Incentives to Discover Talent

Tobias Brünner
University of Lincoln

Guido Friebel
Goethe University Frankfurt

Richard Holden*
UNSW Sydney and AGORA Centre for Market Design

Suraj Prasad
University of Sydney

We study an agent's incentives to discover where her talents lie before putting them to productive use. In our setting, an agent can *specialize* and learn about the same type of talent repeatedly, or *experiment* and learn about different types of talent. While experimentation is efficient for a range of distributions of talent and initial signals, labor-market institutions play a crucial role for individual incentives to experiment. Institutions that give the agent sufficiently large bargaining power, provide incentives for experimentation, but for weak bargaining power, agents specialize. We also look at how competition in the labor market, human capital accumulation, and correlation across talents affect incentives to experiment. (*JEL* codes: D83; J24; J42)

1. Introduction

The idea that there are gains from the *division of labor* with people specializing their efforts across tasks is an old one–dating to around 2400 years ago in Plato's *Republic*. It was, of course, expanded into one of the cornerstones of modern economics by Adam Smith in *The Wealth of Nations* where he emphasized the benefits of breaking down tasks in his hypothetical pin factory.

---

*UNSW Sydney and AGORA Centre for Market Design, 1905/101 Bathurst Street, NSW 2052, Sydney. Email: richard.holden@unsw.edu.au. We are very grateful to the Editor, Michael Powell, and to three anonymous referees for their comments and suggestions. We would also like to thank Murali Agastya, Peter Exterkate, Chiaki Hara, Jonathan Newton, Masatoshi Tsumagari, and seminar audiences at Keio University and Kyoto University.

The Journal of Law, Economics, and Organization, Vol. 00, No. 0
doi:10.1093/jleo/ewab004

Less emphasized by Smith, though equally relevant in the modern economy, is that gains from division of labor also arise from people having different talents (Arrow1974, p. 19). These gains among people with different talents remain a fundamental consideration in fields from labor economics to international trade. But far less attention has been paid to how people come to discover their talents. Sometimes talents are apparent, but more often they must be discovered.

In an insightful book, *Range: Why Generalists Triumph in a Specialized World*, Epstein (2019) explores different paths to discovering talent. Some, like the golfer Tiger Woods, *specialize* by starting an activity early on and sticking to it. Others, like the tennis player Roger Federer, *experiment* by dabbling in many activities before honing in on one. Drawing on examples from sport, music, education, and careers, Epstein observes that a significant proportion of top performers choose a path of experimentation—a view that is at odds with the widespread belief that only high-frequency repetition makes top performers.[1]

The discovery of talents (through specialization or experimentation) and the critical role that labor markets and other institutions play in providing (implicit or explicit) incentives for this discovery are the topic of this paper.

The workhorse model for analyzing the role of talent in economics is the now classic *career concerns* model of Holmström (1999) in which the market and a worker symmetrically learn about the worker's innate ability. By now there is a substantial literature that applies aspects of the career concerns framework to issues of institutional and organizational design: public sector management (Dewatripont et al. 1999), team management and compensation (Jeon 1996; Auriol et al. 2002; Ortega 2003), job design (Meyer 1994; Ortega 2001; Kaarboe and Olsen 2006), and compensation design (Gibbons and Murphy 1992; Meyer and Vickers 1997).

We use this framework to ask a simple question: what incentives do people have to discover their talents? In particular, we focus on the role of incentives as provided by the institutional structure of the labor market and the organizations that reside within them.

In our model, there are two sectors with at least one firm in each sector and an agent who can choose to work in either sector. The agent's sector-specific talent is unknown and production depends on an agent's talent.

There are two phases: *learning* and *working*. Prior to working, the agent can get a signal about her talents by *sampling*, but she can only sample one type of talent per period of learning. In the working phase, the agent shares the surplus generated from the employment relationship via bargaining.

The learning phase is readily interpretable as education with the sampling taking place at the level of a course, a major, or a degree in a certain

1. See Colvin (2008) for an example of this view which emphasizes "deliberate practice"—that is, a designed activity often with a teacher's help that can be frequently repeated with continuous feedback—over talent.

field, and with signals in the form of grades or references. As Schultz (1968) put it, one of the "three major functions of higher education . . . [is] . . . the discovery of talent." Thus, our model speaks to the incentives that education systems and labor markets provide for discovering talent. Alternatively, learning may also take place within the context of an internal labor market where new hires or interns in an organization are assigned to different jobs to see what they are good at. We elaborate on these two modes of learning in Section 7.

Our first central result is that when the two sectors are symmetric in the sense that talents in both sectors have the same mean and the same variance, and are normally distributed, experimentation is efficient—regardless of the initial draw of talent.

To understand the result, consider the choice of the agent after having sampled, say, sector A in the first period. Experimentation involves sampling sector B in period 2, whereas specialization involves sampling sector A again. By symmetry of the sectors, a first signal from sector B is Blackwell more informative than a second signal from sector A. So there is more to learn from experimentation. But there is also nothing to lose. If the agent learns that she has little aptitude for sector B she can always switch to sector A.

Our second central result is that labor-market institutions do not always provide incentives for efficient experimentation. In particular, when the agent's bargaining power is low and labor markets in each sector monopsonistic, the agent's wage is forced down to her next best talent in another sector. Consequently, she (inefficiently) specializes or does not sample her talent. By contrast, when the agent's bargaining power is high, so that she is a residual claimant to the returns from her talent, she efficiently experiments.

To see the intuition for this result, note that sampling a talent involves risk (over the posterior mean) with experimentation being more risky than specialization as less is known about a talent sampled for the first time. When the agent is a residual claimant, the risk is all upside, leading to a convex payoff profile that encourages risk taking through experimentation. By contrast, when the agent's wage equals her next best option in an outside sector, the risk is all downside, leading to a concave payoff profile where the agent prefers the safer option of specialization or not sampling a talent.

The key feature of our model driving the inefficiency result above—namely that an agent's options outside a sector matter for wage determination in monopsonistic labor markets—is consistent with evidence in Schubert et al. (2020). Using a database of online vacancy postings, they find that the negative effect of labor market concentration on wages is stronger for occupations with lower outward mobility to another occupation, thus establishing a link between outside options and wages.[2] It is

2. Their result, however, does not pin down the exact mechanism underlying the link between outside options and wages, nor does it shed light on incentives to experiment.

also worth pointing out that there are other labor-market institutions besides an agent's high bargaining power which encourage residual claimancy: these include competition in labor markets (which we analyze in Section 6.4), entrepreneurship, and high-powered incentive pay.

A strength of our model is that we can clearly see to what extent our result that experimentation is efficient relies on our assumptions of symmetry, normality, and the absence of human capital accumulation. Relaxing these assumptions highlights circumstances in which there is an efficiency role for specialization. When talents are asymmetric, sampling the talent with the higher variance repeatedly can be optimal because there is more to learn from that talent. To understand the role normality plays, we consider a setting where talents have a *t*-distribution. Unlike the case where talents are normally distributed, the posterior variance now depends on the realization of the signal, which once again makes specialization efficient for extreme draws (negative or positive) of the talent sampled. Finally, when sampling a talent is associated with the accumulation of sector-specific human capital, specialization once again has a role to play for high draws of the initial talent sampled.

Our study shares features with the multi-armed bandit literature.[3] As in the multi-armed bandit problem, our setup includes experimentation and exploitation; but whereas in the former an agent can repeatedly switch between experimentation and exploitation over an infinite horizon, in our setup, experimentation can only happen in the learning phase (periods 1 and 2) and exploitation occurs in the working phase (period 3). In other words, we emphasize a different, yet fundamental, type of (pre work) sampling that complements these theories above. Modeling the labor market as a multi-armed bandit problem, Miller (1984) shows that, ceteris paribus, an agent chooses the job with the highest information value. Since working in a job reduces uncertainty about the job-specific match value, agents have an incentive to switch between jobs (and occupations). This result is consistent with our finding that agents who are residual claimants experiment during the learning phase. Following Jovanovic (1979), Miller (1984) considers many firms that compete for a single agent, who is, therefore, able to extract the entire match value. In contrast, Felli and Harris (1996) assumes that there are only two firms, one for each job, such that the agent's wage from its current employer equals the agent's match value with the other firm. In this setting, Felli and Harris (1996) show that the agent experiments efficiently, in the sense that total surplus is maximized. This contrasts with our finding that inefficiencies can arise in the sampling stage.[4]

3. See Bergemann and Välimäki (2008) for a review of the multi-armed bandit problem and its applications in economics.

4. The key difference that leads to these contrasting results regarding the efficiency of the sampling strategy is in terms of the sequence of decisions. In Felli and Harris (1996), firms make their wage offers before the agent chooses a firm and then learns about her talent on the job. In this setup, firms choose their wage offers to incentivize the agent to sample

In our model, talent has multiple dimensions. This feature appears in other learning contexts as well: job design (Meyer 1994; Ortega 2001) and education systems (Malamud 2010, 2011). But in these models, the decision about which type of talent to learn about is *fixed* upfront and cannot depend on new information. By contrast, this decision in our framework is *flexible* and is allowed to vary with new information.

The idea that learning can be inefficient when the costs of discovering talent cannot be adequately compensated also plays a role in other settings. In Terviö (2009) and Pallais (2014), firms bear the opportunity cost of trying out a worker to discover her talent but when talents are general they cannot recoup the benefits. As a result, firms underinvest in hiring new workers with upside potential. Using a field experiment in an online marketplace, Pallais (2014) finds evidence consistent with this result. The difference in our paper is the source of these opportunity costs of learning and the compensation for them. With multiple talents, the opportunity cost of sampling a talent is not being able to sample another, which in turn interacts with an agent's bargaining power to determine compensation.

The trade-off between experimentation and specialization is related to the comparison of breadth versus depth in Geng et al. (2018). In their baseline model, Geng et al. (2018) consider an agent who chooses between N products which each have N attributes. The value of a product is the sum of its i.i.d. attributes. Before choosing a product, the agent decides whether to learn the values of all N attributes of one product (depth) or pick one attribute, say attribute j, and learn the values of attribute j for all N products (breadth). Geng et al. (2018) show that when $N=2$ and the distribution of attributes is symmetric, the agent is indifferent between breadth and depth, whereas we find that the agent prefers experimentation (breadth). The key difference that leads to the divergent findings is that in Geng et al. (2018), all attribute values are i.i.d., whereas in our model, the second signal from sector A is less informative than the first signal from sector B.

Finally, our paper is related to an extension in Holmström (1999) where he shows how a risk averse agent may inefficiently choose projects to reduce her exposure to risk from inferences of her talent. In our setting, the agent avoids risk even though she is risk neutral because of the concave payoff profile induced from her bargaining power and a monopsonistic labor market. Another distinction is that an agent who is a residual claimant always makes efficient choices in our framework whereas this is not the case in Holmström (1999).

2. The Model

2.1 Environment and Production

There are two sectors: sector A and sector B. Associated with each sector is one risk-neutral firm so that the labor market in a sector is

efficiently. In our model, however, firms make wage offers after the sampling stage and therefore cannot influence the agent's sampling decision.

monopsonistic. There is a risk-neutral agent who chooses to work in either sector A or sector B.

The agent's talent in sector i, $i \in \{A, B\}$ is given by η^i. This talent is unknown and is distributed normally with mean 0 and variance $\sigma_\eta^2 > 0$. Talents across sectors are independent of one another. We relax this assumption of independence later in Section 6. Production depends on the talent of the agent in the sector. An agent who works in sector i produces an output η^i.

If the agent does not work in one of the two sectors, her reservation utility is minus infinity. The firm's reservation utility is 0.

2.2 Sampling Talents

Prior to working, the agent can *sample* (or learn about) her talents in a sector over two periods. An agent who samples sector i in period t, $t \in \{1, 2\}$, draws an informative signal $s_t^i = \eta^i + \epsilon_t^i$ at the end of the period, where ϵ_t^i is an idiosyncratic error term which is normally distributed with mean 0 and variance $\sigma_\epsilon^2 > 0$. The error terms are independent across periods.

The key constraint that the agent faces is that she can sample at most one type of talent per period. If the agent samples the same type of talent over both periods, we say that she *specializes*. If the agent samples different types of talents over both periods, we say that she *experiments*.

We assume that there are costs associated with not sampling a talent in a period. In the first period, we assume that this cost is prohibitively large. In the second period, the cost is $\phi > 0$. These costs can be thought of as costs to *access* a labor market. For instance, norms in a labor market may make it more difficult for an agent who has not sampled her talent to work in the market.

2.3 Bargaining Power and Incentives

An agent's wage is determined in the following way. With probability $\mu \in [0, 1]$, the agent makes a take it or leave it offer to the firm she chooses to work in, and with probability $1 - \mu$, the firms simultaneously make offers to the agent, who chooses between one of them.[5] The parameter μ can thus be interpreted as the agent's *bargaining power* to claim the returns from her talent. As we will see later, in equilibrium an agent with the power to make an offer extracts all of the surplus from her talent. A high μ thus corresponds to a setting where the agent is more likely to be a *residual claimant* of the returns to her talent. At the other end, a smaller μ exposes the agent to competition across firms in a labor market as in Holmström (1999). But the key difference here is that competition is between monopsonistic firms across sectors. We discuss some of these institutions in greater detail in Section 7.

5. If both firms offer the same wage, the agent randomly chooses either firm with equal probability.

2.4 Timing and Information Structure

There are three periods in the model: two sampling periods followed by a working period. Given the symmetry in the distribution of both talents, we assume without loss of generality that the agent samples sector A in the first period. Thus, for all of our analysis, we treat the realized signal in period 1 for talent A, s_1^A, as an exogenous parameter. The timing and information structure of the model then is as follows.

The agent samples sector A at the start of period 1. At the end of this period, she draws a publicly observable signal s_1^A. Conditional on this realized signal, the agent decides which sector (if any) to sample at the start of the second period. At the end of the second period, the signal s_2^i, where $i \in \{A, B\}$, is realized when sector i is sampled. In the beginning of period 3, the agent decides which sector to work in and wages are determined. Finally, at the end of period 3, production takes place. Figure 1 depicts the timing of the model.

The solution concept we use is subgame-perfect equilibrium. Given that the wage-offer subgame, where firms make simultaneous offers to the agent, has multiple Nash equilibria when expected talents across sectors differ, we use a refinement that selects equilibria with undominated strategies (Blume 2003; Kartik 2011).

3. Efficiency

We start in this section by characterizing the efficient sampling strategy—that is, the strategy that maximizes expected output—before turning to the agent's incentives in the following section. In particular, we compare the expected surplus (output) from specializing versus experimenting, given the realization of the first-period signal s_1^A.[6] We first sketch the total surplus functions associated with specialization and experimentation. We then compare the expected surplus across these two sampling strategies.

Consider the surplus function associated with specialization first. To convey the intuition for our results clearly, it is useful to work with a transformation of the second-period signal in sector A. In particular, define $\hat{s}_2^A = s_2^A - \lambda_1 s_1^A$, where $\lambda_1 = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_\epsilon^2}$. This normalized signal $\hat{s}_2^A$ has a mean of 0 and the same variance as the signal s_2^A. Let F^A and F^B be the distribution functions for $\hat{s}_2^A$ and s_2^B, respectively.

Because the agent can pick which sector to work in after sampling talents, the surplus from specialization is given by

$$\begin{aligned} \mathrm{TS}_S &= \max\{E(\eta_A | s_1^A, \hat{s}_2^A), E(\eta_B | s_1^A, \hat{s}_2^A)\} \\ &= \max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}, \end{aligned}$$

where $\lambda_2 = \frac{\sigma_\eta^2}{2\sigma_\eta^2 + \sigma_\epsilon^2}$.

6. Given the agent can choose to work in the sector which maximizes her talent, not sampling a talent is never efficient.

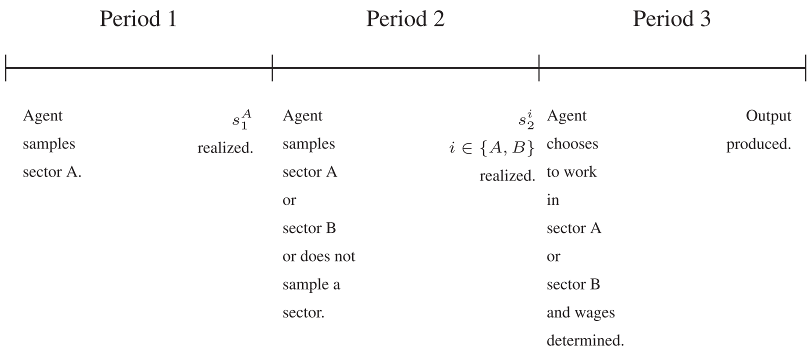

Figure 1. Timeline.

Similarly, the surplus from experimentation is given by

$$\begin{aligned}\mathrm{TS}_E &= \max\{E(\eta_A|s_1^A, s_2^B), E(\eta_B|s_1^A, s_2^B)\}\\ &= \max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}.\end{aligned}$$

The expected surplus from specialization, V_S, is then given by

$$V_S = E_{\hat{s}_2^A}[\mathrm{TS}_S] = E_{\hat{s}_2^A}\left[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}\right]$$

and the expected surplus from experimentation, V_E, is given by

$$V_E = E_{s_2^B}[\mathrm{TS}_E] = E_{s_2^B}\left[\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}\right].$$

Figures 2 and 3 plot the surplus from experimentation and specialization as a function of the realization of the second-period signal. Looking at these figures, it is not clear which of the two sampling strategies yields a higher expected surplus. Notice that the surplus functions overlap. Also, expectations are taken with respect to different random variables: $\hat{s}_2^A$ and s_2^B. Our main result in this section is that experimentation yields a higher expected surplus relative to specialization regardless of the initial draw of talent.

But first we state a useful Lemma.

Lemma 1. Let x be a normally distributed random variable with mean 0. Let a be a positive real number and let c and d be real numbers. Then $E_x[\max\{ax + c, d\}] = E_x[\max\{ax + d, c\}]$.

The proof of Lemma 1 is in Appendix A. The Lemma says that when a random variable is normally distributed with a mean of zero, then interchanging intercepts across components of the max function does not change the expected value of the max function. It is worth pointing out that the lemma above holds not just for a normal distribution but for any symmetric distribution with mean 0.

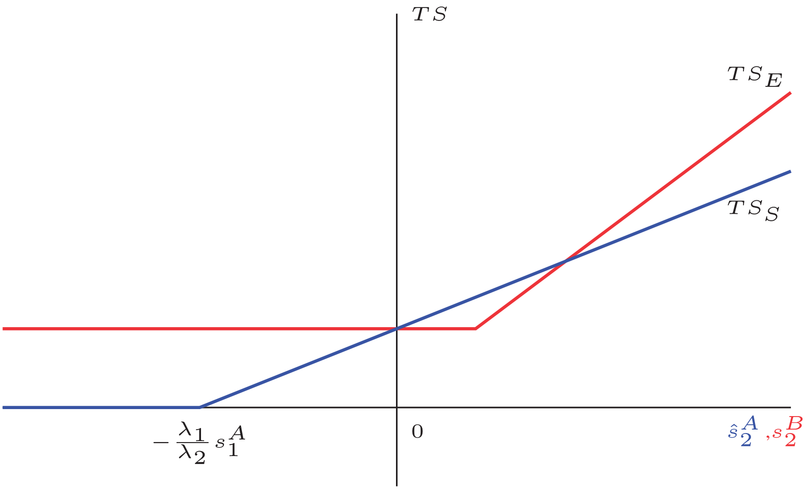

Figure 2. Total Surplus Functions: $s_1^A > 0$.

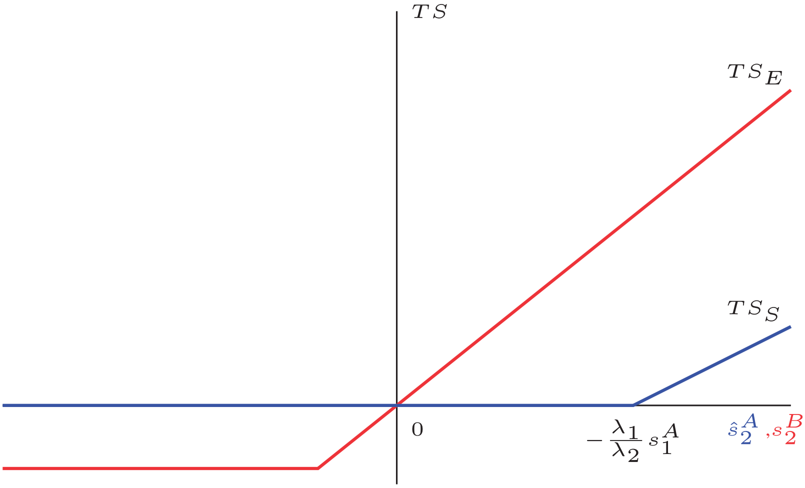

Figure 3. Total Surplus Functions: $s_1^A < 0$.

We now turn to our main result in this section.

Proposition 1. Experimentation, where the agent samples different sectors in each period, is efficient.

Proof. We split the proof into three claims.

Claim 1: $E_{\hat{s}_2^A}[\mathrm{TS}_S] = E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}] \leq E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_2 s_2^B, 0\}]$.

The distribution of signal $\hat{s}_2^A$ given s_1^A is $N(0, (1-\lambda_1^2)(\sigma_\eta^2 + \sigma_\epsilon^2))$. The distribution of signal s_2^B is $N(0, \sigma_\eta^2 + \sigma_\epsilon^2)$. Therefore, the two random variables $\hat{s}_2^A$ and s_2^B have the same mean but the former has smaller variance than the latter. Thus, $\hat{s}_2^A$ second order stochastically dominates s_2^B. Since the max function is convex, $E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}] \leq E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_2 s_2^B, 0\}]$.

Claim 2: $E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_2 s_2^B, 0\}] < E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_1 s_2^B, 0\}]$.

Consider two possible cases.

First, suppose $s_1^A \leq 0$. As $\lambda_1 > \lambda_2$, it follows that $\max\{\lambda_1 s_1^A + \lambda_2 s_2^B, 0\} \leq \max\{\lambda_1 s_1^A + \lambda_1 s_2^B, 0\}$ with the inequality strict for s_2^B sufficiently large. Thus, $E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_2 s_2^B, 0\}] < E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_1 s_2^B, 0\}]$.

Second, suppose $s_1^A > 0$. Then, $\max\{\lambda_1 s_1^A, \lambda_2 s_2^B\} \leq \max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}$ with the inequality strict for s_2^B sufficiently large. From Lemma 1, it follows that $E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_2 s_2^B, 0\}] = E_{s_2^B}[\max\{\lambda_1 s_1^A, \lambda_2 s_2^B\}] < E_{s_2^B}[\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}] = E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_1 s_2^B, 0\}]$.

Claim 3: $E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_1 s_2^B, 0\}] = E_{s_2^B}[\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}] = E_{s_2^B}[\mathrm{TS}_E]$.
This claim follows from Lemma 1.

Taking all three claims together, the result holds. ■

To understand the intuition for this result it helps to take a closer look at the surplus functions in Figures 2 and 3. In particular, notice that there is an *upside* effect: a high signal in the second period increases the posterior mean of the sampled talent and thus increases surplus, whereas a low signal entails no cost because the agent can switch to the non-sampled sector. It turns out that the upside effect is stronger in the case of experimentation for the following two reasons.

First, since the agent's talent in sector B is sampled for the first time in the case of experimentation, the weight placed on this signal is larger relative to the weight placed on the signal in the specialization case ($\lambda_1 > \lambda_2$). This is because a signal drawn for the first time is more informative about talent.

Second, both the signals $\hat{s}_2^A$ and s_2^B have the same mean of 0, but the signal in sector B, which is drawn for the first time, has larger variance.[7] Or put differently, the signal $\hat{s}_2^A$ second-order stochastically dominates the signal s_2^B. This is because less is known about a talent which is sampled for the first time.

To summarize, there is *more to learn* from experimentation: the weight placed on signal B when updating beliefs is stronger ($\lambda_1 > \lambda_2$) and extreme values of signal B are more likely ($\hat{s}_2^A$ second-order stochastically dominates s_2^B). As a result, the upside effect is larger for experimentation. This larger upside effect combined with the symmetry of the normal distribution ensures that experimentation yields a higher expected surplus relative to specialization.

7. The signal $\hat{s}_2^A \sim N(0, (1-\lambda_1^2)(\sigma_\eta^2 + \sigma_\epsilon^2))$, whereas the signal $s_2^B \sim N(0, \sigma_\eta^2 + \sigma_\epsilon^2)$.

Given that experimentation always does better than specialization, we now look at how the difference in the expected surplus across both of these cases varies as we vary parameters in our model.

Proposition 2. $V_E - V_S$ is

(i) strictly increasing in σ_η^2 and
(ii) strictly decreasing in $|s_1^A|$ and tends to 0 as $|s_1^A|$ tends to infinity.

The proofs of all the propositions that follow are in Appendix A. The intuition for the first part of this proposition is clear. As the variance of talents gets larger, there is more to learn from experimentation, which makes it more valuable relative to specialization. Part (ii) of the proposition is less obvious and says that the gains from experimentation are the largest for intermediate draws of talent, and that in the limit (for very good or very bad draws of talent) these gains disappear.

To see why the second part of Proposition 2 holds, notice from Lemma 1 that we can rewrite $E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}] = E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A, \lambda_2 \hat{s}_2^A\}]$. Comparing this expression with the expected surplus from experimentation, which is $E_{s_2^B}[\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}]$, we see that two things matter: the floor of the total-surplus function $\lambda_1 s_1^A$ which is common across both expressions, and the inferences drawn from the second-period signal across sectors ($\lambda_2 \hat{s}_2^A$ versus $\lambda_1 s_2^B$). Ignoring the floor, we see that positive second-period signals favor experimentation, with its steeper upside and larger variance of the signal in sector B. Negative second-period signals favor specialization with a flatter downside and lower variance of the signal in sector A. When $s_1^A > 0$, the floor curtails some of the upside benefit from experimentation. And when $s_1^A < 0$, the floor accommodates some downside cost from experimenting. It is at $s_1^A = 0$, where all of the upside benefits from experimentation are realized without any downside costs, where the difference in the value across experimentation and specialization is the largest.

4. Incentives to Sample Talents

We now turn to incentives that the agent has to sample her talents. When she has the power to make an offer to the firm (this occurs with probability μ), she chooses to work in the sector where her expected talent is larger and she asks for, and gets, a wage that equals her expected talent in that sector. When she does not have the power to make the offer (this occurs with probability $1 - \mu$), there are multiple equilibria in the wage-offer subgame when expected talents differ across sectors. Drawing on Kartik (2011), we use a refinement that selects equilibria with undominated strategies when there are multiple equilibria. Any of these equilibria lead to the outcome where the agent works for the firm where she is more productive and where the equilibrium wage equals her expected talent in the next best sector.[8,9]

Thus, her expected utility from specializing, experimenting, and not sampling a talent in the second period is given by

$$\mathrm{EU}_S = E_{\hat{s}_2^A}[w_S] = E_{\hat{s}_2^A}\left[\mu\max\left\{\lambda_1 s_1^A + \lambda_2\hat{s}_2^A, 0\right\} + (1-\mu)\min\left\{\lambda_1 s_1^A + \lambda_2\hat{s}_2^A, 0\right\}\right],$$
$$\mathrm{EU}_E = E_{s_2^B}[w_E] = E_{s_2^B}\left[\mu\max\left\{\lambda_1 s_1^A, \lambda_1 s_2^B\right\} + (1-\mu)\min\left\{\lambda_1 s_1^A, \lambda_1 s_2^B\right\}\right],$$
$$\mathrm{EU}_N = w_N - \phi = \mu\max\left\{\lambda_1 s_1^A, 0\right\} + (1-\mu)\min\left\{\lambda_1 s_1^A, 0\right\} - \phi,$$

where w_S, w_E, and w_N are the expected wages from specialization, experimentation, and not sampling a talent, respectively.

Let $V_S(s_1^A)$ be the agent's expected value from specialization as a function of s_1^A. The proposition below shows the agent's optimal sampling strategy.

Proposition 3

(i) Let $\mu > \frac{1}{2}$. Then the agent experiments.
(ii) Let $\mu = \frac{1}{2}$. Then the agent is indifferent between experimenting and specializing.
(iii) Let $\mu < \frac{1}{2}$. When $(1-2\mu)V_S(0) \le \phi$ the agent specializes. Otherwise, there exist cutoffs, $\underline{s}_1^A < 0$ and $\bar{s}_1^A > 0$, such that the agent specializes when $s_1^A \le \underline{s}_1^A$ or $s_1^A \ge \bar{s}_1^A$, and she does not sample a talent when $s_1^A \in (\underline{s}_1^A, \bar{s}_1^A)$. Furthermore, $|\underline{s}_1^A|$ and $|\bar{s}_1^A|$ are decreasing in the agent's bargaining power μ and the cost of not sampling a talent ϕ.

Proposition 3 says that an agent experiments as long as her bargaining power is sufficiently high. The cutoff that induces experimentation ($\mu = \frac{1}{2}$) lies right in the middle of the two polar cases: $\mu = 1$, where the agent is the residual claimant, and $\mu = 0$, where the agent's wage is her next best option.

To gain intuition for Proposition 3, it is useful to make two separate comparisons: one comparison between experimentation and specialization and the other between sampling and not sampling a talent.

Consider experimentation versus specialization first, and focus on the two extremes of bargaining power. When $\mu = 1$, the agent's wage for a given realization of the signal in the second period, s_2, is the maximum of

8. For the argument in Kartik (2011) to work in our setting, we need to assume that firms face a constraint on setting wages too low for a given set of posterior means of talents.

9. An example of an equilibrium in undominated strategies is where the more productive firm (where the worker has a higher talent) offers a wage equal to $\min\{\hat{\eta}^A, \hat{\eta}^B\}$ and where the less productive firm randomizes uniformly over the interval $[\min\{\hat{\eta}^A, \hat{\eta}^B\} - \delta, \min\{\hat{\eta}^A, \hat{\eta}^B\}]$, where $\hat{\eta}^i$ is the posterior mean of the agent's talent in sector i at the start of period 3, and where $\delta > 0$ is sufficiently small. This example is from Blume (2003).

her talents across sectors. Because the posterior means are linear (from normally distributed talents and signals) and because the max function is convex, the agent's wage as a function of s_2 is convex. In fact, the wages from specialization $w_S = \max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}$ and experimentation $w_E = \max\{\lambda_1 s_1^A + \lambda_1 s_2^B, 0\}$ (modified using Lemma 1) resemble call options with strike prices of $-\frac{\lambda_1 s_1^A}{\lambda_2}$ and $-s_1^A$, respectively. For this case with convex payoffs, experimentation—with its larger variance when the talent is sampled for the first time and a steeper upside—dominates specialization.

When $\mu = 0$, the agent's wage is concave in s_2 so that this function can be thought of as a put option. Receiving a high signal in the second period yields no benefit as wages are capped above by the non-sampled sector, while low signals entail a downside because they reduce the agent's next best option. For this case with concave payoffs, specialization—which involves a lower residual variance from sampling a talent for the second time and a flatter downside—dominates experimentation.

For an intermediate $\mu \in (0, 1)$, the agent's expected payoff is a convex combination of these two polar cases: for $\mu > \frac{1}{2}$, expected wages are piecewise linear and convex in the second-period signal so that experimentation dominates specialization, for $\mu < \frac{1}{2}$, wages are piecewise linear and concave in the second-period signal so that specialization dominates experimentation, and finally for $\mu = \frac{1}{2}$, expected wages are linear with the agent being indifferent between both strategies.

Figures 4 and 5 sketch wages for the two polar cases, $\mu = 1$ and $\mu = 0$, and expected wages for the intermediate case of $\mu = \frac{1}{2}$ for a positive first-period signal.

Next, let's turn to the comparison between sampling a talent or not in period 2. While specialization slows down learning relative to experimentation, not sampling a talent in period 2 completely brings learning to a halt. Thus, this strategy has more value when learning by the market is used to penalize the agent for a bad draw relative to rewarding the agent for a good one. Or put differently, this strategy is useful when the agent's wage is piecewise linear and concave in s_2 with $\mu < \frac{1}{2}$. The last part of Proposition 3 confirms this intuition. In addition, the proposition above says that not sampling can only be optimal for intermediate draws of talent. For a low signal, the wage from not learning is low to begin with so that the downside matters less. For a high signal, once again, the downside from sampling a talent matters less, in this case because the downside is less likely.

5. A Role for Specialization

Our main result in Proposition 1 is that experimentation is efficient regardless of the initial draw of talent. This result, however, relies on several assumptions: talents are symmetrically and normally distributed, and there is no human capital acquired during the learning phase. In this

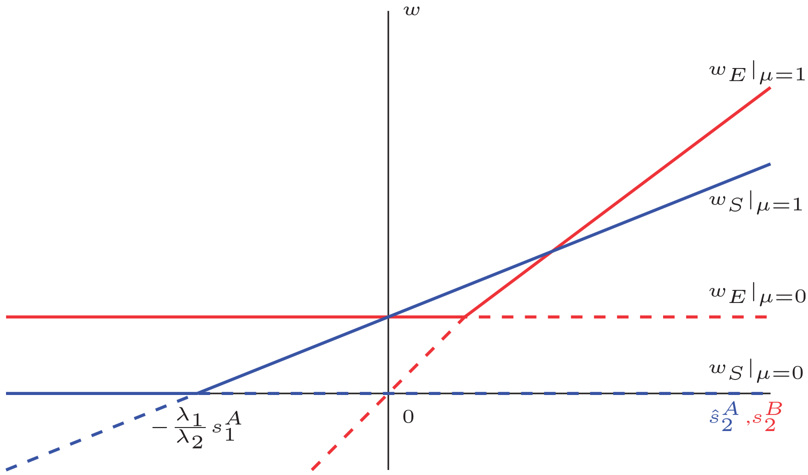

Figure 4. Bargaining Power and Wages: $\mu \in \{0, 1\}$, $s_1^A > 0$.

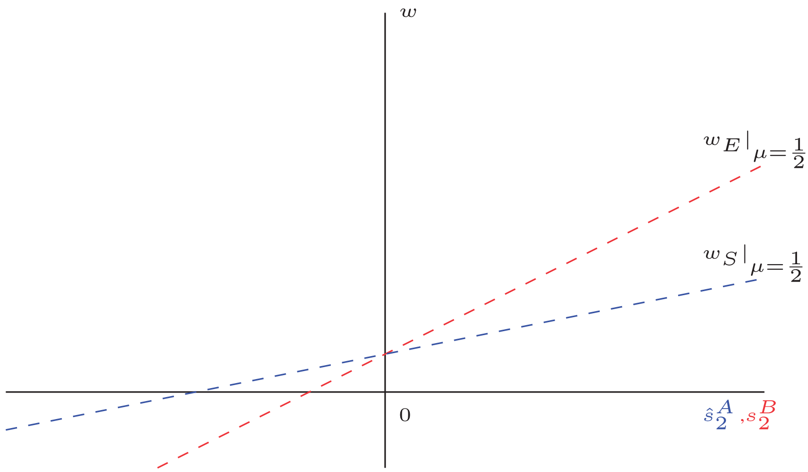

Figure 5. Bargaining Power and Wages: $\mu = \frac{1}{2}$, $s_1^A > 0$.

subsection, we show that specialization has a role to play from an efficiency viewpoint when these assumptions are relaxed.

5.1 Human Capital Accumulation

We now introduce human capital accumulation into our analysis. When an agent samples a sector, she does not just get a signal of her talent; she also acquires human capital $H > 0$. Output in each sector is the agent's talent plus her human capital. Given this specification, we can rewrite the surplus functions as:

$$\mathrm{TS}_S^H = \max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A + 2H, 2\alpha H\}$$

and

$$\mathrm{TS}_E^H = \max\{\lambda_1 s_1^A + (1+\alpha)H, \lambda_1 s_2^B + (1+\alpha)H\},$$

where $\alpha \in [0, 1]$ is a spillover parameter that captures the extent to which the human capital is specific or general across sectors.[10] When $\alpha = 1$, human capital is fully general across sectors, and when $\alpha = 0$, it is fully specific to a sector. The value from specialization for this case with human capital is defined as $V_S^H \equiv E_{\hat{s}_2^A}[\mathrm{TS}_S^H]$, and the value from experimentation as $V_E^H \equiv E_{s_2^B}[\mathrm{TS}_E^H]$.

The following proposition characterizes the efficient sampling strategy with human capital.

Proposition 4

(i) When $\alpha = 1$, so that human capital is fully general across sectors, experimentation is efficient.

(ii) For any $\alpha < 1$, there exist cutoffs, s_1^{*A} and s_{*1}^A, such that specialization is efficient when $s_1^A \geq s_1^{*A}$, and experimentation is efficient when $s_1^A \leq s_{*1}^A$.

In this setup with human capital, the difference in value between experimentation and specialization $V_E^H - V_S^H$ can be decomposed into the sum of two parts: a difference in value from learning, $V_E - V_S$, and a difference in value from allocating human capital. With specialization, human capital is *concentrated* in sector A whereas experimentation *spreads* human capital equally across both sectors.

When human capital in one sector completely spills over to the other sector, so that it is fully general across sectors ($\alpha = 1$), experimentation and specialization generate the same surplus in terms of allocating human capital. Thus, experimentation, which leads to more efficient learning from Proposition 1, is efficient.

With spillovers across sectors ($\alpha < 1$), the initial signal plays an important role. When s_1^A is high, the probability that the agent eventually works in sector A is large, so that concentrating human capital in sector A through specialization leads to a more efficient allocation of human capital. By contrast, a low s_1^A increases the likelihood of working in sector B so that experimentation allocates human capital efficiently.

10. Alternatively we could model human capital accumulated in a sector as a sum of two terms: the innate talent in the sector and a fixed term αH. The results for this case are qualitatively similar to Proposition 4.

The second part of Proposition 4 reflects a tradeoff between allocating human capital in an efficient way and learning. When the initial signal is very large, human capital plays a more important role relative to learning. Indeed, in the limit, when s_1^A tends to infinity, the probability of working in sector A approaches 1, so that the human capital gain in value from specializing is $(1-\alpha)H$, whereas the gains from learning via experimentation vanish in the limit (from Proposition 2). As a result, specialization is efficient for this case. For a very low signal, experimentation does better, both in terms of learning and human capital, so that it is efficient.

Under some additional conditions, the cutoffs in the second part of Proposition 4 coincide so that there is a unique threshold above which specialization is efficient and below which experimentation is efficient. Furthermore, this cutoff is increasing in α, and as $\alpha \to 1$, the cutoff tends to infinity.[11]

5.2 Asymmetric Model

So far in our model, sectors are symmetric: talents in both sectors have the same mean and the same variance. In this section, we allow for asymmetries across sectors.

Let $\eta^A \sim N(0, \sigma_\eta^2)$ and $\eta^B \sim N(b, v\sigma_\eta^2)$ where $v > 0$ and where b is any real number. As before, let $\epsilon_t^i \sim N(0, \sigma_\epsilon^2)$ for $i \in \{A, B\}$, $t \in \{1, 2\}$. Thus, the unconditional distribution of the signal s_2^B is normal with mean b and variance $v\sigma_\eta^2 + \sigma_\epsilon^2$. Given these asymmetries across sectors, the agent's choice of which sector to sample initially is made endogenous.

The following proposition characterizes conditions under which experimentation is efficient.

Proposition 5. Experimentation is efficient if and only if $\frac{\sigma_\eta^2}{\sigma_\epsilon^2} \geq \frac{|v-1|}{v}$. Otherwise, specialization in sector A is efficient when $v < 1$, and specialization in sector B is efficient when $v > 1$.

Notice that the parameter b plays no role in the necessary and sufficient condition for experimentation—only the variances of talents matter. Second, the condition in the proposition for efficient experimentation has a simple interpretation. The left-hand side $\frac{\sigma_\eta^2}{\sigma_\epsilon^2}$ is the signal-to-noise ratio. The right-hand side $\frac{|v-1|}{v}$ reflects the degree of asymmetry across talents as v moves away from 1. The condition above then says that experimentation is efficient if and only if the signal-to-noise ratio is sufficiently large relative to the degree of asymmetry across sectors. In particular, a larger variance in a sector may lead to that sector being sampled repeatedly as there

11. The reason the cutoff may not be unique is that $V_E - V_S$ is non-monotone in s_1^A. A sufficient condition that guarantees a unique cutoff is that H is small enough so that $(V_E - V_S)|_{s_1^A = \frac{2(1-\alpha)H}{\lambda_2}} \geq \int_{\frac{-2(1-\alpha)H}{\lambda_2}}^{0} 2\lambda_2 \hat{s}_2^A (1-\alpha) H dF^A$. This condition ensures that experimentation is efficient for all negative first-period signals so that there is a unique cutoff that is positive. Using the implicit function theorem, the sign of the derivative of the cutoff with respect to α is positive if and only $g_2(s_1^A, \alpha) = \int_{\frac{-(\lambda_1 s_1^A + 2(1-\alpha)H)}{\lambda_2}}^{0} dF^A$ is positive which holds for any positive s_1^A.

is more to learn: for a small v it is efficient to specialize in sector A and for a large v it is efficient to specialize in sector B.

5.3 How Sensitive Is the Result to the Normal–Normal Model?

The fact that the result that experimentation is more efficient than specialization is independent of the realization of the signal drawn in the first period is surprising. We conjecture that the independence on the first-period signal is specific to the normal–normal model and, more specifically, to the property that the variance of the updated normal distribution is independent of the first-period signal.

To explore this conjecture we analyze an alternative but closely related information structure. We assume that both the mean and the variance of the signal distributions are unknown. Let τ^i denote the precision (or inverse of the variance) of talent in sector $i \in \{A, B\}$. The agent has a normal-Gamma prior where, conditional on τ^i, η^i follows a normal distribution with mean zero and variance $(\frac{\sigma_\epsilon^2}{\sigma_\eta^2}\tau^i)^{-1}$, and τ^i follows the Gamma distribution $\mathrm{Ga}(\tau^i | \frac{\nu}{2}, \frac{\nu-2}{2}\sigma_\epsilon^2)$. With these assumptions the agent's talent, η^i, in sector i follows a Student's t-distribution with $\nu > 2$ degrees of freedom, a mean of zero, and scale parameter of $\frac{\nu-2}{\nu}\sigma_\eta^2$, that is,

$$\eta^i \sim t_\nu\left(0, \frac{\nu-2}{\nu}\sigma_\eta^2\right).$$

Conditional on η^i and τ^i, signals are normally distributed with mean η^i and variance $(\tau^i)^{-1}$. With normal-Gamma prior distributions and normally distributed signals, the posterior distributions are also normal-Gamma distributions and the marginal posteriors of the agent's talents are t-distributions (DeGroot 1970). The posterior means of the agent's talents are the same as for the normal–normal model.

Similarly, the unconditional distribution of the first signal and the conditional distribution of the second signal given the first signal follow Student's t-distributions. More specifically,

$$s_2^B \sim t_\nu\left(0, \frac{\nu-2}{\nu}\left(\sigma_\eta^2 + \sigma_\epsilon^2\right)\right)$$

and

$$\hat{s}_2^A | s_1^A \sim t_{\nu+1}\left(0, \left(\frac{\nu-2}{\nu+1} + \frac{1}{\nu+1}\frac{(s_1^A)^2}{\sigma_\eta^2 + \sigma_\epsilon^2}\right)\left(1 - \lambda_1^2\right)\left(\sigma_\eta^2 + \sigma_\epsilon^2\right)\right).$$

These posterior distributions are very similar to the ones obtained in the normal–normal model. Posterior means are identical and, as $\nu \to \infty$, the variances and distributions converge to those of the normal–normal model.

The crucial difference to the normal–normal model is that the posterior variance of the second signal in sector A depends on the first signal. The greater the magnitude of the first signal, the greater the posterior variance of the second signal. If s_1^A is very high or very low, the posterior variance of the second signal from sector A can get larger than the unconditional variance of the signal from sector B. In this case, it can be efficient to sample from sector A again and, thus, specialize.

This point is illustrated in Figure 6, which shows the difference in expected surplus from experimentation and specialization. Here, $\sigma_\eta^2 = \sigma_\epsilon^2 = 1, \lambda_1 = 0.5, \lambda_2 = 0.33$ and $\nu = 3$. Figure 6 confirms our conjecture: The result that experimentation is more efficient than specialization *for all* realizations of the first signal relies on the normal–normal model. In Figure 6, specialization is more efficient when the first signal is below -3.4 or above 3.4. Note that for the parameter values used in Figure 6, the signal s_1^A follows a *t*-distribution with three degrees of freedom, mean zero, and standard deviation of $\sqrt{2}$. Given this distribution, the probability of choosing a signal below -3.4 or above 3.4 is less than 10%. Thus, for the majority of realized signals, experimentation is still efficient.

6. Richer Labor Market Settings

We now modify our framework to consider richer labor market settings. The objective is to show that our main results and insights are robust in these alternative settings.

6.1 A Third Sector

In our model, the agent's reservation utility from not working in either sector A or sector B is minus infinity. In this subsection, we allow for a third sector, sector C, with a firm in the sector that the agent can work in. The agent's talent in this sector is given by η^C and is normally distributed with mean 0 and variance $\mathrm{Var}(\eta^C) \in [0, \sigma_\eta^2]$ and is independent of the talents in the other sectors. Because the variance of talent is lower in sector C, it can be interpreted as a *safe* sector.

As before, the agent can sample at most one sector in each period and there are costs to not sampling a talent in each period.

Given that sector C has a lower variance of talent relative to the other two sectors, there is no loss of generality in assuming that the agent does not sample sector C, and that the agent starts by sampling sector A.[12]

The total surplus functions from specialization and experimentation are then

$$\mathrm{TS}_S = \max\{E(\eta_A|s_1^A, \hat{s}_2^A), E(\eta_B|s_1^A, \hat{s}_2^A), E(\eta_C|s_1^A, \hat{s}_2^A)\}$$

$$= \max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\},$$

and

12. This follows from the same reasoning used in Propositions 1 and 2.

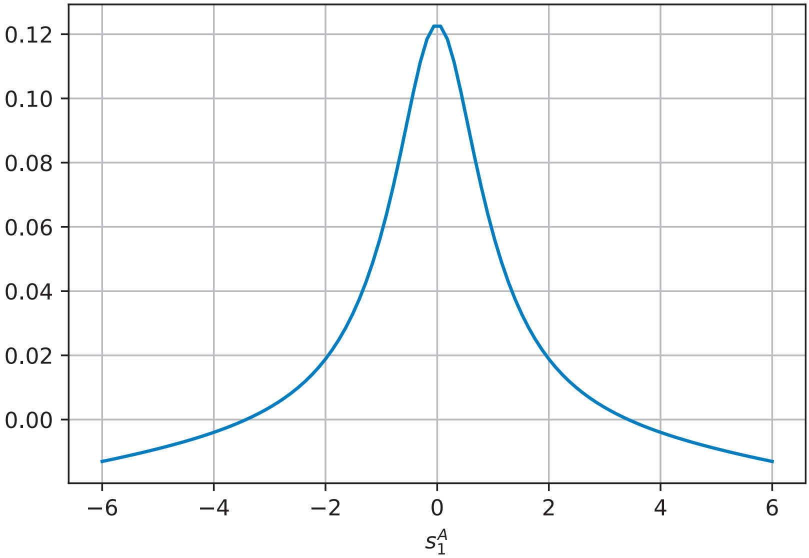

Figure 6. The difference in expected surplus from experimentation over specialization ($V_E - V_S$).

$$\begin{aligned}\mathrm{TS}_E &= \max\{E(\eta_A|s_1^A, s_2^B), E(\eta_B|s_1^A, s_2^B), E(\eta_C|s_1^A, s_2^B)\}\\ &= \max\{\lambda_1 s_1^A, \lambda_1 s_2^B, 0\}.\end{aligned}$$

Notice that the surplus function for specialization is exactly the same as the two-sector case—the presence of the third sector adds no value as a fallback option as sector B (with the same mean of zero) already plays that role. Under experimentation, however, the agent can fallback on sector C when the draws in the other two sectors are bad. This raises the value of experimentation relative to specialization. Thus, experimentation is once again efficient so that our main result in Proposition 1 goes through.

To examine incentives to sample talent, we consider the two polar cases in terms of the agents bargaining power: $\mu = 1$ where the agent has all of the bargaining power and $\mu = 0$ where competition across sectors leads to the agent's wage equaling her talent in her next-best sector. The following proposition characterizes the agent's optimal choice for these two polar cases.

Proposition 6. Suppose the agent has access to a third sector, sector C, with $E[\eta^C] = 0$ and $\mathrm{Var}(\eta^C) \in [0, \sigma_\eta^2]$.

(i) Let $\mu = 1$. Then the agent experiments.
(ii) Let $\mu = 0$. Then the agent experiments if and only if $s_1^A \geq 0$. Otherwise she specializes.

The first part of the proposition above is intuitive and in line with our earlier result in Proposition 3. When an agent has all the bargaining power ($\mu = 1$), she is the residual claimant, giving her incentives to sample efficiently by experimenting. Where the results are different from Proposition 3 is for the case where the agent has no bargaining power. For this case, unlike in Proposition 3, experimentation is optimal for positive initial draws in sector A.

To see the intuition for this proposition, it is useful to note that the agent's wage—which is her talent in her next-best sector—lies between the means of the two non-sampled sectors in period 2.[13] The lower of these two means forms a wage floor, and the higher of the two, a wage ceiling. Under specialization, the symmetry in the prior means across sectors implies that the floor and the ceiling coincide at 0, so that the agent's wage is always 0. Under experimentation, a positive initial signal in sector A imposes a positive ceiling from sector A and a floor of 0 from sector C so that on average the wage is positive. And for a negative initial signal in sector A, the floor from sector A is negative and the ceiling from sector C is 0 so that the expected wage is negative. Taken together, the agent experiments if and only if the initial signal is non-negative. Otherwise, she specializes and gets a wage of 0.

In this section, we assume that sector C has a lower variance of talent. In addition to this, we can also think of sector C as having a lower prior mean for the agent's talent relative to the other sectors so that it is *dominated* by them. It is straightforward to show that experimentation, where the agent samples sectors A and B, is efficient in this case. And, consistent with the second part of Proposition 6, there are inefficiencies in the agent's sampling strategy when she has no bargaining power. The optimal sampling strategies are, however, more complicated. In particular, the agent may sample sector C to raise the floor associated with the agent's wage.

When the prior mean of talent in sector C is low and the variance of talent in this sector is close to that of the other sectors, sampling sector C in the first period, followed by specialization for low initial signals and experimentation otherwise is optimal. This strategy props up the wage floor, and guarantees that wages are non-negative.

By contrast, when the prior mean of talent in sector C is high and the variance of talent in this sector is low relative to the other sectors, sampling sector A in the first period, followed by specialization for low initial draws of the signal, experimentation with sector C for intermediate draws

13. The main argument in Kartik (2011), that undominated equilibria lead to the conventional outcome where the agent works in the firm where she is most productive for a wage that equals the expected value of her next best talent, continues to hold when there are more than two firms, except for the case where the agent's maximum expected talent is the same across two sectors. But for this case, any Nash equilibrium in the wage-offer subgame yields the conventional outcome above.

of the initial signal, and experimentation with sector B for large initial signals, is optimal. This strategy allows the agent to capture the gains from experimentation by sampling the sector with a higher variance of talent and by raising the ceiling associated with wages.

6.2 Correlated Talents

In our main model, talents are independent of one another. In many real world settings, however, talents can be correlated. We consider this case with correlation now.

Let talents in sectors A and B be correlated with correlation coefficient $\rho \in [-1, 1]$. The surplus from specialization is then given by

$$\begin{aligned}\mathrm{TS}_S^{\rho} &= \max\{E(\eta_A | s_1^A, \hat{s}_2^A), E(\eta_B | s_1^A, \hat{s}_2^A)\}\\ &= \max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, \rho(\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A)\}.\end{aligned}$$

Similarly, the surplus from experimentation is given by

$$\begin{aligned}\mathrm{TS}_E^{\rho} &= \max\{E(\eta_A | s_1^A, s_2^B), E(\eta_B | s_1^A, s_2^B)\}\\ &= \max\{\kappa_1 s_1^A + \kappa_\rho s_2^B, \kappa_1 s_2^B + \kappa_\rho s_1^A\},\end{aligned}$$

where $\kappa_1 = \lambda_1 \frac{1-\lambda_1\rho^2}{1-\lambda_1^2\rho^2}$ and $\kappa_\rho = \rho \frac{\lambda_1(1-\lambda_1)}{1-\lambda_1^2\rho^2}$ (DeGroot 1970, p. 175). For $\rho = 0, \kappa_1 = \lambda_1$ and $\kappa_\rho = 0$ and, thus, we are back at our baseline model.

Proposition 7. If talents in the two sectors are correlated with correlation coefficient $\rho \in (-1, 1)$, experimentation, where the agent samples different sectors in each period, is efficient. If $\rho = -1$ or $\rho = 1$, experimentation and specialization are equally efficient.

To see the intuition for the proposition, let's start with the two extreme cases where talents are perfectly correlated with $\rho = -1$ or $\rho = 1$. For both of these cases $\kappa_1 = |\kappa_\rho| = \frac{\lambda_1}{1+\lambda_1} = \lambda_2$ so that the weights put on the signals s_1^A and s_2^i, $i \in \{A, B\}$, are the same regardless of the sector sampled in period 2.[14] Furthermore, the posterior variance of talent at the end of period 1 is the same across sectors and equals $(1-\lambda_1)\sigma_\eta^2$. Put differently, when talents are perfectly correlated, the informational content of a signal is the same for both sectors so that specialization and experimentation are equally efficient.

For intermediate cases of correlation $\rho \in (-1, 1)$, the intuition comes across most clearly when the initial signal $s_1^A = 0$. For this case, the difference in values across experimentation and specialization $V_E^{\rho} - V_S^{\rho}$ can be expressed as

14. Note that $\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A = \lambda_2 s_1^A + \lambda_2 s_2^A$.

$$\begin{aligned} V_E^{\rho} - V_S^{\rho}|_{s_1^A=0} &= E_{s_2^B}\left[\max\{\kappa_\rho s_2^B, \kappa_1 s_2^B\}\right] - E_{\hat{s}_2^A}\left[\max\{\lambda_2 \hat{s}_2^A, \rho\lambda_2 \hat{s}_2^A\}\right] \\ &= E_{s_2^B}\left[\max\left\{(\kappa_1-\kappa_\rho)s_2^B, 0\right\}\right] - E_{\hat{s}_2^A}\left[\max\{(1-\rho)\lambda_2 \hat{s}_2^A, 0\}\right] \\ &= E_{s_2^B}\left[\max\left\{\frac{(1-\rho)\lambda_1 s_2^B}{1-\lambda_1\rho}, 0\right\}\right] - E_{\hat{s}_2^A}\left[\max\left\{\frac{(1-\rho)\lambda_1 \hat{s}_2^A}{1+\lambda_1}, 0\right\}\right], \end{aligned}$$

where the second line follows from the symmetry of the normal distribution. As in Proposition 1, more weight is placed on a signal when sector B is sampled for the first time: $\frac{(1-\rho)\lambda_1}{1-\lambda_1\rho} > \frac{(1-\rho)\lambda_1}{1+\lambda_1}$ when $\rho \in (-1,1)$. And similarly, the signal $\hat{s}_2^A$ second-order stochastically dominates the signal s_2^B: the signal $\hat{s}_2^A \sim N(0,(1-\lambda_1^2)(\sigma_\eta^2+\sigma_\epsilon^2))$, whereas the signal $s_2^B \sim N(0,(1-\lambda_1^2\rho^2)(\sigma_\eta^2+\sigma_\epsilon^2))$. Taken together, experimentation is more efficient than specialization even when talents are correlated. The proof in Appendix A shows that this reasoning continues to hold when the initial signal is not zero.

Interestingly, the gains from experimenting over specializing are not maximized when talents are independent across sectors. To see why, it is useful to consider two effects: the effect of a signal on its *own* sector (captured by the coefficients κ_1 and λ_2 above) and the effect of a signal on the *other* sector (captured by the coefficients κ_ρ and $\rho\lambda_2$ above). The gains from experimenting over specializing that arise from the own-sector effects are symmetric around $\rho = 0$. The gains from the other-sector effects are, however, asymmetric. For a positive ρ, experimentation, with a lower downside effect for the other sector (as $\kappa_\rho < \rho\lambda_2$), dominates specialization. For a negative ρ, specialization, with its larger upside effect for the other sector (as $|\rho\lambda_2| > |\kappa_\rho|$), dominates experimentation. There is thus an asymmetry in the overall gains from experimenting over specializing which is maximized for some positive ρ.[15]

Finally, it is also useful to note that the difference in expected utilities across experimentation and specialization can be written as $(2\mu-1)\left(V_E^{\rho} - V_S^{\rho}\right)$ so that the results from Proposition 3 continue to hold for this case where talents are correlated.

6.3 Endogenous Bargaining Power

So far, the bargaining power of an agent to extract surplus μ is exogenous and independent of the sampling profile. In some settings, bargaining power may be endogenous. For instance, if we think of sampling vertically across degrees in the case of education, then it seems plausible that a specialist may have more bargaining power relative to an agent with a mixed background.

In this subsection, we allow for bargaining power to depend on the sampling profile and on the sector that the agent works in. The way we do this

15. Numerical examples suggest that this asymmetric effect where the gains are maximized for some positive ρ continues to hold for $s_1^A \neq 0$.

is to assume that the agent's bargaining power in a sector is larger when she samples it twice. With this specification, the bargaining power is constant across sectors when the agent experiments: $\mu_E^A = \mu_E^B = \mu$. But for specialization, the agent's bargaining power is larger in sector A: $\mu_S^B = \mu$ whereas $\mu_S^A = \mu + \delta$, with $0 \leq \delta \leq 1 - \mu$.

Proposition 8. Suppose $\delta > 0$.

(i) Let $\mu > \frac{1}{2}$. For s_1^A sufficiently large, the agent prefers specialization over experimentation.
(ii) Let $\mu \leq \frac{1}{2}$. Then the agent prefers specialization over experimentation.

The proposition above says that endogenous bargaining power that favors a specialist makes inefficiencies in sampling talents worse. In particular, the agent prefers specialization over experimentation for a range of initial signals even when $\mu > \frac{1}{2}$ which is in contrast to Proposition 3.

6.4 Labor Market Competition

In our baseline model, we assume that the labor market is monopsonistic with one firm in each sector. In this subsection, we allow for more firms in each sector to introduce labor market competition within a sector. In particular, suppose there are $N \geq 2$ firms in each sector. Once again, with probability μ the agent offers a take it or leave it offer to the firm she chooses to work in. With probability $1 - \mu$, a set of *active* firms in the labor market makes offers simultaneously to the agent. Suppose that with probability $\gamma \in [0, 1]$ this set of active firms includes more than one firm in each sector whereas with probability $1 - \gamma$ the agent, as before, gets offers from only one firm in each sector. The parameter γ can be thought of as how thick labor markets in a sector are.

With this setup, with probability $\hat{\mu} = \mu + (1 - \mu)\gamma$ the agent is a residual claimant and with probability $1 - \hat{\mu} = (1 - \mu)(1 - \gamma)$ the agent's wage is her next best talent in another sector. The parameter $\hat{\mu}$ can be thought of as the agent's effective bargaining power, and is increasing in the probability of her making a take it or leave it offer μ and in γ which measures market thickness and thus competition in the labor market.

With more firms in a sector, Proposition 3 continues to hold, but with μ replaced by the agent's effective bargaining power $\hat{\mu}$. The key point to note is that as long γ and μ are small enough, the agent does not efficiently sample her talents.

7. Discussion

Two key features underpin our model. First, talents are *publicly* sampled *prior* to the agent working. Second, an agent's bargaining power and the structure of the labor market affects her incentives to learn about her talents. In this section, we highlight economic settings where these features play a prominent role. In particular, we provide details of the sampling process, and spell out the sources from which incentives arise in these settings. This allows us to apply the main results from our model to these settings and derive policy implications.

The most natural setting where we believe talents are sampled prior to work is education. Experimentation in this context can involve broader course work which is the defining feature of higher education in countries like the USA.[16] Or it could be thought of as a vertical switch in fields across degrees: for example, a switch from an undergraduate degree in engineering to a graduate degree in business or law.

With grades and references publicly observable, incentives arise from an *external* labor market.[17] Our main policy implication within this context is that an education system that encourages experimentation is not enough to lead to efficient choices at the individual level. These systems have to go hand in hand with labor market institutions that encourage residual claimancy and thus induce risk taking through experimentation. Monopsonistic labor markets in a sector hinder efficient risk taking.

A necessary feature of labor markets for our argument above to work is that there is cross-hiring across sectors with firms in one sector being able to interpret grades and references in fields related to other sectors. We believe that cross-hiring and interpreting signals of talent across sectors is plausible in many labor market settings, but especially for fields where talents are positively correlated. For example, economists are hired for business-related jobs and vice-versa, and signals of talent are quite easily interpretable across these sectors. Other examples of occupational pairs with positively correlated talents and the potential to cross-hire include finance and physics, psychology and education, and computer science and math. Cross-hiring is also a common feature in team settings where individuals from different fields collaborate with each other.

Recent evidence from Malamud (2010, 2011) suggests that learning does play an important role in an education context. In these companion papers, Malamud compares the education system in England, where students specialize *early*, with the system in Scotland, where students specialize *late*. While there do not seem to be substantial differences in residual claimancy structures between England and Scotland, both of which have sufficiently integrated labor markets, his empirical findings with respect to education and labor markets are broadly consistent with our model of learning.[18]

Malamud (2010) finds that applicants to jobs outside of their field of study earn lower wages. While his interpretation is that the lack of specialization lowers an applicant's productivity, our model, where the non-

16. For example, with a liberal arts structure students have the flexibility to divide their courses across a major and other electives (or minors), or even choose a dual major.

17. There may be informative private signals of talent as well, such as interviews. Our model with publicly observable signals abstracts from this problem of private information and its implications for incentives.

18. The reasons for the differences in these education systems are likely to lie outside the scope of the model.

specialized sector is a fallback option for low realizations of a signal, offers an explanation that is different but consistent with this finding. Malamud (2011) also finds that early specialization is associated with a larger probability of working in a different field from the one the student has specialized in. This is once again consistent with our model if we modify it to compare early versus late specialization. In particular, if we assume that early specialization is associated with a larger variance of talent, then the likelihood of getting a low value of the signal relative to the non-specialized sector and the weight attached to this signal are larger, both of which increase the switch probability.

The logic of our results extends to the talent sampling that takes place inside an organization. Many firms maintain trainee programs where new recruits or interns are assigned to different jobs (or tasks) prior to choosing their career path in the organization. Job rotation (which corresponds to experimentation in our framework) is a common feature in "talent factories" like General Electric (GE), but has been criticized because it results in a lack of specialized human capital, leading GE to reconsider high-intensity job rotation.[19]

Assuming, for simplicity, that talents are specific to a firm or that signals are only observable inside an organization, the results can be used to highlight the incentive problems created by job rotation in an *internal* labor market. If we think of the sectors (and the corresponding firms) in our model as divisions or functions that use different types of talent, it becomes clear that even if job rotation is optimal for learning purposes (dominating the potential losses in specialized human capital), a firm needs to think about a workers incentives to engage in such a rotation program. In particular, if moving an agent from one division to another—and thus, drawing an informative signal of talent in the new division—requires the consent of the worker, then incentives have to be designed to induce the worker to experiment. One way to do this is to credibly commit—say by developing a reputation—to share the surplus from talent that the agent generates. Another way, is to *centralize* wage determination in the organization to mute competition for workers across divisions or functions in the organization. Offering the worker a flat wage, for instance, that is independent of talent, ensures efficient sampling outcomes.

8. Conclusion

People's talents are the driving force for innovation and growth. But in many cases, these talents are unknown. Various institutions impact how individuals learn about their talents, and consequently, it is important that these are designed to provide incentives for individuals to learn about their talents in an efficient way. Our paper is a deliberately abstract

19. The Wall Street Journal (March 7, 2012) came up with the headline New GE Way: Go Deep, Not Wide. After several decades, GE ended the practice of job rotation—or job hopping every two years—for future top executives or high potentials.

attempt to make this link between institutions—particularly those governing competition in the labor market—and incentives to discover talents. We develop a tractable model to compare the relative merits of experimentation (where different types of talent are sampled) and specialization (where the same type of talent is sampled repeatedly). We also find that while competitive labor markets induce efficient learning of talents, monopsonistic labor markets move incentives toward specialization.

Our focus has mainly been on incentives provided by labor-market institutions. But there are other institutions that matter for discovering talent: education systems, regulations that entrepreneurs are subject to, access to finance, and taxation are all arguably important. We hope that our framework will serve as a useful starting point to better understand the role that these institutions play in discovering talent.

Funding

This work was supported by the German Academic Exchange Service & Group of Eight (project nr 56266902).

Conflict of interest statement. None declared.

Appendix A

Proof of Lemma 1. Let F denote the distribution function of x. Since the normal distribution is symmetrical around zero, $F(x) = 1 - F(-x)$. Then

$$E_x[\max\{ax + c, d\}] - E_x[\max\{ax + d, c\}] = F\left(\frac{d-c}{a}\right)d + a\int_{\frac{d-c}{a}}^{\infty} xdF$$
$$+\left(1 - F\left(\frac{d-c}{a}\right)\right)c - F\left(\frac{c-d}{a}\right)c - a\int_{\frac{c-d}{a}}^{\infty} xdF - \left(1 - F\left(\frac{c-d}{a}\right)\right)d$$
$$= a\int_{\frac{d-c}{a}}^{\infty} xdF - a\int_{\frac{c-d}{a}}^{\infty} xdF$$
$$= 0,$$

where the last step again follows from the symmetry of the normal distribution. ■

Proof of Proposition 2. We first prove that:

$$V_E - V_S = \int_{\frac{|s_1^A|}{\sigma_B}}^{\frac{\lambda_1|s_1^A|}{\lambda_2\sigma_A}} (\lambda_1\sigma_B z - \lambda_1|s_1^A|)f_z dz + \int_{\frac{\lambda_1|s_1^A|}{\lambda_2\sigma_A}}^{\infty} (\lambda_1\sigma_B z - \lambda_2\sigma_A z)f_z dz,$$

where z is distributed normally with mean 0 and variance 1, σ_A is the standard deviation of the random variable $\hat{s}_2^A$, and σ_B is the standard deviation of the random variable s_2^B.

Consider two cases. Suppose $s_1^A \geq 0$. Then:

$$\begin{aligned}
V_E - V_S &= E_{s_2^B}[\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}] - E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}] \\
&= E_{s_2^B}[\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}] - E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A, \lambda_2 \hat{s}_2^A\}] \\
&= E_{s_2^B}[\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\} - \lambda_1 s_1^A] - E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A, \lambda_2 \hat{s}_2^A\} - \lambda_1 s_1^A] \\
&= \int_{\frac{|s_1^A|}{\sigma_B}}^{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}} \left(\lambda_1 \sigma_B z - \lambda_1 s_1^A\right) f_z dz + \int_{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}}^{\infty} \left(\lambda_1 \sigma_B z - \lambda_1 s_1^A\right) f_z dz \\
&\quad - \int_{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}}^{\infty} \left(\lambda_2 \sigma_A z - \lambda_1 s_1^A\right) f_z dz \\
&= \int_{\frac{|s_1^A|}{\sigma_B}}^{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}} \left(\lambda_1 \sigma_B z - \lambda_1 |s_1^A|\right) f_z dz + \int_{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}}^{\infty} (\lambda_1 \sigma_B z - \lambda_2 \sigma_A z) f_z dz,
\end{aligned}$$

where the second line above follows from Lemma 1, and the fourth line makes use of the transformation $s_2^B = \sigma_B z$ and $\hat{s}_2^A = \sigma_A z$.

Next, suppose $s_1^A < 0$. Then:

$$\begin{aligned}
V_E - V_S &= E_{s_2^B}[\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}] - E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}] \\
&= E_{s_2^B}[\max\{\lambda_1 s_1^A + \lambda_1 s_2^B, 0\}] - E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}] \\
&= \int_{-\frac{s_1^A}{\sigma_B}}^{-\frac{\lambda_1 s_1^A}{\lambda_2 \sigma_A}} \left(\lambda_1 \sigma_B z + \lambda_1 s_1^A\right) f_z dz + \int_{-\frac{\lambda_1 s_1^A}{\lambda_2 \sigma_A}}^{\infty} \left(\lambda_1 \sigma_B z + \lambda_1 s_1^A\right) f_z dz \\
&\quad - \int_{-\frac{\lambda_1 s_1^A}{\lambda_2 \sigma_A}}^{\infty} \left(\lambda_2 \sigma_A z + \lambda_1 s_1^A\right) f_z dz \\
&= \int_{\frac{|s_1^A|}{\sigma_B}}^{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}} \left(\lambda_1 \sigma_B z - \lambda_1 |s_1^A|\right) f_z dz + \int_{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}}^{\infty} (\lambda_1 \sigma_B z - \lambda_2 \sigma_A z) f_z dz.
\end{aligned}$$

Now consider the comparative static results with respect to σ_η^2 and $|s_1^A|$, respectively.

(i)

$$\begin{aligned}
\frac{\partial(V_E - V_S)}{\partial \sigma_\eta^2} &= \frac{\partial \frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}}{\partial \sigma_\eta^2} \left(\frac{\lambda_1 \sigma_B - \lambda_2 \sigma_A}{\lambda_2 \sigma_A}\right) \lambda_1 |s_1^A| f_z \left(\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}\right) - \frac{\partial \frac{|s_1^A|}{\sigma_B}}{\partial \sigma_\eta^2} (0) \\
&\quad + \int_{\frac{|s_1^A|}{\sigma_B}}^{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}} \lambda_1 \frac{\partial \sigma_B}{\partial \sigma_\eta^2} z f_z dz \\
&\quad + \int_{\frac{|s_1^A|}{\sigma_B}}^{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}} \left(\frac{\partial \lambda_1}{\partial \sigma_\eta^2} \left(\sigma_B z - |s_1^A|\right)\right) f_z dz
\end{aligned}$$

$$-\frac{\partial \frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}}{\partial \sigma_\eta^2}\left(\frac{\lambda_1 \sigma_B - \lambda_2 \sigma_A}{\lambda_2 \sigma_A}\right)\lambda_1 |s_1^A| f_z\left(\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}\right)$$

$$+\int_{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}}^{\infty}\left(\lambda_1 \frac{\partial \sigma_B}{\partial \sigma_\eta^2} + \frac{\partial \lambda_1}{\partial \sigma_\eta^2}\sigma_B - \lambda_2 \frac{\partial \sigma_A}{\partial \sigma_\eta^2} - \frac{\partial \lambda_2}{\partial \sigma_\eta^2}\sigma_A\right) z f_z dz.$$

Notice that the first and fourth lines in the expression above cancel each other out. The third line is positive since $z \geq \frac{|s_1^A|}{\sigma_B}$. Also, since $\lambda_1 > \lambda_2$, $\sigma_B > \sigma_A$, $\frac{\partial \lambda_1}{\partial \sigma_\eta^2} > \frac{\partial \lambda_2}{\partial \sigma_\eta^2}$ and $\frac{\partial \sigma_B}{\partial \sigma_\eta^2} > \frac{\partial \sigma_A}{\partial \sigma_\eta^2}$, the last line is positive. Thus, $\frac{\partial (V_E - V_S)}{\partial \sigma_\eta^2} > 0$.

(ii)

$$\frac{\partial (V_E - V_S)}{\partial |s_1^A|} = \frac{\lambda_1^2 |s_1^A|}{\lambda_2^2 \sigma_A^2}(\lambda_1 \sigma_B - \lambda_2 \sigma_A) f_z\left(\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}\right)$$

$$-\frac{\lambda_1 |s_1^A|}{\sigma_B}(0) f_z\left(\frac{|s_1^A|}{\sigma_B}\right)$$

$$-\int_{\frac{|s_1^A|}{\sigma_B}}^{\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}} \lambda_1 f_z dz$$

$$-\frac{\lambda_1^2 |s_1^A|}{\lambda_2^2 \sigma_A^2}(\lambda_1 \sigma_B - \lambda_2 \sigma_A) f_z\left(\frac{\lambda_1 |s_1^A|}{\lambda_2 \sigma_A}\right).$$

Notice that the first and fourth lines cancel each other out. Thus, $\frac{\partial (V_E - V_S)}{\partial |s_1^A|} < 0$.

Proof of Proposition 3. We can rewrite the agent's expected utility from specialization, experimentation, and not sampling a talent in the following way:

$$\begin{aligned}
\mathrm{EU}_S &= E_{\hat{s}_2^A}[\mu \max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\} + (1-\mu)\min\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}] \\
&= (2\mu - 1) E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\}] + (1-\mu) E_{\hat{s}_2^A}[\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A] \\
&= (2\mu - 1) V_S + (1-\mu)\lambda_1 s_1^A,
\end{aligned}$$

where the second line follows from the fact that $\max\{x, y\} + \min\{x, y\} = x + y$, and third line follows from $\hat{s}_2^A$ having a mean of zero.

$$\begin{aligned}\mathrm{EU}_E &= E_{s_2^B}[\mu\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\} + (1-\mu)\min\{\lambda_1 s_1^A, \lambda_1 s_2^B\}] \\ &= (2\mu - 1)E_{s_2^B}[\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}] + (1-\mu)E_{s_2^B}[\lambda_1 s_1^A + \lambda_1 s_2^B] \\ &= (2\mu - 1)V_E + (1-\mu)\lambda_1 s_1^A,\end{aligned}$$

where the second line, once again, follows from the fact that $\max\{x,y\} + \min\{x,y\} = x + y$, and the third line follows from s_2^B having a mean of zero. And finally,

$$\mathrm{EU}_N = \begin{cases} \max\{\mu\lambda_1 s_1^A, (1-\mu)\lambda_1 s_1^A\} - \phi & \text{if} \mu \geq \dfrac{1}{2} \\ \min\{\mu\lambda_1 s_1^A, (1-\mu)\lambda_1 s_1^A\} - \phi & \text{if} \mu < \dfrac{1}{2}. \end{cases}$$

To prove parts (i) and (ii), observe that when $\mu \geq \frac{1}{2}$,

$$\begin{aligned}\max\{\mathrm{EU}_S, \mathrm{EU}_E\} &= (2\mu - 1)\max\{V_S, V_E\} + (1-\mu)\lambda_1 s_1^A \\ &> \max\{\mu\lambda_1 s_1^A, (1-\mu)\lambda_1 s_1^A\} - \phi \\ &= \mathrm{EU}_N,\end{aligned}$$

where the inequality in the second line holds because $\max\{V_S, V_E\} > \max\{\lambda_1 s_1^A, 0\}$. Thus, it is optimal for the agent to sample her talent in period 2. Since $\mathrm{EU}_E - \mathrm{EU}_S = (2\mu - 1)(V_E - V_S)$, the agent experiments when $\mu > \frac{1}{2}$ and is indifferent between experimenting and specializing when $\mu = \frac{1}{2}$.

To prove part (iii), notice that $\mathrm{EU}_S > \mathrm{EU}_E$ when $\mu < \frac{1}{2}$. Thus, if the agent samples a talent in period 2, she specializes. To see when specialization is optimal, observe that

$$\begin{aligned}\mathrm{EU}_S - \mathrm{EU}_N &= (2\mu - 1)V_S + (1-\mu)\lambda_1 s_1^A - \left(\min\{\mu\lambda_1 s_1^A, (1-\mu)\lambda_1 s_1^A\} - \phi\right) \\ &= (2\mu - 1)V_S - \min\{0, (2\mu - 1)\lambda_1 s_1^A\} + \phi.\end{aligned}$$

Consider two possible cases. Suppose $s_1^A \leq 0$. Then $\min\{0, (2\mu - 1)\lambda_1 s_1^A\} = 0$. Since $\lim_{s_1^A \to -\infty} V_S = 0$ it follows that $\lim_{s_1^A \to -\infty} \mathrm{EU}_S - \mathrm{EU}_N = \phi$. Also, since V_S is strictly increasing in s_1^A, $\mathrm{EU}_S - \mathrm{EU}_N$ is strictly decreasing in s_1^A for this case.

Next, suppose $s_1^A > 0$. Then $\min\{0, (2\mu - 1)\lambda_1 s_1^A\} = (2\mu - 1)\lambda_1 s_1^A$. Since $\lim_{s_1^A \to \infty} V_S - \lambda_1 s_1^A = 0$ it follows that $\lim_{s_1^A \to \infty} \mathrm{EU}_S - \mathrm{EU}_N = \phi$ with

$$\begin{aligned}(\mathrm{EU}_S - \mathrm{EU}_N)'(s_1^A) &= (2\mu - 1)\left(V_S - \lambda_1 s_1^A\right)'(s_1^A) \\ &= (2\mu - 1)\left(\int_{\frac{-\lambda_1 s_1^A}{\lambda_2}}^{\infty} \lambda_2 \hat{s}_2^A dF^A\right)'(s_1^A). \\ &> 0\end{aligned}$$

Thus, $\mathrm{EU}_S - \mathrm{EU}_N$ is quasiconvex and is minimized at $s_1^A = 0$ with a value of $(2\mu - 1)V_S(0) + \phi$. When $(1 - 2\mu)V_S(0) \leq \phi$, the minimum value of $\mathrm{EU}_S - \mathrm{EU}_N$ (at $s_1^A = 0$) is non-negative so that specialization is optimal for all realizations of the first-period signal (i.e., $\underline{s}_1^A = \overline{s}_1^A = 0$). When $(1 - 2\mu)V_S(0) > \phi$, on the other hand, the minimum value of $\mathrm{EU}_S - \mathrm{EU}_N$ (at $s_1^A = 0$) is negative so that $\mathrm{EU}_S - \mathrm{EU}_N \geq 0$ if and only if $s_1^A \leq \underline{s}_1^A < 0$ and $s_1^A \geq \overline{s}_1^A > 0$. ■

Proof of Proposition 4. Observe that

$$V_E^H = E_{s_2^B}[\mathrm{TS}_E^H] = V_E + (1 + \alpha)H.$$

Similarly, we can write

$$V_S^H = E_{\hat{s}_2^A}[\mathrm{TS}_S^H] = V_S + (1 + \alpha)H + g\left(s_1^A, \alpha\right),$$

where $g\left(s_1^A, \alpha\right) = \int_{\frac{-\left(\lambda_1 s_1^A + 2(1-\alpha)H\right)}{\lambda_2}}^{0} 2(1 - \alpha)HdF^A + \int_{\frac{-\left(\lambda_1 s_1^A + 2(1-\alpha)H\right)}{\lambda_2}}^{\frac{-\lambda_1 s_1^A}{\lambda_2}} \left(\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A\right) dF^A$.

Thus, the expected gain in surplus from experimenting over specializing is given by $V_E^H - V_S^H = V_E - V_S - g\left(s_1^A, \alpha\right)$.

Note that $\lim_{s_1^A \to \infty} g\left(s_1^A, \alpha\right) = (1 - \alpha)H$ and $\lim_{s_1^A \to -\infty} g\left(s_1^A, \alpha\right) = -(1 - \alpha)H$. Also observe that

$$\begin{aligned} g_1\left(s_1^A, \alpha\right) = {} & \frac{2(1 - \alpha)H\lambda_1}{\lambda_2} f^A\left(\frac{-\left(\lambda_1 s_1^A + 2(1 - \alpha)H\right)}{\lambda_2}\right) \\ & + \lambda_1 \int_{\frac{-\left(\lambda_1 s_1^A + 2(1-\alpha)H\right)}{\lambda_2}}^{\frac{-\lambda_1 s_1^A}{\lambda_2}} dF^A \\ & - \frac{2(1 - \alpha)H\lambda_1}{\lambda_2} f^A\left(\frac{-\left(\lambda s_1^A + 2(1 - \alpha)H\right)}{\lambda_2}\right) \\ & > 0. \end{aligned}$$

(i) Let $\alpha = 1$. Then $g\left(s_1^A, \alpha\right) = 0$. Thus, $V_E^H - V_S^H = V_E - V_S$ which is positive from Proposition 1.

(ii) Let $\alpha < 1$. From Proposition 2 and the limits above, $\lim_{s_1^A \to \infty} V_E - V_S - g\left(s_1^A, \alpha\right) = -(1 - \alpha)H < 0$. Since $V_E - V_S - g\left(s_1^A, \alpha\right)$ is continuous in s_1^A, there exists some $s_1^{*A} > 0$ for which $V_E^H - V_S^H < 0$. Because $\frac{\partial(V_E - V_S)}{\partial s_1^A} < 0$ for $s_1^A > 0$ and $g_1\left(s_1^A, \alpha\right) > 0$, it follows that $V_E^H - V_S^H < 0$ for $s_1^A \geq s_1^{*A}$. Thus, specialization is efficient when $s_1^A \geq s_1^{*A}$.

(iii) To show that experimentation is efficient below a threshold level of the initial signal, note that $\lim_{s_1^A \to -\infty} -g(s_1^A, \alpha) = \lim_{s_1^A \to -\infty} V_E - V_S - g(s_1^A, \alpha) = (1-\alpha)H > 0$ (from Propositions 2 and the limits above). Since $V_E - V_S$ and g are continuous in s_1^A, there exists some s_{*1}^A for which $V_E^H - V_S^H > 0$ and $g(s_1^A, \alpha) < 0$. As $V_E - V_S > 0$ and $g_1(s_1^A, \alpha) > 0$, it follows that $V_E^H - V_S^H > 0$ for $s_1^A \le s_{*1}^A$. Thus, experimentation is efficient when $s_1^A \le s_{*1}^A$.

Proof of Proposition 5. Define $\lambda_1^B \equiv \frac{v\sigma_\eta^2}{v\sigma_\eta^2 + \sigma_\epsilon^2}$ and $\lambda_2^B \equiv \frac{v\sigma_\eta^2}{2v\sigma_\eta^2 + \sigma_\epsilon^2}$. Also note that since we can switch the order of integration, experimentation yields the same expected surplus regardless of which sector the agent samples first.

Let $v > 1$. Consider two cases. First, suppose the agent samples sector A in period 1. Let $\hat{s}_2^B = s_2^B - b$. Then,

$$\begin{aligned} E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, b\}] &\le E_{\hat{s}_2^B}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^B, b\}] \\ &< E_{\hat{s}_2^B}[\max\{\lambda_1 s_1^A + \lambda_1^B \hat{s}_2^B, b\}] \\ &= E_{s_2^B}[\max\{\lambda_1 s_1^A, (1-\lambda_1^B)b + \lambda_1^B s_2^B\}], \end{aligned}$$

where the first line follows from the fact that $\hat{s}_2^A$ second-order stochastically dominates $\hat{s}_2^B$, the second line from the fact that $\lambda_1^B > \lambda_2$, and finally, the last line from Lemma 1. Thus, experimentation dominates specializing in sector A.

Second, suppose the agent samples sector B in the first period. Define $\hat{s}_2^{B\prime} \equiv s_2^B - ((1-\lambda_1^B)b + \lambda_1^B s_1^B)$. Then,

$$\begin{aligned} E_{s_2^A}[\max\{\lambda_1 s_2^A, (1-\lambda_1^B)b + \lambda_1^B s_1^B\}] &\gtreqless E_{\hat{s}_2^{B\prime}}[\max\{\lambda_1 \hat{s}_2^{B\prime}, (1-\lambda_1^B)b + \lambda_1^B s_1^B\}] \\ &\gtreqless E_{\hat{s}_2^{B\prime}}[\max\{\lambda_2^B \hat{s}_2^{B\prime}, (1-\lambda_1^B)b + \lambda_1^B s_1^B\}] \\ &= E_{s_2^B}[\max\{0, \lambda_2^B s_2^B + (1-\lambda_2^B)((1-\lambda_1^B)b + \lambda_1^B s_1^B)\}]. \end{aligned}$$

Observe that,

$$\frac{\sigma_\eta^2}{\sigma_\epsilon^2} \gtreqless \frac{v-1}{v} \iff \sigma_\eta^2 \gtreqless \frac{v\sigma_\eta^2\sigma_\epsilon^2}{v\sigma_\eta^2 + \sigma_\epsilon^2} = (1-\lambda_1^B)v\sigma_\eta^2 \iff \lambda_1 \gtreqless \lambda_2^B.$$

Thus, using similar arguments based on the convexity of the max function and second-order stochastic dominance, the inequalities in lines 1 and 2 above hold if and only if $\frac{\sigma_\eta^2}{\sigma_\epsilon^2} \gtreqless \frac{v-1}{v}$.

Now let $v < 1$. Consider two cases. First, suppose the agent samples sector B in period 1. Then

$$\begin{aligned} &E_{s_2^B}[\max\{0, \lambda_2^B s_2^B + (1-\lambda_2^B)((1-\lambda_1^B)b + \lambda_1^B s_1^B)\}] \\ &= E_{\hat{s}_2^{B\prime}}[\max\{\lambda_2^B \hat{s}_2^{B\prime}, (1-\lambda_1^B)b + \lambda_1^B s_1^B\}] \\ &< E_{\hat{s}_2^{B\prime}}[\max\{\lambda_1 \hat{s}_2^{B\prime}, (1-\lambda_1^B)b + \lambda_1^B s_1^B\}] \\ &\leq E_{s_2^A}[\max\{\lambda_1 s_2^A, (1-\lambda_1^B)b + \lambda_1^B s_1^B\}], \end{aligned}$$

where the first line follows from Lemma 1, the second line from the fact that $\lambda_1 > \lambda_2^B$, and the third line from the fact that $\hat{s}_2^{B\prime}$ second-order stochastically dominates s_2^A. Thus, experimentation dominates specializing in sector B.

Next, suppose the agent samples sector A in period 1. Then,

$$\begin{aligned} E_{\hat{s}_2^A}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, b\}] &\lesseqgtr E_{\hat{s}_2^B}[\max\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^B, b\}] \\ &\lesseqgtr E_{\hat{s}_2^B}[\max\{\lambda_1 s_1^A + \lambda_1^B \hat{s}_2^B, b\}] \\ &= E_{s_2^B}[\max\{\lambda_1 s_1^A, (1-\lambda_1^B)b + \lambda_1^B s_2^B\}]. \end{aligned}$$

Observe that,

$$\frac{\sigma_\eta^2}{\sigma_\epsilon^2} \gtrless \frac{1-v}{v} \iff v\sigma_\eta^2 \gtrless \frac{\sigma_\eta^2 \sigma_\epsilon^2}{\sigma_\eta^2 + \sigma_\epsilon^2} = (1-\lambda_1)\sigma_\eta^2 \iff \lambda_1^B \gtrless \lambda_2..$$

Thus, using similar arguments based on the convexity of the max function and second-order stochastic dominance, the inequalities in lines 1 and 2 above hold if and only if $\frac{\sigma_\eta^2}{\sigma_\epsilon^2} \gtrless \frac{1-v}{v}$.

Proof of Proposition 6.

(i) As the agent is the residual claimant when $\mu = 1$, her optimal choice is the efficient one, which is to experiment.

(ii) Suppose the agent specializes. There are two cases to consider. First, suppose $\hat{s}_2^A \geq -\frac{\lambda_1}{\lambda_2} s_1^A$. Then expected output is the largest for sector A and the agent's wage is the mean of her talent in the next best sector which is 0. Second, suppose $\hat{s}_2^A < -\frac{\lambda_1}{\lambda_2} s_1^A$, then expected output is the largest for sector B or sector C, both of which equal 0. The agent's wage is thus 0, once again, for this case. Thus, the agent's expected utility from specialization is 0. Note that specialization dominates not sampling a talent where the agent gets a wage of 0 and has to incur the cost ϕ.

(iii) Next, suppose the agent experiments. Consider two possible cases. First, suppose $s_1^A > 0$. Then the agent's output is $\max\{\lambda_1 s_1^A, \lambda_1 s_2^B\}$ and the agent's wage is max$\{\min\{\lambda_1 s_1^A, \lambda_1 s_2^B\}, 0\}$. Thus, the agent's wage is non-negative (and positive for $s_2^B > 0$). As a result, the agent's expected utility is positive and so experimentation is optimal.

(iv) Second, suppose $s_1^A \leq 0$. Then the agent's output is max$\{\lambda_1 s_2^B, 0\}$ and the agent's wage is max$\{\min\{\lambda_1 s_2^B, 0\}, \lambda_1 s_1^A\}$. Thus, the agent's wage is non-positive so that specialization does at least as well as experimentation. In fact, when $s_1^A < 0$, the agent's wage is negative for $s_2^B \leq 0$ so that specialization does strictly better than experimentation. ■

Proof of Proposition 7. Note first, that the expected surplus from specialization can be written as

$$V_S^\rho = E_{\hat{s}_2^A}\left[\rho\left(\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A\right) + \max\{(1-\rho)\left(\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A\right), 0\}\right]$$
$$= \rho\lambda_1 s_1^A + E_{\hat{s}_2^A}\left[\max\{(1-\rho)\left(\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A\right), 0\}\right].$$

Let $\hat{s}_2^B = s_2^B - \rho\lambda_1 s_1^A$, such that the distribution of signal $\hat{s}_2^B$ is $N(0, (1-\lambda_1^2\rho^2)(\sigma_\eta^2 + \sigma_\epsilon^2)$. The expected surplus from experimentation can then be expressed as

$$V_E^\rho = E_{\hat{s}_2^B}\left[\max\left\{\lambda_1 s_1^A + \kappa_\rho \hat{s}_2^B, \rho\lambda_1 s_1^A + \kappa_1 \hat{s}_2^B\right\}\right]$$
$$= E_{\hat{s}_2^B}\left[\rho\lambda_1 s_1^A + \kappa_\rho \hat{s}_2^B + \max\left\{(1-\rho)\lambda_1 s_1^A + (\kappa_1 - \kappa_\rho)\hat{s}_2^B, 0\right\}\right].$$
$$= \rho\lambda_1 s_1^A + E_{\hat{s}_2^B}\left[\max\left\{(1-\rho)\lambda_1 s_1^A + (\kappa_1 - \kappa_\rho)\hat{s}_2^B, 0\right\}\right]$$

The expected advantage of experimentation over specialization is then

$$V_E^\rho - V_S^\rho = (1-\rho)\lambda_1\left(E_{\hat{s}_2^B}\left[\max\left\{s_1^A + \frac{\kappa_1 - \kappa_\rho}{(1-\rho)\lambda_1}\hat{s}_2^B, 0\right\}\right] - E_{\hat{s}_2^A}\left[\max\left\{s_1^A + \frac{\lambda_2}{\lambda_1}\hat{s}_2^A, 0\right\}\right]\right)$$
$$= (1-\rho)\lambda_1\left(E_{\hat{s}_2^B}\left[\max\left\{s_1^A + \frac{\hat{s}_2^B}{1-\lambda_1\rho}, 0\right\}\right] - E_{\hat{s}_2^A}\left[\max\left\{s_1^A + \frac{\hat{s}_2^A}{1+\lambda_1}, 0\right\}\right]\right).$$

Clearly, $V_E^\rho - V_S^\rho = 0$, for $\rho = 1$.

For $\rho < 1$, using a second-order stochastic dominance argument, $V_E^\rho - V_S^\rho > 0$, if $\mathrm{Var}(\frac{\hat{s}_2^B}{1-\lambda_1\rho}) > \mathrm{Var}(\frac{\hat{s}_2^A}{1+\lambda_1})$. The difference between these variances is

$$\mathrm{Var}\left(\frac{\hat{s}_2^B}{1-\lambda_1\rho}\right) - \mathrm{Var}\left(\frac{\hat{s}_2^A}{1+\lambda_1}\right) = \frac{1+\rho}{1-\lambda_1\rho}\frac{2\lambda_1}{1+\lambda_1}\left(\sigma_\eta^2 + \sigma_\epsilon^2\right).$$

This difference is zero for $\rho = -1$ and positive for all $\rho > -1$.

Finally, note that $(1-\rho)$ is strictly decreasing in ρ and $\mathrm{Var}(\frac{\hat{s}_2^B}{1-\lambda_1\rho}) - \mathrm{Var}(\frac{\hat{s}_2^A}{1+\lambda_1})$ is strictly increasing in ρ. Thus, experimentation dominates specialization for all $\rho \in (-1,1)$, and experimentation and specialization are equally efficient for $\rho \in \{-1,1\}$. ■

Proof of Proposition 8. Let EU_E^δ and EU_S^δ be the value from experimenting and specializing when bargaining weights are endogenous. We first show that

$$\mathrm{EU}_E^\delta - \mathrm{EU}_S^\delta = \mathrm{EU}_E - \mathrm{EU}_S - \delta V_S = (2\mu-1)(V_E - V_S) - \delta V_S.$$

When the agent experiments, bargaining weights stay are the same across sectors and equal μ. From the proof of Proposition 3, it follows that $\mathrm{EU}_E^\delta = \mathrm{EU}_E = (2\mu-1)V_E + (1-\mu)\lambda_1 s_1^A$.

When the agent specializes,

$$\mathrm{EU}_S^\delta = (\mu+\delta)\int_{-\frac{\lambda_1 s_1^A}{\lambda_2}}^{\infty}\left(\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A\right)dF^A + (1-\mu)\int_{-\infty}^{-\frac{\lambda_1 s_1^A}{\lambda_2}}\left(\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A\right)dF^A$$

$$= (\mu+\delta)\left(\int_{-\frac{\lambda_1 s_1^A}{\lambda_2}}^{\infty}\max\left\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\right\}dF^A + \int_{-\infty}^{-\frac{\lambda_1 s_1^A}{\lambda_2}}\max\left\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\right\}dF^A\right)$$

$$+(1-\mu)\left(\int_{-\infty}^{-\frac{\lambda_1 s_1^A}{\lambda_2}}\min\left\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\right\}dF^A + \int_{-\frac{\lambda_1 s_1^A}{\lambda_2}}^{\infty}\min\left\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\right\}dF^A\right)$$

$$= E_{\hat{s}_2^A}\left[(\mu+\delta)\max\left\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\right\} + (1-\mu)\min\left\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\right\}\right]$$

$$= E_{\hat{s}_2^A}\left[\mu\max\left\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\right\} + (1-\mu)\min\left\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\right\}\right]$$

$$+\delta E_{\hat{s}_2^A}\left[\max\left\{\lambda_1 s_1^A + \lambda_2 \hat{s}_2^A, 0\right\}\right]$$

$$= (2\mu-1)V_S + (1-\mu)\lambda_1 s_1^A + \delta V_S,$$

where the last line follows from the fact that $\max\{x,y\}+\min\{x,y\}=x+y$ and $\hat{s}_2^A$ having a mean of zero. Thus, $\mathrm{EU}_E^\delta - \mathrm{EU}_S^\delta = \mathrm{EU}_E - \mathrm{EU}_S - \delta V_S = (2\mu-1)(V_E - V_S) - \delta V_S$.

(i) Suppose $\mu > \frac{1}{2}$. Since $\lim_{s_1^A \to \infty}(2\mu-1)(V_E - V_S) - \delta V_S = -\infty$, and since V_E and V_S are continuous in s_1^A, the agent prefers specialization over experimentation for s_1^A sufficiently large.

(ii) Suppose $\mu \le \frac{1}{2}$. Since $V_E - V_S > 0$ from Proposition 1, it follows that $(2\mu-1)(V_E - V_S) \le 0$. Thus, $\mathrm{EU}_E^\delta - \mathrm{EU}_S^\delta \le -\delta V_S < 0$. As a result, specialization dominates experimentation for the agent.

References

Arrow, Kenneth J. 1974. *The Limits of Organization*. New York, NY: W. W. Norton & Company.

Auriol, Emmanuelle, Guido Friebel, and Lambros Pechlivanos. 2002. "Career Concerns in Teams," 20 *Journal of Labor Economics* 289–307.

Bergemann, Dirk, and Välimäki. Juuso 2008. "Bandit Problems," in Steven N. Durlauf and Lawrence E. Blume, eds., *The New Palgrave Dictionary of Economics*, 2nd edition, Vol. 1, 336–340. Palgrave Macmillan.

Blume, Andreas. 2003. "Bertrand without Fudge," 78 *Economics Letters* 167–168.

Colvin, Geoff. 2008. *Talent Is Overrated: What Really Separates World-Class Performers from Everybody Else*. Boston: Nicholas Brealey Publishing Group.

DeGroot, Maurice H. 1970. *Optimal Statistical Decisions*. New Jersey: John Wiley and Sons.

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole. 1999. "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies," 66 *Review of Economic Studies* 199–217.

Epstein, David. 2019. *Range: Why Specialists Triumph in a Specialized World*. New York: Riverhead Books.

Felli, Leonardo, and Christopher Harris. 1996. "Learning, Wage Dynamics, and Firm-Specific Human Capital," 104 *Journal of Political Economy* 838–868.

Geng, Sen, Pejsachovicz Leonardo, and Richter. Michael 2018. "Breadth versus Depth." Working paper.

Gibbons, Robert, and Kevin J. Murphy. 1992. "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," 100 *Journal of Political Economy* 468–505.

Holmström, Bengt. 1999. "Managerial Incentive Problems—A Dynamic Perspective," 66 *Review of Economic Studies* 169–182.

Jeon, Seonghoon. 1996. "Moral Hazard and Reputational Concerns in Teams: Implications for Organizational Choice," 14 *International Journal of Industrial Organization* 297–315.

Jovanovic, Boyan. 1979. "Job Search and the Theory of Turnover," 87 *Journal of Political Economy* 972–990.

Kaarboe, Oddvar M., and Trond E. Olsen. 2006. "Career Concerns, Monetary Incentives and Job Design," 108 *The Scandinavian Journal of Economics* 299–316.

Kartik, Navin. 2011. "A Note on Undominated Bertrand Equilibria," 11 *Economics Letters* 125–126.

Malamud, Ofer. 2010. "Breadth versus Depth: The Timing of Specialization in Higher Education," 24 *LABOUR* 359–390.

——— 2011. "Discovering One's Talent: Learning from Academic Specialization," 62 *Industrial and Labor Relations Review* 372–405.

Meyer, Margaret. 1994. "The Dynamics of Learning with Team Production: Implications for Task Assignment," 109 *Quarterly Journal of Economics* 1157–1184.

Meyer, Margaret, and John Vickers. 1997. "Performance Comparisons and Dynamic Incentives," 105 *Journal of Political Economy* 547–581.

Miller, Robert A. 1984. "Job Matching and Occupational Choice," 92 *Journal of Political Economy* 1086–1120.

Ortega, Jaime. 2001. "Job Rotation as a Learning Mechanism," 47 *Management Science* 1361–1370.

——— 2003. "Power in the Firm and Managerial Career Concerns," 12 *Journal of Economics and Management Strategy* 1–29.

Pallais, Amanda. 2014. "Inefficient Hiring in Entry-Level Labor Markets," 104 *American Economic Review* 3565–3599.

Schubert, Gregor, Stansbury Anna, and Taska. Bledi 2020. "Monopsony and Outside Options." Working paper.

Schultz, Theodore. 1968. "Resources for Higher Education: An Economist's View," 67 *Journal of Political Economy* 327–347.

Terviö, Marco. 2009. "Superstars and Mediocrities: Market Failure in the Discovery of Talent," 76 *Review of Economic Studies* 829–850.