ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Müller, Daniel

Working Paper

On Horns and Halos: Confirmation Bias and Job Rotation

Bonn Econ Discussion Papers, No. 05/2010

Provided in Cooperation with:
Bonn Graduate School of Economics (BGSE), University of Bonn

Suggested Citation: Müller, Daniel (2010) : On Horns and Halos: Confirmation Bias and Job Rotation, Bonn Econ Discussion Papers, No. 05/2010, University of Bonn, Bonn Graduate School of Economics (BGSE), Bonn

This Version is available at:
https://hdl.handle.net/10419/38813

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

Bonn Econ Discussion Papers

Discussion Paper 05/2010

On Horns and Halos: Confirmation Bias and Job Rotation

by

Daniel Müller

April 2010

Bonn Graduate School of Economics
Department of Economics
University of Bonn
Kaiserstrasse 1
D-53113 Bonn

Financial support by the
Deutsche Forschungsgemeinschaft (DFG)
through the
Bonn Graduate School of Economics (BGSE)
is gratefully acknowledged.

Deutsche Post World Net is a sponsor of the BGSE.

On Horns and Halos: Confirmation Bias and Job Rotation

Daniel Müller

Bonn Graduate School of Economics, University of Bonn

Adenauerallee 24-42, D-53113 Bonn Germany

E-mail address: daniel.mueller@uni-bonn.de[1]

[1]Acknowledgements: In preparing this paper I have greatly benefited from comments made by Meike Ahrends, Fabian Herweg, Matthias Kräkel, and Patrick Schmitz.

Abstract: *Confirmation bias, which refers to unintentional and unknowing selectivity in the use of evidence, belongs to the major problems faced by organizations. In this article, we discuss job rotation as a natural solution to this problem. In a nutshell, adopting job rotation provides an organization that is plagued by confirmation bias with a more reliable informational footing upon which to base its decisions. Job rotation, however, also comes with a cost, e.g. a loss of productivity or a disruption of work flows. We study this trade-off and identify conditions under which job rotation and specialization are each optimal.*

JEL classification:

Keywords: Confirmation Bias; Job Rotation; Work Organization

1 Introduction

"If one were to attempt to identify a single problematic aspect of human reasoning that deserves attention above all others, the confirmation bias would have to be among the candidates for consideration."

– Raymond S. Nickerson

Confirmation bias refers to unintentional and unknowing selectivity in the acquisition and use of evidence. Ample empirical evidence supports the view that once one has come to believe in a position on an issue, one's primary purpose becomes that of justifying or defending that position.[2] In consequence, regardless of whether treatment of evidence was evenhanded before the position was taken, it can become highly biased afterward. Though confirmation bias is considered as one of the most widely accepted notions of inferential errors, as suggested by the above quote by Nickerson (1999), its implications for organizational design have not been subject of thorough formal investigation.[3] This is surprising because in organizations there seems to be ample room for confirmation bias to arise and in consequence to adversely affect intra-organizational decision processes and organizational performance. In this paper, we aim at making a first step toward drawing out potential responses of organizational design to confirmation bias and its effects.

One aspect of organizational life where confirmation bias has major impact immediately comes to mind: performance appraisal. In the community practicing performance appraisal, confirmation bias is also known as the horns-and-halo effect, which refers to supervisors' tendency to judge employees as either good or bad, and then to seek evidence that supports

[2]See Nickerson (1999) for an excellent survey.

[3]Other behavioral biases have been considered in the literature on organizational theory: leniency, favoritism, or centrality bias on the side of supervisors, reference-dependent preferences, inequity aversion, or violation of procedure-neutrality on the side of employees, just to name a few. Surveys regarding the former and the latter kind of biases in the context of organizations are found in Prendergast and Topel (1993) and Camerer and Malmendier (2009), respectively.

that opinion.[4] Many, if not most performance measures regarding a firm's employees are subjective rather than objective in nature.[5] This makes performance appraisal a process by which humans judge other humans, thereby opening the door for behavioral biases and inferential errors to enter and – more importantly – to distort this process. Raters' bias in performance appraisal is considered a severe problem in practice. According to Brian Davis, executive vice president of Personnel Decisions International, "[t]he problem with rater-bias is that it takes away the organization's ability to objectively use data from performance evaluations with any validity. [...] [Y]ou can't count on the objectivity or accuracy of a performance assessment, and you have no differentiating data that allows you to make confident decisions about promotions, training, or leadership development."[6] In consequence, with bad promotion decisions having dire consequences, the biggest of which are lower employee morale, decreased productivity, and lost customer share, organizations have a vested interest in identifying the right person for a job because the cost of getting it wrong is high.[7]

In this paper we argue that organizational design provides a tool which is capable of thwarting confirmation bias not only in performance appraisal but also in other situations:

---

[4] The horns-and-halo effect, in turn, is one possible explanation for the so-called Matthew effect, which suggests that no matter how hard an employee strives, their past appraisal records will prejudice their future attempts to improve. For more on this, see http://www.performance-appraisal.com/bias.htm.

[5] For papers emphasizing this point, see, for example, Prendergast (1999) and MacLeod (2003).

[6] See http://www.management-issues.com/2007/6/7/research/bias-blights-performance-reviews.asp. Further information about Personnel Decisions International (PDI), a Minneapolis-based consultancy firm, can be found at http://www.personneldecisions.com/.

[7] According to a survey of 444 organizations throughout North America conducted by Right Management, a globally operating career transition and organizational consulting firm, the average cost of coping with an employee who does not work out is 2.5 times his salary. According to Rick Smith, Senior Vice President of Right Management, "[t]here is a smaller margin for error today in selection and promoting people into key positions, and a greater need to target development efforts to ensure that they really make a difference." For the corresponding press release from 04/11/2006, see http://phx.corporate-ir.net/phoenix.zhtml?c=65255&p=irol-newsArticle&ID=849080&highlight=.

job rotation.[8] Under confirmation bias the outcome of a judgment process often is determined by early pieces of evidence which color all subsequently received pieces of information, i.e., first impressions matter. By its very nature, in many situations job rotation creates "multiple first impressions" – and thus unbiased evaluations – by regularly breaking up the matches of the judging person and the situation to be judged. The work practice of job rotation, however, commonly is acknowledged to be associated with some sort of cost, e.g. a serious loss of productivity caused by a disruption of work flows or the sacrifice of job-specific human capital. We show that, when organizational members are subject to confirmation bias, incuring this cost for implementing job rotation may well be worthwhile for an organization in order to obtain a more accurate probability assessment upon which to base its decisions.

In Section 2, we briefly review some of the many forms that confirmation bias can take, survey some (mostly psychological) evidence for these phenomena, and finally present the model of confirmation bias proposed by Rabin and Schrag (1999), which we are going to apply throughout the paper.

Inspired by the anecdotal evidence presented above, in Section 3 we turn to the most immediate situation one can think of when pondering where confirmation bias might take effect in organizations: promotion decisions based on the evaluation of workers by their supervisors. We present a simple model in which, with different types of jobs being available, the efficient allocation of a worker depends on his ability, which is assumed to be commonly unknown. If the firm wants to base these decisions on a more solid informational footing by gathering additional information, due to an exogenously given need to delegate some tasks, it has to rely on supervisors to do so. Since under the assumptions we impose no

[8] Job rotation refers to a job practice which assigns an employee not to a single specific task but to a set of several tasks (associated with a meaningful change in job content) among which he rotates with some frequency. For evidence on job rotation being used by a significant and increasing number of companies in the United States and other OECD countries, see Osterman (1994, 2000), Gittleman et al. (1998), and OECD (1999).

incentive-compatibility or truthful-revelation complications arise, supervisors are happy to truthfully report their observations to the firm. The only friction that we allow for is that supervisors are subject to confirmation bias. The firm can choose between two types of work design, specialization or job rotation. If the firm opts for specializing the worker, he remains in one and the same division which leads to an increase in his productivity in this field of activity. Under specialization, however, the worker is evaluated by this division's supervisor exclusively. When supervisors succumb to confirmation bias, this leads to later evaluations being biased due to earlier established beliefs. If the firm decides to implement job rotation, on the other hand, the worker is placed in various of the firm's divisions and becomes a generalist who is less productive than a specialist. Under job rotation, however, the firm regularly breaks up the matches of supervisors and their subordinates, thereby creating multiple unbiased evaluations of many supervisors regarding one particular employee. We show that preventing confirmation bias from affecting supervisors' judgment can indeed outweigh the loss of productivity due to implementing job rotation. Moreover, we show that job rotation is more likely to be the optimal form of work design the stronger the degree of supervisors' confirmation bias is.

After discussing our modeling assumptions in Section 4, in Section 5 we provide an alternative interpretation of our model in order to emphasize its applicability to situations different from supervisor-worker relationships. We consider an employee who has to evaluate where a productive asset might be put to use most profitably. In contrast to the supervisor-worker setting, here job rotation does not sever the link between the judging person and the situation to be judged. In consequence an unbiased evaluation in this case probably is not to be obtained. Nevertheless, empirical evidence documents that preferential treatment of information supporting existing beliefs as well as overconfidence in one's own judgment can be reduced by forcing people to evaluate their own views, especially when that includes providing reasons against their current opinion.[9] By placing them in various positions, by

[9]See Perkins et al. (1991), Fischhoff (1977), Hoch (1984, 1985), Koriat et al. (1980), Tetlock and Kim

its very nature job rotation forces employees to look at their field of activity from different perspectives, thereby most likely broadening their view and making them less susceptible for one-sided treatment of evidence. By showing that there is scope for the firm to benefit from the resulting more reliable probability assessment even when confirmation bias is merely reduced but not fully eliminated, we provide an explanation for the often found statement that firms prefer "well-rounded employees", which neither relies on the folk wisdom that future managers should be equipped with a broad view of the entire firm, nor on the need of multi-skilled workers in order to cope efficiently with technological change.[10]

Section 6 concludes by briefly summarizing our results, relating our findings to alternative theories of job rotation, and drawing out potential implications for empirical analysis.

Related Literature When it comes to naming potential costs of implementing job rotation, there almost seems to be unanimity in the theoretical literature: Transferring individuals to new jobs sacrifices job-specific human capital, and frequent job rotation may in consequence entail a serious loss of productivity. With regard to benefits of this particular kind of work design, on the other hand, over the years many explanations have been put forth why it may be worthwhile to incur the afore-mentioned loss in productivity. One of these explanations, formalized in Cosgel and Miceli (1999), posits that workers dislike monotonous jobs. In consequence, regular job transfers increase employees' motivation and overall satisfaction by reducing their boredom and keeping them interested in their jobs, which in turn allows firms to economize on wages. A large part of the theoretical literature, however, focuses on the effects of job rotation on firm learning by placing firms and their employees on very unequal informational footing, with the firm being in a disadvantageous position. In a framework where the firm can neither observe workers' effort nor the productivity of the jobs the workers are placed in (which subsequently is observed by the respective workers),

(1987).

[10]See, for example, Schaeffer (1983) and Koike (1993) on the former argument, and Carmichael and MacLeod (1993) on the latter.

both Ickes and Samuelson (1987) and Arya and Mittendorf (2004) show that job transfers alleviate the ratchet effect.[11,12] Abstracting from any moral hazard problems, Ortega (2001) finds that a firm can benefit from implementing job rotation in order to optimally match employees to jobs when there is uncertainty about both the profitability of different jobs and the productivity of different persons at different jobs. Eguchi (2005) considers a multi-task situation where, next to regular work activities, the worker can engage in influence activities which become more profitable for the worker the longer he is in his current position. It is shown that when the firm is harmed by this rent-seeking behavior of its employees but cannot use incentive payment schemes effectively due to difficulties in measuring workers' performance, frequent job transfers are useful to limit these influence activities. Finally, when the firm faces workers of different but unobservable ability, Arya and Mittendorf (2006) argue that implementing optional job rotation programs can help firms to better match pay to an employee's true worth by achieving a self-selection of the workers: When undertaking different tasks is costly for workers but less costly for highly talented employees than for employees of low talent, the former opt for the job transfer program in order to prove their versatility, whereas the latter refrain from doing so because it is too costly.

We see this paper as complementing the existing theoretical literature on job rotation in the following sense: We abstract from any hidden action problems (e.g. hidden gaming by supervisors or shirking and influence abilities by workers) and we also remove any informational disadvantage in the afore-mentioned sense, which organizational members might profitably exploit (e.g. private information of workers with respect to their own ability or workplace productivity). Moreover, the preferences of the organizational members are com-

[11] The ratchet effect refers to workers' shirking in order to disguise the productivity of their jobs and to prevent an increase in performance standards.

[12] Arguing that under the tie-breaking rule used by Ickes and Samuleson (1987) there exists a second equilibrium in which both agents shirk in the productive job and thus are overall better off, Ma (1988) proposes an alternative compensation mechanism which uniquely implements the second-best identified by Ickes and Samuleson.

pletely standard with no inherent taste for diversity in their field of activity. The only friction that we allow for is that organizational members are subject to confirmation bias. We show that – even in the absence of any informational asymmetries – the firm may benefit from incuring the cost for implementing job rotation in order to obtain a more accurate probability assessment upon which to base its decisions. This observation adds a new item, which is based upon psychological foundations, to the list of benefits associated with job rotation as work design.

2 Confirmation Bias

Empirical Evidence Confirmation bias, which refers to unwitting selectivity in both the acquisition and evaluation of evidence, comes along in many guises. When seeking or interpreting information, people display the tendency to give greater weight to evidence that is supportive to beliefs they hold dear than to information that is counter indicative of those established opinions. Empirical evidence for this preferential treatment of evidence, also referred to as my-side bias, is provided by Baron (1991, 1995), Perkins et al. (1983), Perkins et al. (1991), and Kuhn (1989).[13] Another well-documented phenomenon is the primacy effect, which refers to the finding that when information is gathered and integrated over time, evidence acquired in the early stages is likely to carry more weight than evidence acquired later in the process. In consequence, opinions are formed early in the process and subsequently acquired information evaluated in a way that is partial to that opinion.[14] The primacy effect, which can be seen as possible manifestation of belief persistence,[15] can also

[13]Even if there is no "my side", i.e., even when people have no vested interest in the truth of a particular hypothesis, they appear to seek confirmatory information regarding this hypothesis. See, for example, Maynatt et al. (1977), Schwartz (1982), Zuckerman et al. (1995).

[14]See, for example, Nisbett and Ross (1980), Lingle and Ostrom (1981), Sherman et al. (1983).

[15]Belief persistence refers to the resistance of once established opinions to change even when faced with compelling disconforming evidence. See, for example, Ross et al. (1975), Ross (1977), Ross and Lepper (1980).

lead to a biased evaluation and interpretation of evidence that is subsequently acquired: people tend to question conflicting information more willingly than information supportive of preexisting beliefs (Ross and Anderson, 1982), to see ambiguous evidence more likely as supporting rather than disconfirming an established opinion (Lord et al., 1979; Darley and Gross, 1983), to explain away events that are inconsistent with a held position (Henrion and Fischhoff, 1986), and even to interpret evidence that should count against a hypothesis as counting in favor of it (Pitz et al., 1967).

The explanations that have been put forth to account for confirmation bias are numerous, ranging from "the desire to believe" over pragmatism and error avoidance to educational effects. At this point, however, we take the occurrence of this phenomenon as given.

A Formal Model In order to formally draw out the implications of confirmation bias for organizational design, we adopt the model of confirmation bias and belief formation proposed by Rabin and Schrag (1999). There are two exhaustive and mutually exclusive states of the world, $\theta \in \{\theta_L, \theta_H\}$. A priori, an individual considers both states of the world equiprobable, i.e., $\text{prob}(\theta = \theta_L) = \text{prob}(\theta = \theta_H) = 0.5$. In every period $t \in \{1, 2, 3, \ldots\}$ the person receives a signal, $s_t \in \{L, H\}$, that is correlated with the true state of the world. Signals received over time are independently and identically distributed with $\text{prob}(s_t = L|\theta = \theta_L) = \text{prob}(s_t = H|\theta = \theta_H) = \mu$, for some $\mu \in (0.5, 1)$. After receiving each signal, the individual updates her belief about the relative likelihood of $\theta = \theta_L$ and $\theta = \theta_H$.

When subject to confirmation bias, the person may misinterpret signals that contradict her currently held belief about which state of the world is more likely. Formally, in each period $t \in \{1, 2, 3, \ldots\}$ the individual perceives a signal $\sigma_t \in \{h, l\}$. When the person perceives signal $\sigma_t = l$ she believes that she actually received signal $s_t = L$, and if she perceives signal $\sigma_t = h$ she believes that she actually received signal $s_t = H$. Given her (possibly erroneous) perception of the information she is receiving, the individual each period updates her beliefs according to Bayes' Rule. With probability $q \in (0, 1)$ the individual misreads a signal s_t

that conflicts with her current belief about the true state of the world, which is based on the sequence of perceived signals $\sigma^{t-1} = (\sigma_1, \ldots, \sigma_{t-1})$. Signals that are supportive of the currently held belief, on the other hand, are always interpreted correctly. So, for example, if the person currently believes that θ_H is more likely, then with certainty she interprets a signal $s_t = H$ as $\sigma_t = h$, but with probability q she misinterprets a signal $s_t = L$ as $\sigma_t = h$. In order to summarize the distribution of a person's perceived signal σ_t more concisely, let $\mu^*(q)$ denote the probability that the person perceives a signal confirming her belief that one state is more likely when in fact the other state is the true state of the world. Analogously, let $\mu^{**}(q)$ denote the probability that the person perceives a signal confirming her belief that one state is more likely when in fact it is the true state of the world. Formally,

$$\begin{aligned}
\mu^*(q) &= \text{prob}\,(\sigma_t = h \,|\, \text{prob}(\theta = \theta_H|\sigma^{t-1}) > 0.5,\, \theta = \theta_L) \\
&= \text{prob}\,(\sigma_t = l \;\,|\, \text{prob}(\theta = \theta_L\,|\sigma^{t-1}) > 0.5,\, \theta = \theta_H) \\
&= (1-\mu) + q\mu
\end{aligned}$$

and

$$\begin{aligned}
\mu^{**}(q) &= \text{prob}\,(\sigma_t = h \,|\, \text{prob}(\theta = \theta_H|\sigma^{t-1}) > 0.5,\, \theta = \theta_H) \\
&= \text{prob}\,(\sigma_t = l \;\,|\, \text{prob}(\theta = \theta_L\,|\sigma^{t-1}) > 0.5,\, \theta = \theta_L) \\
&= \mu + q(1-\mu).
\end{aligned}$$

Note that $\mu^{**}(q) > \mu^*(q)$ for all $q \in (0,1)$ and $\mu \in (0.5, 1)$.

3 Supervision and Job Allocation

Suppose a firm hires a worker who has two periods of active work life. Both the firm and the worker are assumed to be risk neutral. The firm's objective is to maximize overall output over the two periods. The worker's ability is either high or low, $\theta \in \{\theta_L, \theta_H\}$ with $\theta_H > 0$. In the first period, neither the firm nor the worker know the worker's ability. It is

common knowledge, however, that both types of workers are equally likely among the overall population, prob$(\theta = \theta_H) = 0.5$.

The firm comprises of two divisions. At the outset, the firm commits to one of two possible types of job design, specialization or job rotation. If the firm opts for specialization of the new worker, he is placed in one of these divisions and stays there for at least the first period. If the firm implements job rotation, the worker spends the first half of the first period in one division and the second half of the first period in the other division. Thus, under job rotation the worker becomes a generalist in the sense that he learns as much about one division as he learns about the other. Let $r \in \{1, 2\}$ denote the number of divisions that the worker is placed in during the first period, i.e., $r = 1$ corresponds to specialization and $r = 2$ to job rotation.[16]

We abstract from any moral hazard problems: presence of the worker is enough for the firm to benefit from his input, i.e., no costly effort from the worker is needed. In his first period with the firm, the worker has to be trained and has to learn work flows, organizational design, and communication channels. Since each new worker faces these basic tasks regardless of his talent or his work place, first-period output is assumed to be independent of both his ability and the division he is placed in. Moreover, first-period output is independent of the type of work design, job rotation or specialization. We normalize first-period output to zero. The worker's second-period output, on the other hand, depends on both his ability and the type of job he is allocated to in the following way: There are two types of jobs for the worker, $j \in \{A, B\}$, that the firm can install in the second period in any division the worker visited during the first period. Let y_j denote the worker's second-period output in job j. Output in job A is independent of the worker's ability, $y_A = \bar{y} > 0$. In job B, on the other hand,

[16]The assumption that the worker switches divisions only once under job rotation is shared with most contributions to the extant theoretical literature on job rotation, e.g. Ickes and Samuelson (1987), Cosgel and Miceli (1999), Ortega (2001), and Arya and Mittendorf (2004).

output depends on the worker's ability as follows:

$$y_B = \begin{cases} \bar{y} + k(r)\theta_H & \text{for} \quad \theta = \theta_H \\ 0 & \text{for} \quad \theta = \theta_L \end{cases},$$

where $0 < k(2) < k(1) = 1$. More vividly spoken, job A might be thought of as a back-office job where the worker has to do (possibly tedious but nevertheless straightforward) paperwork. Job B, on the other hand, could be that of a product designer or marketing manager, where skills like creativity or analytical thinking are important for success. With the impact of high talent on output being decreasing in the degree of rotation, $1 - k(2)$ represents the benefits of specialization. Let $\theta_H < \bar{y}$, which implies that the firm would place the worker in job A even under specialization if it had to rely on its prior beliefs when allocating the worker to a job in period 2.

For the sake of simplicity, we assume that once the worker starts working for the firm, he stays with that firm for both periods. Thus, all the firm has to do is to compensate the worker for his (discounted lifetime) reservation utility, which we assume to be zero.

Under these assumptions the only remaining decision the firm has to take is in what type of job to place the worker at the beginning of period 2. With $\theta_H \in (0, \bar{y})$, in order to allow for a meaningful analysis, the firm must be able to gather information about the worker's ability. Due to some exogenously given need for delegation, the firm itself cannot observe this information about the worker's ability, but has to rely on the divisions' supervisors for doing so. We assume that over his first period with the firm, there are two evaluation periods of equal length in each of which the worker is evaluated by the supervisor of the division in which he is currently placed.[17] Under specialization the worker is evaluated twice in one and the same division by this division's supervisor, whereas under job rotation he is evaluated exactly once in each division he is placed in and thus by two different supervisors. In each evaluation period, the current supervisor of the worker receives a signal $s_t \in \{L, H\}$, $t = 1, 2$,

[17] The (admittedly) ad hoc restriction to only two evaluation periods will be discussed at length in the following section.

about the worker's ability. This signal represents, for example, the realization of some set of (at least to some extent) subjective performance measures. Let $\mu = \text{prob}(s_t = H|\theta = \theta_H) = \text{prob}(s_t = L|\theta = \theta_L) \in (0.5, 1)$.

Supervisors are risk neutral and we abstract from any incentives for supervisors to lie about the signals they perceive, e.g. disutility from handing out bad evaluations. Moreover, we assume that supervisors costlessly observe signals and that the informativeness of the signals is independent of any costly effort of the supervisors. Under these assumptions, an arbitrarily small incentive to identify the true ability of the worker, e.g. an arbitrarily small stake in the firm's profits, will lead to the supervisors reporting truthfully. The only friction we allow for is that supervisors are subject to confirmation bias. As described in the previous section, with probability $q \in (0, 1)$ a supervisor misinterprets signals that contradict her current hypothesis about the worker's ability as supporting her hypothesis. Let the supervisor's perception of signal $s_t \in \{L, H\}$ be denoted by $\sigma_t \in \{l, h\}$. We assume that all supervisors share the same (common knowledge) prior about the worker being of high talent, $\text{prob}(\theta = \theta_H) = 0.5$, and that there is no communication among supervisors. In consequence, confirmation bias will only affect a supervisor's judgment under specialization when she receives subsequent signals about the same worker. Last, we assume that the firm is aware of the supervisors being subject to confirmation bias. If this was not the case, there would be no reason for the firm to implement anything else but specialization.

With regard to information transmission, at the end of each evaluation period, a supervisor reports her perceived signal immediately to the firm, where this information is stored. Thus, at the end of period 1, the firm is faced with a tuple of reports, $(\sigma_1, \sigma_2) \in \{(h, h), (h, l), (l, h), (l, l)\} \equiv \mathcal{M}$. We assume that both the content and the date of reception of these reports are verifiable, and that in consequence, when choosing the type of job design at the outset, the firm can commit to an allocation rule based on the content and the order of the supervisors' reports.[18] For a given job design which places the worker in $r \in \{1, 2\}$

[18]This assumption allows us to sidestep the issue whether the firm itself is subject to confirmation bias. We

divisions during the first period, this allocation rule $\mathcal{B}_r$ prescribes for which pairs of reports the worker is allocated to job B. Formally, either $\mathcal{B}_r \subseteq \mathcal{M}$ or $\mathcal{B}_r = \emptyset$, where the latter refers to the worker being allocated to job A no matter what.[19] Clearly, the optimal allocation rule depends on the updated posterior belief of the worker being of high talent, which in turn depends on the type of job design implemented. The timing of events is summarized in Figure 1.

In a first-best situation, i.e., when the worker's ability is known to the firm, the firm would place the worker in job A when $\theta = \theta_L$ and in job B when $\theta = \theta_H$, where in the latter case the worker stays in one and the same division over the first period in order to capitalize on the benefits of specialization. When the worker's talent is unknown to the firm, it has to rely on the reports of the supervisors when allocating the worker to a job in period 2.

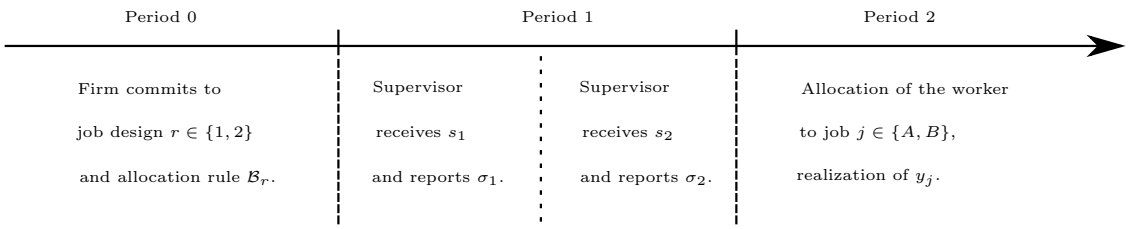

Figure 1: Timing of events.

Allocation under Specialization First, suppose the firm decides to reap the benefits of specialization and does not implement job rotation. Under specialization, after two evaluation periods the worker will be allocated to job B if and only if, given the updated posterior belief that the worker is highly talented, the expected output in job B exceeds the ability-

will comment on this assumption in the next section.

[19] More precisely, an allocation rule for a job design with $r \in \{1, 2\}$ is a mapping $\mathcal{B}_r : \mathcal{M} \to \{A, B\}$, which prescribes for each pair of possible reports $(\sigma_1, \sigma_2) \in \mathcal{M}$ in which job the worker is placed in period 2. The above "operationalization" of such an allocation rule, however, will turn out to be quite convenient.

independent output in job A, or equivalently, if and only if the firm's posterior belief about the worker being of high talent exceeds

$$\bar{p} := \frac{\bar{y}}{\bar{y} + \theta_H}.$$

Note that $\bar{p} \in (0.5, 1)$ due to our assumptions that $\theta_H \in (0, \bar{y})$. With supervisors being subject to confirmation bias, if a supervisor receives in the second evaluation period a signal which contradicts her current opinion about the worker's ability, with probability $q \in (0, 1)$ she misinterprets that signal as supporting her current opinion. When forming its updated posterior belief based on the supervisor's report at the end of the first period, the firm has to take into account the supervisor's possible misperception of the signals she received. In consequence, the order in which signals are received, or more precisely perceived, is important. Suppose, for example, the supervisor reports that she has observed two h signals. The firm now has to take into account that the supervisor, after having received an H signal in the first evaluation period, at the beginning of the second evaluation period considered the agent more likely to be of high ability than of low ability. Therefore, since with probability q she misinterprets an L signal as supporting her opinion, the probability that the supervisor perceived a second h signal is higher than the probability that she actually received a second H signal.[20] Let $p(\sigma_1, \sigma_2; q) := \text{Prob}(\theta = \theta_H | \sigma_1, \sigma_2; q)$ denote the firm's posterior believe about the worker being of high ability after the supervisor reports $(\sigma_1, \sigma_2) \in \mathcal{M}$ under specialization. Then, according to Bayes' rule,

$$p(h, h; q) = \frac{\mu\mu^{**}(q)}{\mu\mu^{**}(q) + (1 - \mu)\mu^{*}(q)}.$$

Analogously we obtain $p(h, l; q) = p(l, h; q) = 0.5$, and $p(l, l; q) = (1-\mu)\mu^*(q)/[(1-\mu)\mu^*(q) + \mu\mu^{**}(q)]$. It is readily verified that $\mu > 0.5$ implies $p(l, l; q) < 0.5 < p(h, h; q)$ for all $q \in (0, 1)$. From above we know that the firm will allocate the worker to job B only if the ex post belief

[20] While the probability of receiving a second H signal when $\theta = \theta_H$ is μ, the probability that the supervisor perceives a second h signal is $\mu^{**}(q) > \mu$. Analogously, while the probability of receiving a second H signal when $\theta = \theta_L$ is $1-\mu$, the probability that the supervisor perceives a second h signal is $\mu^*(q) > 1-\mu$.

about the worker being of high ability exceeds $\bar{p} > 0.5$. Thus, under specialization, the worker will be placed in job B only if the supervisor reports two h signals and $p(h, h; q) \geq \bar{p}$.

Lemma 1: *If* $p(h, h; q) \geq \bar{p}$*, then* $\mathcal{B}_1 = \{(h, h)\}$*. Otherwise,* $\mathcal{B}_1 = \emptyset$*.*

Allocation under Job Rotation Under our assumptions on intra-organizational information transmission, job rotation helps the firm to get rid of the supervisors' confirmation bias. Each evaluation period the worker is evaluated by a different supervisor, and each of these supervisors shares the common prior about the worker's ability since she encounters the worker for the first time. Thus, job rotation creates multiple unbiased "first impressions", which in turn allows the firm to derive a more accurate probability assessment about the worker's talent. Clearly, in this situation the order in which signals are observed is of no importance for the updated posterior belief. Formally, let $p(n_h, n_l) = \text{Prob}(\theta = \theta_H | n_h, n_l)$ denote the firm's updated posterior belief about the agent being highly talented, where n_h and n_l are the overall number of h signals and l signals, respectively, reported by the supervisors over the two evaluation periods. According to Bayes' rule we have

$$p(2, 0) = \frac{\mu^2}{\mu^2 + (1 - \mu)^2}.$$

Analogously, we obtain $p(0, 2) = (1 - \mu)^2/[\mu^2 + (1 - \mu)^2]$ and $p(1, 1) = 0.5$. Since $\mu > 1/2$, we have $p(0, 2) < 0.5 < p(2, 0)$. Removing the distortion due to confirmation bias, however, comes at the cost of sacrificing the benefits of specialization since $k(2) < 1$. Under job rotation, the worker will be placed in job B only if the posterior belief about the worker being of high talent exceeds

$$\bar{\bar{p}} := \frac{\bar{y}}{\bar{y} + k(2)\theta_H}.$$

Since $k(2) \in (0, 1)$, we have $0.5 < \bar{p} < \bar{\bar{p}} < 1$. Thus, under job rotation, the worker will be allocated to job B if and only if two h signals have been reported and $p(2, 0) \geq \bar{\bar{p}}$.

Lemma 2: *If* $p(2, 0) \geq \bar{\bar{p}}$*, then* $\mathcal{B}_2 = \{(h, h)\}$*. Otherwise,* $\mathcal{B}_2 = \emptyset$*.*

Comparison of Job Designs The question of interest is whether job allocation under job rotation can outperform job allocation under specialization in terms of ex-ante expected output. So far we know that the allocation rule under specialization depends on whether or not $p(h, h; q)$ exceeds $\bar{p}$, whereas under job rotation it depends on whether or not $p(2, 0)$ exceeds $\bar{\bar{p}}$. Since $\mu > 0.5$ and $q > 0$, we have $p(h, h; q) < p(2, 0)$, which reflects that the firm is more confident that the worker is highly talented after two h signals being reported under job rotation than under specialization due to a more accurate probability assessment. With $k(2) < k(1) = 1$, on the other hand, we have $\bar{p} < \bar{\bar{p}}$, which accounts for the loss of productivity under job rotation. Thus, we have to distinguish the following cases:

(a) $\bar{p} \leq p(h, h; q)$ and $\bar{\bar{p}} \leq p(2, 0)$;

(b) $\bar{p} \leq p(h, h; q)$ and $p(2, 0) < \bar{\bar{p}}$;

(c) $p(h, h; q) < \bar{p}$ and $\bar{\bar{p}} \leq p(2, 0)$;

(d) $p(h, h; q) < \bar{p}$ and $p(2, 0) < \bar{\bar{p}}$.

Obviously, case (d) is of little interest since under both forms of job design the worker will always be allocated to job A, $\mathcal{B}_1 = \mathcal{B}_2 = \emptyset$, which yields output $\bar{y}$ with certainty. In case (a), the allocation rule is identical under both types of job design, since the worker is allocated to job B whenever two h signals are reported, and to job A otherwise, $\mathcal{B}_1 = \mathcal{B}_2 = \{(h, h)\}$; ex-ante expected output, however, may differ under both types of job design due to a different probability assessment on the one hand, and the benefit of specialization on the other hand. Cases (b) and (c) obviously give rise to different allocation rules: In case (b), while the worker is always placed in job A under job rotation, $\mathcal{B}_2 = \emptyset$, he is allocated to job B if two h signals are reported under specialization, $\mathcal{B}_1 = \{(h, h)\}$. In case (c), allocation rules are vice versa.

In order to compare job designs in cases (a)-(c), we first characterize these cases in terms of the underlying model parameters μ and $k(2)$ for given values $\bar{y}$, θ_H and $q \in (0,1)$.[21] It is readily verified that $p(2,0) \geq \bar{\bar{p}}$ if and only if $k(2) \geq \bar{k}$, where

$$\bar{k} := \frac{(1-\mu)^2}{\mu^2} \frac{\bar{y}}{\theta_H}.$$

Since $k(2) < 1$, for $k(2) \geq \bar{k}$ to be possible we must have $\bar{k} < 1$. Regarding $\bar{k}$ as a function of μ, we find that $\bar{k} < 1$ if and only if $\mu > \bar{\mu}$, where

$$\bar{\mu} := \frac{\sqrt{\bar{y}}}{\sqrt{\bar{y}} + \sqrt{\theta_H}}.$$

By the assumption that $\theta_H \in (0, \bar{y})$ we have $\bar{\mu} \in (0.5, 1)$. Next, note that $p(h,h;q) < \bar{p}$ if and only if $\mu < \bar{\bar{\mu}}(q)$, where

$$\bar{\bar{\mu}}(q) := \frac{2\bar{y} - q(\bar{y} - \theta_H) - \sqrt{q^2(\bar{y} - \theta_H)^2 + 4\bar{y}\theta_H}}{2(1-q)(\bar{y} - \theta_H)}.$$

In the appendix we show that $\lim_{q \to 0} \bar{\bar{\mu}}(q) = \bar{\mu}$ and that, for all $q \in (0,1)$, $d\bar{\bar{\mu}}(q)/dq > 0$ and $\bar{\bar{\mu}}(q) < 1$, which implies that $\bar{\bar{\mu}}(q) \in (\bar{\mu}, 1)$ for $q \in (0,1)$. The fact that $\bar{\bar{\mu}}(q)$ is increasing in q reflects that if the distortion through confirmation bias becomes stronger, for the firm to be willing to allocate the worker to the ability-dependent job B under specialization the signal itself must become more reliable. Taken together, these observations allow us to establish the following lemma.

Lemma 3: *Given $\bar{y}$, θ_H, and $q \in (0,1)$, we have*

(a)	$\bar{p} \leq p(h,h;q)$, $\bar{\bar{p}} \leq p(2,0)$	*iff*	$\mu \in [\bar{\bar{\mu}}(q), 1)$ *and* $k(2) \geq \bar{k}$;
(b)	$\bar{p} \leq p(h,h;q) < p(2,0) < \bar{\bar{p}}$	*iff*	$\mu \in [\bar{\bar{\mu}}(q), 1)$ *and* $k(2) < \bar{k}$;
(c)	$p(h,h;q) < \bar{p} < \bar{\bar{p}} \leq p(2,0)$	*iff*	$\mu \in (\bar{\mu}, \bar{\bar{\mu}}(q))$ *and* $k(2) \geq \bar{k}$.

Proof: See Appendix.

[21] For details, see the proof of Lemma 3 in the Appendix.

Note that μ has to be sufficiently large ($\mu > \bar{\mu}$) to allow for the possibility of job rotation being the optimal choice of work design. Intuitively, if the correlation of the (unbiased) signal with the true state of the world is too low per se, it does not pay off for the firm to incur the cost of job rotation in order to prevent this bad signal from becoming somewhat more distorted.

To compare job rotation and specialization in terms of ex-ante expected output, we introduce one further piece of notation. Let $P(r)$ denote the probability of two h signals being reported when the number of divisions the worker is placed in equals r. Then $P(1) = (1/2)(\mu\mu^{**}(q) + (1-\mu)\mu^{*}(q))$ and $P(2) = (1/2)(\mu^2 + (1-\mu)^2)$. Moreover, let $\mathbb{E}[y|r]$ denote the ex-ante expected output under a job design with $r \in \{1, 2\}$.

Case (a): Under both specialization and job rotation the same allocation rule is implemented, $\mathcal{B}_1 = \mathcal{B}_2 = \{(h, h)\}$. Thus, $\mathbb{E}[y|2] > \mathbb{E}[y|1]$ if and only if

$$P(2)p(2,0)(\bar{y} + k(2)\theta_H) + (1 - P(2))\bar{y} > P(1)p(h,h;q)(\bar{y} + \theta_H) + (1 - P(1))\bar{y},$$

or equivalently, if and only if $k(2) > \bar{\bar{k}}(q)$, where

$$\bar{\bar{k}}(q) := 1 - \frac{1-\mu}{\mu} q \left[\frac{\bar{y}}{\theta_H} - 1 \right].$$

First, note that $\bar{\bar{k}}(q) < 1$ for all $q \in (0, 1)$. Moreover, it is readily verified that $\bar{\bar{k}}(q) \geq \bar{k}$ if and only if $\mu \geq \bar{\bar{\mu}}(q)$. Thus, in case (a), we have $0 < \bar{k} \leq \bar{\bar{k}}(q) < 1$.

Case (b): While the allocation rule under job rotation is $\mathcal{B}_2 = \emptyset$, under specialization we have $\mathcal{B}_1 = \{(h, h)\}$. Thus, $\mathbb{E}[y|2] \leq \mathbb{E}[y|1]$ if and only if

$$\bar{y} \leq P(1)p(h,h;q)(\bar{y} + \theta_H) + (1 - P(1))\bar{y},$$

or equivalently, if and only if $\mu \geq \bar{\bar{\mu}}(q)$. Since this last inequality is satisfied in case (b), specialization unconditionally outperforms job rotation. This result follows more immediately from the fact that under specialization the firm prefers to implement allocation rule $\mathcal{B}_1 = \{(h, h)\}$ instead of $\mathcal{B}_1 = \emptyset$.

Case (c): Under specialization the allocation rule is $\mathcal{B}_1 = \emptyset$, whereas under job rotation we have $\mathcal{B}_2 = \{(h, h)\}$. Thus, $\mathbb{E}[y|2] > \mathbb{E}[y|1]$ if and only if

$$P(2)p(2,0)(\bar{y} + k(2)\theta_H) + (1 - P(2))\bar{y} > \bar{y},$$

or equivalently, if and only if $k(2) > \bar{k}$. Since this last inequality is satisfied in case (c), job rotation unconditionally outperforms specialization. This result follows more immediately from the fact that under job rotation the firm prefers to implement allocation rule $\mathcal{B}_2 = \{(h, h)\}$ instead of $\mathcal{B}_2 = \emptyset$.

We summarize the above observations in the following proposition.

Proposition 1: *Given $\bar{y}$, θ_H, $q \in (0, 1)$, job rotation strictly outperforms specialization, $\mathbb{E}[y|2] > \mathbb{E}[y|1]$, if and only if (i) $\mu \in [\bar{\bar{\mu}}(q), 1)$ and $k(2) > \bar{\bar{k}}(q)$, or (ii) $\mu \in (\bar{\mu}, \bar{\bar{\mu}}(q))$ and $k(2) > \bar{k}$.*

Thus, given that the benefits of specialization are sufficiently small, there are two reasons for job rotation being superior compared to specialization. First, in case (c), there are different allocation rules implemented under the different types of job design. Under specialization, confirmation bias is so strong that the worker will always be placed in the ability-independent job A because the firm is (justifiedly) pessimistic – even if two h signals are reported – about the worker's talent.[22] Under job rotation, in contrast, with an unbiased probability assessment, the firm dares to place the worker in job B when two h signals are reported, which ex ante generates higher expected profits. Secondly, in case (a), both types of job design nominally implement the same allocation rule, i.e., the worker is allocated to job B if two h signals are reported and to job A otherwise. Under job rotation, however, due to unbiased reports, the probability of actually facing a highly-talented worker is higher than under specialization, which, again, leads to ex ante higher expected profits.

Having characterized the circumstances where job rotation outperforms specialization and vice versa, allows us to establish the following comparative static result.

[22] Formally, given $\bar{y}$, θ_H, and μ, q is sufficiently large such that $\mu < \bar{\bar{\mu}}(q)$, and in turn, $p(h, h; q) < \bar{p}$.

Proposition 2: *Given $\bar{y}$, θ_H, μ, and q such that $\mu \in [\bar{\bar{\mu}}(q), 1)$ and $k(2) \in [\bar{k}, \bar{\bar{k}}(q)]$. An increase in the degree of confirmation bias from q to $q' > q$ makes it more likely that job rotation strictly outperforms specialization.*

The intuition for this result is straightforward. In the original situation, a subcase of case (a) in Lemma 3, under both types of job design the worker is allocated to job B if two h signals are reported and to job A otherwise. According to Proposition 1, however, specialization outperforms job rotation in terms of ex ante expected output because the benefits of specialization are large. Under specialization, an increase in the degree of confirmation bias, q, reduces the posterior belief about the worker being highly talented after two h signals have been reported. The posterior belief under job rotation, in contrast, is unaffected by an increase in q. There are two reasons why this might lead to job rotation becoming the optimal form of job design. First, if the posterior belief under specialization is lowered sufficiently, the firm will adopt a different allocation rule under specialization and place the worker in job A no matter what, in which case job rotation unconditionally becomes superior. Formally, the increase in q raises $\bar{\bar{\mu}}(q)$. Letting $q < q'$, if $\bar{\bar{\mu}}(q) \leq \mu < \bar{\bar{\mu}}(q')$, then the shift from q to q' leads to a transition from case (a) to case (c) in Lemma 3. Secondly, even if the firm sticks to the original allocation rule, since the reliability of the supervisor's report decreases under specialization, the threshold which the cost of implementing job rotation must not exceed in order for job rotation to be optimal, becomes less stringent, $\bar{\bar{k}}(q') < \bar{\bar{k}}(q)$. If $\bar{\bar{k}}(q') < k(2) \leq \bar{\bar{k}}(q)$, job rotation becomes the optimal form of job design. Thus, when confirmation bias becomes a more severe problem, the stronger distortion of the supervisor's reports under specialization is more likely to outweigh the loss in productivity that comes along with job rotation.

Before we move on to a discussion of our modeling assumptions, we want to relate the above analysis to a statement found in Ickes and Samuelson (1987). There we read that " [w]hile uncertainty about employee productivity may be important, job transfers can optimally arise only if there is also uncertainty about the productivity of the job. Allowing uncertainty about

employee characteristics [...] cannot serve as an alternative explanation for job transfers." As we have seen, however, when we allow for another type of friction in form of confirmation bias of supervisors, job rotation may be the optimal form of work design even if there is no uncertainty regarding job characteristics but only regarding employee characteristics.

4 Discussion

Multiple periods While we stripped our model bare of hidden action and hidden information problems on purpose, the restriction to two evaluation periods is not that voluntarily but imposed by Rabin and Schrag (1999)'s model of confirmation bias. To illustrate, suppose the firm comprises of three divisions, implements three evaluation periods, and rotates the worker three times when opting for job rotation. The allocation rule under specialization depends on how the firm's posterior belief compares to $\bar{p}$, whereas the allocation rule under job rotation depends on how the firm's posterior belief compares to $\bar{\bar{p}} = \bar{y}/(\bar{y} - k(3)\theta_H)$, where $k(3) < 1$ represents the cost of rotating the worker three times compared to specialization. Under both types of work design, with $0.5 < \bar{p} < \bar{\bar{p}}$, a necessary condition for the worker to be placed in the ability-dependent job B is that the firm's posterior belief about the worker being of high talent exceeds 0.5. Application of Bayes' rule and straightforward calculations reveal that the firm's posterior belief exceeds 0.5 if and only if at least two H signals have been reported. More precisely, $p(3,0) > p(h,h,h;q) > p(2,1) = p(h,l,h;q) = p(l,h,h;q) = \mu > p(h,h,l;q) > 0.5$. Thus, in order to compare specialization and job rotation in terms of expected output, we have to distinguish ten different cases, as illustrated in Table 1.

While dealing with that many cases clearly would be tedious enough, we run into further problems when characterizing these cases in terms of the underlying model parameter μ. For example, in order to determine the values of μ for which $p(h,h,h;q) < \bar{p}$, we have to figure

	$\mathcal{B}_1$	$\mathcal{B}_2$
$p(h,h,h;q) < \bar{p},\ p(3,0) < \bar{\bar{p}}$	$\emptyset$	$\emptyset$
$\mu < \bar{p} \leq p(h,h,h;q),\ p(3,0) < \bar{\bar{p}}$	$\{(h,h,h)\}$	$\emptyset$
$p(h,h,l;q) < \bar{p} \leq \mu,\ p(3,0) < \bar{\bar{p}}$	$\{(h,l,h),(l,h,h),(h,h,h)\}$	$\emptyset$
$0.5 < \bar{p} \leq p(h,h,l;q),\ p(3,0) < \bar{\bar{p}}$	$\{(h,h,l),(h,l,h),(l,h,h),(h,h,h)\}$	$\emptyset$
$p(h,h,h;q) < \bar{p},\ \mu < \bar{\bar{p}} \leq p(3,0)$	$\emptyset$	$\{(h,h,h)\}$
$\mu < \bar{p} \leq p(h,h,h;q),\ \mu < \bar{\bar{p}} \leq p(3,0)$	$\{(h,h,h)\}$	$\{(h,h,h)\}$
$p(h,h,l;q) < \bar{p} \leq \mu,\ \mu < \bar{\bar{p}} \leq p(3,0)$	$\{(h,l,h),(l,h,h),(h,h,h)\}$	$\{(h,h,h)\}$
$0.5 < \bar{p} \leq p(h,h,l;q),\ \mu < \bar{\bar{p}} \leq p(3,0)$	$\{(h,h,l),(h,l,h),(l,h,h),(h,h,h)\}$	$\{(h,h,h)\}$
$p(h,h,l;q) < \bar{p} \leq \mu,\ 0.5 < \bar{\bar{p}} \leq \mu$	$\{(h,l,h),(l,h,h),(h,h,h)\}$	$\{(h,h,l),(h,l,h),(l,h,h),(h,h,h)\}$
$0.5 < \bar{p} \leq p(h,h,l;q),\ 0.5 < \bar{\bar{p}} \leq \mu$	$\{(h,h,l),(h,l,h),(l,h,h),(h,h,h)\}$	$\{(h,h,l),(h,l,h),(l,h,h),(h,h,h)\}$

Table 1: Allocation rules for three evaluation periods.

out when

$$\frac{\mu(\mu + q(1-\mu))^2}{\mu(\mu + q(1-\mu))^2 + (1-\mu)((1-\mu) + q\mu)^2} < \frac{\bar{y}}{\bar{y} - \theta_H}, \tag{1}$$

which basically boils down to finding the zeros of a polynomial of third order. With both the number of cases to consider and the degree of the polynomials characterizing the threshold levels for μ increasing in the number of evaluations, the necessary calculations soon become infeasible. Thus, in order to inquire into interesting questions of optimal rotation intervals or the optimal number of evaluation periods a more tractable model of confirmation bias is needed.

Commitment to an allocation rule The assumption that the firm can commit to an allocation rule at the outset can be dropped when the firm itself is not subject to confirmation bias when making the allocation decision based on the reports it received. With the firm's perception being unbiased, its updated posterior from report (σ_1, σ_2) is the same irrespective of whether it is calculated ex ante, before actually receiving this report, or ex post, i.e., after

having received (and thus after having evaluated) this report. Therefore, if commitment to an allocation rule is not possible, after receiving the reports there is no incentive for the firm to deviate from the allocation rule which is optimal if commitment is possible before receiving the reports.[23]

The assumption that the head of the firm is not subject to confirmation bias does not seem that far-fetched. As we already mentioned in Section 2, for confirmation bias to arise, pieces of evidence have to be received and evaluated subsequently, and there also often needs to be some degree of vagueness that leaves room for misinterpretation. Both can be imagined not to be the case for the head of the organization: First, when deciding in which job to place a worker, the head of the firm refers to the reports and recommendations of the supervisors that evaluated this worker. With forced distribution rankings or specified evaluation schemes often being used in practice, there probably is less room for arbitrariness in interpretation of these reports than in the original evaluation of the worker by his supervisor(s). Secondly, with bosses and CEOs almost always having got their hands full, the head of the firm will not spend days over days in advance pondering where to place the worker, but more likely he will focus on the decision shortly before it is due with most (all) relevant information available. Though only one evaluation report or memo can be read at a time, for the head of the firm this removes the sequential character of information acquisition at least to some degree. Third, while probably only one supervisor at a time is responsible for evaluation of the worker, when making the allocation decision, the head of the organization most likely comprises of several actors, like the personnel manager, the managers in whose division the worker might be placed in, and so on. One might imagine that preferential treatment of evidence might be less likely to occur when discussing pros and cons of a decision with other equally skilled people. Last, the assumption of a rational head of the organization

[23]The firm's updated posterior beliefs ex ante and ex post, however, would not coincide if the firm also succumbs to confirmation bias when evaluating the reports it received. In this case, the ability to commit to an allocation rule clearly makes a difference.

also is in line with the largest part of the literature on behavioral industrial organization, where rational firms/principals interact with behaviorally biased consumers/agents: while the former know about the biases of the latter, the latter often are assumed to be naive and do not know about their own bias.[24]

5 An Alternative Interpretation

In this section we want to emphasize applicability of the above analysis to situations different from supervisor-worker relationships. In order to do so, we provide an alternative interpretation of our model. Apart from very few exceptions, we basically just relabel variables. Therefore, as long as there is no danger of confusion, we will make use of notation and definitions introduced before without explicitly saying so.

Suppose a risk neutral firm, which comprises of several divisions, e.g. production and marketing, faces two opportunities where to deploy an asset in the second period, project A or project B. For example, the asset might be a machine used in production, and the two projects represent the production of different products. While the return of project A is assumed to be riskless, $R_A = \bar{R} > 0$, the return of project B depends on the true state of the world, $\theta \in \{\theta_L, \theta_H\}$, as follows:

$$R_B = \begin{cases} \bar{R} + \Delta & \text{for} \quad \theta = \theta_H \\ 0 & \text{for} \quad \theta = \theta_L \end{cases},$$

where $\Delta \in (0, \bar{R})$. Thus, in the above example, we can think of project A as the production of a product which is already established in the market and for which the firm is familiar with the production process. Project B, on the other hand, can be thought of as the production of a newly developed product, the success of which depends on factors like market acceptance or whether there will be complications in the production process. In this sense, θ_H can be

[24] See, for example, Eliaz and Spiegler (2006), DellaVigna and Malmendier (2004), Gilpatric (2008), and Gabaix and Laibson (2006). For a survey on bounded rationality in industrial organization, see Ellison (2006).

interpreted as a situation, i.e., a constellation of market characteristics and technological circumstances, in which the launch of the new product will be successful, whereas it will be a failure if the state is θ_L. A priori, the two states are equally likely. Thus, without any further information, the firm allocates the asset to project A.

In order to obtain further information about the true state of the world, the firm can hire a risk neutral worker in period 1, who lives and stays with the firm for two periods. The worker's (discounted lifetime) reservation utility equals zero. Once again, we abstract from hidden information or hidden action problems: there is no uncertainty about the worker's talent, and the worker's output is (for expositional purposes only) equal to zero in both periods. The only meaningful task the firm can assign to the worker is to gather information about the true state of the world in order to improve the decision where to place the productive asset.

Over the first period, the worker costlessly receives two subsequent signals about the true state of the world, $s_t \in \{L, H\}$ with $t = 1, 2$. These signals are identically and independently distributed according to $\mu \in (0.5, 1)$. Due to confirmation bias, however, the signal the worker receives may differ from the signal he actually perceives: once he comes to believe that one state of the world is more likely than the other, with probability $q \in (0, 1)$ the agent misinterprets a signal which contradicts this hypothesis as actually supporting this hypothesis. Let the worker's perception of signal $s_t \in \{L, H\}$ be denoted by $\sigma_t \in \{l, h\}$. We assume that the firm itself does not receive the signals, and therefore has to rely on the reports of the worker. The firm, however, is aware of the worker's confirmation bias. With the gathering of information being costless and the worker being risk neutral, an arbitrarily small incentive to identify the true state of the world, e.g. an arbitrarily small stake in the return of the second-period project, will lead to the worker reporting truthfully what signals he has perceived. Both the content and the order of these reports, which we can identify with the worker's perceived signals $(\sigma_1, \sigma_2) \in \mathcal{M}$, are verifiable.

At the outset, the firm can commit to a particular type of job design, specialization or job

rotation. Letting $r \in \{1, 2\}$ denote the number of divisions the worker is placed in during the first period, $r = 1$ corresponds to specialization and $r = 2$ to job rotation. If the firm opts for implementing job rotation and makes the worker switch divisions during his first period with the firm, it incurs a cost $c > 0$.[25] Moreover, the firm can commit to an allocation rule, which prescribes for which reports of the worker the asset is placed in project B. Thus, for a job design with $r \in \{1, 2\}$, this allocation rule is either $\mathcal{B}_r = \emptyset$ or $\mathcal{B}_r \subseteq \mathcal{M}$.

We do not believe it to be too far-fetched to assume that job rotation will reduce confirmation bias in this scenario as well. Suppose the firm opts for specialization and worker is placed in the production division throughout his first period with the firm. Though he might be able to get all the relevant information from the marketing division, it is easy to imagine that he will see all bits of information he gathers through the eyes of a production engineer, thus attaching too much weight to technological aspects and too little weight to market related data.[26] If the firm implements job rotation, on the other hand, during the first period the worker switches from production into marketing, and thus basically is forced to open his eyes more widely with respect to the market-related data as well. Thus, while the agent already holds some belief about the true state of the world when being placed in the marketing division, the new perspective from which he now has to assess the problem might make him more willing to let go of this hypothesis. However, since one and the same worker evaluates one and the same problem, it is likely that confirmation bias will merely be

[25] Basically $c > 0$ may reflect any cost possibly associated with job rotation. Campion et al. (1994), for example, identify productivity losses and disruption of work flows for both the department gaining a rotating employee and the department losing the employee as potential costs of job rotation, resulting from training requirements in the first case and from having a vacancy in the second case. Also Burke and Moore (2000) draw attention to reverberating negative effects of job rotation on nonrotaters' perception of organizational justice.

[26] Evidence for experts being more confident than justified when making judgments in their own areas of expertise is provided by Kidd (1970), Loftus and Wagenaar (1988), Oskamp (1965) regarding engineers, attorneys, and psychologists, respectively.

reduced by job rotation but not fully eliminated. Thus, letting q_r denote the probability that the agent misinterprets a contradicting signal under a job design which places the worker in $r \in \{1, 2\}$ divisions during the first period, we assume $0 < q_2 < q_1 < 1$.

Defining

$$\bar{c} := \frac{\mu(1-\mu)(q_1 - q_2)(\bar{R} - \Delta)}{2}$$

and

$$\bar{\bar{c}} =: \frac{\mu(\mu + q_2(1-\mu))\Delta - (1-\mu)((1-\mu) + q_2\mu)\bar{R}}{2},$$

and following the lines of the analysis in Section 3, we obtain the following result.

Proposition 3: *Given $\bar{R}$, Δ, q_1, q_2, job rotation strictly outperforms specialization, $\mathbb{E}[R|2] > \mathbb{E}[R|1]$ if and only if (i) $\mu \in [\bar{\bar{\mu}}(q_1), 1)$ and $c < \bar{c}$, or (ii) $\mu \in (\bar{\bar{\mu}}(q_2), \bar{\bar{\mu}}(q_1))$ and $c < \bar{\bar{c}}$.*

Proof: See Appendix.

The above result is familiar by now: if the benefit of specialization is sufficiently small, job rotation may be superior to specialization for two reasons. First, in case (ii), with confirmation bias being strong under specialization, the firm implements a very conservative allocation rule under specialization which places the asset always in project A, whereas under job rotation the asset is placed in project B if two h signals are reported and in project A otherwise. This more "daring" allocation rule, which is based on a more reliable probability assessment, yields higher expected profits. In case (i), on the other hand, even though allocation rules are identical under both types of work design, under job rotation, due to unbiased reports, the probability of the worker being of high talent after two h signals have been reported is higher than under specialization, which yields higher expected profits. Moreover, if the degree of confirmation bias under specialization increases or job rotation becomes more efficient in reducing the employees degree of confirmation bias, i.e., if q_1 increases or q_2 decreases, it becomes more likely that job rotation is the optimal form of

work design. This follows from the fact that the threshold which the cost of implementing job rotation must not exceed in order for job rotation to be optimal becomes less stringent, $d\bar{c}/q_2 = -d\bar{c}/q_1 < 0$ and $d\bar{\bar{c}}/dq_2 < 0$.

As we briefly mentioned in the introduction, the literature on management development and employee learning recommends job rotation in order to endow the managers-to-be with a deeper understanding of more aspects of business, which they will need as they move up to broader jobs, or to help employees to cope better with uncertainty and technological change. The above analysis suggests, however, that a firm may have an incentive to provide its employees with a broader view of the organization even in the absence of such considerations: if knowledge of different organizational aspects makes employees less susceptible to confirmation bias, and if the firm has to base some of its decisions on its employees' judgments, job rotation may help to provide a more profound informational footing for the firm's decisions.

6 Conclusion

In this paper we examined a setting in which an organization is faced with its members being subject to confirmation bias, i.e., the tendency to treat subsequent information partially after an initial position has been taken. Given that job rotation is able (i) to sever the link between the judge and the situation to be judged, or (ii) to force the judge to be more open-minded for contradicting evidence, we have shown that implementing this particular form of job design may be profitable for the organization, even if it comes along with certain costs. The reason is that job rotation leads to a more reliable informational footing for the organization's decision making. We do not, however, obtain a call for a universal mandate for job rotation, but we find that optimality of job rotation is circumstance specific. In particular, the higher the degree of confirmation bias and the lower the cost associated with the implementation of job rotation, the more likely it is that job rotation is superior to specialization.

As briefly mentioned in the introduction, there are three major approaches to explain why work place organization may take the particular form of job rotation: employee moti-

vation, employee learning, and employer learning. The employee motivation theory posits that job rotation helps to make work more interesting, thereby in particular providing motivation for so-called "plateaued" employees, i.e., employees with limited promotion prospects. The employee learning theory, on the other hand, contents that job rotation is an effective way to develop employees' abilities and to improve organizational knowledge in order to help prepare junior employees to become top managers or to better cope with uncertainty. Last, according to the employer learning theory, job rotation improves job assignments by providing the employer with information about the employee's abilities, both general and job-specific, and also job-specific factors unrelated to the employee. Though we do not see an immediate connection to the first of these approaches, the two tales told in this paper suggest that the presence of confirmation bias in organizations might interact with the two latter explanations, employee and employer learning. As for employer learning, our model about employees' evaluation by supervisors indicates that job rotation may become an even more valuable learning device for the firm when confirmation bias is an issue because it may prevent distortion of the signal that the employer receives. The alternative interpretation of our model, on the other hand, in a sense links employer and employee learning theory: though the ultimate goal of the employer is to learn where best to deploy the asset, when confirmation bias is present this may be achieved most profitably by making the employee learn to know the different building blocks of the organization in order to broaden his view and make him less susceptible for partial treatment of information.

In particular this last observation might be relevant for empirical analysis. In a rigorous test of the afore-mentioned explanations for the practice of job rotation, Eriksson and Ortega (2006) find "only very limited support for the employee motivation hypothesis, [but that] statistical evidence is more amenable to the employee learning hypothesis and employer learning hypothesis."[27] This is correct in the sense that a number of the hypothesized re-

[27] Arguing that a satisfactory test of the three major theories of job rotation should combine a representative sample of establishments with data on employee characteristics, Eriksson and Ortega (2006) merge a

lationships between job rotation and the set of relevant variables were found to be in the predicted direction at a statistically significant level, e.g. a positive correlation between the use of job rotation and firm size or the number of hierarchical levels, which is consistent with both employee and employer learning theory. Regarding hypotheses for which the two learning theories predict different directions, however, there is no clear-cut result which theory better explains the data. For example, the finding that firms that spend more to train their employees are more likely to use job rotation schemes is favorable to the employee learning hypothesis but contradicts the employer learning hypothesis. Tenure in the industry not having a statistically significant effect on rotation, on the other hand, is consistent with the employer learning theory but contradicts the employee learning theory. In the light of our second story, we believe that these two theories sometimes cannot be treated separately but have to be seen as interwoven with each other. Therefore, in order to obtain even sharper predictions, it might be insightful to differentiate cases where the ultimate goal of employee learning is firm learning from cases of pure employee learning.

Last, we want to point out a more directly testable implication of this paper. We have seen that the stronger the degree of confirmation bias, the more likely is job rotation the optimal form of workplace design. In consequence, we should expect to find rotation arrangements more often in firms where there is more scope for confirmation bias to arise. While the degree of confirmation bias might be quite difficult to measure per se, there might be several ways to operationalize its measurement. For example, based on the observation that for confirmation bias to arise there needs to be some room for misinterpretation of evidence, the extent to which evaluation of employees is based upon subjective performance measures, which (by their very nature) are more vague and thus more susceptible to misinterpretation than

representative survey of Danish firms with the employer-employee linked panel constructed by Statistics Denmark, which provides data on each employee at the sampled firms. The resulting database is richer than most surveys of establishments and provides more representative evidence than do single-firm case studies.

objective performance measures, might serve as an indicator for the presence and strength of confirmation bias.

A Appendix: Proofs of Lemmas and Propositions

Proof of Lemma 3

In order to give the proof some structure, we proceed in several steps.

CLAIM 1: $p(2,0) > \bar{p}$ iff $k(2) > \bar{k}$.

PROOF: Follows immediately from rearranging. ||

CLAIM 2: $\bar{k} < 1$ iff $\mu > \bar{\mu}$.

PROOF: Rearranging yields

$$\bar{k} < 1 \quad \Longleftrightarrow \quad \mu^2 - \frac{2\bar{y}}{(\bar{y} - \theta_H)}\mu + \frac{\bar{y}}{(\bar{y} - \theta_H)} < 0.$$

Define

$$f(\mu) := \mu^2 - \frac{2\bar{y}}{(\bar{y} - \theta_H)}\mu + \frac{\bar{y}}{(\bar{y} - \theta_H)}.$$

Straight-forward differentiation reveals that $f(\mu)$ is a strictly convex function, $f''(\mu) > 0$, which reaches its minimum at $\mu = \bar{y}/(\bar{y} - \theta_H)$. The zeros of $f(\mu)$ are obtained by solving

$$f(\mu) = 0 \iff \bar{\mu}_{1,2} \quad = \quad \frac{\sqrt{\bar{y}}(\sqrt{\bar{y}} \pm \sqrt{\theta_H})}{(\sqrt{\bar{y}} + \sqrt{\theta_H})(\sqrt{\bar{y}} - \sqrt{\theta_H})}.$$

Let

$$\bar{\mu}_1 = \frac{\sqrt{\bar{y}}}{(\sqrt{\bar{y}} + \sqrt{\theta_H})} \quad \text{and} \quad \bar{\mu}_2 = \frac{\sqrt{\bar{y}}}{(\sqrt{\bar{y}} - \sqrt{\theta_H})}.$$

Obviously, $\bar{\mu}_2 > 1$, which allows us to focus on $\bar{\mu}_1$ because we are interested only in values of μ from the interval $(0.5, 1)$. Since, by assumption, $\theta_H \in (0, \bar{y})$, we have $\bar{\mu}_1 \in (0.5, 1)$. Thus, with $f(\mu)$ being a strictly convex function which is strictly decreasing for $\mu < 1$, letting $\bar{\mu} = \bar{\mu}_1$ concludes the proof. ||

Claim 3: $p(h,h;q) < \bar{p}$ iff $\mu < \bar{\bar{\mu}}$.

Proof: Rearranging yields

$$p(h,h;q) < \bar{p} \iff \mu^2 - \frac{2\bar{y} - q(\bar{y} - \theta_H)}{(1-q)(\bar{y} - \theta_H)}\mu + \frac{\bar{y}}{(1-q)(\bar{y} - \theta_H)} > 0$$

Define

$$g(\mu) := \mu^2 - \frac{2\bar{y} - q(\bar{y} - \theta_H)}{(1-q)(\bar{y} - \theta_H)}\mu + \frac{\bar{y}}{(1-q)(\bar{y} - \theta_H)}.$$

Differentiation with respect to μ reveals that $g(\mu)$ is a strictly convex function, $g''(\mu) > 0$, which reaches its minimum at $\mu = (2\bar{y} - q(\bar{y} - \theta_H))/2(1-q)(\bar{y} - \theta_H) > 1$. The zeros of $g(\mu)$ are obtained by solving

$$g(\mu) = 0 \iff \bar{\bar{\mu}}_{1,2} = \frac{2\bar{y} - q(\bar{y} - \theta_H) \pm \sqrt{q^2(\bar{y} - \theta_H)^2 + 4\theta_H\bar{y}}}{2(1-q)(\bar{y} - \theta_H)}$$

Once again, we are interested in values of μ from the interval $(0.5, 1)$. Since for all $q \in (0.5, 1)$ we have

$$\bar{\bar{\mu}}_2 = \frac{2\bar{y} - q(\bar{y} - \theta_H) + \sqrt{q^2(\bar{y} - \theta_H)^2 + 4\theta_H\bar{y}}}{2(1-q)(\bar{y} - \theta_H)} > 1$$

we can focus on

$$\bar{\bar{\mu}}_1 = \frac{2\bar{y} - q(\bar{y} - \theta_H) - \sqrt{q^2(\bar{y} - \theta_H)^2 + 4\theta_H\bar{y}}}{2(1-q)(\bar{y} - \theta_H)}.$$

Straightforward calculations reveal that for $q \in (0,1)$ we have $\bar{\bar{\mu}}_1 \in (0.5, 1)$. Thus, with $g(\mu)$ being a strictly convex function which is strictly decreasing for $\mu < 1$, letting $\bar{\bar{\mu}}(q) = \bar{\bar{\mu}}_1$ concludes the proof. ||

Claim 4: $d\bar{\bar{\mu}}(q)/dq > 0$.

Proof: First, note that $\bar{\bar{\mu}}(q)$ is continuously differentiable with respect to q for all $q \in (0,1)$. By definition of $\bar{\bar{\mu}}(q)$, the following identity holds:

$$g(\bar{\bar{\mu}}(q)) = \bar{\bar{\mu}}(q)^2 - \frac{2\bar{y} - q(\bar{y} - \theta_H)}{(1-q)(\bar{y} - \theta_H)}\bar{\bar{\mu}}(q) + \frac{\bar{y}}{(1-q)(\bar{y} - \theta_H)} \equiv 0$$

Rearranging and differentiation with respect to q yield

$$\frac{d\bar{\bar{\mu}}(q)}{dq} = \frac{(\bar{y} - \theta_H)\bar{\bar{\mu}}(q)(\bar{\bar{\mu}}(q) - 1)}{2\bar{y}(\bar{\bar{\mu}}(q) - 1) - 2\theta_H\bar{\bar{\mu}}(q) - q(\bar{y} - \theta_H)(2\bar{\bar{\mu}}(q) - 1)}.$$

In the proof of Claim 3 we established that $\bar{\bar{\mu}}(q) \in (0.5, 1)$ for $q \in (0, 1)$, which immediately implies $d\bar{\bar{\mu}}(q)/dq > 0$. ||

Claim 5: $\forall q \in (0, 1)$, $0.5 < \bar{\mu} < \bar{\bar{\mu}}(q) < 1$.

Proof: In Claim 2 and 3 we have already established that $\bar{\mu} \in (0.5, 1)$ and $\bar{\bar{\mu}}(q) \in (0.5, 1)$. It remains to show that $\bar{\mu} < \bar{\bar{\mu}}(q)$ for all $q \in (0, 1)$. Note that $\bar{\bar{\mu}}(q)$ is a continuous and continuously differentiable function on the interval $(-\infty, 1)$, thus $\lim_{q\to 0} \bar{\bar{\mu}}(q)$ exists and is given by

$$\lim_{q\to 0} \bar{\bar{\mu}}(q) = \frac{2\bar{y} - \sqrt{4\bar{y}\theta_H}}{2(\bar{y} - \theta_H)} = \frac{\sqrt{\bar{y}}}{\sqrt{\bar{y}} + \sqrt{\theta_H}} = \bar{\mu}.$$

From Claim 4, we know that $d\bar{\bar{\mu}}(q)/dq > 0$ for $q \in (0, 1)$, which establishes the result. ||

Combining Claims 1-5 establishes the desired result. ■

Proof of Proposition 3

Let $p(\sigma_1, \sigma_2; q_r)$ denote the firm's updated posterior belief about the true state of the world being $\theta = \theta_H$ after receiving report $(\sigma_1, \sigma_2) \in \mathcal{M}$ from the worker under a job design which places the worker in $r \in \{1, 2\}$ divisions during the first period. The expected return from allocating the asset to project B exceeds the expected return from allocating the asset to project A if and only if $p(\sigma_1, \sigma_2; q_r)$ exceeds

$$\bar{p} = \frac{\bar{R}}{\bar{R} + \Delta} \in (0.5, 1).$$

Since $\mu \in (0.5, 1)$ implies that $p(l, l; q_r) < p(l, h; q_r) = p(h, l; q_r) = 0.5 < p(h, h; q_r)$ for $r \in \{1, 2\}$, the following observation follows immediately.

Lemma 4: *For $r \in \{1, 2\}$, if $p(h, h; q_r) \geq \bar{p}$, then $\mathcal{B}_r = \{(h, h)\}$. Otherwise $\mathcal{B}_r = \emptyset$.*

It is readily verified, that $\mu^*(q)/\mu^{**}(q)$ is increasing in q, which implies that $p(h,h;q)$ is decreasing in q. With the question of interest being whether the firm can benefit from implementing job rotation compared to specialization, this observation renders the case where $p(h,h;q_2) < \bar{p}$ uninteresting. In this case, $\mathcal{B}_1 = \mathcal{B}_2 = \emptyset$, i.e., under both types of job design the asset is allocated to the riskless project A irrespective of the worker's report. Thus, job rotation can never be optimal because it comes along with additional costs without providing any benefit. This leaves us with two cases in which there is scope for job rotation to outperform specialization due to a more accurate probability assessment. In the first of these cases, $\bar{p} \leq p(h,h;q_1)$, the allocation rule is the same under both types of job design, $\mathcal{B}_1 = \mathcal{B}_2 = \{(h,h)\}$. In the second case, $p(h,h;q_1) < \bar{p} \leq p(h,h;q_2)$, allocation rules differ, $\mathcal{B}_1 = \emptyset$ and $\mathcal{B}_2 = \{(h,h)\}$.

It can be shown that $p(h,h;q) > \bar{p}$ if and only if $\mu > \bar{\bar{\mu}}(q)$, where

$$\bar{\bar{\mu}}(q) = \frac{2\bar{R} - q(\bar{R}-\Delta) - \sqrt{q^2(\bar{R}-\Delta)^2 + 4\bar{R}\Delta}}{2(1-q)(\bar{R}-\Delta)}$$

with $\bar{\bar{\mu}}'(q) > 0$ and $\bar{\bar{\mu}}(q) \in (0.5, 1)$ for all $q \in (0,1)$. These properties of $\bar{\bar{\mu}}(q)$ follow immediately from the proof of Lemma 3. The following observation then is immediate.

Lemma 5: *Given $\bar{R}$, Δ, and $0 < q_2 < q_1 < 1$, we have*

(a′) $\quad \bar{p} \leq p(h,h;q_1) < p(h,h;q_2) \qquad$ *iff* $\qquad \mu \in [\bar{\bar{\mu}}(q_1), 1)$;

(b′) $\quad p(h,h;q_1) < \bar{p} \leq p(h,h;;q_2) \qquad$ *iff* $\qquad \mu \in [\bar{\bar{\mu}}(q_2), \bar{\bar{\mu}}(q_1))$.

In both cases (a′) and (b′), job rotation can outperform specialization if the cost of job rotation is sufficiently small. To formally establish this result, let $P(q)$ denote the probability of two h signals being reported for a given q. Moreover, let $\mathbb{E}[R|r]$ denote the firm's ex-ante expected return from asset allocation under a job design with $r \in \{1,2\}$.

Case (a′): Both types of job design lead to the same allocation rule, $\mathcal{B}_1 = \mathcal{B}_2 = \{h,h\}$. Thus, $\mathbb{E}[R|2] > \mathbb{E}[R|1]$ if and only if

$$P(q_2)p(h,h;q_2)(\bar{R}+\Delta) + (1-P(q_2))\bar{R} - c > P(q_1)p(h,h;q_1)(\bar{R}+\Delta) + (1-P(q_1))\bar{R},$$

or equivalently, if and only if $c < \bar{c}$, where

$$\bar{c} = \frac{\mu(1-\mu)(q_1 - q_2)(\bar{R} - \Delta)}{2}.$$

Case (b′): Under specialization we have $\mathcal{B}_1 = \emptyset$, whereas under job rotation the allocation rule is $\mathcal{B}_2 = \{h, h\}$. Thus, $\mathbb{E}[R + y|2] > \mathbb{E}[R + y|1]$ if and only if

$$P(q_2)p(h, h; q_2)(\bar{R} + \Delta) + (1 - P(q_2))\bar{R} - c > \bar{R},$$

or equivalently, if and only if $c < \bar{\bar{c}}$, where

$$\bar{\bar{c}} = \frac{\mu(\mu + q_2(1-\mu))\Delta - (1-\mu)((1-\mu) + q_2\mu)\bar{R}}{2}.$$

It is readily verified that $\bar{\bar{c}} > 0$ whenever $\mu \in (\bar{\bar{\mu}}(q_2), \bar{\bar{\mu}}(q_1))$.

This establishes the desired result. ■

References

[1] **Arya, A.** and **B. Mittendorf (2004):** Using Job Rotation to Extract Employee Information, *Journal of Law, Economics, and Organization*, Vol. 20, 400-414.

[2] **Arya, A.** and **B. Mittendorf (2006):** Using Optional Job Rotation to Gauge On-the-Job Learning, *Journal of Institutional and Theoretical Economics*, Vol. 162, 505-515.

[3] **Baron, J. (1991):** Beliefs about Thinking, in *Informal Reasoning and Education*, J.F. Voss, D.N. Perkins, and J.W. Segal (eds.), Hillsdale, NJ, Erlbaum, 169-186.

[4] **Baron, J. (1995):** Myside Bias in Thinking about Abortion, *Thinking and Reasoning*, Vol. 7, 221-235.

[5] **Burke, L.A.** and **J.E. Moore (2000):** The Reverberating Effects of Job Rotation: A Theoretical Exploration of Nonrotaters' Fairness Perceptions, *Human Resource Management Review*, Vol. 10, 127-152.

[6] **Camerer, C.** and **U. Malmendier (2007):** Behavioral Economics of Organizations, in *Behavioral Economics and Its Applications*, P. Diamond and H. Vartiainen (eds.), Princeton University Press, 235-290.

[7] **Campion, M.A.**, **L. Cheraskin**, and **M.J. Stevens (1994):** Career-Related Antecedents and Outcomes of Job Rotation, *Academy of Management Journal*, Vol. 37, 1518-1542.

[8] **Carmichael, L.** and **B. MacLeod (1993):** Multiskilling, Technical Change and the Japanese Firm, *Economic Journal*, Vol. 103, 142-160.

[9] **Cosgel, M.M.** and **T.J. Miceli (1999):** Job Rotation: Cost, Benefits, and Stylized Facts, *Journal of Theoretical and Institutional Economics*, Vol. 155, 301-320.

[10] **Darley, J.M.** and **P.H. Gross (1983):** A Hypothesis-Confirming Bias in Labelling Effects, *Journal of Personality and Social Psychology*, Vol. 44, 20-33.

[11] **DellaVigna, S.** and **U. Malmendier (2004):** Contract Design and Self-Control: Theory and Evidence, *Quarterly Journal of Economics*, Vol. 119, 353-402.

[12] **Ellison, G. (2006):** Bounded Rationality in Industrial Organization, *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Blundell, Newey and Persson (eds.), Cambridge University Press.

[13] **Eguchi, K. (2005):** Job Transfer and Influence Activity, *Journal of Economic Behavior and Organization*, Vol. 56, 187-197.

[14] **Eliaz, K.** and **R. Spiegler (2006):** Contracting with Diversely Naive Agents, *Review of Economic Studies*, Vol. 73, 689-714.

[15] **Eriksson, T** and **J. Ortega (2006):** The Adoption of Job Rotation: Testing the Theories, *Industrial and Labor Relations Review*, Vol. 59, 653-666.

[16] **Fischhoff, B. (1977):** Perceived Informativeness of Facts, *Journal of Experimental Psychology: Human Perception and Performance*, Vol. 3, 349-358.

[17] **Gabaix, X.** and **D. Laibson (2006):** Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets, *Quarterly Journal of Economics*, Vol. 121, 505-540.

[18] **Gilpatric, S.M. (2008):** Present-Biased Preferences, Self-Awareness and Shirking, *Journal of Economic Behavior and Organization*, Vol. 67, 735-754.

[19] **Gittleman, M.**, **M. Horrigan** and **M. Joyce (1998):** Flexible' Workplace Practices: Evidence from a Nationally Representative Survey, *Industrial and Labor Relations Review*, Vol. 52 , 99-115.

[20] **Henrion, M.** and **B. Fischhoff (1986):** Assessing Uncertainty in Physical Constants, *American Journal of Physics*, Vol. 54, 791-798.

[21] **Hoch, S.J. (1984):** Availability and Inference in Predictive Judgment, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol. 10, 649-662.

[22] **Hoch, S.J. (1985):** Counterfactual Reasoning and Accuracy in Predicting Personal Events, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol. 11, 719-731.

[23] **Ickes, B.W.** and **L. Samuelson (1978):** Job Transfers and Incentives in Complex Organizations: Thwarting the Ratchet Effect, *RAND Journal of Economics*, Vol. 18, 275-286.

[24] **Kidd, J.B. (1970):** The Utilization of Subjective Probabilities in Production Planning, *Acta Psychologica*, Vol. 34, 338-347.

[25] **Koriat, A.**, **S. Lichtenstein** and **B. Fischhoff (1980):** Reasons for Confidence, *Journal of Experimental Psychology: Human Learning and Memory*, Vol. 6, 107-118.

[26] **Kuhn, D. (1989):** Children and Adults as Intuitive Scientists, *Psychological Review*, Vol. 96, 674-689.

[27] **Lingle, J.H.** and **T.M. Ostrom (1981):** Principles of Memory and Cognition in Attitude Formation, in *Cognitive Responses in Persuasive Communications: A Text in Attitude Change*, R.E. Petty, T.M. Ostrom, and T.C. Brock (eds.), Hillsdale, NJ, Erlbaum, 399-420.

[28] **Loftus, E.F.** and **Wagenaar, W.A. (1988):** Lawyers' Predictions of Success, *Jurimetrics Journal*, Vol. 28, 437-453.

[29] **Lord, C.**, **L. Ross** and **(M.R. Lepper (1979):** Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence, *Journal of Personality and Social Psychology*, Vol. 37, 2098-2109.

[30] **Ma, C.A. (1988):** Implementation in Dynamic Job Transfers, *Economics Letters*, Vol. 28, 391-395.

[31] **MacLeod, B. (2003):** Optimal Contracting with Subjective Evaluation, *American Economic Review*, Vol. 93, 216-240.

[32] **Mynatt, C.R.**, **M.E. Doherty** and **R.D. Tweney (1977):** Confirmation Bias in a Simulated Research Environment: An Experimental Study of Scientific Inferences, *Quarterly Journal of Experimental Psychology*, Vol. 29, 85-95.

[33] **Nickerson, R.S. (1998):** Confirmation Bias: A Ubiquitous Phenomenon in Many Guises, *Review of General Psychology*, Vol. 2, 175-220.

[34] **Nisbett, R.E.** and **L. Ross (1980):** Human Inference: Strategies and Shortcomings of Social Judgment, Englewood Cliffs, NJ, Prentice Hall.

[35] **OECD (1999):** OECD Employment Outlook, Paris, Organization for Economic Cooperation and Development.

[36] **Ortega, J. (2001):** Job Rotation as a Learning Mechanism, *Management Science*, Vol. 47, 1361-1370.

[37] **Oskamp, S. (1965):** Overconfidence in Case Study Judgments, *Journal of Consulting Psychology*, Vol. 29, 261-265.

[38] **Osterman, P. (1994):** How Common Is Workplace Transformation and Who Adopts It?, *Industrial and Labor Relations Review*, Vol. 47, 173-188.

[39] **Osterman, P. (2000):** Work Reorganization in an Era of Restructuring: Trends in Diffusion and Effects on Employee Welfare, *Industrial and Labor Relations Review*, Vol. 53, 179-196.

[40] **Perkins, D.N.**, **R. Allen** and **J. Hafner (1983):** Difficulties in Everyday Reasoning, in *Thinking: The Frontier Expands*, W. Maxwell (ed.), Philadelphia, Franklin Institute Press, 83-105.

[41] **Perkins, D.N.**, **M. Farady** and **B. Bushey (1991):** Everyday Reasoning and the Roots of Intelligence, in *Informal Reasoning and Education*, J.F. Voss, D.N. Perkins, and J.W. Segal (eds.), Hillsdale, NJ, Erlbaum, 83-106.

[42] **Pitz, G.F.**, **L. Downing** and **H. Reinhold (1967):** Sequential Effects in the Revision of Subjective Probabilities, *Canadian Journal of Psychology*, Vol. 21, 381-393.

[43] **Prendergast, C. (1999):** The Provision of Incentives in Firms, *Journal of Economic Literature*, Vol. 37, 7-63.

[44] **Prendergast, C.** and **R. Topel (1993):** Discretion and Bias in Performance Evaluation, *European Economic Review*, Vol. 37, 355-365.

[45] **Rabin, M.** and **J.L. Schrag (1999):** First Impressions Matter: A Model of Confirmatory Bias, *Quarterly Journal of Economics*, Vol. 114, 37-82.

[46] **Ross, L. (1977):** The Intuitive Psychologist and his Shortcomings: Distortions in the Attribution Process, in *Advances in Experimental Social Psychology*, L. Berkowitz (ed.), Orlando, FL, Academic Press, Vol. 10, 174-221.

[47] **Ross, L.** and **C. Anderson (1982):** Shortcomings in the Attribution Process: On the Origins and Maintenance of Erroneous Social Assessments, in *Judgments under Uncertainty: Heuristics and Biases*, D. Kahneman, P. Slovic, and A. Tversky (Eds.), Cambridge, Cambridge University Press, 129-152.

[48] **Ross, L.** and **M.R. Lepper (1980):** The Perseverance of Beliefs: Empirical and Normative Considerations, in *New Directions for Methodology of Social and Behavioral Science: Fallible Judgment in Behavioral Research*, R. Shweder and D. Fiske (eds.), San Francisco, Jossey-Bass, Vol. 4, 17-36.

[49] **Ross, L.**, **M. Lepper** and **M. Hubbard (1975):** Perseverance in Self Perception and Social Perception: Biased Attributional Processes in the Debriefing Paradigm, *Journal of Personality and Social Psychology*, Vol. 32, 880-892.

[50] **Schaeffer, R.G. (1983):** Staffing Systems: Managerial and Professional Jobs, Elsevier Science Publisher, Amsterdam.

[51] **Schwartz, B. (1982):** Reinforcement-Induced Behavioral Stereotype: How Not to Teach People to Discover Rules, *Journal of Experimental Psychology: General*, Vol. 111, 23-59.

[52] **Sherman, S.J.**, **K.S. Zehner**, **j. Johnson** and **E.R. Hirt (1983):** Social Explanation: The Role of Timing, Set, and Recall on Subjective Likelihood Estimates, *Journal of Personality and Social Psychology*, Vol. 44, 1127-1143.

[53] **Tetlock, P.E.** and **J.I. Kim (1987):** Accountability and Judgment Processes in a Personality Prediction Task, *Journal of Personality and Social Psychology*, Vol. 52, 700-709.

[54] **Zuckerman, M.**, **R. Knee**, **H.S. Hodgins** and **K. Miyake (1995):** Hypothesis Confirmation: The Joint Effect of Positive Test Strategy and Acquiescence Response Set, *Journal of Personality and Social Psychology*, Vol. 68, 52-60.