

Non-Common Priors, Incentives, and Promotions: The Role of Learning*

Matthias Fahn[†] and Nicolas Klein[‡]

May 29, 2023

Abstract

We analyze a repeated principal-agent setting in which the principal cares about the agent's verifiable effort as well as an extra profit that can be generated only if the agent is talented. The agent is overconfident about his talent and updates beliefs using Bayes' rule. An exploitation contract in which the agent is only compensated for his effort if the extra profit materializes maximizes the principal's profits. In this optimal contract, the agent's principal-expected compensation decreases over time and learning exacerbates his exploitation, unless he has been revealed to be talented. Therefore, the principal's profits may increase with failures, and the agent may only be employed if his perceived talent is sufficiently low. As an application of these results, we analyse a firm's optimal promotion policy, and show that promotion to a new job may optimally be based on the agent being successful in a previous job, even if the agent's talent across jobs is

*We thank Miguel Ballester, Gabriel Carroll, Raphael Godefroy, Tilman Klumpp, Rumen Kostadinov, Aditya Kuvalekar, Daniel Krämer, Stephan Lauer mann, Takeshi Murooka, Georg Nöldeke, Michel Poitevin, Sven Rady, Heiner Schumacher, Peter Schwardmann, Masha Titova, and Hannu Vartiainen, and seminar audiences at Stony Brook 2022 and the Society for the Advancement of Economic Theory (SAET), for helpful comments.

[†]JKU Linz, CESifo and IZA, matthias.fahn@jku.at.

[‡]University Paris-Panthéon-Assas, Université de Montréal and CIREQ, klein-nic@yahoo.com.

entirely uncorrelated. This provides a novel explanation for the so-called “Peter Principle”, for which Benson et al., 2019 have recently provided evidence in a setting with verifiable performance and highly confident workers.

1 Introduction

Humans systematically overestimate their abilities. Many think they are better drivers than the average, more intelligent, or better at predicting political outcomes (Myers, 2010; Bondt and Thaler, 1995; see Meikle et al., 2016 or Santos-Pinto and de la Rosa, 2020 for excellent overviews). Recent evidence points towards the prevalence of such “overconfidence” also in the workplace – among managers (Malmendier and Tate, 2005; Malmendier and Tate, 2015; Huffman et al., 2022) as well as non-executives (Hoffman and Burks, 2020).

We are just beginning to understand the extent and persistence of workers’ overconfidence, and how it may affect the structure of long-term employment relationships. Whereas some studies argue that it can be cheaper for firms to hire overconfident workers who overestimate their chances of achieving a successful outcome (Santos-Pinto, 2008; de la Rosa, 2011; Sautmann, 2013), their focus is on one-shot interactions. But the relevance of such “exploitation contracts” relies on their ongoing use over an extended period of time. If workers learn and update their assessments (as studies such as Grossman and Owens, 2012 or Yaouanq and Schwardmann, 2022 indicate), their exploitation may quickly become infeasible.

In this paper, we show that a firm’s exploitation of a worker’s overconfidence about his talent can *intensify* over time, even though he incorporates informative signals and updates beliefs using Bayes’ rule. This implies that the firm’s expected profits can go up as bad signals about the worker’s talent accumulate and firm and worker become increasingly pessimistic. Then, employing a worker may only be profitable if he is believed to be sufficiently *unproductive*. We apply these results to a firm’s promotion decision and demonstrate that it can be optimal to base a promotion on success in the

current job, even if the task requirements in the current and the new job are entirely unrelated. The reason is that a success reduces the uncertainty about the worker's ability, and a subsequent promotion re-instates belief divergence and consequently exploitation possibilities. Thereby, we provide a microfoundation for the so-called Peter Principle, according to which past successes are a bigger driver of promotion decisions than what naively appears to be optimal. In contrast to the prevailing alternative theoretical explanations, our approach does not rely on (parts of) the worker's performance being unverifiable, and is thus able to rationalize recent evidence by Benson et al., 2019 for the existence of the Peter Principle among highly confident sales agents whose performance can easily be verified.

Our results are derived in a continuous-time setting, where a risk-neutral principal can hire a risk-neutral agent to work on a task. The agent's value to the principal is given by his (costly) effort and his talent (or match quality), which might either be high or low. If talent is high, the agent's effort generates an extra profit to the principal with some probability at each instant in time. If talent is low, the extra profit is never generated. The agent's talent is initially uncertain, and both players adjust their beliefs using Bayes' rule: Once the extra profit materializes for the first time, beliefs of the agent being talented jump to 1. Otherwise, beliefs go down. The agent is *overconfident* about his talent, i.e., his starting belief of being talented exceeds the principal's.

The agent's effort as well as the realization of the extra profit of high talent are verifiable. Therefore, it is possible to incentivize the agent by deterministically compensating him for his effort cost. Because of the agent's overconfidence, however, this does not maximize the principal's profits. Instead, as long as there is uncertainty (i.e., until a first success has been realized), the principal finds it optimal to only pay the agent conditionally on success. The reason is that, with such an arrangement the expected payment from the principal's perspective is below the effort cost, whereas the agent's overconfidence makes him believe that his costs are covered.

The principal-expected compensation paid to the agent and thus the extent

of the agent's exploitation depend on the ratio between the principal's and the agent's belief. This ratio – and with it the agent's principal-expected compensation – decreases over time as long as there is no success. The reason is that, whereas his overconfidence diminishes in absolute terms (and converges to zero), the agent's overconfidence increases in relative terms.

This implies that, even though the agent learns and becomes increasingly pessimistic, the principal's profits from exploiting the agent due to the provision of incentives go up as long as there has been no success. However, the total profits from hiring the agent also contain the extra profit in case he is talented, and this component goes down in expectation without a success. Still, the former effect may dominate the latter; this is the case if the extra profit is small or occurs only rarely, or if the agent is initially very overconfident. Then, as long as no success has been realized, the principal's expected profits *increase with failures*, i.e., as the principal and the agent become more and more pessimistic about the latter's talent. If the principal's outside option is so attractive that hiring an agent known to be talented is not profitable, it can be optimal to hire him only with *sufficiently pessimistic* beliefs concerning his talent. Thus, we show that receiving information about one's talent does not necessarily reduce the extent of the possible exploitation of an agent's overconfidence.

As an application of these results, we take into account that the principal's outside option may not only correspond to a termination of the employment relationship, but could also involve her value of promoting the agent to a different position. Then, for the case in which the principal's profits increase with failures, it might be optimal to promote the agent after a success in the original job, even if the agent's talents in both jobs are entirely unrelated. In general, the agent's overconfidence lets the principal put less weight on the agent's inherent ability for the new job than would be optimal based on fundamentals. This result is further exacerbated if the agent is also overconfident in the second job. Then, a first-job success wipes out the principal's exploitation opportunities there. Promoting him to the second job again introduces uncertainty regarding his talent, thus creating new room to exploit

his overconfidence. Moreover, a worker who is currently not successful but who is expected to be talented in the second stage may instead not be promoted because his continued lack of success increases the firm's profits by exploiting him. Also if the agent overestimates a potential positive correlation between his talent in both jobs, a success in the first job might increase the agent's overconfidence and render a promotion particularly profitable.

This mechanism encourages a promotion policy in which it is not necessarily the best-suited agent who will have the most stellar career. Thereby, we provide a micro-foundation for the Peter Principle¹ for which Benson et al. (2019) have recently provided evidence. Different from alternative explanations, our approach can generate the Peter Principle even if the agent's performance is verifiable, which indeed seems to be the case in the setting described by Benson et al. (2019). They demonstrate that the promotion of sales workers is to a larger extent determined by their verifiable sales than would be justified by their fit for a managerial position. Moreover, this link between sales and promotion is especially strong for so-called "lone wolves" who are highly self-confident but whose fit for managerial positions is particularly poor because of a lack of willingness to collaborate with others.

Related Literature

We relate to the theoretical literature on the "Peter Principle," according to which promotion decisions are based on success in the current, instead of (perceived) ability in the new, job. We argue that the previous explanations are insufficient to explain this phenomenon in a setting with sales agents as observed by Benson et al. (2019).

For example, one explanation having been proposed is that employees may value the signalling role of promotions. Waldman (1984) and DeVaro and

¹The *Peter Principle* states that a person performing well in his job tends to be promoted until he reaches *Peter's Plateau*, a level of responsibility for which he is incompetent. Although named for Canadian management scientist Laurence Peter, the principle was already enunciated half a century earlier by Spanish philosopher José Ortega y Gasset, in 1910.

Waldman (2012) set up models in which firms privately observe workers' abilities for the "new" job. Because a promotion provides a signal about this ability to the market, it has to come with a steep wage increase to fend off counteroffers, and the ability threshold above which someone is promoted is higher than without private information. Yet these theories do not predict that the wrong people are promoted, but instead only the very best. In the setting explored by Benson et al. (2019), a promotion indicates a sales worker's ability for his current job, rather than managerial talent. Employees may also value the social status that comes with a promotion. DellaVigna and Pope (2018) find that social comparisons can increase incentives to exert effort; however, the effects are negligible compared to even small monetary incentives. Cullen and Perez-Truglia (2022) identify negative motivation effects of earning less than one's colleagues, but only if they are at the same level and not in a higher position. Therefore, promotion-based incentives could reduce the negative consequences of pay inequality. Nevertheless, whereas Breza et al. (2017) confirm the negative effects of pay inequality, they demonstrate that these effects disappear if workers observe their higher-paid colleagues to be more productive than themselves – and relative productivity can certainly be observed among sales agents. As a further explanation for the extensive use of promotions as a motivation device, firms may use promotions instead of monetary bonuses because the latter are more prone to influence activities by workers (Milgrom and Roberts, 1988), an idea formally modelled by Fairburn and Malcomson (2001). These models rely on an effort dimension that is *not* objectively measurable and can therefore be misreported by supervisors. By the same token, in Lazear (2004), firms do observe but a noisy signal of agents' talent. In expectation, a high observation will correspond to a high noise term. Firms anticipate this, but, given the information they have access to, they cannot avoid the Peter principle. In Koch and Nafziger (2012), by contrast, agents' talents are observable, while their efforts are not. Firms thus trade off the most productive allocation of jobs to agents with their desire to minimize the agency costs stemming from Moral Hazard. Assuming that success in the more difficult job is a stronger statistical signal for high effort, they show that, for certain parameter values, the principal's latter

desideratum optimally determines the task assignment, leading to a “Peter principle”-allocation, by which some productive efficiency is optimally sacrificed for a reduction in agency costs. Such a trade-off between productive efficiency and rent extraction is also at the heart of our mechanism. In addition, our model derives a dynamic link between past performance and the likelihood of being promoted, and that this link is particularly strong for (over-)confident individuals.

We therefore conclude that, although these theories are able to rationalize the incentive roles of promotions, they are insufficient to explain the observations made by Benson et al. (2019), which are based on an easily verifiable task and highly confident individuals. Instead, we argue that their results might be the consequence of the optimal exploitation of overconfident workers.

We also contribute to the literature on incentive contracts with overconfident agents. DellaVigna and Malmendier (2004) and Heidhues and Kőszegi (2010) provide early work on how to design incentive contracts when consumers are overconfident, in this case about their future self control. They show that exploitation is optimal and feasible. In a static employment setting with a risk-neutral principal and a risk-averse agent, Santos-Pinto (2008) and de la Rosa (2011) demonstrate that implementing effort can be cheaper if the agent is overconfident about his ability. Moreover, exploitation contracts can emerge, in which an agent’s overconfidence gives him a realized expected utility that is smaller than anticipated by himself. Schumacher and Thysen (2022) explore the consequences of an agent having misspecified beliefs that pertain to the consequences of his actions off the equilibrium path. This can also make it cheaper to provide incentives for a risk-averse agent who underestimates the benefits of shirking.

Goel and Thakor (2008) and Gervais et al. (2011) analyze the investment decisions of risk-averse and overconfident managers. They demonstrate that overconfidence makes managers less conservative in their project choices and can increase firm value; however this link is not monotonic.

There also is evidence for the existence of exploitation contracts, in the lab as well as in the field. In the lab, Sautmann (2013) finds that agents who are

overconfident about their abilities overestimate their expected payoffs and consequently are worse off than underconfident agents. Larkin et al., 2012 observe that participants who overestimate their performance in a standard multiplication task are more likely to select convex (instead of linear) incentive schemes that offer generous rewards for levels of performance they are unlikely to attain.

Evidence from the field is mostly based on executive compensation, where overconfident managers receive incentive-heavy compensation contracts (Humphery-Jenner et al., 2016). Firms benefit from these arrangements because overconfident CEOs receive fewer bonus payments and smaller stock option grants than their peers and therefore ultimately receive less total compensation (Otto, 2014).

Although the mechanisms underlying such “exploitation contracts” seem well understood, their benefits for firms depend on whether they can repeatedly be applied over a sufficiently long time horizon. Thus, it is important to understand how employees assess the feedback they receive about their performance. If they learn and update their assessments (such as in Yaouanq and Schwardmann, 2022), one might expect their exploitation to quickly become infeasible. We show, however, that learning about the source of the underlying overconfidence can actually exacerbate the agent’s exploitation. Moreover, even if complete learning is achieved, firms may re-instate uncertainty – and consequently overconfidence – by promoting the agent.

Existing dynamic models with overconfident agents either rely on environments of misspecified learning in which success has several determinants and the agent is overconfident about one of them (Heidhues et al., 2018; Heidhues et al., 2021; Hestermann and Yaouanq, 2021; Murooka and Yamamoto, 2021), or assume that the agent assigns probability 1 to one state of the world and therefore does not update when receiving new information (Englmaier et al., 2020).

When deciding whether to hire the agent, the principal is facing a *one-armed bandit problem*. As a stylized formalization of the trade-off between experimentation and exploitation, the bandit problem goes back to Thompson

(1933) and Robbins (1952). Gittins (1974) showed the structure of the optimal policy; Presman (1991) calculated the *Gittins Index* for the case in which the underlying uncertainty is modeled by a Poisson process.

2 Model

A principal and an agent interact in continuous time over an infinite horizon. Both parties discount future payoffs at the rate of $r > 0$. At each instant $t \in \mathbb{R}_+$, the principal can either hire the agent or produce himself. If he produces himself in $[t, t + dt)$, he receives a profit flow of $\bar{\pi}dt \geq 0$. If the agent is hired at instant t , the agent chooses his effort level at instant t , $e_t \in \{0, 1\}$; we impose that e_t be càglàd in t . Choosing an effort of $e_\tau = 1$ ($e_\tau = 0$) for a.a. $\tau \in [t, t + dt)$ entails a cost of $cdt > 0$ (0) to the agent whose outside utility flow is normalized to zero. The choice of effort is observable and contractible. The agent's time-invariant talent $\theta \in \{0, 1\}$ determines, together with the agent's effort choice, the principal's profit flow over those time intervals in which the agent is hired.

Indeed, if the agent is hired at a flow wage of $w \in \mathbb{R}_+$, and exerts effort e , over a time interval $[t, t + dt)$, the principal's profit flow over that period is given by $(e - w)dt + \eta$ with probability θadt , and $(e - w)dt$ with the counterprobability, for some $a > 0$ and $\eta > 0$. Thus, a talented agent (i.e., one with $\theta = 1$) yields the principal an extra profit of η at an instantaneous rate of adt . The principal initially believes that the agent is talented with probability $p_0^P \in (0, 1)$; the agent initially believes that he is talented with probability $p_0^A \in [p_0^P, 1)$. We thus assume that $p_0^A \geq p_0^P$, i.e., the agent is *over-confident*. Both players update their respective beliefs according to Bayes' rule: as soon as an extra profit has been observed, both players' beliefs jump to 1, and stay there. If no extra profit has arrived by period t , party i 's belief can be written as $p_t^i = \frac{p_0^i e^{-a \int_0^t e_\tau d\tau}}{p_0^i e^{-a \int_0^t e_\tau d\tau} + 1 - p_0^i}$. In what follows, we shall write beliefs in the

form of the odds ratio $x_t^i = p_t^i / (1 - p_t^i) = x_0^i e^{-a \int_0^t e_\tau d\tau}$. Thus,

$$\frac{x_t^P}{x_t^A} = \Psi \in (0, 1]$$

is constant over time; Ψ is an inverse measure of the agent's over-confidence, with $\Psi = 1$ corresponding to the case of common priors. In the following, we shall call x_t (Ψx_t) the agent's (principal's) *belief* at instant t .

Contracts, Information, and Equilibrium Only spot contracts are possible. These specify the agent's instantaneous wage payment as a function of the agent's current effort and the principal's current profit, which is assumed to be verifiable. The agent is protected by limited liability; i.e., these wage payments must be non-negative at all times, after any history. The agent's belief is common knowledge; our results are not affected by whether the agent is aware of the principal's belief. For this, it is important that the agent's belief, and his overconfidence, are not affected by the contract offered by the principal. In our setting, the form of contract offered by the principal is (weakly) optimal for all levels of overconfidence $\Psi \leq 1$. For example, both might agree to disagree. We solve for a perfect Bayesian equilibrium (PBE) that maximizes the principal's profits (given her beliefs).

3 Results

First, we derive the optimal compensation structure. Clearly, it is not optimal to pay the agent anything he does if not exert effort. Since effort is verifiable, it is possible to offer a spot contract that pays the agent c for his effort, independently of beliefs about θ . The agent would be willing to accept this contract offer, which would allow the principal to extract the whole rent from effort. However, with $\Psi < 1$, i.e., with $p_0^A > p_0^P$, it is optimal for the principal to exploit the agent's overconfidence and only to pay him conditionally on his exerting effort *and* producing the extra profit η for the principal. The reason is that the agent's belief of being talented and thus of receiving the payment

is higher than the principal's, so that both players gain by engaging in a side-bet on the arrival of the extra profit. The risk-neutral agent is willing to accept any contract that at least covers his effort cost in expectation, $\frac{c}{ap_t^A} = \frac{1+x_t}{ax_t}c$. In a profit-maximizing equilibrium it is suboptimal to leave the agent a rent, thus the lump-sum wage payment paid after a success is $W_t = \frac{1+x_t}{ax_t}c$. The principal-expected cost of hiring the agent then amounts to

$$\frac{p_t^P}{p_t^A}c = a\Psi \frac{x_t}{1 + \Psi x_t} W_t = \frac{1 + x_t}{1 + \Psi x_t} \Psi c,$$

which is smaller than c .² For a given x_t , this amount is increasing in Ψ . Thus, the greater the agent's overconfidence, the greater the scope for side-bets between the players (the limited-liability constraint notwithstanding), and therefore the lower the amount the principal expects to pay the agent for his services.

This structure is (strictly) optimal (for $\Psi < 1$) as long as there has been no success. Once the extra profit has been realized and both players' beliefs jump to 1, this contract generates the same profits as one in which the principal just pays a flow of c irrespectively of whether η materializes or not.

3.1 The Cost of Learning

Now, we explore how the agent's expected compensation evolves over time. Clearly, after a success, beliefs jump to 1 and stay there forever thereafter, which implies that expected hiring cost then also become time-invariant. As long as no success has been realized, though, these expected costs decrease as time passes.

Lemma 1 *W_t is decreasing in x_t and hence increasing in time t if there is no success.*

²The optimality of such side-bets is widely known in settings with non-common priors, see Eliaz and Spiegler (2006), or Grubb (2015) for an overview.

The principal-expected cost of hiring the agent,

$$\frac{p_t^P}{p_t^A}c = a\Psi \frac{x_t}{1 + \Psi x_t} W_t = \frac{1 + x_t}{1 + \Psi x_t} \Psi c,$$

is increasing in x_t ; it tends to Ψc as $x \rightarrow 0$, and to c as $x \rightarrow \infty$. It is a martingale on the principal's information filtration; in case of a success, it jumps up to c , and is decreasing in time t if there is no success.

Without any success, both the principal's and the agent's beliefs go down and eventually approach zero. Because $p_t^P < p_t^A$, though, Bayes' rule indicates that the relative reduction of the principal's belief is more pronounced than the agent's:

$$\frac{dp_t^-}{p_t} = -a(1 - p_t)dt,$$

thus $|-a(1 - p_t^P)dt| = |a(1 - p_t^P)dt| > |a(1 - p_t^A)dt| \Leftrightarrow p_t^P < p_t^A$.

This implies that learning does not necessarily benefit the agent. If no success is observed and negative signals accumulate, the agent's (principal-)expected compensation goes down. Therefore, even if agents update their beliefs about the underlying source of their overconfidence using Bayes' rule (for which there is evidence, see Yaouanq and Schwardmann, 2022), their exploitation need not vanish in the long run – to the contrary, it may even exacerbate.

Note that this result does not rely on time being continuous but also holds if time is discrete. We use continuous time because it allows us to explicitly characterize value functions.

3.2 The Optimal Hiring and Firing Decision

The principal's strategy thus boils down to, at each instant, choosing whether to hire the agent as a function of the previous history. Formally, the principal's hiring decisions are a process $\{\chi_t\}_{t \in \mathbb{R}_+}$ that is predictable with respect to the public information, where $\chi_t = 1$ if the agent is hired at instant t , and

$\chi_t = 0$ otherwise. Clearly, it is without loss to restrict the principal to choosing a Markov strategy, i.e., an effort process $\{\chi_t\}_{t \in \mathbb{R}_+}$ such that $\chi_t = \chi(x_t)$ for all $t \in \mathbb{R}_+$, where $\chi : \mathbb{R}_+ \cup \{\infty\} \rightarrow \{0, 1\}$ is a time-invariant function of beliefs.³ In summary, the principal chooses a Markov strategy so as to maximize

$$\begin{aligned} \Pi(x) = & \mathbb{E} \left[\int_0^\infty r e^{-rt} \left(1 - \frac{\Psi x_0}{1 + \Psi x_0} \left(1 - e^{-a \int_0^t e(x_\tau) d\tau} \right) \right) \chi(x_t) \right. \\ & \left. \times \left(1 - \bar{\pi} - \frac{1 + x_t}{\frac{1}{\Psi} + x_t} c + \frac{\Psi x_t}{1 + \Psi x_t} a \left(\eta + \max \left\{ 0, \frac{1 - \bar{\pi} + a\eta - c}{r} \right\} \right) \right) dt \middle| x_0 = x \right], \end{aligned} \quad (1)$$

where the expectation is with respect to the process $\{x_t\}_{t \in \mathbb{R}_+}$.

Bellman Equation

We now set up the Bellman equation for the problem. It is given by

$$V^*(x) = \max_{\chi \in \{0,1\}} \chi [\mathcal{B}(x, V^*) + \mathcal{M}(x)],$$

where

$$\mathcal{M}(x) := \left[1 - \bar{\pi} - \frac{1 + x}{1 + \Psi x} \Psi c + \frac{\Psi x}{1 + \Psi x} a \eta \right],$$

is the principal's myopic payoff from hiring the agent, given his belief is Ψx . A myopic principal (i.e., one whose discount rate $r \rightarrow \infty$) would hire the agent at x if and only if $\mathcal{M}(x) \geq 0$. The same policy would be optimal if the principal did not update her belief regarding the agent's talent (e.g., because the agent's talent is continuously drawn anew). Clearly, $\mathcal{M}(x) \geq 0$ for all x if $\min\{1 - \bar{\pi} - c\Psi, 1 - \bar{\pi} + a\eta - c\} \geq 0$; in this case, a myopic principal would always hire the agent. By the same token, $\mathcal{M}(x) \leq 0$ if $\max\{1 - \bar{\pi} - c\Psi, 1 - \bar{\pi} + a\eta - c\} \leq 0$; in this case, a myopic principal would never hire the agent. If $1 - \bar{\pi} - c\Psi < 0 < 1 - \bar{\pi} + a\eta - c$, $\mathcal{M}(x) \geq 0$, and a myopic principal would thus hire the agent, if and only if $x \geq -\frac{1 - \bar{\pi} - \Psi c}{\Psi(1 - \bar{\pi} + a\eta - c)} =: x^m$.

³Our payoff-maximizing perfect Bayesian equilibrium will thus be a *Markov perfect equilibrium (MPE)* with players' beliefs as a state variable.

If, however, $1 - \bar{\pi} + a\eta - c < 0 < 1 - \bar{\pi} - c\Psi$, a myopic principal would hire the agent if and only if $x \leq x^m$. We note that $x^m \in (0, \infty)$ in both these cases.

Yet, a principal that is not myopic also takes the learning benefit of employing the agent into account. This learning benefit amounts to $\frac{1}{r}$ times the infinitesimal generator of the process of posterior beliefs applied to the value function V , and can be written as

$$\mathcal{B}(x, V) := \frac{xa}{r} \left[\frac{\Psi}{1 + \Psi x} (\max\{0, 1 - \bar{\pi} + a\eta - c\} - V(x)) - V'(x) \right].$$

We write $V^*(x) = \max\{0, V(x)\}$, where V satisfies the ODE

$$\begin{aligned} & ax(1 + \Psi x)V'(x) + (r + \Psi x(r + a))V(x) \\ &= r[(1 + \Psi x)(1 - \bar{\pi}) - (1 + x)\Psi c + \Psi xa\eta] + \Psi xa \max\{0, 1 - \bar{\pi} + a\eta - c\}, \end{aligned}$$

which is solved by

$$V(x) = 1 - \bar{\pi} + \frac{\Psi x}{1 + \Psi x} a\eta - c\Psi \frac{1 + x}{1 + \Psi x} - \mathbb{1}_{\{1 - \bar{\pi} + a\eta - c < 0\}} \frac{a}{a + r} \frac{\Psi x}{1 + \Psi x} (1 - \bar{\pi} + a\eta - c) + C \frac{x^{-\frac{r}{a}}}{1 + \Psi x},$$

with C denoting a constant of integration. We furthermore note that⁴

$$\lim_{x \downarrow 0} V(x) = 1 - \bar{\pi} - \Psi c;$$

$$\lim_{x \rightarrow \infty} V(x) = (1 - \bar{\pi} + a\eta - c) \left(1 - \mathbb{1}_{\{1 - \bar{\pi} + a\eta - c < 0\}} \frac{a}{a + r} \right);$$

in what follows, we shall write $V(0)$ and $V(\infty)$ respectively for these limits.

If $V(0)$ and $V(\infty)$ have the same sign, the principal's hiring decision under (almost) perfect information will be the same, independently of whether that almost perfect information is positive or negative regarding the agent's talent.

⁴There is a discontinuity in payoffs at $x = 0$, which stems from the fact that, at $x = 0$, the contract we are looking at (payments contingent on success) ceases to be possible. As our contract continues to be possible, and (weakly) optimal, when $p^A = p^P = 1$, there is no such discontinuity at $x = \infty$.

It is thus no surprise that the principal will make the same hiring decision for all beliefs, and hence the learning benefit $\mathcal{B} = 0$ in this case, as the following proposition shows.

Proposition 1 *The following cases describe the conditions for always or never hiring the agent being optimal.*

- [1.]If $\min\{1 - \bar{\pi} - c\Psi, 1 - \bar{\pi} + a\eta - c\} \geq 0$, $\chi(x) = 1$ for all $x \in \mathbb{R}_+ \cup \{\infty\}$ is optimal. The value function is given by $V^*(x) = 1 - \bar{\pi} + \frac{\Psi x}{1 + \Psi x} a\eta - \frac{1+x}{1+\Psi x} c\Psi$. If $a\eta > (1 - \Psi)c$, it is strictly increasing and strictly concave; if $a\eta < (1 - \Psi)c$, it is strictly decreasing and strictly convex. If $a\eta = (1 - \Psi)c$, $V^*(x) = 1 - \bar{\pi} - c\Psi$.
- [2.]If $\max\{1 - \bar{\pi} - c\Psi, 1 - \bar{\pi} + a\eta - c\} \leq 0$, $\chi(x) = 0$ for all $x \in \mathbb{R}_+ \cup \{\infty\}$ is optimal. The value function is $V^* = 0$ in this case.

Proofs for our results rely on standard verification arguments; please refer to the Appendix for details.

In the following propositions, we shall show that, in the cases not covered by Proposition 1, the principal's learning benefit will be strictly positive, and that his hiring decision will admit of a simple cutoff structure. First, if $1 - \bar{\pi} - c\Psi < 0 < 1 - \bar{\pi} + a\eta - c$, i.e., if the extra profit is important to the principal, meaning that η is large, and the initial disagreement regarding the agent's talent is not too severe, i.e., Ψ is not too low, the principal will hire the agent if and only if he is optimistic enough about his talent, as the following proposition shows.

Proposition 2 *If $1 - \bar{\pi} - c\Psi < 0 < 1 - \bar{\pi} + a\eta - c$, $\chi = \mathbb{1}_{(x^*, \infty]}$, with $x^* = \frac{r}{r+a}x^m$, is optimal. The value function is C^1 and given by*

$$V^*(x) = \mathbb{1}_{(x^*, \infty]}(x) \left[\frac{x^{-\frac{r}{a}} C}{1 + \Psi x} + 1 - \bar{\pi} + \frac{\Psi x}{1 + \Psi x} a\eta - \frac{1+x}{1 + \Psi x} c\Psi \right],$$

where $C = -x^{*\frac{r}{a}}(1 + \Psi x^*) \left[1 - \bar{\pi} + \frac{\Psi x^*}{1 + \Psi x^*} a\eta - \frac{1+x^*}{1 + \Psi x^*} c\Psi \right]$ is a constant of integration determined by value matching at $x = x^*$. On (x^*, ∞) , V^* is

strictly increasing, and strictly convex (concave) on (x^*, \tilde{x}) ((\tilde{x}, ∞)), for some inflection point $\tilde{x} \in (x^*, \infty)$.

In this case, the principal will either never hire the agent if $x_0 \leq x^*$, or, if $x_0 > x^*$, he will initially hire the agent and keep hiring him till the time τ at which the belief $x_\tau = x^*$; the agent is fired for good at this time τ . The firing time $\tau = \tau^*$, where $\tau^* := \frac{1}{a} \ln(x_0/x^*)$, if the agent produces no extra profit η in $[0, \tau^*]$; otherwise, $\tau = \infty$, i.e., the agent is hired forever. This case is thus equivalent to a standard one-armed Poisson bandit problem, in which the risky arm is pulled whenever the decision maker is optimistic enough about its quality. The value function in this case is smooth, verifying the usual *smooth pasting* property. In our case, a success is fully revealing, so that the risky arm will be used forever after a success. In the absence of a success, optimism about its quality wanes continuously; the risky arm will be abandoned forever when beliefs hit a threshold (or we start out below this threshold). The principal's learning benefit shows up in the fact that she will hire the agent below the myopic cutoff x^m ; indeed, on $(\frac{1}{1+\frac{a}{r}}x^m, x^m)$, she is hiring the agent, even though her current payoffs would be higher if she produced herself. The concept of forgoing current payoffs in exchange for information that is then parlayed into better decisions in the future is what the literature commonly refers to as *experimentation*. The extent of experimentation in our model is governed by the discounted arrival rate of information $\frac{a}{r}$; it vanishes as the principal becomes myopic ($r \rightarrow \infty$), and becomes large as information arrives quickly (a large).

If, however, $1 - \bar{\pi} + a\eta - c < 0 < 1 - \bar{\pi} - c\Psi$, i.e., if η and Ψ are relatively small, the opposite dynamics obtain. In this case, the extra profit is relatively unimportant to the principal, and the initial disagreement concerning the agent's talent is large. Then, the principal will hire the agent if and only if he is *pessimistic* enough about his talent, as the following proposition details.

Proposition 3 *If $1 - \bar{\pi} + a\eta - c < 0 < 1 - \bar{\pi} - c\Psi$, $\chi = \mathbb{1}_{[0, \tilde{x}]}$, with $\tilde{x} = \frac{a+r}{r}x^m$, is optimal. The value function in this case is given by $V^*(x) = \mathbb{1}_{[0, \tilde{x}]}(x) [1 - \bar{\pi} + \frac{\Psi x}{1+\Psi x}a\eta - \frac{1+x}{1+\Psi x}c\Psi - \frac{a}{a+r} \frac{\Psi x}{1+\Psi x}(1 - \bar{\pi} + a\eta - c)]$; it*

is C^1 , except for a convex kink at \check{x} , flat on $[\check{x}, \infty)$, and strictly decreasing and strictly convex on $(0, \check{x})$.

In this case, the principal will either never hire the agent if $x_0 > \check{x}$, or, if $x_0 \leq \check{x}$, she will hire the agent until he produces the extra profit, at which time she will fire him forever. In this case, the stopping boundary is not a *regular* boundary, as beliefs can only move away from, rather than toward, the boundary \check{x} , over the course of time. As in Keller and Rady (2015), therefore, *smooth pasting* fails, and the value function admits a kink at the boundary. As in the previous case, the extent of experimentation is increasing in the ratio $\frac{a}{r}$, with $\check{x} = \left(\frac{a}{r} + 1\right) x^* = \left(\frac{a}{r} + 1\right) x^m$.

The following remark indicates that no matter if $V^*(x)$ is increasing or decreasing, the principal experiments less if her beliefs are closer to agent's, i.e., if the extent of the latter's overconfidence is lower.

Remark 1 *More similar beliefs (higher Ψ), and therefore fewer exploitation opportunities, lead to less experimentation. Thus:*

- In Proposition 2, $\frac{\partial x^*}{\partial \Psi} > 0$.
- In Proposition 3, $\frac{\partial \check{x}}{\partial \Psi} < 0$.

Finally, the following remark, which follows from Propositions 2 and 3, collects the conditions for the monotonicity of the value function.

Remark 2 *The value function V^* is monotonically increasing if and only if $a\eta \geq (1 - \Psi)c$; it is constant if and only if $a\eta = (1 - \Psi)c$. It is monotonically decreasing if and only if $a\eta \leq (1 - \Psi)c$.*

4 Application – Optimal Promotion Policies and the Peter Principle

We have shown that the principal benefits from a discrepancy between her and the agent's belief. Once a success has revealed the agent to be good,

beliefs are aligned and the principal can no longer benefit from an exploitation contract. If the discrepancy has been large, this can lead to a situation where the principal consumes her outside option after a success.

Now, we argue that, instead of corresponding to a complete termination of the relationship, triggering the outside option may also consist of moving the agent to another position. If this move is labelled a promotion, it can be optimal to promote the agent conditionally on his being successful in his previous position, even if such a success does not indicate fitness for the new job. Indeed, Benson et al. (2019) provide causal evidence for this link, which has been widely known as the “Peter Principle.”

In the following, we will take a closer look at how the agent’s overconfidence can generate a promotion policy that is based on success in previous jobs; in Section 4.5, we relate it to the evidence provided by Benson et al. (2019). For this purpose, we will first assume that the outside option $\bar{\pi}$ is the agent’s value in the higher position; we will later explore potential features of that job (such as the agent being overconfident there as well), which generate this value. This first exercise allows us to demonstrate that, even if the agent’s talent in both jobs is unrelated, his overconfidence can make a promotion after a success in the first job more likely.

4.1 Outside Option as the Value of a Promotion

Assume the agent starts out in the first job, which is as described in Section 2. At a time of her choosing, the principal can promote the agent to a second job where his value to the principal is $\bar{\pi}$; not promoting him thus entails a flow opportunity cost of $\bar{\pi}$, as before. It is, however, not possible to move the agent back again to the first job. For simplicity, we set the value of firing, or temporarily not employing, the agent in the first job to 0. Importantly, there is no correlation between the jobs regarding the agent’s talent for either, and he is (weakly) over-confident concerning the first. Potential overconfidence in the second job is explored in the subsequent Section 4.2.

Clearly, the principal will promote the agent at time $\tau^* = \inf \{t \geq 0 : V^*(x_t) < 0\}$,

where $V^*(x_t)$ is the value of employing the agent in the first job net of the value of the outside option $\bar{\pi}$. Generally, our results will depend on whether $a\eta$ is larger or smaller than $(1 - \Psi)c$, i.e., whether $V^*(x)$ is increasing or decreasing (see Remark 2).

As a benchmark, we first analyze the principal’s promotion decision for an agent who is not overconfident, i.e., $\Psi = 1$. Then, the agent is more likely to be promoted after no success has been observed in the first job. Indeed, if $\Psi = 1$, $a\eta > (1 - \Psi)c$, and V^* is increasing in x . Thus, the agent will either be promoted after a long enough history of failures in the first job – or right away or never. This is because the longer history of failures makes the opportunity costs of promoting the agent less severe. As V^* is monotone for common priors, the following Lemma is immediate:

Lemma 2 *With common priors ($\Psi = 1$), there is a cutoff $\bar{\pi}(x)$ such that the agent is promoted iff $\bar{\pi} > \bar{\pi}(x)$; moreover, $\bar{\pi}(x)$ is increasing.*

The result that a promotion is more likely after the agent has been sufficiently *unsuccessful* in the first job may seem unintuitive but follows from our assumption that the jobs are uncorrelated; i.e., a success (failure) in the first job does not indicate (lack of) fitness for the second job. Then, the tradeoff is between the agent’s value in the second job (which is constant prior to a promotion) and the opportunity costs of losing him in the first job, which are increasing in x .

Moreover, the opportunity costs of promoting able workers are indeed affecting real-world promotion decisions. For example, there is evidence for “talent hoarding” in firms, meaning that middle managers who benefit from having good employees in their teams suppress the promotion opportunities of those they value the most (Haegele, 2022).

Next, assume that the agent is overconfident, i.e., $\Psi < 1$. Then, the results derived in Propositions 2 and 3 can be used to show

Proposition 4 *There is a cutoff $\bar{\pi}(x, \Psi)$ such that the agent is promoted iff $\bar{\pi} > \bar{\pi}(x, \Psi)$. $\bar{\pi}(x, \Psi)$ is strictly increasing in x if and only if $a\eta > (1 - \Psi)c$,*

strictly decreasing in x if and only if $a\eta < (1 - \Psi)c$, and constant in x if and only if $a\eta = (1 - \Psi)c$.

For all $x < \infty$, $\bar{\pi}(x, \Psi)$ is strictly decreasing in the players' belief alignment Ψ , when the principal's belief $x \cdot \Psi$ is held constant.

If $a\eta < (1 - \Psi)c$, $V^*(x)$ is decreasing and the agent's value goes up over time as long as no success is observed. Once a success occurs, the principal's value of keeping the agent in the first job falls because of the eliminated exploitation opportunities. Then, the resulting value reduction increases the relative benefits of a promotion (i.e., the cutoff $\bar{\pi}(x, \Psi)$ drops) even though the success is not informative of the agent's talent in the second job – and the Peter Principle is indeed the manifestation of optimal firm policy. In these cases, the agent is promoted after a success simply because, after his type has been revealed, the value of keeping him in the first job is too low.

If $a\eta > (1 - \Psi)c$, $V^*(x)$ is increasing and the general pattern is the same as with common priors ($\Psi = 1$). Either the agent is immediately (or never) promoted, or after many failures in the first job have sufficiently reduced the opportunity costs of a promotion. Still, the threshold $\bar{\pi}(x, \Psi)$ is higher than with $\Psi = 1$ because the exploitation opportunities in the first job decrease in Ψ (holding the principal's belief $x \cdot \Psi$ constant).

Finally, the last result of Proposition 4 illustrates the fact that the higher the agent's overconfidence the more valuable he is to the principal in the first job.

4.2 Endogenizing $\bar{\pi}$

Now, we endogenize the agent's value in the second job and assume that his overconfidence can extend to it. Assume that the second job also has the features described in Section 2; there is still no correlation between the agent's talent across both jobs. The details can be found below, in Section 6.2 in the Appendix. As before, the agent's value in the second job remains constant

as long as he is not promoted. Therefore, the same effects as in Section 4.1 obtain, while introducing the agent’s overconfidence in the second job allows for additional comparative statics. The reason is that a promotion after a success in the first job re-instates uncertainty and overconfidence, and thus again allows the principal to exploit the agent. Therefore, a lower Ψ in the second job (holding the principal’s belief there constant) makes it *ceteris paribus* more likely that the agent is promoted after a first-job success.

4.3 Correlated Jobs and Endogenous Overconfidence

We have demonstrated that the agent’s overconfidence can affect promotion decisions in a way that is reminiscent of the Peter Principle. We have assumed that the agent’s talent across both jobs is not correlated. However, even with a positive correlation between jobs, the agent’s overconfidence induces the principal to put less weight on the agent’s talent for the second job. Indeed, while a success in the first job then increases players’ beliefs concerning the agent’s talent for the second job, this increase is less pronounced than the increase in the belief about his talent for the first job (unless correlation was perfect). In the absence of divergent beliefs, therefore, such a success should make a promotion less likely. With divergent beliefs, however, as the success also eliminates exploitation opportunities in the first job, promotion will become more likely whenever the belief divergence is important enough (i.e., whenever Ψ is low enough).

Moreover, a success in the first job could also by itself increase the agent’s overconfidence. For example, assume that the agent overestimates the *correlation* between talent across both jobs. This could be the result of an inherent bias,⁵ or of the principal’s subterfuge. Then, our results would only require the agent to naively believe the principal’s claim that being successful in the first job is indicative of his potential in the second job. In this case,

⁵For example, the widely observed self-attribution bias, in which people attribute their success to their own abilities instead of just being lucky (see Daniel et al., 1998 or Billett and Qian, 2008 for evidence in the context of managers), could be a factor leading to the agent’s attribution of a first-job success to a general skill that also transfers to other realms.

promoting an agent who has proven to be talented in the first job would again create the additional benefit of being able to exploit his overconfidence in the second job. Importantly, this result would not require the agent to be inherently or initially overconfident – instead his overconfidence would endogenously emerge from a wrong belief that talent in one domain transfers to talent in another.

4.4 Two Agents

Finally, we argue that employing overconfident agents may also lead to the principal’s putting less weight on an agent’s perceived value in the second job when making a promotion decision, compared to the case where agents are not overconfident. Assume there is some time T by which the principal wants to promote one out of two agents, $i \in \{1, 2\}$. As in Section 4.1, the principal’s value of promoting agent i , $\bar{\pi}_i$, is solely given by his (expected) inherent talent in the second job. Without loss, we assume that $\bar{\pi}_1 \geq \bar{\pi}_2$. To isolate the role of an agent’s overconfidence on the principal’s promotion policy and abstract from differences in the opportunity costs of a promotion, we focus on cases in which the principal’s belief $\Psi_i x_i$ is the same for both agents, while only their Ψ_i might differ. If agent i has produced a success, he is commonly known to be talented, i.e., $\Psi_i = 1$. Now, if both agents have produced a success, or agents are not inherently overconfident, i.e., $\Psi_1 = \Psi_2 = 1$, agent 1 is promoted because the principal bases the promotion decision solely on the agent’s perceived value in the second job. Otherwise, agent 1 is not necessarily promoted: Assume $\Psi_1 < \Psi_2 \leq 1$, so that, due to his greater overconfidence, agent 1’s value is also higher in the first job. Then, if the difference between $\bar{\pi}_1$ and $\bar{\pi}_2$ is small relative to the difference between Ψ_2 and Ψ_1 , the principal will indeed promote agent 2 although he is less well suited for the second job.

Putting together this outcome as well as the insights from Section 4.1 and holding the agent’s perceived value in the second job constant, we can conclude that an overconfident agent is more likely to be promoted if he has

proven to be talented in the first job, whereas he is less likely to be promoted otherwise – compared to the benchmark case of common beliefs. Thus, we would predict a positive correlation between current performance and a promotion, even if the requirements for both jobs are entirely unrelated.

4.5 Evidence

Using microdata on sales workers, Benson et al. (2019) find evidence for promotion policies putting too much weight on current performance, as opposed to perceived fit for the new job. Although sales clearly are a *verifiable* performance measure, high sales are not only rewarded with cash compensation, but also increase a salesperson’s chances of being promoted to a managerial position. This policy disregards managerial potential and is costly because it reduces managerial quality (measured as value added to subordinate sales) by 30% compared to a counterfactual where the ones with the highest managerial potential would be promoted. Benson et al. (2019) discuss a number of potential theoretical explanations for these outcomes which, however, we argue cannot fully rationalize their observations, which are based on an easily verifiable task (see the Related Literature Section above). Instead, we argue that it is not the nature of the job that renders the promotion of successful sales agents (instead of those with the best fit) optimal, but their personal characteristics. Indeed, there is evidence that sales agents are particularly prone to being overconfident. Sevy (2016), in a Forbes blog, argues that, because of the availability of clear performance indicators, sales is an environment that attracts people who want to prove their ability. Those who go for sales care about personal advancement and not about helping a team thrive; this is different in sales *management*, where holding back one’s ego and letting others shine is important.

Moreover, whereas Benson et al. (2019) find that collaboration experience is indicative of better *managerial* performance, so-called “lone wolves,” who never collaborate and are known to be highly self confident (Dixon and Adamson, 2011) are significantly more likely to be promoted to a managerial position.

5 Conclusion

We have shown that, in a model in which the principal benefits both from the agent's work effort (in a deterministic fashion) and (stochastically) from his talent, the monotonicity of the principal's value function depends on the agent's overconfidence Ψ . If the agent's appraisal of his talent is close to that of the principal, i.e., if Ψ is close to 1, the principal's value function is increasing in her belief, and the agent is fired after a long enough streak of failures. If, however, the agent is very overconfident, i.e., if Ψ is low, the principal's value function is *decreasing* in her belief, and the agent is fired *after a success*. As in our model *firing* can be interpreted as promotion to a second, unrelated, job, we provide a novel explanation for the well-documented *Peter principle*: As the agent's type becomes common knowledge after a success, a success makes exploitation contracts impossible; thus, if exploiting the agent's overconfidence is an important part of the principal's objective, she will not want to hire the agent in the current job any longer after a success there, preferring to promote him to another job instead.

We have assumed that a success fully reveals the agent's type, i.e., an untalented agent never produces a success. While this makes our model analytically tractable, we expect our main qualitative conclusions to continue to obtain in a setting in which an untalented agent may also at times, albeit less frequently, produce a success.

6 Appendix

6.1 Proofs

We shall write

$$\hat{V}(x) = 1 - \bar{\pi} + \frac{\Psi x}{1 + \Psi x} a\eta - c\Psi \frac{1+x}{1+\Psi x} - \mathbb{1}_{\{1-\bar{\pi}+a\eta-c < 0\}} \frac{a}{a+r} \frac{\Psi x}{1+\Psi x} (1-\bar{\pi}+a\eta-c)$$

for the principal's payoff of never firing the agent in the absence of a success.

In all four cases, the proposed policy χ implies a well-defined law of motion of the belief x , and the closed-form expression for V^* is the payoff function associated with the policy χ . To prove optimality of χ , it suffices to show that $\mathcal{B}(x, V^*) \geq -\mathcal{M}(x)$ ($\mathcal{B}(x, V^*) \leq -\mathcal{M}(x)$) whenever $\chi = 1$ ($\chi = 0$) on some open subset of \mathbb{R}_+ .

For Proposition 1, Case (1.), direct computation shows that $\mathcal{B}(x, \hat{V}) \geq -\mathcal{M}(x)$ for all $x \geq 0$. Moreover, $\hat{V}' > 0 > \hat{V}''$ if $a\eta > (1 - \Psi)c$, $\hat{V}' < 0 < \hat{V}''$ if $a\eta < (1 - \Psi)c$, and $\hat{V} = 1 - \bar{\pi} - c\Psi$ if $a\eta = (1 - \Psi)c$.

In Case (2.), $\mathcal{B}(x, V^*) = \mathcal{B}(x, 0) = 0$, for all $x \geq 0$. Thus, all that remains to be shown is that $\mathcal{M}^*(x) \leq 0$ for all $x \geq 0$. As \mathcal{M} is increasing, this is equivalent to $\lim_{x \rightarrow \infty} \mathcal{M}(x) = 1 - \bar{\pi} - c + a\eta \leq 0$, which holds by the definition of Case (2.).

Let us turn to Proposition 2. For $x < x^*$, $V^*(x) = 0$ and $\mathcal{B}(x, V^*) = \frac{\Psi x a}{r(1+\Psi x)}(1 - \bar{\pi} + a\eta - c)$. Direct computation shows that $\mathcal{B}(x, V^*) \leq -\mathcal{M}(x)$ for $x < x^*$. For $x > x^*$, one shows by direct computation that $\mathcal{B}(\cdot, V^*) > -\mathcal{M}(\cdot)$ in this range. Thus, $\chi = \mathbb{1}_{(x^*, \infty]}$ is optimal. Direct computation furthermore shows that $\lim_{x \downarrow x^*} V^{*'}(x) = 0$ and $V^{*'}(x) > 0$ for all $x > x^*$. By the same token, direct computation shows that $\lim_{x \downarrow x^*} V^{*''}(x) > 0$, $\lim_{x \rightarrow \infty} V^{*''}(x) < 0$, while $V^{*'''}|_{(x^*, \infty)} < 0$.

We now turn to Proposition 3. For $x > \check{x}$, $V^*(x) = \mathcal{B}(x, V^*) = 0$. By the same token, $\mathcal{M}(x) \leq 0$ if and only if $x \geq x^m = \frac{r}{a+r}\check{x}$. For $x < \check{x}$, one shows by direct computation that $\mathcal{B}(\cdot, V^*) > -\mathcal{M}(\cdot)$ in this range. Thus, $\chi = \mathbb{1}_{(0, \check{x}]}$

is optimal. Direct computation furthermore shows that $V^{*\prime\prime}|_{(0,\tilde{x}]} > 0$, and that $\lim_{x \uparrow \tilde{x}} V^{*\prime}(x) < 0$.

6.2 Microfoundation for Second Job

The purpose of this appendix is to show how to extend the model so as explicitly to incorporate the second job. Specifically, we shall denote $x_0 \in (0, \infty)$ ($\Psi_x x_0$) the agent's (principal's) belief (measured in odds ratios, as before) that the agent is talented for the first job, and hence produces the extra profit $\eta_x > 0$ at the rate $a_x > 0$ in the first job. By the same token, we shall write $y_0 \in (0, \infty)$ ($\Psi_y y_0$) for the agent's (principal's) belief that the agent is talented for the second job, and hence produces the extra profit $\eta_y > 0$ at the rate $a_y > 0$ in the second job. Flow effort costs in either job are $c_x > 0$, and $c_y > 0$, respectively.

We continue to assume that the agent is (weakly) overconfident regarding both jobs, i.e., that $\Psi_x \leq 1$ and $\Psi_y \leq 1$. Since talent across jobs is uncorrelated, we have $y_t = y_0$ for all times t at which the agent is employed in the first job. Both parties discount future payoffs at the rate $r > 0$. After the agent has been promoted to the second job, the principal, as before, receives a flow payoff of $\bar{\pi}_y \geq 0$ if she does not hire the agent. Before the agent is promoted, the principal receives a flow payoff of $\bar{\pi}_x \geq 0$ if she does not hire the agent. We shall write V_x^* for the agent's value to the principal in the first job, ignoring the possibility of promotion to the second job. Clearly, the principal will promote the agent at time $\tau^* = \inf \{t \geq 0 : \bar{\pi}_x + V_x^*(x_t) < \bar{\pi}_y + V_y^*(y_0)\}$.

The value functions V_x^* and V_y^* are computed as above. Before the agent is promoted, y_t , and therefore $V_y^*(y_t) \equiv V_y^*(y_0)$, remain constant, while x_t , and hence $V_x^*(x_t)$, evolve as described above. The key to our subsequent analysis is the monotonicity of the value function, which we have noted in Remark 2. In particular V_i^* ($i \in \{x, y\}$) is strictly increasing (decreasing) if and only if $a_i \eta_i > (1 - \Psi_i)c_i$ ($a_i \eta_i < (1 - \Psi_i)c_i$), and constant if and only if $a_i \eta_i = (1 - \Psi_i)c_i$.

As before a promotion, y_t , and hence $V_y^*(y_t)$, remain constant, only the mono-

tonicity of V_x^* , and hence the properties of the first job, matter for the dynamics. In particular, for arbitrary parameters for the second job:

- If $a_x \eta_x > (1 - \Psi_x)c_x$, the agent is promoted after a long enough dearth of lump sums $[0, \tau^*]$, with $\tau^* \in [0, \infty]$;
- if $a_x \eta_x < (1 - \Psi_x)c_x$, the agent is promoted either right away, never, or at the arrival time of the first lump sum in the first job;
- if $a_x \eta_x = (1 - \Psi_x)c_x$, the agent is either promoted right away or never.⁶

Promotion dynamics thus depend only on the characteristics of the first job. In particular, the agent is promoted after a long enough streak of failures if $a_x \eta_x > (1 - \Psi_x)c_x$. If $a_x \eta_x = (1 - \Psi_x)c_x$, his performance in the first job does not matter; he either stays in the first job forever, or is immediately affected to the second job. If $a_x \eta_x < (1 - \Psi_x)c_x$, Peter-principle dynamics apply: the agent is promoted after a success in the first job.

Thus, if $a_x \eta_x \leq (1 - \Psi_x)c_x$, the agent is either promoted right away or never in the absence of a success. If $a_x \eta_x > (1 - \Psi_x)c_x$, however, the agent is never promoted after a success, but, in the absence of a success, may be promoted at any time $\tau^* \in [0, \infty]$, the exact realization of which depends on the precise parameter values.

References

- BENSON, A., D. LI, AND K. SHUE (2019): “Promotions and the Peter Principle,” *Quarterly Journal of Economics*, 134, 2085–2134, 10.1093/qje/qjz022. (document), 1, 1, 4, 4.5
- BILLETT, M. T., AND Y. QIAN (2008): “Are Overconfident CEOs Born or Made? Evidence of Self-Attribution Bias from Frequent Acquirers,” *Management Science*, 54, 1037–1051, 10.1287/mnsc.1070.0830. 5

⁶This is neglecting the knife-edge case where $V_x^* = 1 - \bar{\pi}_x - \Psi_x c_x = V_y^*(y_0)$; in this case, the principal is indifferent over all promotion times in $[0, \infty]$, independently of the history.

- BONDT, W. F. D., AND R. H. THALER (1995): “Chapter 13 Financial decision-making in markets and firms: A behavioral perspective,” in *Handbooks in Operations Research and Management Science*: Elsevier, 385–410, 10.1016/s0927-0507(05)80057-x. 1
- BREZA, E., S. KAUR, AND Y. SHAMDASANI (2017): “The Morale Effects of Pay Inequality,” *Quarterly Journal of Economics*, 133, 611–663, 10.1093/qje/qjx041. 1
- CULLEN, Z., AND R. PEREZ-TRUGLIA (2022): “How Much Does Your Boss Make? The Effects of Salary Comparisons,” *Journal of Political Economy*, 130, 766–822, 10.1086/717891. 1
- DANIEL, K., D. HIRSHLEIFER, AND A. SUBRAHMANYAM (1998): “Investor Psychology and Security Market Under- and Overreactions,” *The Journal of Finance*, 53, 1839–1885, 10.1111/0022-1082.00077. 5
- DELLAVIGNA, S., AND U. MALMENDIER (2004): “Contract Design and Self-Control: Theory and Evidence,” *Quarterly Journal of Economics*, 119, 353–402, 10.1162/0033553041382111. 1
- DELLAVIGNA, S., AND D. POPE (2018): “What Motivates Effort? Evidence and Expert Forecasts,” *Review of Economic Studies*, 85, 1029–1069, 10.1093/restud/rdx033. 1
- DEVARO, J., AND M. WALDMAN (2012): “The Signaling Role of Promotions: Further Theory and Empirical Evidence,” *Journal of Labor Economics*, 30, 91–147, 10.1086/662072. 1
- DIXON, M., AND B. ADAMSON (2011): *The Challenger Sale: Taking Control of the Customer Conversation*: Penguin Publishing Group, <https://books.google.de/books?id=pioPC90iMdMC>. 4.5
- ELIAZ, K., AND R. SPIEGLER (2006): “Contracting with Diversely Naive Agents,” *Review of Economic Studies*, 73, 689–714, 10.1111/j.1467-937x.2006.00392.x. 2

- ENGLMAIER, F., M. FAHN, AND M. A. SCHWARZ (2020): “Long-Term Employment Relations when Agents are Present Biased,” Working Paper. 1
- FAIRBURN, J. A., AND J. M. MALCOMSON (2001): “Performance, Promotion, and the Peter Principle,” *Review of Economic Studies*, 68, 45–66, 10.1111/1467-937x.00159. 1
- GERVAIS, S., J. B. HEATON, AND T. ODEAN (2011): “Overconfidence, Compensation Contracts, and Capital Budgeting,” *Journal of Finance*, 66, 1735–1777, 10.1111/j.1540-6261.2011.01686.x. 1
- GITTINS, J. (1974): “A dynamic allocation index for the sequential design of experiments,” *Progress in statistics*. 1
- GOEL, A. M., AND A. V. THAKOR (2008): “Overconfidence, CEO Selection, and Corporate Governance,” *Journal of Finance*, 63, 2737–2784, 10.1111/j.1540-6261.2008.01412.x. 1
- GROSSMAN, Z., AND D. OWENS (2012): “An unlucky feeling: Overconfidence and noisy feedback,” *Journal of Economic Behavior and Organization*, 84, 510–524, 10.1016/j.jebo.2012.08.006. 1
- GRUBB, M. D. (2015): “Overconfident consumers in the marketplace,” *Journal of Economic Perspectives*, 29, 9–36. 2
- HAEGELE, I. (2022): “Talent Hoarding in Organizations,” Working Paper. 4.1
- HEIDHUES, P., AND B. KŐSZEGI (2010): “Exploiting Naivete about Self-Control in the Credit Market,” *American Economic Review*, 100, 2279–2303. 1
- HEIDHUES, P., B. KŐSZEGI, AND P. STRACK (2018): “Unrealistic Expectations and Misguided Learning,” *Econometrica*, 86, 1159–1214, 10.3982/ecta14084. 1

- HEIDHUES, P., B. KOSZEGI, AND P. STRACK (2021): “Convergence in models of misspecified learning,” *Theoretical Economics*, 16, 73–99, 10.3982/te3558. 1
- HESTERMANN, N., AND Y. L. YAOUANQ (2021): “Experimentation with Self-Serving Attribution Biases,” *American Economic Journal: Microeconomics*, 13, 198–237, 10.1257/mic.20180326. 1
- HOFFMAN, M., AND S. V. BURKS (2020): “Worker overconfidence: Field evidence and implications for employee turnover and firm profits,” *Quantitative Economics*, 11, 315–348, 10.3982/qe834. 1
- HUFFMAN, D., C. RAYMOND, AND J. SHVETS (2022): “Persistent Overconfidence and Biased Memory: Evidence from Managers,” *American Economic Review*, 112, 3141–3175, 10.1257/aer.20190668. 1
- HUMPHERY-JENNER, M., L. L. LISIC, V. NANDA, AND S. D. SILVERI (2016): “Executive overconfidence and compensation structure,” *Journal of Financial Economics*, 119, 533–558, 10.1016/j.jfineco.2016.01.022. 1
- KELLER, G., AND S. RADY (2015): “Breakdowns,” *Theoretical Economics*, 10, 175–202, 10.3982/te1670. 3.2
- KOCH, A. K., AND J. NAFZIGER (2012): “Job Assignments under Moral Hazard: The Peter Principle Revisited,” *Journal of Economics and Management Strategy*, 21, 1029–1059, 10.1111/j.1530-9134.2012.00347.x. 1
- LARKIN, I., L. PIERCE, AND F. GINO (2012): “The psychological costs of pay-for-performance: Implications for the strategic compensation of employees,” *Strategic Management Journal*, 33, 1194–1214, 10.1002/smj.1974. 1
- LAZEAR, E. P. (2004): “The Peter Principle: A Theory of Decline,” *Journal of Political Economy*, 112, S141–S163, 10.1086/379943. 1
- MALMENDIER, U., AND G. TATE (2005): “CEO Overconfidence and Corporate Investment,” *Journal of Finance*, 60, 2661–2700, <http://www.jstor.org/stable/3694800>. 1

- (2015): “Behavioral CEOs: The Role of Managerial Overconfidence,” *Journal of Economic Perspectives*, 29, 37–60, 10.1257/jep.29.4.37. 1
- MEIKLE, N. L., E. R. TENNEY, AND D. A. MOORE (2016): “Overconfidence at work: Does overconfidence survive the checks and balances of organizational life?” *Research in Organizational Behavior*, 36, 121–134, 10.1016/j.riob.2016.11.005. 1
- MILGROM, P., AND J. ROBERTS (1988): “An Economic Approach to Influence Activities in Organizations,” *American Journal of Sociology*, 94, S154–S179, 10.1086/228945. 1
- MUROOKA, T., AND Y. YAMAMOTO (2021): “Misspecified Bayesian Learning by Strategic Players: First-Order Misspecification and Higher-Order Misspecification,” *Working Paper*. 1
- MYERS, D. G. (2010): *Social psychology*: McGraw-Hill, 610. 1
- OTTO, C. A. (2014): “CEO optimism and incentive compensation,” *Journal of Financial Economics*, 114, 366–404, 10.1016/j.jfineco.2014.06.006. 1
- PRESMAN, E. L. (1991): “Poisson Version of the Two-Armed Bandit Problem with Discounting,” *Theory of Probability and Its Applications*, 35, 307–317, 10.1137/1135038. 1
- ROBBINS, H. (1952): “Some aspects of the sequential design of experiments.” 1
- DE LA ROSA, L. E. (2011): “Overconfidence and moral hazard,” *Games and Economic Behavior*, 73, 429–451, 10.1016/j.geb.2011.04.001. 1, 1
- SANTOS-PINTO, L. (2008): “Positive Self-image and Incentives in Organisations,” *Economic Journal*, 118, 1315–1332, 10.1111/j.1468-0297.2008.02171.x. 1, 1

- SANTOS-PINTO, L., AND L. E. DE LA ROSA (2020): “Overconfidence in labor markets,” *Handbook of Labor, Human Resources and Population Economics*, 1–42. 1
- SAUTMANN, A. (2013): “Contracts for Agents with Biased Beliefs: Some Theory and an Experiment,” *American Economic Journal: Microeconomics*, 5, 124–156, <http://www.jstor.org/stable/43189633>. 1, 1
- SCHUMACHER, H., AND H. C. THYSEN (2022): “Equilibrium contracts and boundedly rational expectations,” *Theoretical Economics*, 17, 371–414, [10.3982/te4231](https://doi.org/10.3982/te4231). 1
- SEVY, B. (2016): “Why Great Salespeople Make Terrible Sales Managers,” October, <https://www.forbes.com/sites/groupthink/2016/10/10/why-great-sales-people-make-terrible-sales-managers/>. 4.5
- THOMPSON, W. R. (1933): “On the Likelihood That one Unknown Probability Exceed Another in View of the Evidence of Two Samples,” *Biometrika*, 25, 285–294, [10.1093/biomet/25.3-4.285](https://doi.org/10.1093/biomet/25.3-4.285). 1
- WALDMAN, M. (1984): “Job Assignments, Signalling, and Efficiency,” *The RAND Journal of Economics*, 15, 255–267, [10.2307/2555679](https://doi.org/10.2307/2555679). 1
- YAOUANQ, Y. L., AND P. SCHWARDMANN (2022): “Learning About One’s Self,” *Journal of the European Economic Association*, 20, 1791–1828, [10.1093/jeea/jvac012](https://doi.org/10.1093/jeea/jvac012). 1, 1, 3.1