance, and the manager starts period $t = 14$ with an expected discounted payoff $w^y = 12$. After another poor cashflow realization, his promise falls at a low point of $w^q = 6.60$. Inefficient termination is looming, and case (a-i) of Proposition 6 applies. If a growth opportunity arrived, the agent would be dismissed with certainty with zero severance; on the other hand, with a contingent continuation promise w_N^ℓ raised to 7.33 the manager has a higher chance of surviving the dismissal stage in case no growth opportunity arises, i.e., the most inefficient form of turnover is made less likely. In our example, no growth opportunity arises in period $t = 14$, but the agent's promise is still below the dismissal threshold $\underline{w}_N = 8.41$, and therefore he is at risk of being fired with no severance (with 13% chance). The challenged manager survives the dismissal threat and finds his promise increased to the no-growth dismissal threshold.

³¹As noted above, under these parameters replacement is more inefficient ex-post in the absence of a growth opportunity ($\delta_N > \delta_G$), and part (a) of Proposition 6 applies.

³²When the firm takes up a growth opportunity with an incumbent manager, his expected discounted payoff remains unchanged. The adjustment to the promise, which appears in Eq. (28), is merely due to the fact that the end-of-period promise w^c is scaled by next-period size, whereas w_G^ℓ is scaled by current size.

The firm performance in the next period ($t = 15$) is not good enough for the manager to be sure to keep his position ($w^q < qw_G + (1 - q)\underline{w}_N = 8.49$). Case (a-ii) of Proposition 6 now applies since $w^q > (1 - q)\underline{w}_N = 7.57$. If no growth opportunity had materialized in that period, the manager would have been safe ($w_N^l = \underline{w}_N$). However, given the realization of a second growth opportunity, he is again at risk of being fired (with 43% chance). The manager is dismissed and leaves the firm without severance pay after a tenure of 15 periods. A new manager is hired to run the firm at a size of 1.21 (indeed χ^n is relatively low, and it is therefore more beneficial for the firm to take up growth with its new manager than passing it up).

To summarize the insights from Section 7, in our model extended to allow for the endogenous choice of whether an incumbent or a new manager will grow the firm, we first find that when the cost of growing with the incumbent manager (χ^i) are sufficiently high, we recover all the results of the benchmark model where we assumed only a new manager is able to grow the firm (Proposition 4). Conversely, when the costs of growing with the incumbent are sufficiently low, then the incumbent may grow the firm, but only if his past performance has been sufficiently good (Proposition 5).

8 Conclusion

In this paper we explore the relationship between managerial incentive provision and growth in a dynamic agency model of the firm. In contrast with previous studies, we consider a long-lived firm with stochastic growth opportunities run by a sequence of managers over time. In this setting managerial turnover may occur not only to discipline management but also to facilitate growth.

Our framework produces new insights on managerial compensation, turnover, and firm growth. We show how optimal contracts in firms with growth opportunities can be implemented with a system of deferred compensation credit and bonuses that are similar to those found in practice. Firms with very good growth prospects tend to rely less on back-loading of compensation than firms with poor growth prospects. When growth opportunities are contractible, granting severance pay to dismissed managers is (at least weakly) suboptimal under the optimal contract because it has no agency cost-reducing benefit once the agent leaves the firm. However growth-induced turnover can result in positive severance if the principal needs to incentivize the manager to truthfully report the arrival of a growth opportunity. The growth trajectory of a firm depends on the severity of agency problems as well as the quality of its growth opportunities. When growth entails a change of management, growth can be forsaken after periods of good performance. When instead incumbent managers are able to implement growth, they only do so when past performance has been sufficiently good. When past performance has been poor, firms grow with new managers. Finally, we identify a new component of agency costs which relates exclusively to managerial turnover, which is due to the spillover effect of the length of an existing managerial contract onto the present value of all future contracts signed by the firm.

Our framework can be extended in multiple ways to study more extensively how growth opportunities affect managerial turnover and compensation. While in our framework the rate at which growth opportunities arise is exogenous, it would be interesting to explore optimal incentive provision when incumbents need to be given proper incentives to increase the likelihood of growth. Managers may also need to allocate their efforts between producing cashflows from

assets in place and developing new opportunities for growth. There may be a trade-off between these two types of activities in that they may both require top management time but also because they use different management skills.

Another possible extension of our analysis may be to incorporate time-variation in the arrival rate and quality of firms' growth opportunities, so as to better capture firms' life-cycles— from small start-ups to larger more mature firms. This extension may provide further guidance in bringing the model to the data for the study of managerial turnover, CEO compensation and firm growth, by distinguishing firms at different stages of growth.

A Proofs

Proof of Lemma 1: Non scale-adjusted value functions are defined recursively as follows. Given $B_{t+1,G}^\ell(\Phi, w)$ and $B_{t+1,N}^\ell(\Phi, w)$, we have

$$B_{t+1}^q(\Phi, w) = \max_{w_G, w_N \geq 0} qB_{t+1,G}^\ell(\Phi, w_G) + (1-q)B_{t+1,N}^\ell(\Phi, w_N), \quad (30)$$

subject to $qw_G + (1-q)w_N = w$. Then

$$B_{t+1}^y(\Phi, w) = \max_{\{w^q(y)\}_{y \in \mathcal{Y}}} \Phi\mu + \mathbb{E}_y\{B_{t+1}^q[\Phi, w^q(y)]\} \quad (31)$$

subject to promise-keeping condition $\mathbb{E}_y[w^q(y)] = w$, limited liability $w^q(y) \geq 0$, and incentive-compatibility constraint

$$w^q(y) \geq w^q(\tilde{y}) + \lambda(y - \tilde{y}), \quad \forall y \in \mathcal{Y}, \forall \tilde{y} \in [0, y]. \quad (32)$$

Note that the limited liability and incentive-compatibility constraints imply that B_{t+1}^y is only defined for $w \geq \lambda\mu$. Now, given B_{t+1}^y , we can define

$$B_t^e(\Phi, w) = e^{-r} B_{t+1}^y(\Phi, e^\rho w), \quad w \geq e^{-\rho} \lambda\mu \quad (33)$$

Next

$$B_t^c(\Phi, w) = \max_{C, w^e \geq 0} -C + B_t^e(\Phi, w^e) \quad (34)$$

subject to $C + \Phi w^e = \Phi w$. Note that the first argument in functions B^c and B^e is the beginning-of-next-period size, which has already been determined, and cash compensation C is not size-adjusted. We can also define

$$L_{t,N}(\Phi) = e^{-r} B_{t+1}^y(\Phi, w_0) - \kappa\Phi, \quad (35)$$

$$L_{t,G}(\Phi) = e^{-r} B_{t+1}^y((1+\gamma)\Phi, w_0) - (\kappa + \chi)\Phi, \quad (36)$$

and

$$B_{t,\theta}^\ell(\Phi, a) = \max_{p, S, w^c} p(L_{t,\theta}(\Phi) - S) + (1-p)B_t^c(\Phi, w^c) \quad (37)$$

subject to $pS + (1-p)\Phi w^c = \Phi w$, $S \geq 0$, $p \in [0, 1]$, and $w^c \geq e^{-\rho} \lambda\mu$. The homogeneity result and the definition of the scale-adjusted value functions as they appear in Section 3.1 follows directly from the observation that in the last period $B_T^q(\Phi, w) = -\Phi w$. Then given the homogeneity of B_T^q , the homogeneity of B_T^y follows, and homogeneity of earlier value functions obtains recursively. \blacksquare

Proof of Lemma 3: Our goal is to show how the concavity of b_t^e arises for $t < T - 1$. For that purpose, we need to go through the detailed construction of the value functions within period $T - 1$. Our starting point is that in the last period $b_T^y(w) = \mu - w$, for $w \geq \lambda\mu$, which in turn implies $b_{T-1}^e(w) = e^{-r}\mu - e^{\rho-r}w$, for $w \geq e^{-\rho}\lambda\mu$. Since the slope of b_{T-1}^e is strictly below -1 ,

the solution of the constrained maximization problem in (9) involves setting $w^e = e^{-\rho}\lambda\mu$ and $c = w - e^{-\rho}\lambda\mu$. Therefore,

$$b_{T-1}^c(w) = e^{-\rho}\lambda\mu + (1 - \lambda)e^{-r}\mu - w, \quad w \geq e^{-\rho}\lambda\mu. \quad (38)$$

We can now analyze $b_{T-1,N}^\ell$. The relevant continuation value upon replacement is

$$\ell_{T-1,N} = e^{-r}b^y(w_0) - \kappa = e^{-r}\mu - (e^{-r}w_0 + \kappa). \quad (39)$$

Note that $w_0 \geq \lambda\mu$ implies that $\ell_{T-1,N} < e^{-\rho}\lambda\mu + (1 - \lambda)e^{-r}\mu$, which in turn implies that $\delta_{T-1,N} > -1$ and $b_{T-1,N}^\ell$ is piecewise linear and globally concave, with a kink at $\underline{w}_{T-1,N} = e^{-\rho}\lambda\mu$. The same characterization applies to $b_{T-1,G}^\ell$ if $\delta_{T-1,G} > -1$; otherwise $b_{T-1,G}^\ell$ is simply linearly decreasing with slope -1 . Furthermore, note that $\ell_{G,T-1} > \ell_{N,T-1}$ implies $\delta_{T-1,G} < \delta_{T-1,N}$. Consider now the constrained optimization problem in (4). Given our previous characterization of $b_{T-1,N}^\ell$ and $b_{T-1,G}^\ell$, we know the maximum is reached (though not necessarily uniquely) by setting $w_G = 0$ and $w_N = w/(1 - q)$. Therefore we can write

$$b_{T-1}^q(w) = q\ell_{G,T-1} + (1 - q)b_{T-1,N}^\ell\left(\frac{w}{1 - q}\right). \quad (40)$$

This further implies that b_{T-1}^q is piecewise linear and globally concave, with slope $\delta_{T-1,N} > -1$ for $w < (1 - q)\underline{w}_{T-1,N}$ and slope -1 for $w > (1 - q)\underline{w}_{T-1,N}$, with a kink at $(1 - q)\underline{w}_{T-1,N}$. We now turn to the function b_{T-1}^y as defined in 5. Using Lemma 2, we can write

$$b_{T-1}^y(w) = \mu + \int b_{T-1}^q(w + \lambda(y - \mu))dF(y), \quad (41)$$

where F denotes the cumulative probability distribution of size-adjusted cashflows. Consider two promises w_A and w_B greater or equal to $\lambda\mu$, and for $\alpha \in (0, 1)$, define $w_C = \alpha w_A + (1 - \alpha)w_B$. Note that

$$\begin{aligned} & \alpha \int b_{T-1}^q(w_A + \lambda(y - \mu))dF(y) + (1 - \alpha) \int b_{T-1}^q(w_B + \lambda(y - \mu))dF(y) \\ &= \int [\alpha b_{T-1}^q(w_A + \lambda(y - \mu)) + (1 - \alpha)b_{T-1}^q(w_B + \lambda(y - \mu))]dF(y) \\ &\leq \int b_{T-1}^q[\alpha(w_A + \lambda(y - \mu)) + (1 - \alpha)(w_B + \lambda(y - \mu))]dF(y) \\ &= \int b_{T-1}^q[(\alpha w_A + (1 - \alpha)w_B) + \lambda(y - \mu)]dF(y). \end{aligned}$$

Therefore $\alpha b_{T-1}^y(w_A) + (1 - \alpha)b_{T-1}^y(w_B) \leq b_{T-1}^y(w_C)$, and b_{T-1}^y is concave. Further inspection shows that b_{T-1}^y is strictly concave for $w < (1 - q)\underline{w}_{T-1,N} + \lambda\mu$, and decreases linearly with slope -1 above that threshold. The concavity of b_{T-2}^c follows directly. That concavity is preserved in earlier periods can be established using similar arguments. \blacksquare

Proof of Proposition 1: We drop time subscripts for notational convenience. Taking ℓ_θ , $b^c(\cdot)$ and $w_\theta \geq 0$ as given, we consider the constrained optimization problem

$$b_\theta^\ell(w_\theta) = \max_{p_\theta, s_\theta, w_\theta^c} p_\theta(\ell_\theta - s_\theta) + (1 - p_\theta)b^c(w_\theta^c) \quad (42)$$

subject to the promise keeping condition $p_\theta s_\theta + (1 - p_\theta)w_\theta^c = w_\theta$, the limited liability condition $s_\theta \geq 0$, $w_\theta^c \geq e^{-\rho}\lambda\mu$, and $p_\theta \in [0, 1]$.

One candidate solution consists in setting $p_\theta = 1$. Promise keeping then requires that $s_\theta = w_\theta$, and the value taken by the objective in (42) is $\ell_\theta - w_\theta$. Other candidate solutions involve setting $p_\theta < 1$. For $w_\theta < e^{-\rho}\lambda\mu$, the set of possible value for p_θ is bounded below by $1 - w_\theta/(e^{-\rho}\lambda\mu) > 0$. For values of p_θ below this threshold, the promise keeping constraint and the constraints $s_\theta \geq 0$, $w_\theta^c \geq e^{-\rho}\lambda\mu$ cannot all be satisfied. Given $p_\theta \in [\max\{1 - w_\theta/(e^{-\rho}\lambda\mu), 0\}, 1] \equiv \mathcal{P}(w_\theta)$ and for $0 \leq s_\theta \leq [w_\theta - (1 - p_\theta)e^{-\rho}\lambda\mu]/p_\theta$, we define

$$f(s_\theta|w_\theta, p_\theta) \equiv p_\theta(\ell_\theta - s_\theta) + (1 - p_\theta)b^c\left(\frac{w_\theta - p_\theta s_\theta}{1 - p_\theta}\right).$$

The upper bound on s_θ guarantees that $w_\theta^c = (w_\theta - p_\theta s_\theta)/(1 - p_\theta)$ remains above $e^{-\rho}\lambda\mu$. Differentiating $f(\cdot|w_\theta, p_\theta)$ with respect to s_θ shows that the sign of $f'(s_\theta|w_\theta, p_\theta)$ coincides with the sign of $-1 - b^c(w_\theta^c)$. Lemma 4 implies that $f'(0|w_\theta, p_\theta) \leq 0$, and $f'(0|w_\theta, p_\theta) < 0$ if $w_\theta/(1 - p_\theta) < \bar{w}$. When $p_\theta < 1$, it is therefore (at least weakly) optimal to set $s_\theta = 0$ and $w_\theta^c = w_\theta/(1 - p_\theta)$, and the objective in (42) is then equal to

$$p_\theta \ell_\theta + (1 - p_\theta)b^c\left(\frac{w_\theta}{1 - p_\theta}\right) \equiv h(p_\theta|w_\theta).$$

Differentiating with respect to p_θ gives

$$h'(p_\theta|w_\theta) = \ell_\theta - b^c\left(\frac{w_\theta}{1 - p_\theta}\right) + \frac{w_\theta}{1 - p_\theta}b^c\left(\frac{w_\theta}{1 - p_\theta}\right).$$

Whenever the optimal choice of p_θ results in an interior solution, the first-order optimality condition $h'(p_\theta|w_\theta) = 0$ is satisfied, which can be rewritten as

$$b^c\left(\frac{w_\theta}{1 - p_\theta}\right) = \frac{b^c\left(\frac{w_\theta}{1 - p_\theta}\right) - \ell_\theta}{\frac{w_\theta}{1 - p_\theta}}. \quad (43)$$

Now, given ℓ_θ and $b^c(\cdot)$, consider the auxiliary problem

$$\delta_\theta = \sup_{w \geq e^{-\rho}\lambda\mu} \frac{b^c(w) - \ell_\theta}{w}.$$

Graphically, this problem involves finding the value of $w \geq e^{-\rho}\lambda\mu$ for which the slope of the line relating the points with coordinates $(0, \ell_\theta)$ and $(w_\theta, b^c(w_\theta))$ is maximal. The first-order optimality condition is

$$b^c(w) = \frac{b^c(w) - \ell_\theta}{w}. \quad (44)$$

There are three possible configurations.

1. For high values of ℓ_θ , such that $\ell_\theta \geq b^c(\bar{w}) + \bar{w}$, then $(b^c(w) - \ell_\theta)/w < b^c(w)$ for any $w \geq e^{-\rho}\lambda\mu$. There is no finite optimum and $\delta_\theta \leq -1$.

2. For low values of ℓ_θ , such that $[b^c(e^{-\rho}\lambda\mu) - \ell_\theta]/(e^{-\rho}\lambda\mu) > b^c(e^{-\rho}\lambda\mu)$, the optimum is reached at the corner $w = e^{-\rho}\lambda\mu$ and $\delta_\theta > -1$. In this configuration, Eq. (15) implies $\underline{w}_\theta = e^{-\rho}\lambda\mu$.
3. For intermediate values of ℓ_θ , there exists a unique, interior optimum, and $\delta_\theta > -1$. In this configuration, the value \underline{w}_θ defined by Eq. (15) satisfies $e^{-\rho}\lambda\mu \leq \underline{w}_\theta < \bar{w}$ and

$$b^c(\underline{w}_\theta) = \frac{b^c(\underline{w}_\theta) - \ell_\theta}{\underline{w}_\theta}. \quad (45)$$

Graphically, this condition says that the line going through the points with coordinates $(0, \ell_\theta)$ and $(\underline{w}_\theta, b^c(\underline{w}_\theta))$ is tangent to the curve representing the function $b^c(w)$.

Suppose configuration 3 prevails. By inspection of (43) and (45), we see that whenever the optimal choice of p_θ is interior, it satisfies $w_\theta/(1 - p_\theta) = \underline{w}_\theta$, i.e.,

$$p_\theta = 1 - \frac{w_\theta}{\underline{w}_\theta}, \quad (46)$$

$$s_\theta = 0, \quad (47)$$

$$w_\theta^c = \underline{w}_\theta. \quad (48)$$

The value taken by the objective in (42) at this candidate solution is

$$\left(1 - \frac{w_\theta}{\underline{w}_\theta}\right) \ell_\theta + \frac{w_\theta}{\underline{w}_\theta} b^c(\underline{w}_\theta) = \ell_\theta + \frac{b^c(\underline{w}_\theta) - \ell_\theta}{\underline{w}_\theta} w_\theta = \ell_\theta + \delta_\theta w_\theta > \ell_\theta + w_\theta.$$

For $w_\theta \leq \underline{w}_\theta$, $1 - (w_\theta/\underline{w}_\theta)$ belongs to the interior of $\mathcal{P}(w_\theta)$, and since $\ell_\theta + \delta_\theta w_\theta > \ell_\theta + w_\theta$, the solution to (42) involves setting $(p_\theta, s_\theta, w_\theta^c)$ per (46)-(48). For $w_\theta > \underline{w}_\theta$ however, setting $p_\theta = 1 - w_\theta/\underline{w}_\theta$ would violate the condition $p_\theta \geq 0$. The best choice of $p_\theta < 1$ in that case is $p_\theta = 0$. Indeed, $h'(0|w_\theta)$ has the sign of $b^c(w_\theta) - (b^c(w_\theta) - \ell_\theta)/w_\theta$, which is negative. Hence for $w_\theta > \underline{w}_\theta$, the optimization problem in (42) boils down to either setting $p_\theta = 1$ and $s_\theta = w_\theta$, or setting $p_\theta = 0$ and $w_\theta^c = w_\theta$ — i.e., $b_\theta^\ell(w_\theta) = \max\{\ell_\theta - w_\theta, b^c(w_\theta)\}$. Since $\delta_\theta > -1$, we have

$$\begin{aligned} b^c(w_\theta) &= b^c(\underline{w}_\theta) + \int_{\underline{w}_\theta}^{w_\theta} b^c(w)dw = \ell_\theta - \delta_\theta \underline{w}_\theta + \int_{\underline{w}_\theta}^{w_\theta} b^c(w)dw \\ &> \ell_\theta - \underline{w}_\theta + \int_{\underline{w}_\theta}^{w_\theta} b^c(w)dw \\ &> \ell_\theta - w_\theta, \end{aligned}$$

where the last inequality follows from $b^c(w) \geq -1$. It is therefore optimal to set $p_\theta = 0$ and $w_\theta^c = w_\theta$.

Suppose configuration 2 prevails. Consider first values of w_θ such that $w_\theta \leq \underline{w}_\theta$, i.e., $w_\theta \leq e^{-\rho}\lambda\mu$. Recall that p_θ must satisfy $p_\theta \geq 1 - w_\theta/(e^{-\rho}\lambda\mu)$. Moreover the fact $[b^c(e^{-\rho}\lambda\mu) - \ell_\theta]/(e^{-\rho}\lambda\mu) > b^c(e^{-\rho}\lambda\mu)$ implies that $h'(p_\theta)$ evaluated at this corner is strictly negative. The best candidate solution with $p_\theta < 1$ can be written as per (46)-(48), noting that in the present configuration $\underline{w}_\theta = e^{-\rho}\lambda\mu$. Since the value taken by the objective at this candidate solution, $\ell_\theta + \delta_\theta w_\theta$, is greater than $\ell_\theta - w_\theta$, this is the solution to (42). For $w_\theta > \underline{w}_\theta$, the same reasoning as above proves that it is optimal to set $p_\theta = 0$ and $w_\theta^c = w_\theta$.

Suppose configuration 1 prevails. The fact that $(b^c(w) - \ell_\theta)/w < b^c(w)$ for any $w \geq e^{-\rho}\lambda\mu$ implies that for any w_θ and for any $p_\theta \in \mathcal{P}(w_\theta)$, $h'(p_\theta|w_\theta) > 0$. Moreover $h(p_\theta|w_\theta)$ goes to $\ell_\theta - w_\theta$ as p_θ goes to 1. It is therefore optimal to set $p_\theta = 1$ and $s_\theta = w_\theta$, for any value of w_θ .

Part (i) of Proposition 1 summarizes the optimal solution in configurations 2 and 3, whereas part (ii) describes the solution in configuration 1. ■

Proof of Proposition 2: We drop time subscripts for notational convenience and define the function $V_w(w_G) = qb_G^\ell(w_G) + (1 - q)b_N^\ell[w_N(w_G, w)]$, where

$$w_N(w_G, w) = \frac{1}{1 - q}(w - qw_G). \quad (49)$$

For any $w \geq 0$, we consider the problem

$$\max_{w_G \in [0, \frac{w}{q}]} V_w(w_G). \quad (50)$$

Note that $V_w'(w_G)$ has the sign of $b_G^{\ell'}(w_G) - b_N^{\ell'}[w_N(w_G, w)]$. Consider first the case where $\delta_G > -1$ and $\underline{w}_G < \infty$, as depicted in Figure 4. For $w < (1 - q)\underline{w}_G$, $V_w'(0) < 0$; indeed $w_N(0, w) = w/(1 - q) < \underline{w}_G$ and therefore $b_N^{\ell'}(w_N(0, w)) > \delta_G$. Hence we have the corner solution $w_G = 0$ and $w_N = w/(1 - q)$. For $w \geq (1 - q)\underline{w}_G$, the first-order optimality condition $V_w'(w_G) = 0$ is satisfied at $w_G = w$. Indeed $w_N(w, w) = w$, and $b_G^{\ell'}(w) = b_N^{\ell'}(w)$ since b_G^ℓ and b_N^ℓ both coincide with b^c in that range. Setting $w_G = w_N = w$ is the unique solution when $w \in [(1 - q)\underline{w}_G, \bar{w}]$ since V_w is strictly concave over that range. However for $w > \bar{w}$, the maximum of V_w is reached at any $w_G \geq \bar{w}$ such that $w_N(w_G, w) \geq \bar{w}$. This comes from the fact that b^c is linear over that region. Consider now the case where $\underline{w}_G = \infty$ and b_G^ℓ decreases linearly with slope -1 . This case is as depicted in Figure 3. For $w < (1 - q)\bar{w}$, $V_w'(0) < 0$; indeed $w_N(0, w) < \bar{w}$ and therefore $b_N^{\ell'}(w_N(0, w)) > -1$. Hence we have the corner solution $w_G = 0$ and $w_N = w/(1 - q)$. However for $w > (1 - q)\bar{w}$, the maximum of V_w is reached at any $w_G \geq 0$ such that $w_N(w_G, w) \geq \bar{w}$. ■

Proof of Corollary 3: By definition of a high growth firm, part (b) of Proposition 1 applies for $\theta = G$, i.e., $s_G(w_G) = w_G$ and we know from Proposition 3 that $w_G = w$. ■

Proof of Proposition 5: Consider the constrained optimization problem in (24). For given $w > (1 + \gamma)e^{-\rho}\lambda\mu$, the objective function evaluated at the candidate solution $p = 0$, $s = 0$ and $w^c = w/(1 + \gamma)$ is equal to $\hat{b}^c(w)$, where \hat{b}^c is defined in (26). Note that the lower bound of the domain of \hat{b}^c follows directly from the lower bound of the domain of b^c . All achievable payoffs are within the convex hull of $(0, \ell_G)$ and the payoff frontier \hat{b}^c . ■

Proof of Proposition 6: The argument of the proof relies crucially on the slopes of the value functions b_G^ℓ and b_N^ℓ . Under the assumption that χ^i is low so that $b_G^\ell = \hat{b}_G^\ell$, we note that for $w > \underline{w}_G$, $b_G^{\ell'}(w) = \hat{b}^c(w) = b^c(w/(1 + \gamma))$. Then we apply the same logic as in the proof of Proposition 2. ■

References

- Anderson, R.W. and K. Nyborg**, 2011, "Financing and Corporate Growth Under Repeated Moral Hazard," *Journal of Financial Intermediation* 20, 1-24 .
- Biais, B., T. Mariotti, G. Plantin and J.C. Rochet**, 2007, "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications," *Review of Economic Studies* 74, 345-390.
- Biais, B., T. Mariotti, J.C. Rochet and S. Villeneuve**, 2010, "Large Risks, Limited Liability and Dynamic Moral Hazard," *Econometrica* 78, 73-118.
- Casamatta, C. and A. Guembel**, 2010, "Managerial Legacies, Entrenchment and Strategic Inertia," *Journal of Finance*. 65, 2403-2436.
- Clementi G.L. and H. Hopenhayn**, 2006, "A Theory of Financing Constraints and Firm Dynamics," *Quarterly Journal of Economics* 121, 229-265.
- DeMarzo, P. and M. Fishman**, 2007a, "Agency and Optimal Investment Dynamics," *Review of Financial Studies* 22, 859-892.
- DeMarzo, P. and M. Fishman**, 2007b, "Optimal Long-term Financial Contracting," *Review of Financial Studies* 20, 2079-2128.
- DeMarzo, P., M. Fishman, Z. He, and N. Wang**, 2012, "Dynamic Agency Theory meets the Q-Theory of Investment," forthcoming, *Journal of Finance*.
- DeMarzo, P. and Y. Sannikov**, 2006, "Dynamic Agency and the Q-theory of Investment," *Journal of Finance* 61, 2681-2724.
- Eisfeldt, A. and C. Kuhnen**, 2012, "CEO Turnover in a Competitive Assignment Framework," Anderson School of Management, UCLA, working paper.
- Garrett, D. and A. Pavan**, 2012, "Managerial Turnover in a Changing World." Northwestern University, working paper.
- He, Z.**, 2008, "Optimal Executive Compensation when Firm Size Follows Geometric Brownian Motion", *Review of Financial Studies*. 22, 859-892.
- Inderst, R. and H. Mueller**, 2010, "CEO Replacement Under Private Information," *Review of Financial Studies* 23, 2935-2969.
- Jenter, D. and K. Lewellen**, 2011, "CEO Preferences and Acquisitions," mimeo.
- Kaplan, S.N. and B.A. Minton** , 2008, "How Has CEO Turnover Changed?" Chicago Booth, working paper.
- Kaplan, S., B. Sensoy and P. Stromberg**, 2008, "Should Investors Bet on the Jockey or the Horse?," *Journal of Finance* 64(1), 75-115.
- Malenko, A.** 2012, "Optimal Design of Internal Capital Markets," mimeo MIT Sloan.
- Martin, K.J. and J.J. McConnell** 1991 "Corporate Performance, Takeovers and Management Turnover," *Journal of Finance*. 46, 671-687.
- Mikkelson, W.H. and M.M Partch** 1997, "The Decline of Takeovers and Disciplinary Managerial Turnover," *Journal of Financial Economics*. 44, 205-228.

- Murphy, K.J.** 1985 “Corporate Performance and Managerial Remuneration: An Empirical Analysis” *Journal of Accounting and Economics* 7, 11-42.
- Murphy, K.J.** 2001 “Performance Standards in Incentive Contracts,” *Journal of Accounting and Economics*. 30, 245-278.
- Penrose, E.** 1959, *The Theory of the Growth of the Firm*. Oxford: Basil Blackwell.
- Philippon, T. and Y. Sannikov,** 2010, “Real Options in a Dynamic Agency Model,” mimeo NYU.
- Quadrini, V.** 2004, “Investment and Liquidation in Renegotiation-Proof Contracts with Moral Hazard,” *Journal of Monetary Economics* 51, 713-751.
- Roberts, D.J.** 2004, *The Modern Firm: Organizational Design for Performance and Growth, the Inaugural Clarendon Lectures in Management Studies*. Oxford: Oxford University Press.
- Spear, S. and C. Wang,** 2005, “When to Fire a CEO: Optimal Termination in Dynamic Contracts,” *Journal of Economic Theory* 120, 239-256.
- Williamson, O.** 1966, “Profit, Growth and Sales Maximization,” *Economica*. 33, 1-16.
- Yermack, D.** , 2006, “Golden Handshakes: Separation Pay for Retired and Dismissed CEOs,” *Journal of Accounting and Economics* 41, 237-256.

Table 1: Illustration of the optimal contract when the incumbent can grow the firm.

Period t	6	7	8	...	14	15
Size Φ_t	1	1	1.1	—	1.1	1.1
Promise w_t^y	30	36.55	41.58	—	12	9.00
Cashflow y_t	15	13	17	—	4	9
Promise w_t^q	34.50	39.25	47.88	—	6.60	8.10
Growth option θ_t	N	G	N	—	N	G
Promise $w_{t,\theta}^\ell$	34.16	42.75	47.41	—	7.33	5.30
Replacement proba p	0	0	0	—	0.13	0.43
Promise w^c	34.16	38.86	47.41	—	8.41	
Cash compensation c	0	0	2.48	—	0	
Promise w^e	34.16	38.86	44.93	—	8.41	

The table illustrates the tenure of a manager under the optimal contract discussed in Section 7.2, for a specific sequence of cashflows and growth opportunity realizations. Scaled cashflows are uniformly distributed on $\{0, 1, \dots, 20\}$, with $\mu = 10$. Other parameter values are $\lambda = 0.9$, $\rho = \log(1.07)$, $r = \log(1.065)$, $q = 0.2$, $\gamma = 0.1$, $\kappa = 7.5$, $\chi^i = \chi^n = 2$ and $w_0 = 14$.

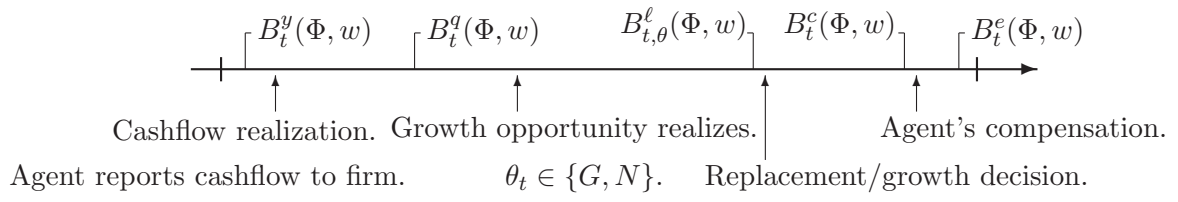


Figure 1: Intra-period timing.

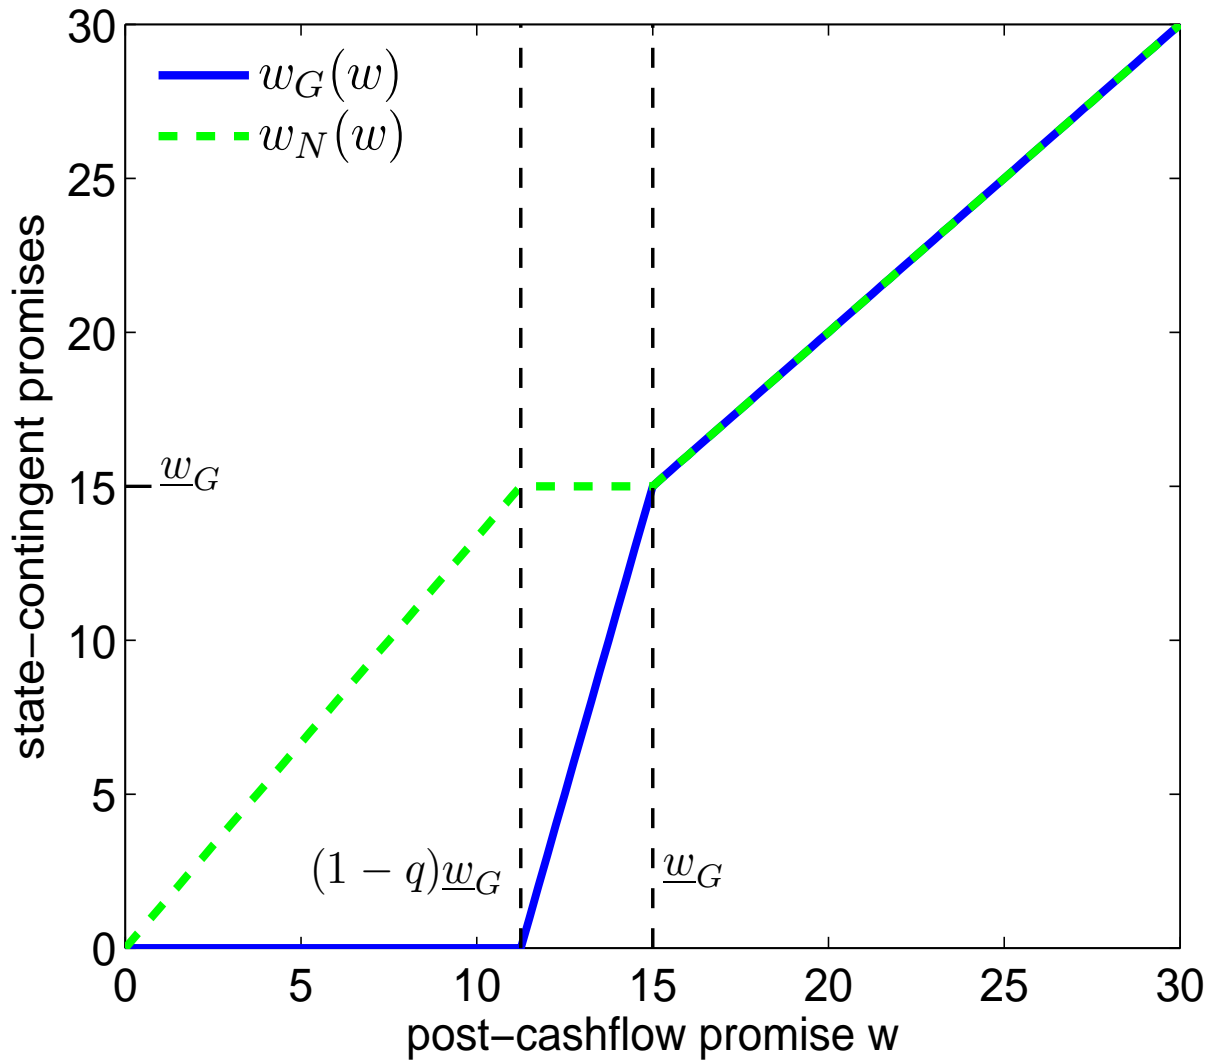


Figure 2: Promises contingent on realization of growth opportunity: low-growth benchmark. The figure is drawn assuming $\underline{w}_G = 15$, $\bar{w} = 25$ and $q = 0.25$.

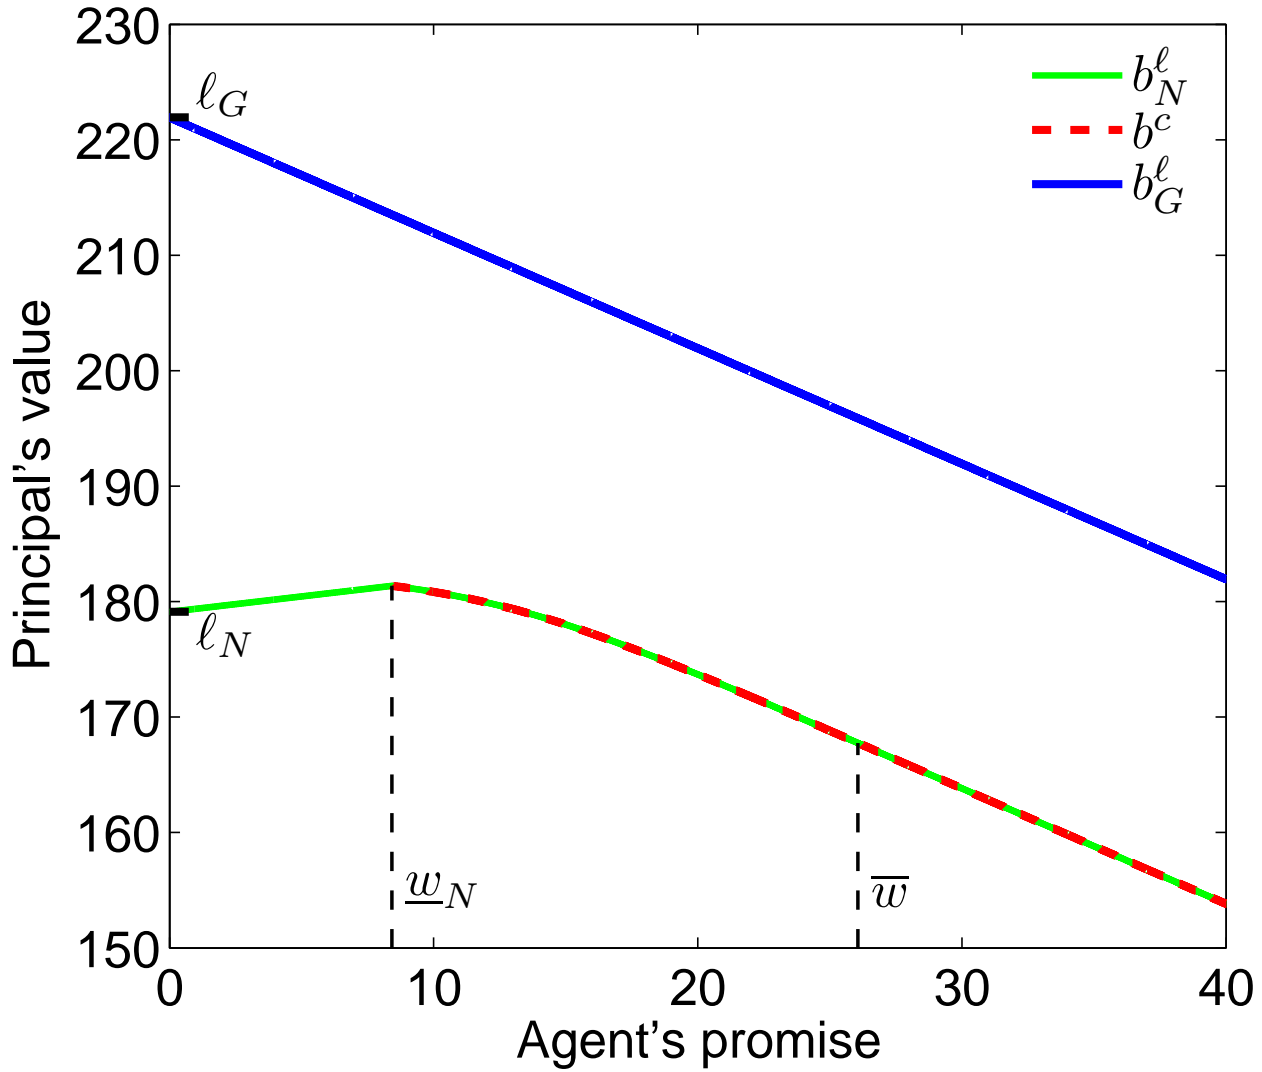


Figure 3: Value functions for high-growth firm benchmark.

Scaled cashflows are uniformly distributed on $\{0, 1, \dots, 20\}$, with $\mu = 10$. Other parameter values are $\lambda = 0.9$, $\rho = \log(1.07)$, $r = \log(1.065)$, $q = 0.1$, $\gamma = 0.25$, $\kappa = 0.2$, $\chi = 2$ and $w_0 = 14$.

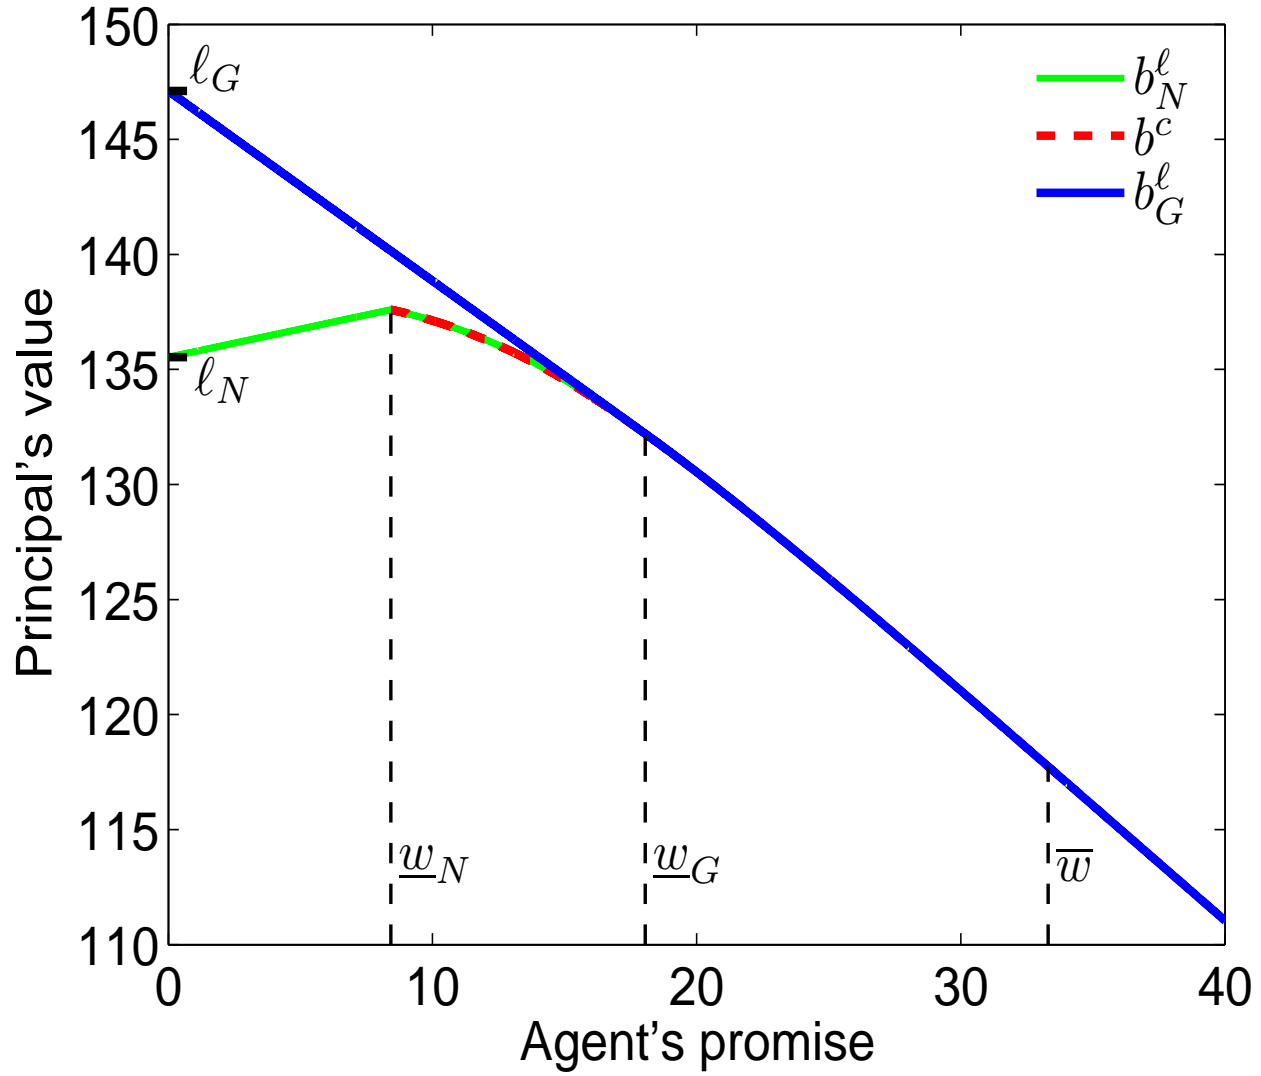


Figure 4: Value functions for low-growth firm benchmark.

Scaled cashflows are uniformly distributed on $\{0, 1, \dots, 20\}$, with $\mu = 10$. Other parameter values are $\lambda = 0.9$, $\rho = \log(1.07)$, $r = \log(1.065)$, $q = 0.1$, $\gamma = 0.1$, $\kappa = 0.2$, $\chi = 2$ and $w_0 = 14$.

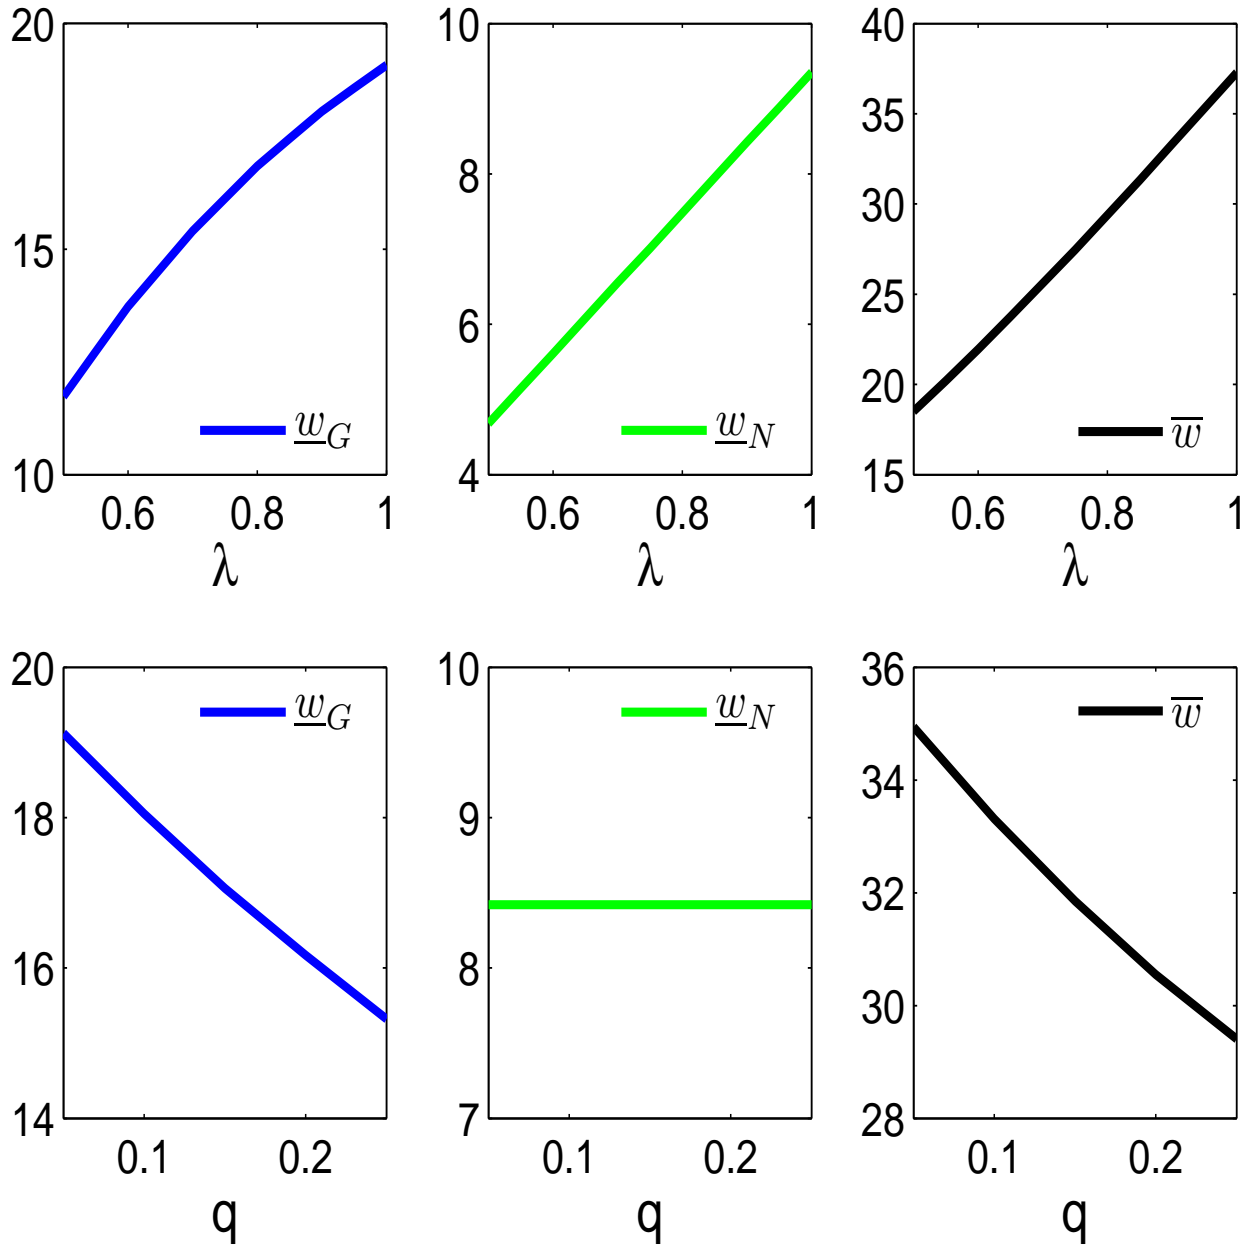


Figure 5: Threshold sensitivities.

Scaled cashflows are uniformly distributed on $\{0, 1, \dots, 20\}$, with $\mu = 10$. Default values for λ and q , when they are not used as variables, are $\lambda = 0.9$ and $q = 0.1$, respectively. Other parameter values are $\rho = \log(1.07)$, $r = \log(1.065)$, $\gamma = 0.1$, $\kappa = 0.2$, $\chi = 2$ and $w_0 = 14$.

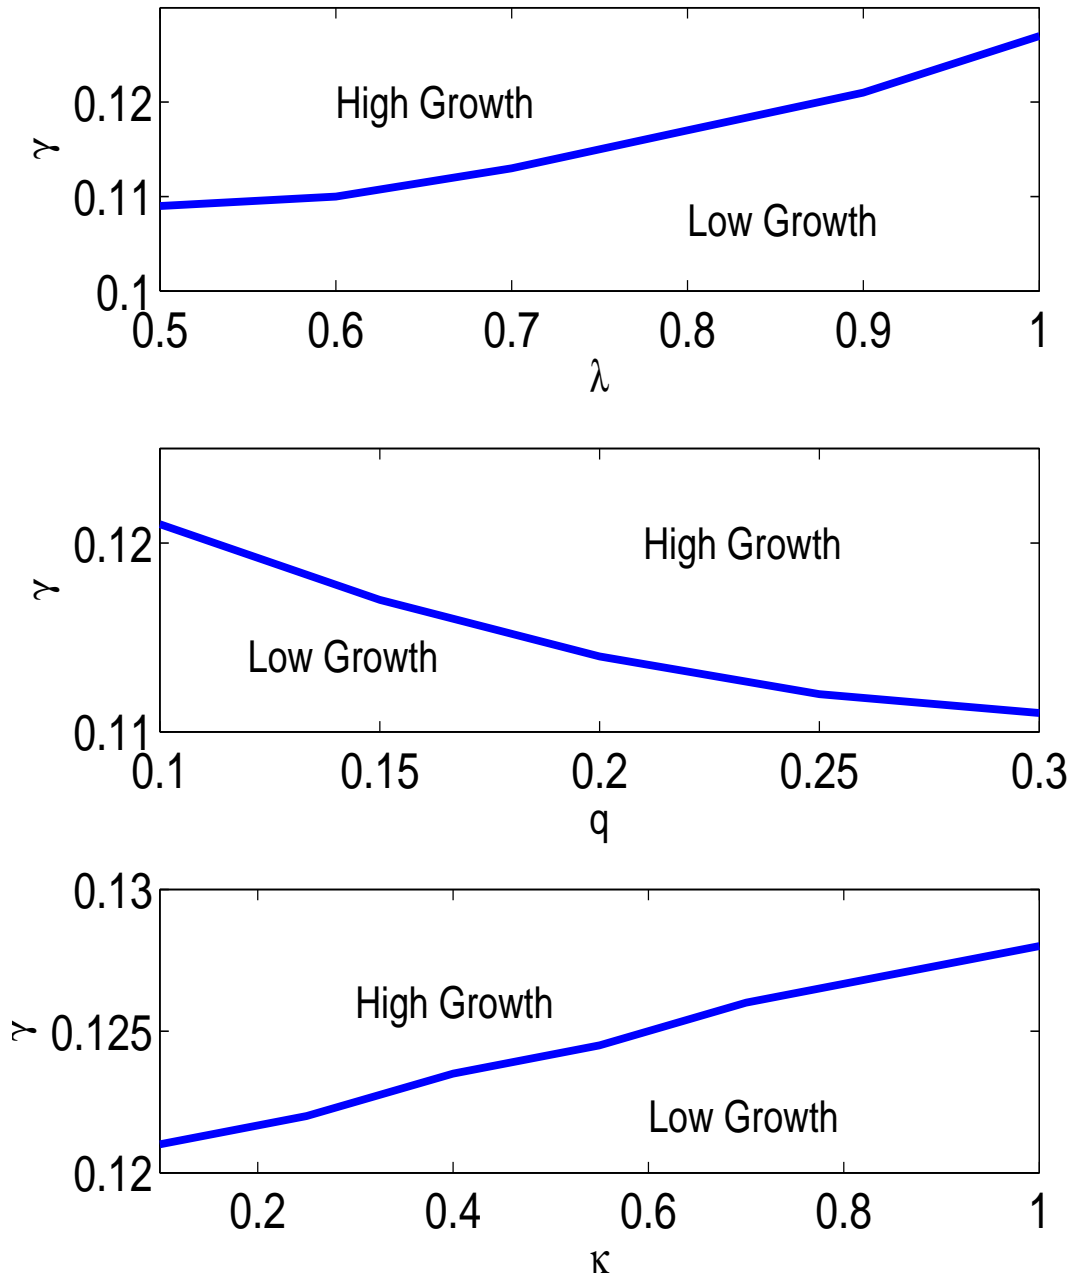


Figure 6: High-growth and low-growth regions of parameter space.

The figure represent the threshold value of γ on the frontier between the high- and low-growth regions, for different values of parameters λ , q and κ . Scaled cashflows are uniformly distributed on $\{0, 1, \dots, 20\}$, with $\mu = 10$. Default values for λ and q and κ are $\lambda = 0.9$, $q = 0.1$ and $\kappa = 0.2$, respectively. Other parameter values are $\rho = \log(1.07)$, $r = \log(1.065)$, $\chi = 2$ and $w_0 = 14$.

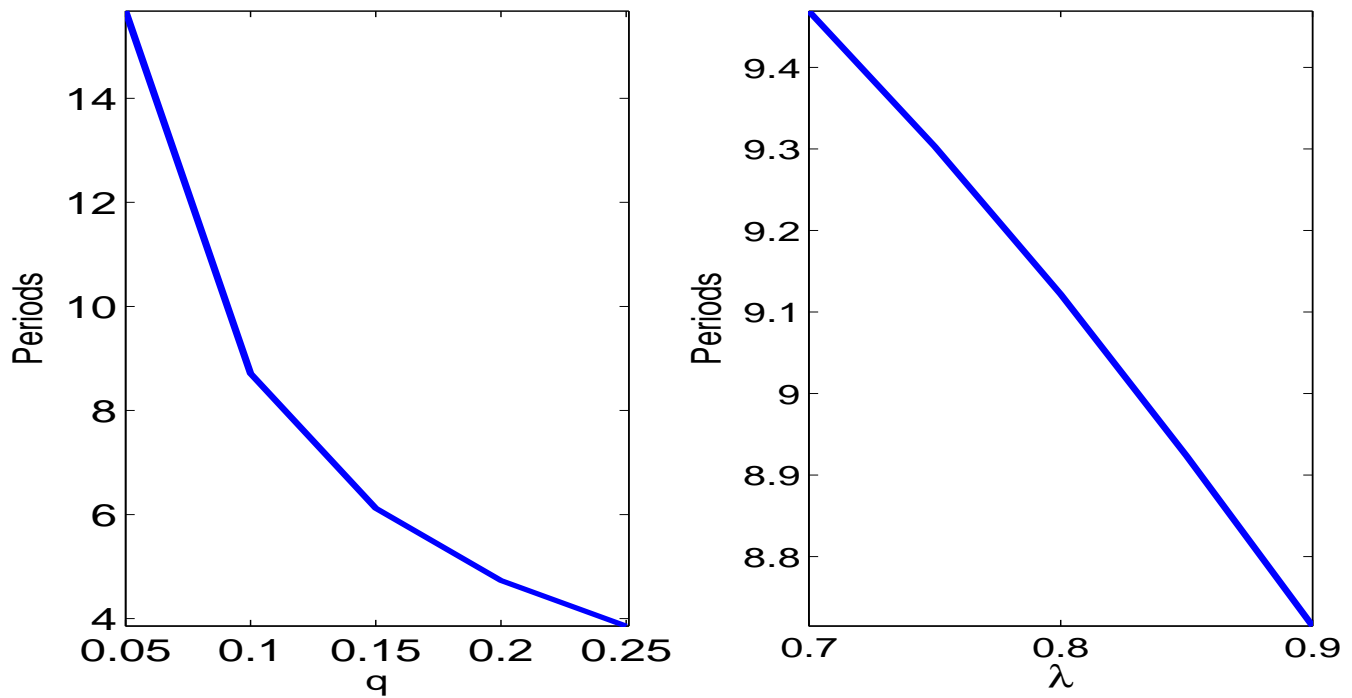


Figure 7: Average tenure in high-growth firms.

The average number of tenure years is obtained from simulating the optimal contract. Scaled cashflows are uniformly distributed on $\{0, 1, \dots, 20\}$, with $\mu = 10$. Default values for λ and q are $\lambda = 0.9$ and $q = 0.1$, respectively. Other parameter values are $\rho = \log(1.07)$, $r = \log(1.065)$, $\gamma = 0.25$, $\kappa = 0.2$, $\chi = 2$ and $w_0 = 14$.

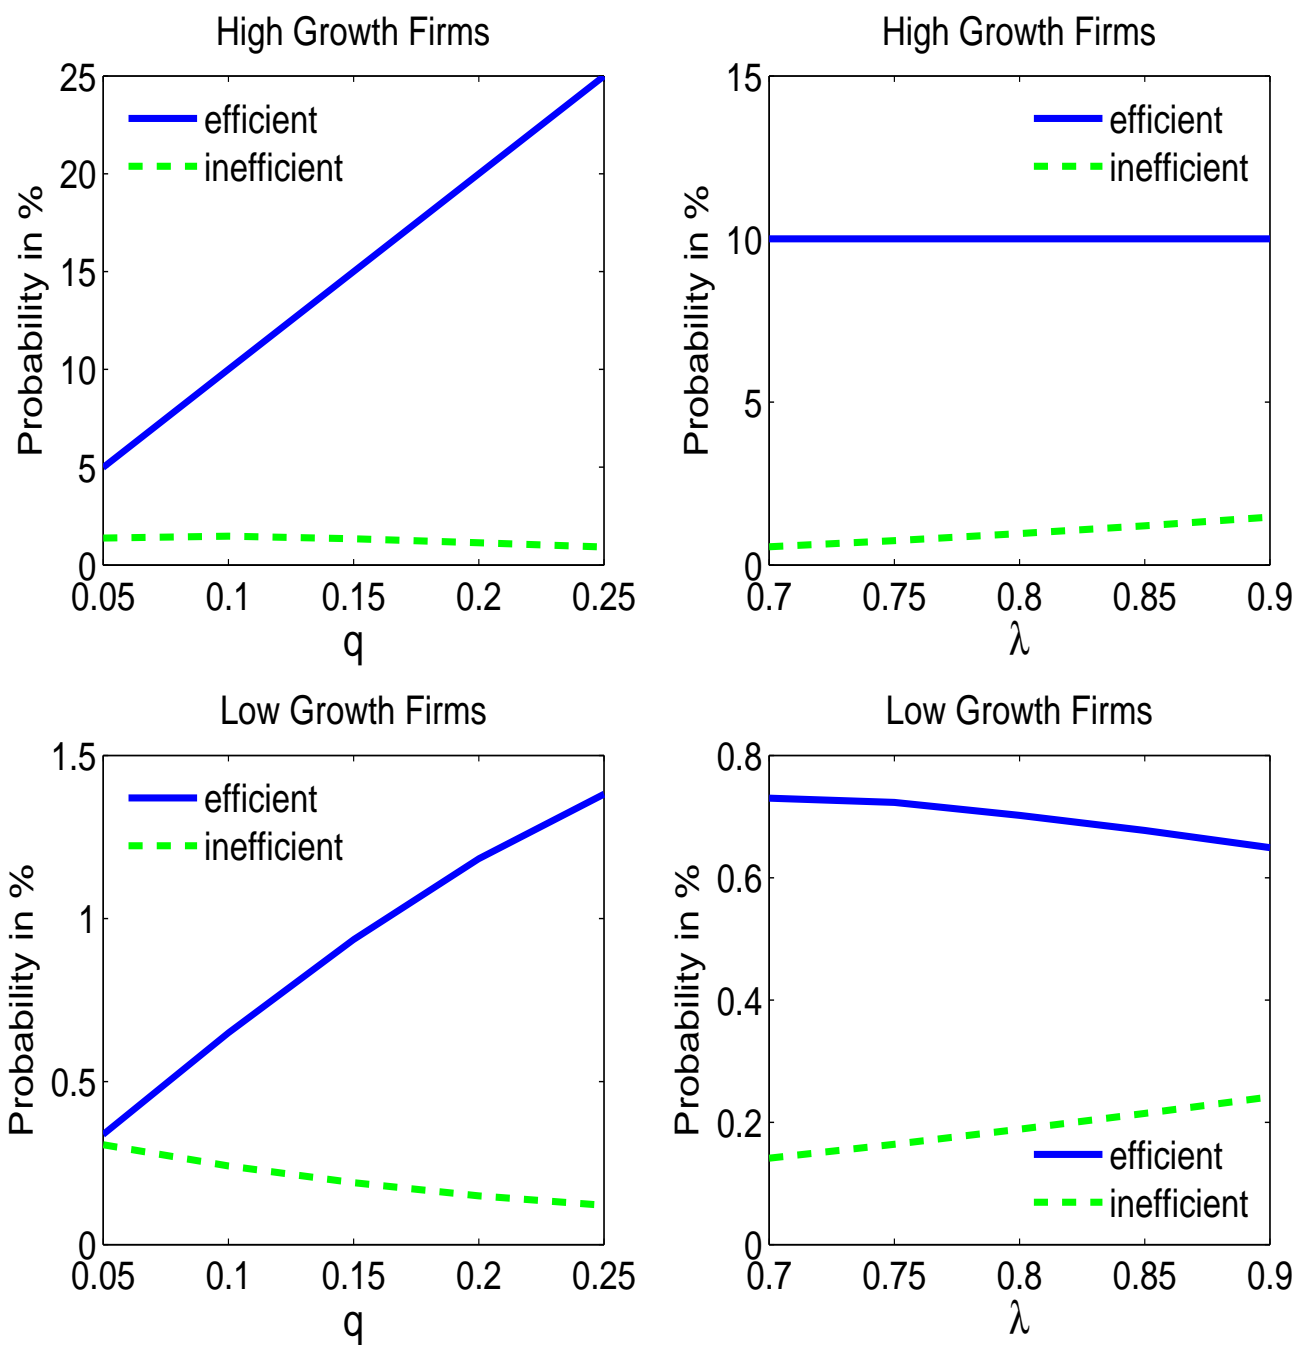


Figure 8: Average dismissal rates.

The unconditional probabilities of efficient and inefficient turnover are obtained from simulating the optimal contract. Scaled cashflows are uniformly distributed on $\{0, 1, \dots, 20\}$, with $\mu = 10$. Default values for λ and q are $\lambda = 0.9$ and $q = 0.1$, respectively. Top (bottom) panel is drawn for $\gamma = 0.25$ ($\gamma = 0.1$). Other parameter values are $\rho = \log(1.07)$, $r = \log(1.065)$, $\kappa = 0.2$, $\chi = 2$ and $w_0 = 14$.

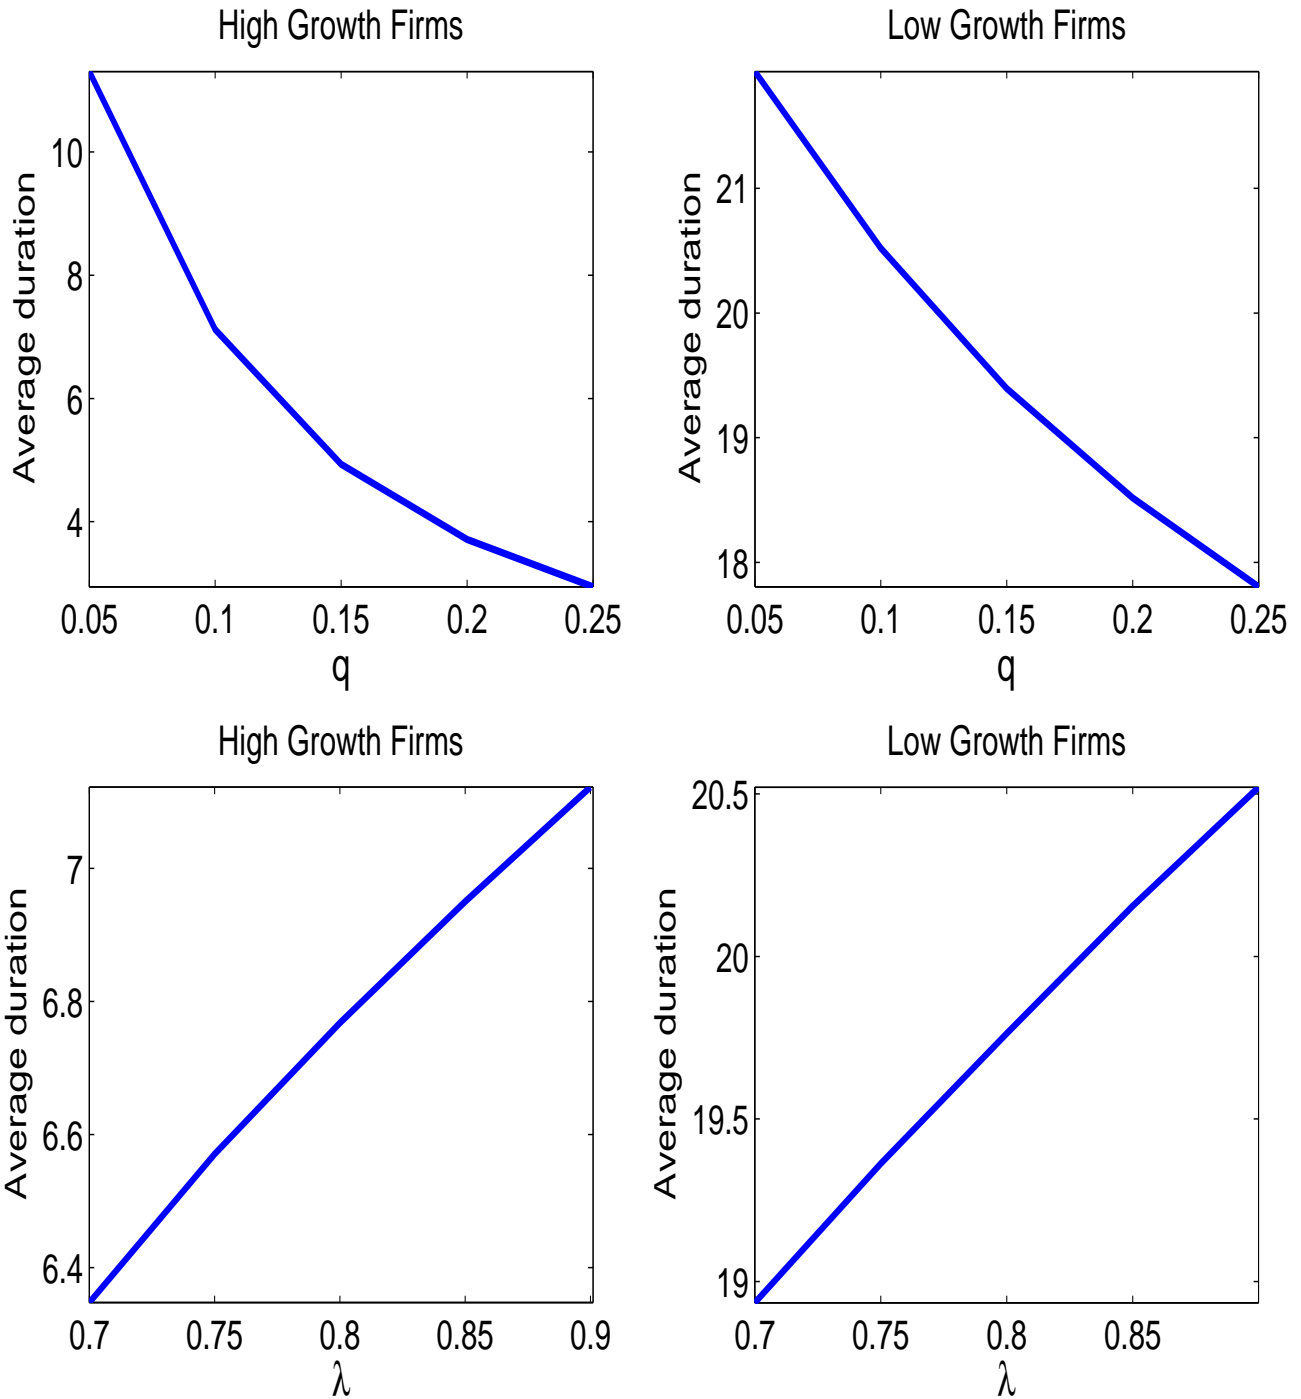


Figure 9: Compensation duration

The average compensation duration is obtained from simulating the optimal contract as explained in the text. Scaled cashflows are uniformly distributed on $\{0, 1, \dots, 20\}$, with $\mu = 10$. Default values for λ and q are $\lambda = 0.9$ and $q = 0.1$, respectively. Left (right) panel is drawn for $\gamma = 0.25$ ($\gamma = 0.1$). Other parameter values are $\rho = \log(1.07)$, $r = \log(1.065)$, $\kappa = 0.2$, $\chi = 2$ and $w_0 = 14$.

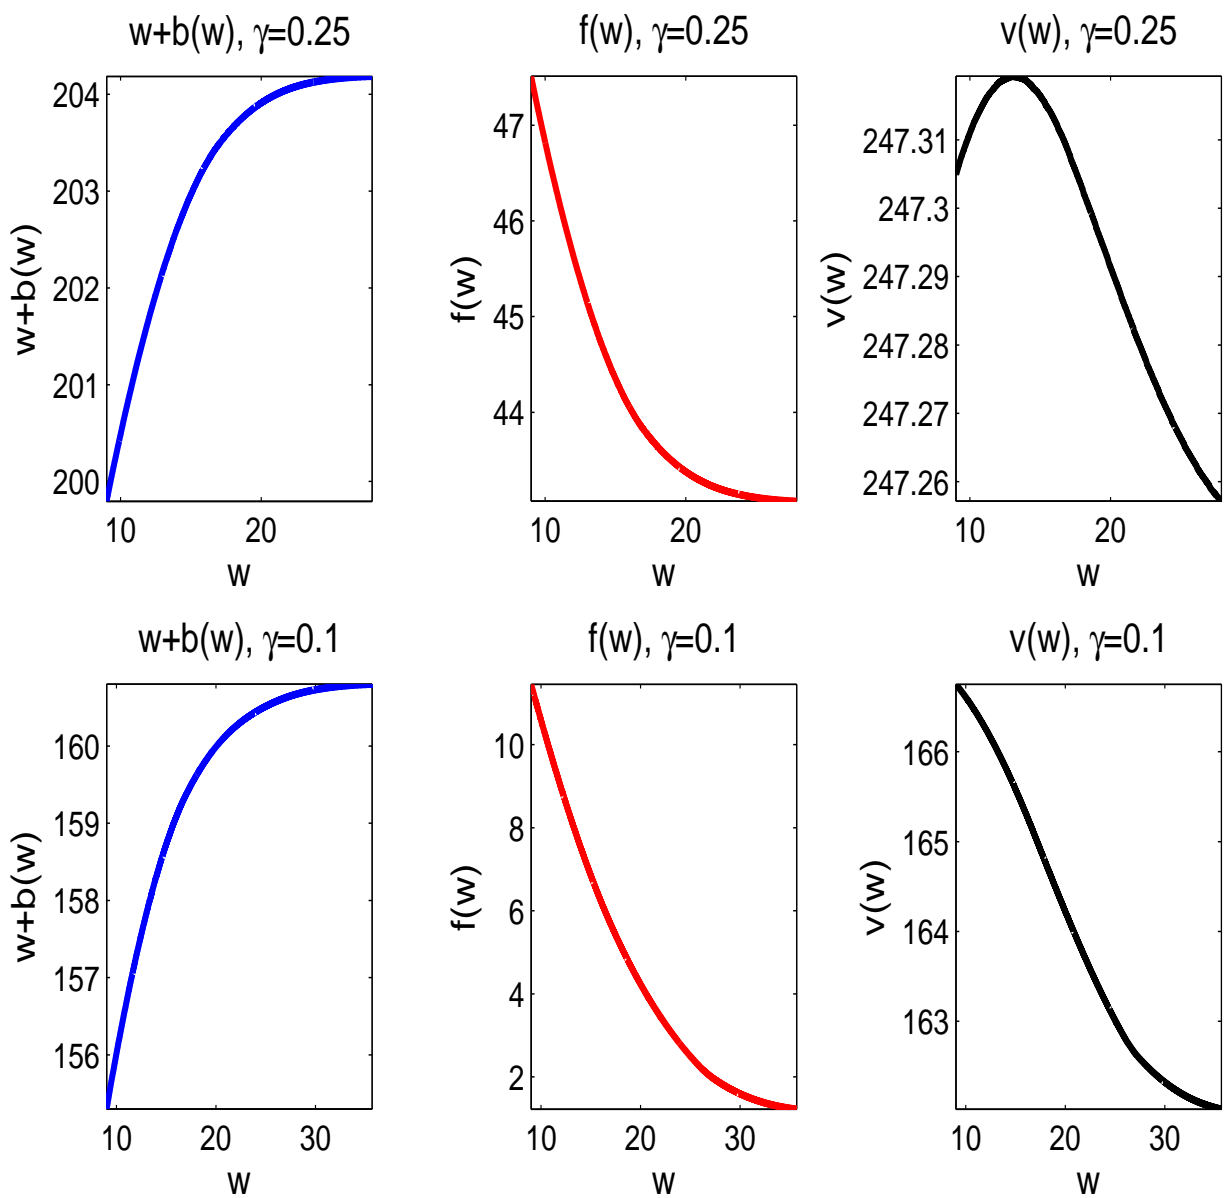


Figure 10: Second-best value for high- (top) and low-growth firms (bottom). The components of second-best firm value are described in the text. Top (bottom) panel is drawn for $\gamma = 0.25$ ($\gamma = 0.1$). Other parameter values are $\lambda = 0.9$, $\rho = \log(1.07)$, $r = \log(1.065)$, $q = 0.1$, $\kappa = 0.2$, $\chi = 2$ and $w_0 = 14$.

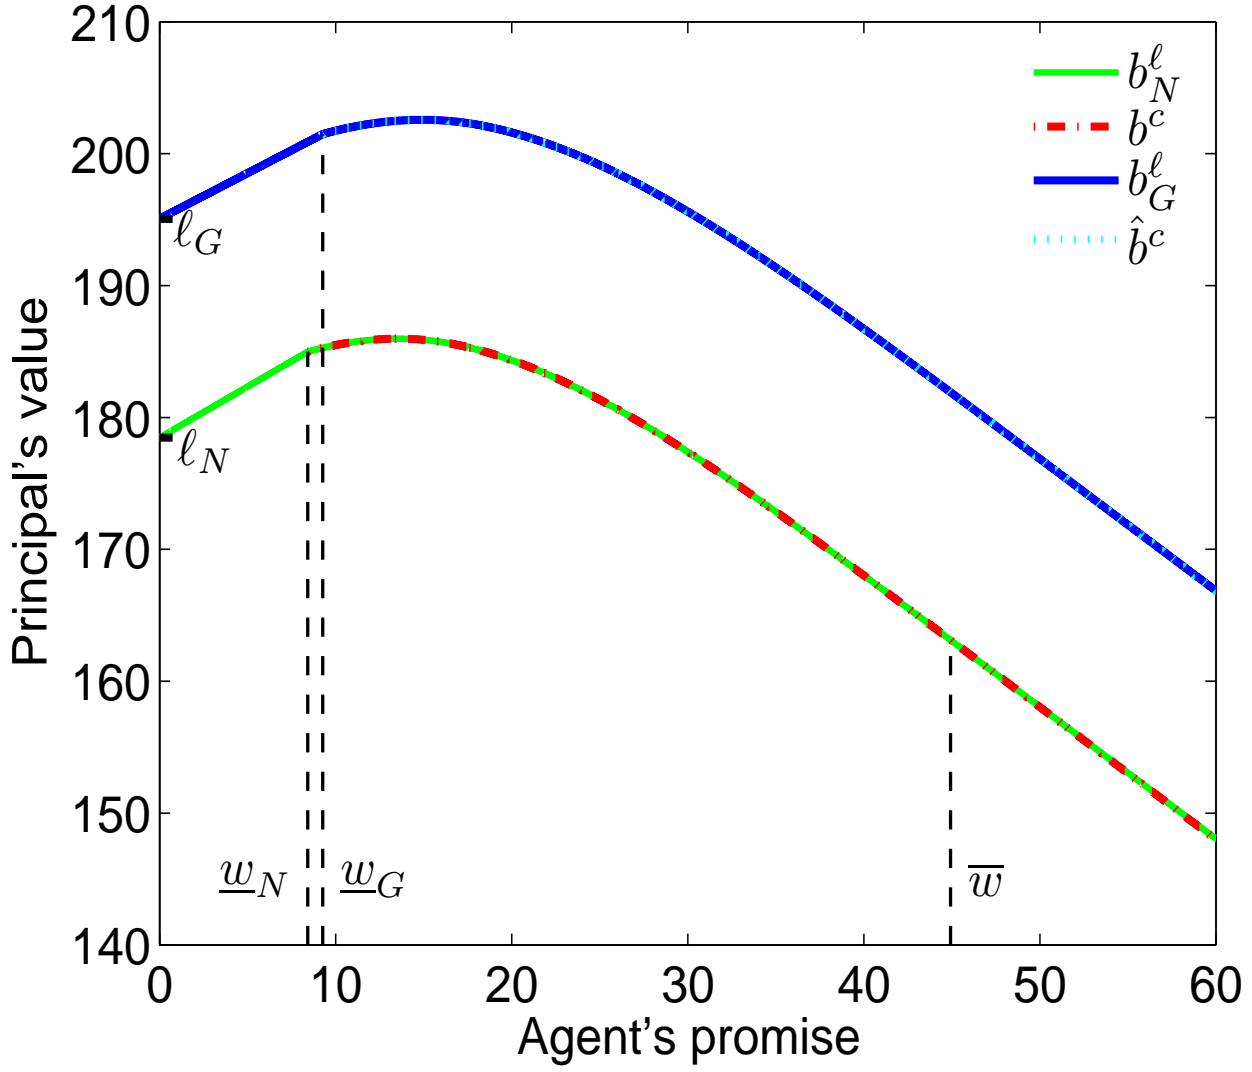


Figure 11: Extension: When the incumbent sometimes grows the firm.

The figure represents the value functions discussed in Section 7.2 in the stationary limit. Scaled cashflows are uniformly distributed on $\{0, 1, \dots, 20\}$, with $\mu = 10$. Other parameter values are $\lambda = 0.9$, $\rho = \log(1.07)$, $r = \log(1.065)$, $q = 0.2$, $\gamma = 0.1$, $\kappa = 7.5$, $\chi^i = \chi^n = 2$ and $w_0 = 14$.

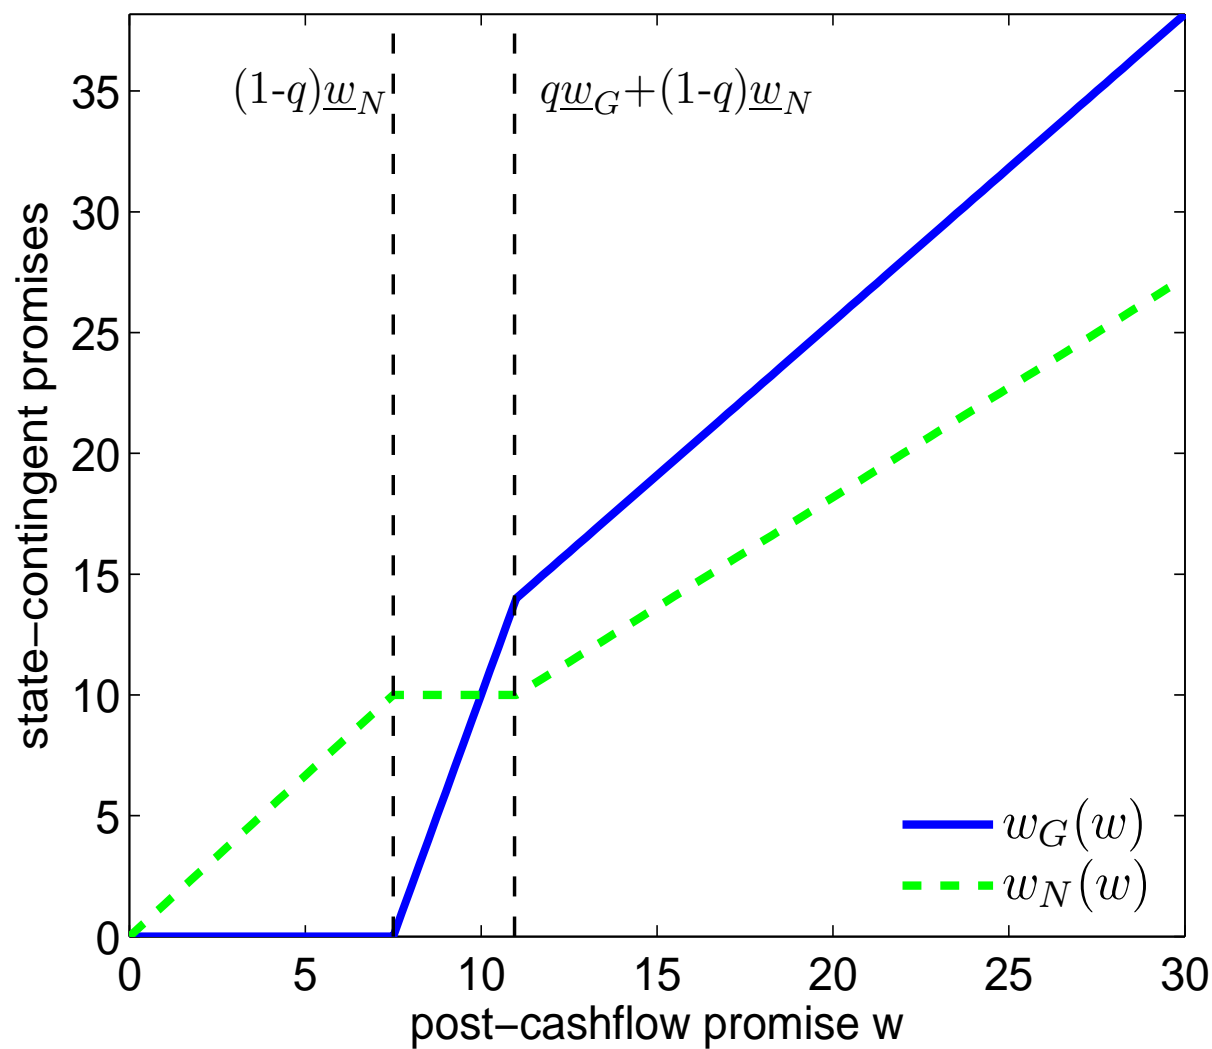


Figure 12: Promises contingent on realization of growth opportunity (case $\delta_N > \delta_G$). The figure illustrates Part (a) of Proposition 6. It is drawn for $\underline{w}_N = 10$, $\underline{w}_G = 14$, $\bar{w} = 25$ and $q = 0.25$.

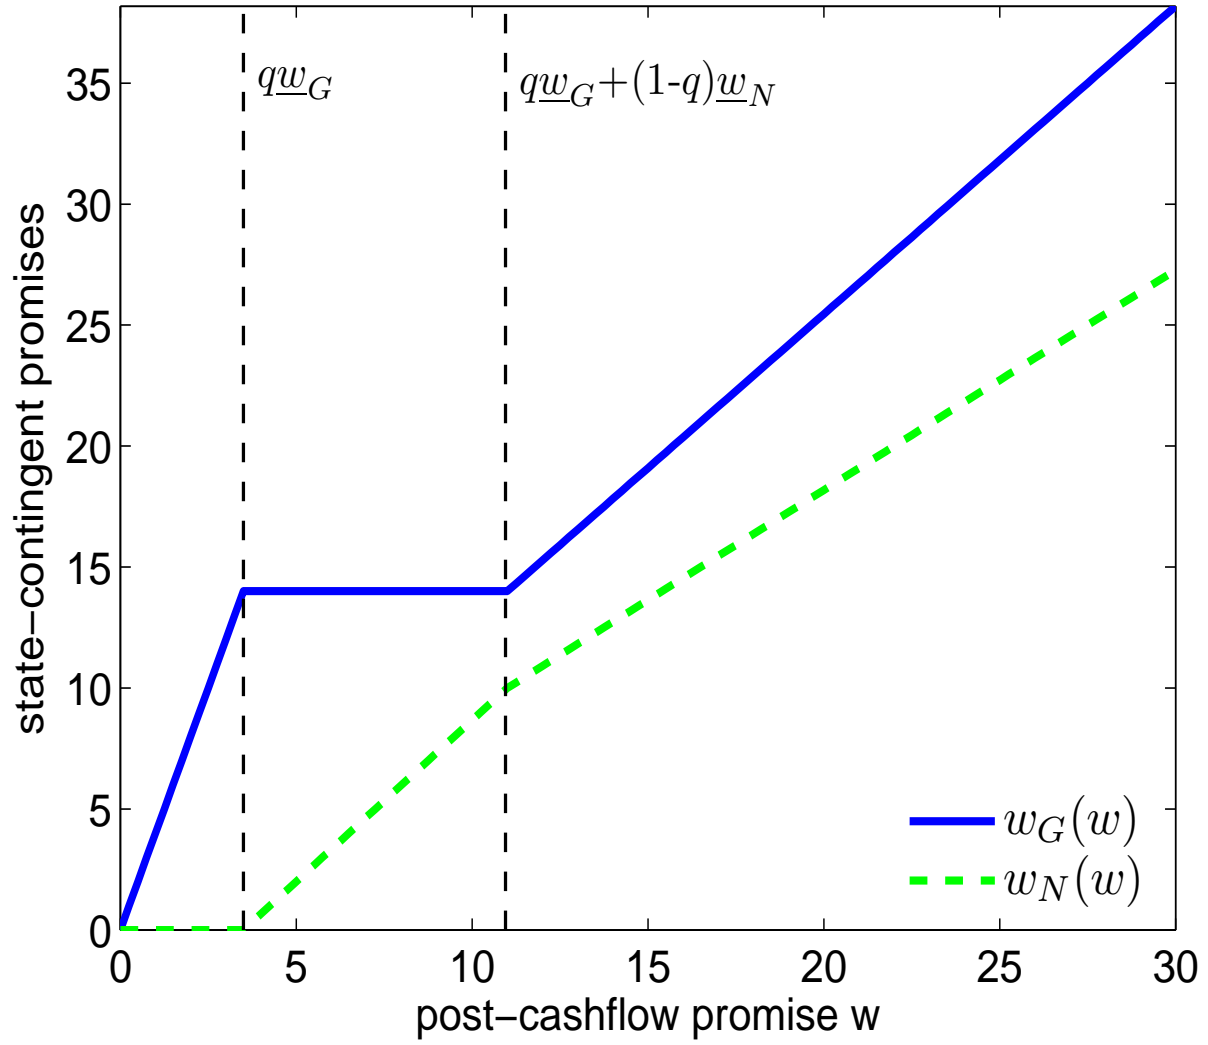


Figure 13: Promises contingent on realization of growth opportunity (case $\delta_N < \delta_G$). The figure illustrates Part (b) of Proposition 6. It is drawn for $\underline{w}_N = 10$, $\underline{w}_G = 14$, $\bar{w} = 25$ and $q = 0.25$.

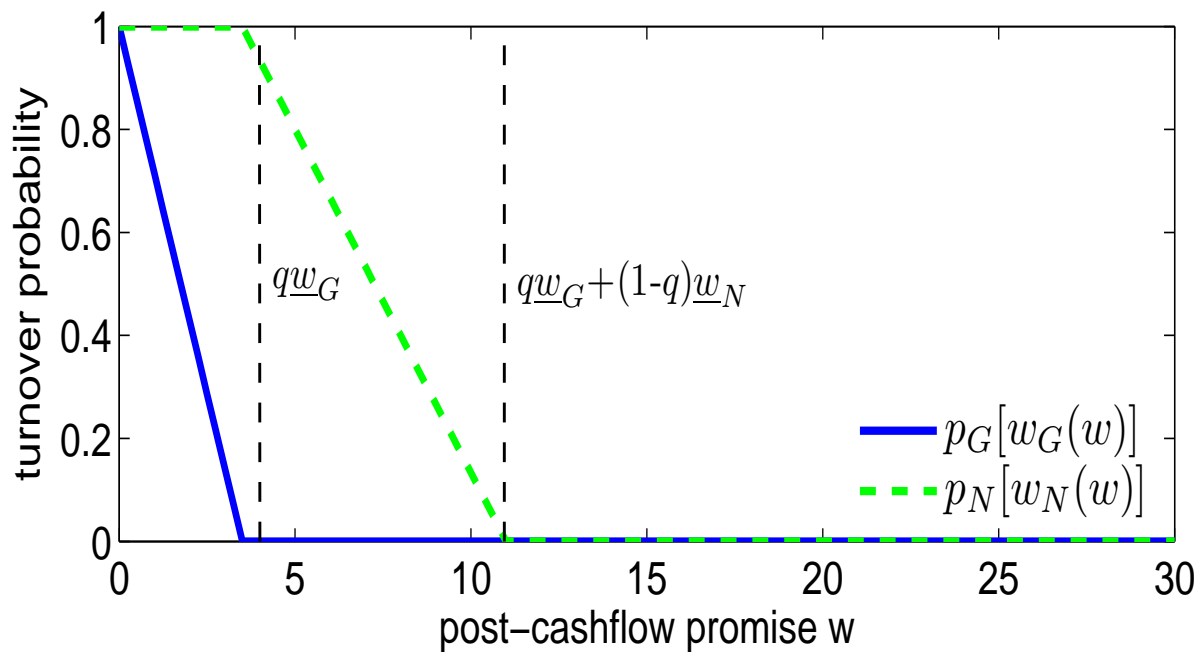
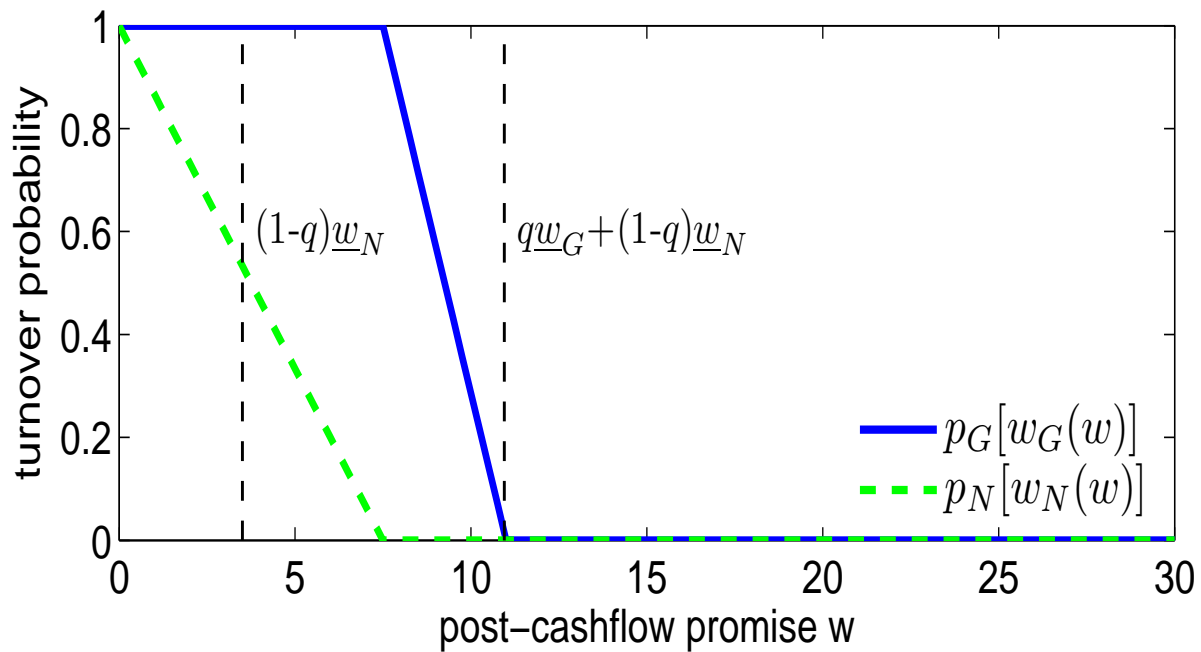


Figure 14: Turnover probability as a function of post-cashflow promise, for $\theta = G$ or $\theta = N$. The top panel corresponds to the case $\delta_N > \delta_G$, while the bottom panel corresponds to the case $\delta_G > \delta_N$. Both panels assume $\underline{w}_N = 10$, $\underline{w}_G = 14$, $\bar{w} = 25$ and $q = 0.25$.