Essays on Contracts and Social Preferences

Robertas Zubrickas

EFI THE ECONOMIC RESEARCH INSTITUTE

Dissertation for the Degree of Doctor of Philosophy, Ph.D
Stockholm School of Economics 2009.

Keywords: Mechanism design, non-pecuniary incentives, student grading, compression of ratings, managerial incentives, soft budget constraint, firm-size effects, inequity aversion, preference evolution.

ISBN 978-91-7258-788-5

Printed by:
Elanders, Vällingby 2009

Author photo:
Cecilia Nordstrand

Distributed by:
EFI, The Economic Research Institute
Stockholm School of Economics
P O Box 6501, SE-113 83 Stockholm
www.hhs.se/efi

To Marsalina

PREFACE

This report is a result of a research project carried out at the Centre of Economics at the Economic Research Institute at the Stockholm School of Economics.

This volume is submitted as a doctor's thesis at the Stockholm School of Economics. As usual at the Economic Research Institute, the author has been entirely free to conduct and present her research in her own ways as an expression of her own ideas.

The institute is grateful for the financial support provided by the Jan Wallander and Tom Hedelius Foundation and Inga och Sixten Holmquists stipendiestiftelse which has made it possible to fulfill the project.

Filip Wijkström
Director of the Economic Research Institute
at the Stockholm School of Economics

Paul Segerstrom
Director of the Centre of Economics
at the Stockholm School of Economics

EFI THE ECONOMIC RESEARCH INSTITUTE

EFI Mission

EFI, the Economic Research Institute at the Stockholm School of Economics, is a scientific institution that works independently of economic, political and sectional interests. It conducts theoretical and empirical research in the management and economic sciences, including selected related disciplines. The Institute encourages and assists in the publication and distribution of its research findings and is also involved in the doctoral education at the Stockholm School of Economics. At EFI, the researchers select their projects based on the need for theoretical or practical development of a research domain, on their methodological interests, and on the generality of a problem.

Research Organization

The research activities at the Institute are organized into 20 Research Centres. Centre Directors are professors at the Stockholm School of Economics.

EFI Research Centre:	*Centre Director:*
Management and Organization (A)	Sven-Erik Sjöstrand
Centre for Entrepreneurship and Business Creation (E)	Carin Holmquist
Public Management (F)	Nils Brunsson
Information Management (I)	Mats Lundeberg
Centre for People and Organization (PMO)	Andreas Werr
Centre for Innovation and Operations Management (T)	Pär Åhlström
Centre for Media and Economic Psychology (P)	Richard Wahlund
Centre for Consumer Marketing (CCM)	Magnus Söderlund
Centre for Information and Communication Research (CIC)	Per Andersson
Marketing, Distribution and Industry Dynamics (D)	Björn Axelsson
Centre for Strategy and Competitiveness (CSC)	Örjan Sölvell
Accounting and Managerial Finance (B)	Johnny Lind
Centre for Financial Analysis and Managerial Economics in Accounting (BFAC)	Kenth Skogsvik
Finance (FI)	Clas Bergström
Centre for Health Economics (CHE)	Magnus Johannesson
International Economics and Geography (IEG)	Mats Lundahl
Economics (S)	Paul Segerstrom
Economic Statistics (ES)	Anders Westlund
Centre for Business Law (RV)	Johnny Herre
Centre for Tax Law (SR)	Bertil Wiman

Chair of the Board: Professor Carin Holmquist
Director: Associate Professor Filip Wijkström

Address

EFI, Box 6501, SE-113 83 Stockholm, Sweden • Website: www.hhs.se/efi/
Telephone: +46(0)8-736 90 00 • Fax: +46(0)8-31 62 70 • E-mail efi@hhs.se

Acknowledgement

This thesis is a product of my doctoral studies at the Department of Economics at the Stockholm School of Economics. It would not have come to light without the help of others. First and foremost, I would like to thank my supervisor Karl Wärneryd. His advice, support, and guidance were very important on my way toward graduation. I would also like to thank Professors Tore Ellingsen, Magnus Johannesson, Jörgen Weibull, and all the teachers of the doctoral program of the Stockholm School of Economics and Stockholm University. The years of my study in Stockholm would have been tedious and difficult if it were not for my fellow students who had been very good friends and colleagues throughout those years. Last, but not the least, my gratitude also extends to the administrative staff of the Department of Economics who helped to make my studies and stay in Stockholm smooth and easy. Financial support from the Stockholm School of Economics, the Jan Wallander and Tom Hedelius Foundation, and Inga och Sixten Holmquists stipendiestiftelse is gratefully acknowledged.

Contents

Summary

In this thesis, I study the problems of optimal grading, employee performance evaluation by unaccountable managers, and the evolution of inequity-averse preferences. I approach these problems with contract-theoretic models and demonstrate how certain patterns observed in related empirical work could have emerged.

In Chapter 1, I analyze a teacher-student relationship as a principal-agent model with a costless reward structure. Based on the predictions of my model, I argue that the stylized fact of a mismatch and low correlation between students' abilities and their grades can be the expected-effort-maximizing outcome of teachers' optimal grading. In Chapter 2, I develop a three-tier model of a firm's economic organization, which is centered on the observation that managers do not fully internalize the payroll expenses they incur. Assuming that the degree of manager accountability varies inversely with firm size, I show that the large-firm wage premium and the inverse relationship between wage dispersion and firm size can be the equilibrium outcomes of the agency problem studied in this chapter. The last chapter of my thesis deals with the empirical regularity that members of a market-integrated society are more pro-social than members of an isolated society. I present an evolutionary argument for this endogeneity of people's preferences by showing that inequity aversion with respect to money distribution could have evolved

as an optimal response to merchants' price discrimination.

Chapter 1: Optimal Grading

In this chapter, I study a principal-agent model in which the principal elicits effort from the agent without bearing any cost of rewarding the agent. I develop the model around the premises of a teacher-student relationship with a grade as a costless reward. Other applications of the model could be a job performance appraisal with ratings, where a manager is not held accountable for the payroll expenses he or she incurs, and a parent-child relationship with praise as a reward. Significantly, the existing empirical evidence supports my modeling assumptions and the theoretical findings I obtain, lending credibility to my modeling framework chosen.

In the model, on the part of a teacher (the principal) I solve for the incentive compatible grade-for-learning-effort allocations that maximize a student's expected learning effort. I obtain that irrespective of distribution for student abilities the teacher finds it optimal to pool some of the most efficient student types subject to the highest grade. In other words and more generally, the "no distortion at the top" property does not hold for optimal allocations with costless rewards. I also get that, in classes with lower ability students, teachers should apply more lenient grading standards, and vice versa. This can lead to heterogeneous distributions of grades among classes different in student abilities, and, in particular, to a mismatch and low correlation between students' grades and their abilities, ubiquitously observed in field studies (Goldman & Widawski (1976)). Based on these findings, I conclude that a teacher-student relationship can be considered as an agency relationship with costless rewards.

Similarly, applying the findings of the model to a job performance appraisal with ratings, I argue that the compression of ratings (Murphy & Cleveland (1995))—which in the model takes the form of a pooling equilibrium among some of the most efficient agent types—can be an average-effort-maximizing outcome as well. I explore this result more closely in Chapter 2 of this thesis.

Chapter 2: Manager Accountability and Employee Wage Schedules

This chapter deals with the empirical regularities of the compression of ratings in job performance appraisals (Murphy & Cleveland (1995)) and firm-size effects on employee wage schedules: the large-firm wage premium and inverse relationship between wage dispersion and firm size (Oi & Idson (1999), Stigler (1962)). I attempt to relate and jointly explain these phenomena with the help of a three-tier agency model of a firm's economic organization. The model is developed around the observation from manager and employee compensation practices that managers do not fully internalize the payroll expenses they incur in evaluating their employees' performance. There is also evidence indicating that the degree of manager accountability varies inversely with firm size: the owner of a smaller firm can more closely align her managers' incentives with the firm's profit maximization, and vice versa.

I show that incorporating these features of vertical managerial relationships into an otherwise standard agency model produces theoretical predictions that offer a good match with the empirical stylized facts on wage patterns. My model predicts that the softer a manager's budget constraint is, the more the owner of a firm limits the manager's discretion by lowering an upper bound on employee rewards. In response, the manager designs a flatter pay-for-effort allocation schedule for his employees. As a result, it leads to a coarse distribution of rewards (or ratings), which, I argue, could be behind the compression of ratings phenomenon. With the problem of a soft budget constraint more aggravated in larger firms, the model predicts that the average wage will increase but wage dispersion will decrease with firm size.

All these findings of the model are consistent with the empirical evidence of employee wage patterns in firms. Therefore, I argue that manager accountability (or rather the lack of it) can be a cause of the compression of ratings and firm-size effects on wages.

Chapter 3: How Exposure to Markets Can Favor Inequity-Averse Preferences

In this chapter, I present an evolutionary argument for the recent empirical finding that members of a market-integrated society are more pro-social than members of an isolated society (Henrich et al. (2004)). The main idea is that non-individualistic preferences can be individual fitness maximizing in the presence of general equilibrium externalities. In particular, I show how inequity aversion with respect to money distribution could have evolved as an optimal response to merchants' price discrimination.

I develop a model—a prototype of a traditional small-scale society—in which people's preferences for money distribution are endogenously determined by the economic environment they live in. In the model, individuals share an endowment among themselves and use the proceeds either for immediate consumption or for the purchase of consumption goods from merchants on the external market if such exists.

When facing a monopolist merchant on the external market, an individual who is selected to distribute the endowment among the members of their society may find it own-consumption maximizing to share with others (instead of keeping the endowment all for himself). The intuition is that by sharing with others one can avoid the merchant's full rent extraction and, consequently, increase the purchasing power of one's own, even diminished, endowment share. Then, assuming that increased consumption means increased individual fitness, I argue that inequity-averse behavior with respect to endowment distribution can be an optimal response to external merchants' price discrimination and lead to the evolution of inequity-averse preferences. At the same time, if in the model a society is isolated from any external trades, selfish preferences should prevail—exactly as the empirical evidence indicates to be the case.

Bibliography

Goldman, Roy D., & Widawski, Mel H. 1976. A Within-Subjects Technique for Comparing College Grading Standards: Implications in the Validity of the Evaluation of College Achievement. *Educational and Psychological Measurement*, 36(2), 355–359.

Henrich, Joseph, Boyd, Robert, Bowles, Samuel, Camerer, Colin, Fehr, Ernst, & Gintis, Herbert (eds). 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford University Press.

Murphy, Kevin R., & Cleveland, Jeanette. 1995. *Understanding Performance Appraisal: Social, Organizational, and Goal-based Perspectives*. SAGE.

Oi, Walter Y., & Idson, Todd L. 1999. Firm Size and Wages. *Pages 2165–2214 of*: Ashenfelter, O., & Card, D. (eds), *Handbook of Labor Economics*, vol. 3. Elsevier.

Stigler, George J. 1962. Information in the Labor Market. *Journal of Political Economy*, 70(5), 94–105.

Chapters

Chapter 1

Optimal Grading*

Abstract

In the framework of static mechanism design games with costless non-pecuniary rewards, we solve for optimal student grading standards and attempt to explain the observed mismatch between students' grades and their abilities. The model predicts that the lower the expectations the teacher holds about her students' abilities, the more lenient the grading standards she should set up. Generally, the "no distortion at the top" property does not hold for optimal contracts with costless rewards. As a result, we also argue that the compression of ratings witnessed in job performance appraisals could be an equilibrium outcome. The theoretical findings presented herein are strongly supported by empirical evidence from related literature in psychological and educational measurement.

1. Introduction

The vast literature on subjective evaluation has long dealt with the phenomenon of the compression of ratings, which concerns raters' shallow differentiation of good from bad performance of their ratees as witnessed, for example, in job performance appraisals. (See Murphy & Cleveland (1995) for a review of related studies from the psychological strand of

*I am indebted to Tore Ellingsen and Magnus Johannesson for their research guidance. I would also like to thank Drew Fudenberg, Wayne Grove, Jean Tirole, Karl Wärneryd, Jörgen Weibull, and seminar participants at the University of California, Santa Barbara. Financial support from the European Network for Advancement of Behavioural Economics (ENABLE) and the Jan Wallander and Tom Hedelius Foundation is gratefully acknowledged. School of Economics. Financial support from the Torsten and Ragnar Söderberg Foundation and the Jan Wallander and Tom Hedelius Foundation is gratefully acknowledged.

the literature, or Prendergast (1999) for economists' account on the issue.) We can also position the phenomenon of mismatch or the low correlation between students' university grades and their abilities in this literature, which similarly raises the question of why teachers turn out to be lenient in grading their students (Goldman & Widawski (1976) and Johnson (2003)). In general, and quite surprisingly, these phenomena persist universally in different settings despite the fact that they seemingly lead to inefficient outcomes of principal-agent relationships. Coarse rating and grading schemes are bound to fail to elicit higher effort levels from more able agents. At the same time, there must be some rationale behind those enduring coarse incentive schemes, and this work attempts to contribute toward an understanding of this.

The approach in the current study is to look at the phenomena raised through the perspective of a static principal-agent model with hidden information, the distinct feature of which is a costless and, in most cases, non-pecuniary reward structure. Examples of such rewards and applications of the model are grades in a teacher-student relationship, ratings in a job performance appraisal, and, less tangibly, praise in a parent-child relationship. For ease of exposition, we build the model around the premises of a teacher's designing grading standards to evaluate her students' performance. But we also apply the theoretical results obtained here to other agency problems—a job performance appraisal with ratings, in particular.

Specifically, we consider a grading standard as an implicit contract between a teacher and a student, which the teacher sets up to assign grades for different observable and verifiable student performance levels (such as test scores).[1] The key assumption of the model is that a grade enters only the student's payoff function, while the teacher herself bears no cost of grading the student. Furthermore, there is no scarcity of rewards (grades) in the model, basically implying that the student's expected utility from a targeted grade is independent of the actual distribution of grades in the class.[2] Assuming a deterministic relation between learning effort and performance, we solve for the optimal grading standards on the part of the teacher, who aims at eliciting the highest expected effort from the student.

Based on the theoretical results obtained, we argue that the empirically observed mismatch between students' grades and their abilities and the compression of ratings can be the optimal solutions to particular agency problems. In the model, we show that if the teacher holds low expectations about her students' abilities, then she should apply more lenient grading standards (in order to elicit on average higher effort levels), and *vice versa*. This can lead to heterogeneous distributions of grades among classes different in student abilities, and, in particular, to a mismatch and low correlation between students' grades and their abilities. Similarly, applying the findings of the model to a job performance appraisal with ratings, we argue that the compression of ratings—which in the model takes the form of a pooling equilibrium among some of the most efficient agent types—can be an average-effort-maximizing outcome.

Significantly, the existing empirical evidence strongly supports the

[1]The assumption of verifiability is in contrast to the modeling framework of MacLeod (2003) and Levin (2003)—two related papers on subjective evaluation—where the results mainly hinge on the fact that performance or effort levels are not verifiable.

[2]This assumption makes the model different from other related models with non-pecuniary rewards, such as status incentives, where those rewards are not valuable *per se*, but are valuable because of their scarcity, see Moldovanu et al. (2007); Besley & Ghatak (2008); or Dubey & Geanakoplos (2005).

findings of the model. For instance, Goldman & Widawski (1976) report a negative correlation between students' Scholastic Aptitude Test scores (which could be seen as a proxy measure of students' abilities) and the grading standards in the classes the students were majoring in. According to this study (conducted at University of California, Riverside), the negative correlation observed is due to the fact that professors in a field consisting of students with high abilities tend to grade more stringently than do professors in a field with lower-ability students—precisely as our model predicts. These empirical findings were confirmed by similar studies conducted at Dartmouth College (Strenta & Elliott (1987)) and at Duke University (Johnson (2003)).

This work has been largely motivated by findings from behavioral and experimental economics literature on contractual relationships. There are a growing number of theoretical and empirical studies arguing that along with pecuniary incentives people also tend to care about non-pecuniary incentives and rewards such as intrinsic motivation, perceived trustworthiness and importance, or elicited praise and esteem (Brennan & Pettit (2004); Frey & Osterloh (2002); Berg et al. (1995); and Falk & Kosfeld (2006)). As a result, the standard principal-agent model with merely monetary incentives may fail to explain the variation in exerted effort levels by employees in fixed-wage jobs with no other present or future pecuniary stimuli, as discussed in Akerlof (1982). The obvious direction for further research is now to enrich the standard model with more elaborate mechanisms aimed at more closely capturing the complexities of a principal-agent relationship (see Bénabou & Tirole (2003); Moldovanu et al. (2007); Besley & Ghatak (2008); or Sliwka (2007)). The current study can also be seen as a part of this more general research line.

The proposed refinement that, unlike the agent, the principal is indifferent to a transfer between them is by no means new in the contract theory literature. It was formally studied, for example, in Guesnerie & Laffont (1984), one of the founding articles on mechanism design aimed at providing an all-encompassing solution to a broadly defined principal-agent problem. In particular, they distinguish between "type A" and "type B" preferences, where with the former preferences the principal's utility does not depend on a transfer, while with the latter (conventional) preferences it does. In their study, however, the "type A" preferences are primarily used to analyze a social planner's problem of social welfare maximization. There, a transfer between the social planner (principal) and the agent is equivalent, figuratively speaking, to distributing money between two pockets of the same jacket, leaving the social welfare intact. Therefore, the framework of Guesnerie & Laffont (1984) does not apply to the problem studied here. In our model, the principal is, in fact, more of "type B", i.e., she cares only about her own utility but does not pay for motivating the agent.

Nor does this study stand alone in designing optimal grading rules.[3] Dubey & Geanakoplos (2005) also target the same problem but from a different perspective: they model a teacher-student relationship as a game of status with private information and stochastic output similar to a tournament. Hence, unlike in our study, they present a multiple-agent model, where an agent's utility from a grade depends on his or her class

[3]However, it needs to be reckoned that not much theoretical work has been done on modeling a teacher-student relationship as a principal-agent model on its own, whereas this relationship has typically been modeled as part of a more global game involving potential employers or university administration (see Ostrovsky & Schwarz (2003)). At the same time, more research has been done on the empirical side of the problem (see Johnson (2003)).

ranking, i.e., status, resulting from the grade rewarded, but not on the grade *per se*. Also given some other modeling differences, we draw different conclusions about optimal grading schemes. Dubey & Geanakoplos (2005) find that teachers should use coarse grading schemes and "pyramid" the allocation of grades: in equilibrium the highest grade would be available to fewer students than the second-highest grade, and so on.[4] Our model similarly predicts that teachers should apply coarse grading schemes, but the level of "coarsening" depends on the distribution of student abilities. Furthermore, we do not find "pyramiding" to be an optimal grade allocation rule, especially, when there is a large mass of less able students in the class.

The remainder of this chapter is organized as follows. Section 2 introduces the model, which is solved in Section 3. Section 4 discusses the main findings of the model and relates them to the phenomena raised in the introduction. Section 5 reviews the existing empirical evidence, and Section 6 concludes the study.

2. Model

This section presents a principal-agent model with costless non-pecuniary rewards, or more specifically, a teacher-student model with grades as a reward. For expositional tractability and whenever it is possible, it closely follows the textbook variants of the related static models with hidden information and the single agent as in, e.g., Bolton & Dewatripont (2004) or Fudenberg & Tirole (1991).

Referring to the empirical findings in Arcidiacono (2004), students select majors in college primarily because of their intrinsic preferences

[4] Moldovanu et al. (2007) makes a similar prediction as well.

for those majors and not so much because of their expected future earnings (which are, on the other hand, correlated with university grades). In our model, therefore, we assume that teachers find students' class selection exogenous with respect to the grading standards to be used in class. Restricting the model's application to university majors only, we shall ignore possible reputational concerns that teachers are thought to face when designing grading standards such as a concern that stricter grading practices may scare some students away. (See, for example, Johnson (1997), which, however, admits that reputational concerns are more prevalent for optional classes than for classes in the students' major.) Neither will the empirical evidence of the model's findings be prejudiced, for the evidence presented in this chapter comes from grading standards in classes in the students' major.

Therefore, consider a teacher who has to set up grading rules for the students enrolled in her class, who come from the population of students with a preference for the subject taught by the teacher. A grading rule assigns grades to students' performance levels, which are observable and verifiable, and are assumed to be perfectly correlated with students' costly learning efforts. The teacher's utility increases in her students' effort levels, but is independent of grades. In contrast, students derive utility only from their own grades[5] and incur disutility from studying (say, because of opportunity costs). By this, we assume that the expected distribution of grades in class does not affect a student's learning effort choice decision (accordingly, his or her *ex ante* utility from a tar-

[5]... due to, for instance, better signals sent to their parents, to friends, or even to themselves about their personal characteristics. Ideally, one would want to think of grades as ability signals to be sent to the labor market, but then the model would need to be closed by introducing one more stage, the recruitment stage (e.g., Ostrovsky & Schwarz (2003)). This extension, however, would eventually require elaborating on the school entry decision as well.

geted grade).[6] Furthermore, even though students are identical with respect to their attempts to save on effort needed for a given grade, they are different in learning abilities, which are their private information, and the teacher knows only the population distribution for student abilities. These assumptions allow us to treat the teacher-student relationship described as a single-agent model with hidden information.[7]

The teacher's goal is to elicit the highest expected effort from the student subject to the individual rationality and incentive compatibility constraints as described below. Using the direct mechanism approach, the teacher designs for the student a grading rule, which is a set of effort-grade allocations $\{x(\theta), r(\theta)\}_{\theta \in [\underline{\theta}, \overline{\theta}]}$, where the parameter θ is the student's ability distributed according to a twice differentiable cumulative distribution function F over the student ability space $[\underline{\theta}, \overline{\theta}]$ with the probability density function f $(f > 0)$. The grade function r maps the ability set $[\underline{\theta}, \overline{\theta}]$ into $[0, 1]$ so that the maximum grade the teacher can offer is 1[8]; the effort function x maps $[\underline{\theta}, \overline{\theta}]$ into a bounded interval $[0, \bar{x}]$, which, when needed, is assumed to be large enough to allow for an interior solution.

From the set of effort-grade allocations $\{x(\theta), r(\theta)\}_{\theta \in [\underline{\theta}, \overline{\theta}]}$ offered by the

[6] An important refinement that students have inter-dependent preferences is left for future research (also, see Dubey & Geanakoplos (2005)).

[7] If a hidden-information framework seems restrictive, as it could be thought of when considering teacher-student relationships, then we can, alternatively, allow that the teacher is able to tell a student's type (similarly to Bénabou & Tirole (2003)). Then, we would require that the teacher cannot discriminate among her students by applying ability-specific grading rules. With this alternative formulation, the optimization problem would basically remain intact as in the case with the hidden-information framework adopted, and, therefore, the latter is retained for its link with the existing literature.

[8] Without an upper bound on the reward function, the optimal grading rule would be to demand the maximal feasible effort from every student type and to reward them whatever abundantly. Putting an upper bound on the reward function not only remedies the implausibility of the solution that arises, but it is also a very natural thing to impose when considering a teacher-student relationship, where there is typically a formal, institutionally set highest grade. Similarly, job performance is also normally appraised on a finite rating scale, and, finally, even praise, which could be thought of as unbounded, may still have only a limited effect on the agent's utility resulting from it.

teacher, the student chooses a type $\hat{\theta} \in [\underline{\theta}, \overline{\theta}]$ to report to the teacher. Then, the assigned allocation $(x(\hat{\theta}), r(\hat{\theta}))$ renders the student of ability θ a net utility of

$$U_A(x(\hat{\theta}), r(\hat{\theta}), \theta) = r(\hat{\theta}) - C(x(\hat{\theta}), \theta),$$

where C is a learning effort cost function with the properties $C_x > 0$, $C_{xx} > 0$, $C_\theta < 0$, and $C_{x\theta} \leq 0$. The student's reservation utility is set to be equal to 0 (for all ability types). The teacher's utility resulting from the implemented allocation $(x(\hat{\theta}), r(\hat{\theta}))$ is equal to

$$U_P(x(\hat{\theta}), r(\hat{\theta})) = V(x(\hat{\theta})),$$

which is concave in effort x. Next, the above utility levels are assumed to satisfy the von Neumann-Morgenstern axioms, and the teacher and the student are both rational utility maximizers.

Applying the Revelation Principle, with respect to direct effort-grade allocations $\{x(\theta), r(\theta)\}_{\theta \in [\underline{\theta}, \overline{\theta}]}$ the teacher maximizes her expected utility

$$\int_{\underline{\theta}}^{\overline{\theta}} V(x(\theta)) f(\theta) d\theta \tag{1.1}$$

subject to

$$r(\theta) - C(x(\theta), \theta) \geq r(\hat{\theta}) - C(x(\hat{\theta}), \theta), \tag{1.2}$$

$$r(\theta) - C(x(\theta), \theta) \geq 0, \tag{1.3}$$

$$0 \leq r(\theta) \leq 1, \text{ for all } \theta \text{ and } \hat{\theta} \text{ in } [\underline{\theta}, \overline{\theta}]. \tag{1.4}$$

In the above, (1.2) is the incentive compatibility constraint, (1.3) is the individual rationality or participation constraint, and the last constraint imposes an upper bound on the teacher's rewards.

The solution to the above problem characterizes the Bayesian-Nash equilibrium of the teacher-student agency problem studied above. The equilibrium consists of the set of effort-grade allocations designed by the teacher and the ability type reported by the student of every type such that the teacher and the student of every type play their best replies given the strategy of the other.

Prior to deriving the conditions that characterize the equilibrium play, we place some further structure on the model to ensure the uniqueness of the problem's solution. First, the single-crossing condition holds (in fact, it has been imposed by the cost function's property $C_{x\theta} \leq 0$). Second, we impose the assumption of the monotone increasing hazard rate, which is defined as $h(\theta) = f(\theta)/(1 - F(\theta))$, and the assumption implies that $h'(\theta) \geq 0$. Finally, we assume that the cost function C is separable in effort x and type θ, and takes the form of $C = g(x)/\theta$, where g is a convex function, implying that $C_{xx\theta} \leq 0$.

3. Solution

The standard principal-agent model with monetary (i.e., costly) transfers is solved using the method attributable to Mirrlees (1971). Its main idea is to obtain a functional equation with one unknown by merging the agent's optimization problem with that of the principal. In our case with costless rewards, we need to modify this solution method. The reason is that, unlike in the standard model, the intercomparison of the agent's and principal's utilities is not possible in our model, for there is no linking term between the two utility functions (as is typical for the transfer function). The existence of an upper bound also makes our model different from the standard model. Therefore, to look into the main properties of the solution to the optimization problem (1.1)–(1.4), we, firstly, approach it through its discrete version, as suggested below. Then, we take the limit of the discrete-case results to arrive at the general solution with the continuum of agent types.[9]

Discretization

We discretize the student ability type space $\left[\underline{\theta}, \overline{\theta}\right]$ into n discrete types $(\theta_1, ..., \theta_i, ..., \theta_n)$, where a student type $\theta_i = \underline{\theta} + (i-1)\partial\theta$, for $i = 1, ..., n$, and $\partial\theta = (\overline{\theta} - \underline{\theta})/n$. Then, we discretize the initial (continuous) distribution F for student types by defining probability weights $p(\theta_i) = \int_{\theta_i}^{\theta_i+\partial\theta} f(\theta)d\theta$ for every θ_i, which is the probability mass of the student types within the interval $[\theta_i, \theta_i + \partial\theta]$. (From this discretization, we later switch to the continuous case by taking the limit $n \to \infty$, or $\partial\theta \to 0$.)

The discrete version of the optimization problem (1.1)–(1.4) is defined below, where the teacher with respect to effort-grade allocations $\{x(\theta_i), r(\theta_i)\}_{i=1}^{n}$ maximizes her expected utility

$$\sum_{i=1}^{n} p(\theta_i)V(x(\theta_i))$$

subject to

$$r(\theta_i) - C(x(\theta_i), \theta_i) \geq 0, \qquad (IR_i)$$

$$r(\theta_i) - C(x(\theta_i), \theta_i) \geq r(\theta_j) - C(x(\theta_j), \theta_i), \qquad (IC_i)$$

$$0 \leq r(\theta_i) \leq 1, \text{ for every } i = 1, ..., n \text{ and } j \neq i.$$

[9]In footnote 7 we suggest that, instead of the hidden-information framework, we could use the perfect-information framework with the condition that all the students irrespective of their abilities are subject to the same grading rules. Under this alternative framework, if there are a finite number of students in the class, then we face a discrete-type problem, as defined in the main text that follows. The limiting continuous case is then just a convenient way of summarizing the properties of the solution to the discrete-type problem.

To solve the above maximization problem, we start with reducing it by getting rid of redundant constraints. Then, from the reduced problem we derive the first-order conditions that characterize the optimal effort-grade allocations. But first, we make the following conjecture.

Conjecture 1.1 *For any partition of the student type space, the equilibrium effort-grade allocations are distinct for every student type.*

Setting up the Lagrangian

As it is standard in principal-agent models with hidden information, among the agent types, who are subject to non-zero allocations, the only individual rationality constraint that binds is that of the least efficient agent. For the rest of the chapter, we assume that the teacher contracts upon all the student types (actually, the incidence of the shutdown of some types is much less likely with costless rewards than it is with costly rewards). Then, it is IR_1 constraint that binds. Furthermore, in the solution the adjacent incentive compatibility constraints need to be downward binding:

$$r(\theta_i) - C(x(\theta_i), \theta_i) = r(\theta_{i-1}) - C(x(\theta_{i-1}), \theta_i), \qquad i = 2, ..., n. \tag{1.5}$$

Supposing that in the solution the effort schedule x is increasing in student type (as also implied by Conjecture 1.1), which has to be checked separately, then due to the single-crossing condition we can ignore the rest of constraints. Next, we make an observation that it must be optimal for the teacher to make at least some type subject to the highest grade of 1 since it costs nothing to the teacher. Then, as follows from Conjecture 1.1, in the equilibrium it is only the most efficient student type that receives the highest grade, i.e., $r(\theta_n) = 1$.

Having said that, we can eliminate the reward allocations r from the optimization problem by combining all the binding constraints together with $r(\theta_n) = 1$ into one optimality constraint:

$$1 - \sum_{i=1}^{n} C(x(\theta_i), \theta_i) + \sum_{i=2}^{n} C(x(\theta_{i-1}), \theta_i) = 0. \tag{1.6}$$

Finally, we set up the Lagrangian of the reduced optimization problem (without the monotonicity constraint on the effort x), which is

$$\begin{aligned} L(\{x(\theta_i)\}_{i=1}^{n}, \mu) \;=\; & \sum_{i=1}^{n} p(\theta_i) V(x(\theta_i)) + \\ & + \mu(1 - \sum_{i=1}^{n} C(x(\theta_i), \theta_i) + \sum_{i=2}^{n} C(x(\theta_{i-1}), \theta_i)), \end{aligned} \tag{1.7}$$

where μ is a Lagrange multiplier on constraint (1.6).

The first-order conditions with respect to effort levels $x(\theta_i)$ for $i = 1, ..., n-1$ are

$$p(\theta_i) V_x(x(\theta_i)) = \mu(C_x(x(\theta_i), \theta_i) - C_x(x(\theta_i), \theta_{i+1})), \tag{1.8}$$

and with respect to $x(\theta_n)$ it is

$$p(\theta_n) V_x(x(\theta_n)) = \mu C_x(x(\theta_n), \theta_n). \tag{1.9}$$

The equilibrium effort-grade allocations need to satisfy the above first-order conditions, which means equating the teacher's gains and losses from a marginal change in the equilibrium allocations.

Dividing the last two first-order conditions (of the two highest contracted effort levels) renders

$$\frac{p(\theta_n) V_x(x(\theta_n))}{p(\theta_{n-1}) V_x(x(\theta_{n-1}))} = \frac{C_x(x(\theta_n), \theta_n)}{C_x(x(\theta_{n-1}), \theta_{n-1}) - C_x(x(\theta_{n-1}), \theta_n)}. \tag{1.10}$$

Multiplying both sides by $\partial\theta$ and then taking the limit $\partial\theta \to 0$ give us the left-hand side of (1.10) which tends to 0 (with $\partial\theta \to 0$, $p(\theta_n)/p(\theta_{n-1}) \to 1$ and $V(x(\theta_n))/V(x(\theta_{n-1})) \to 1$), while the limit of the right-hand side stays positive:

$$C_x(x(\overline{\theta}),\overline{\theta})/(-C_{x\theta}(x(\overline{\theta}),\overline{\theta})) > 0.$$

Therefore, condition (1.10) cannot hold for fine enough partitions of the student type space and, in particular, for the continuum of types θ (i.e., at the limit $\partial\theta \to 0$). Hence, the conjecture that the solution takes the form of a separating equilibrium cannot be universally valid. This implies that, for fine partitions of the type space, the perfect screening of student types (actually, that of the most efficient types) cannot be optimal for the teacher.

Intuitively, the finding that there is no separating equilibrium among the most efficient student types should not be surprising. Suppose it were the case. Then, consider the teacher increasing the effort level demanded of the second-most-efficient student type from $x(\theta_{n-1})$ to $x'(\theta_{n-1})$ against the corresponding grade increase from $r(\theta_{n-1})$ to $r'(\theta_{n-1}) = 1$ so that this change is acceptable to the student (otherwise, he would not report truthfully). Compare the teacher's gains and losses from this change: the loss is the most efficient student type's effort reduction from $x(\theta_n)$ to $x'(\theta_{n-1})$, where it must be that $x(\theta_n) > x'(\theta_{n-1})$. At the same time, the accrued gain is not only the increase in the second-most-efficient student type's effort level by $x'(\theta_{n-1}) - x(\theta_{n-1}) > 0$, but actually it is the whole string of follow-up increases in other effort levels $x(\theta_i) \to x'(\theta_i)$, $i = 1, ..., n-2$, made at no extra cost (due to costless rewards) to fill the slack in the incentive-compatibility constraints arising after the increase in the reward $r(\theta_{n-1})$ to $r'(\theta_{n-1}) = 1$. Hence, unless the probability mass of the most efficient type θ_n is big enough (which in the continuous case is possible only for some irregular distribution F for student abilities), the teacher's gain from the change described is larger than the corresponding loss. Therefore, the teacher can do better by pooling some of the most efficient student types by making them subject to the highest reward of 1 until their probability mass is big enough to offset the gains and losses described above.

"Pooling at the top" interval

When condition (1.10) does not hold—which occurs, as we showed, for fine partitions of the student type space—we continue by gradually increasing the probability mass of student types subject to the highest grade and denote this mass by $P(\theta_m) = \sum_{j=m}^{n} p(\theta_j)$, where $m = n - 1, n-2, \ldots$. We repeat the above solution algorithm for different m (with m replacing n in the above derivations) until we have the optimality conditions met. In particular, for a given m, the first-order condition equivalent to (1.9) is

$$P\left(\theta_m\right) V_x(x(\theta_m)) = \mu C_x(x(\theta_m), \theta_m).$$

Accordingly, the updated condition (1.10) takes the form of

$$\frac{P(\theta_m) V_x(x(\theta_m))}{p(\theta_{m-1}) V_x(x(\theta_{m-1}))} = \frac{C_x(x(\theta_m), \theta_m)}{C_x(x(\theta_{m-1}), \theta_{m-1}) - C_x(x(\theta_{m-1}), \theta_m)}.$$

Leaving the details of the discrete-case problem aside, from now on we revert to the continuous case only. Again, multiplying both sides of the previous expression by $\partial\theta$ and taking the limit $\partial\theta \to 0$ on both sides (and applying L'Hôpital's Rule) give the condition for the pooling equilibrium among the most efficient agents:

$$\frac{1-F(\theta)}{f(\theta)} = -\frac{C_x(x(\theta),\theta)}{C_{x\theta}(x(\theta),\theta)}.$$

Given our assumption that the cost function $C(x,\theta)$ is separable in x and θ, with $C(x,\theta)=g(x)/\theta$, the above condition becomes

$$\frac{1-F(\theta)}{f(\theta)} = \theta. \tag{1.11}$$

Since there may be no type θ in $\left[\underline{\theta},\overline{\theta}\right]$ for which condition (1.11) holds[10], then the starting value of the "pooling at the top" interval is defined as

$$\theta^* = \min\{\theta : (1-F(\theta))/(\theta f(\theta)) \leq 1, \theta \in \left[\underline{\theta},\overline{\theta}\right]\}. \tag{1.12}$$

Note that the expression $(1-F(\theta))/(f(\theta)\theta)$ is monotonically decreasing in θ due to the monotone hazard rate assumption so that the pooling-equilibrium starting point θ^* is uniquely determined. All in all, the pooling-equilibrium effort-grade allocation for every type θ in $\left[\theta^*,\overline{\theta}\right]$ is $(x(\theta^*),1)$, where the effort level $x(\theta^*)$ is determined by the remaining conditions as shown below.

The equilibrium effort-grade allocations

As for the student types in the interval $[\underline{\theta},\theta^*)$, the equilibrium allocations need to satisfy the remaining first-order conditions (1.8), the continuous version of which for any θ in $[\underline{\theta},\theta^*)$ is

$$f(\theta)V_x(x(\theta)) = -\mu C_{x\theta}(x(\theta),\theta). \tag{1.13}$$

The Lagrange multiplier μ, as determined by the updated expressions

[10] The pooling-equilibrium interval comprises the whole agent-type space if, for instance, student types θ are uniformly distributed with $\underline{\theta} \geq \overline{\theta}/2$.

for (1.9) and (1.11), is equal to

$$\mu = -\frac{f(\theta^*)V_x(x(\theta^*))}{C_{x\theta}(x(\theta^*), \theta^*)}. \tag{1.14}$$

From (1.13) and (1.14) one can observe that there is no discontinuity at the equilibrium effort allocations at the point θ^* since $x(\theta) \to x(\theta^*)$ as $\theta \to \theta^*$.

From condition (1.13) and the expression for μ, we can express $x(\theta)$ for every θ in $[\underline{\theta}, \theta^*)$ in terms of the starting point θ^* of the pooling-equilibrium interval, which is determined by (1.12), and the highest demanded effort allocation $x(\theta^*)$, which still needs to be determined. To do so, we consider the continuous version of the optimality condition (1.6), which takes the form of

$$1 - C(x(\theta^*), \theta^*) + \int_{\underline{\theta}}^{\theta^*} C_\theta(x(\theta), \theta)d\theta = 0. \tag{1.15}$$

Then, after plugging in the expression for $x(\theta)$ from (1.13) and (1.14) into the above optimality condition, we can integrate out the parameter θ and obtain the expression for the highest demanded effort allocation. After having determined $x(\theta^*)$, the remaining equilibrium effort levels immediately follow from (1.13). The constraint for the effort function x to be non-decreasing is met, which follows from equation (1.11) and the monotone hazard rate assumption.[11]

Denote the derived equilibrium effort levels by $x^*(\theta)$ for θ in $[\underline{\theta}, \theta^*]$. The equilibrium grade levels $r^*(\theta)$ for θ in $[\underline{\theta}, \theta^*)$ are found from the set

[11]From equation (1.11) it follows that $f(\theta)/(1-F(\theta)) < 1/\theta$ for θ in $[\underline{\theta}, \theta^*)$, and from the monotone hazard rate: $f'(\theta) > -f^2(\theta)/(1 - F(\theta))$. Then, taking the internal derivative of (1.13) together with the above properties renders that $dx/d\theta > 0$.

of the binding incentive-compatibility constraints and are given by

$$r^*(\theta) = C(x^*(\theta), \theta) - \int_{\underline{\theta}}^{\theta} C_\theta(x^*(\tilde{\theta}), \tilde{\theta}) d\tilde{\theta}, \tag{1.16}$$

where $x^*(\theta)$ is the optimal effort allocation for student type θ.

Special case

Below, we illustrate the solution by considering a special case, where the teacher's utility function is linear in effort x, $V(x) = x$, and the student's cost function is $C(x, \theta) = x^2/(2\theta)$.

The pooling-equilibrium condition (1.11) remains intact, while the first-order condition (1.13) becomes for θ in $[\underline{\theta}, \theta^*)$

$$x(\theta) = \frac{x(\theta^*)}{f(\theta^*)\theta^{*2}} f(\theta)\theta^2.$$

Plugging this expression into the optimality constraint (1.15) renders

$$1 - \frac{(x(\theta^*))^2}{2\theta^*} - \frac{(x(\theta^*))^2}{f^2(\theta^*)\theta^{*4}} \int_{\underline{\theta}}^{\theta^*} \frac{f^2(\tilde{\theta})\tilde{\theta}^2}{2} d\tilde{\theta} = 0,$$

from which we can determine the pooling-equilibrium effort level $x(\theta^*)$:

$$x(\theta^*) = \left(\frac{1}{2\theta^*} + \frac{\int_{\underline{\theta}}^{\theta^*} \frac{f^2(\tilde{\theta})\tilde{\theta}^2}{2} d\tilde{\theta}}{f^2(\theta^*)\theta^{*4}}\right)^{-0.5}.$$

Having determined the highest contracted effort level, from the above expression for $x(\theta)$ we can find the remaining equilibrium effort levels. Finally, the equilibrium grade levels follow immediately from (1.16).

The Equilibrium

Proposition 1.1 below characterizes the Bayesian-Nash equilibrium of the teacher-student agency problem studied above.

Proposition 1.1 *Given the assumptions of the model, the solution to the optimization problem (1.1)–(1.4) constitutes the Bayesian-Nash equilibrium, where the teacher's designed effort-grade allocations* $(x^*(\theta), r^*(\theta))$ *and the student's reported type* $\hat{\theta}^*(\theta)$ *for every* θ *in* $\left[\underline{\theta}, \overline{\theta}\right]$ *are such that*

- *for every student type* θ *in* $\left[\theta^*, \overline{\theta}\right]$*, where* θ^* *as in (1.12), the effort-grade allocations are* $r^*(\theta) = 1$*,* $x^*(\theta) = x(\theta^*)$*, where* $x(\theta^*)$ *is determined by (1.13), (1.14) and (1.15); and the student of type* θ *reports* $\hat{\theta}^*(\theta) = \tilde{\theta}$*,* $\tilde{\theta} \in \left[\theta^*, \overline{\theta}\right]$*.*

- *for every* θ *in* $[\underline{\theta}, \theta^*)$*, given* θ^* *and* $x(\theta^*)$ *from above, the effort-grade allocations* $(x^*(\theta), r^*(\theta))$ *are defined by (1.13) and (1.16), respectively; and the student of type* θ *reports* $\hat{\theta}^*(\theta) = \theta$*.*

4. Main Findings and Discussion

Here, we provide an intuitive account of the solution obtained for the agency problem with costless (non-pecuniary) rewards studied earlier. Simultaneously, in this and subsequent sections, we try to relate the predictions of the model to the empirical evidence relevant to the framework examined here.

Compression of Ratings

One of the main theoretical findings of this study is that in a principal-agent model with costless rewards there must be a pooling equilibrium for at least some of the most efficient agents. This result is in sharp contrast to the properties of the standard model with costly rewards, where the "no distortion at the top" property typically holds.[12]

[12]While the "no distortion at the top" property is characteristic of principal-agent models with monetary rewards, see, e.g., Mirrlees (1971) , it has also been shown to hold for agency problems with

Proposition 1.2 *For an agency problem with costless rewards, the "no distortion at the top" property does not hold.*

The proof of this result was provided in the previous section, where we showed that a uniform effort-grade allocation applies to all agent types from the non-empty interval $\left[\theta^{*},\overline{\theta}\right]$. Moreover, the above proposition is silent about the reward structure having an upper bound, for neither the size of this bound nor its mere existence has any impact on the result. (In a principal-agent model with costly rewards, imposing an upper bound on the reward function does not lead to a pooling equilibrium among the most efficient agents as long as this constraint is not binding, i.e., when the upper bound is big enough.)

Bearing this result in mind, we can address the phenomenon of the compression of ratings in job performance appraisals. The incidence of a pooling equilibrium—which, as been noted, can stretch out to comprise the whole agent type space—means that, within the setting studied, if the agent, for any reason, exerts an effort level above the highest contracted effort level $x(\theta^{*})$, he will still receive the reward of 1. Hence, despite variability in effort levels one can still observe compressed rewards, but it could be the outcome of the contractual arrangement optimally set by the principal. As has been analytically argued, it may not be optimal for the principal to differentiate much among her agents in evaluating their performance. Distinguishing the most efficient agents can only be achieved by suppressing the motivation of less efficient agents. But it may have a larger adverse effect on aggregate (or average) performance than the higher contribution achieved from the most efficient agents.

status incentives, see Moldovanu et al. (2007). (However, recently, there have also been papers in which this property does not hold in the optimum, see Levin (2003) or MacLeod (2003).)

Hence, the "leniency bias" rating scheme could, in fact, be optimal.[13]

Furthermore, the compression of ratings phenomenon becomes even more prevalent the more concave the principal's utility function is. Despite the fact that in our model the functional form of the principal's utility function does not affect the size of the pooling-equilibrium interval, as one can see from (1.11), it does, however, affect the range of effort-grade allocations. Intuitively, and as can also be shown analytically, if the teacher, referring back to the teacher-student framework studied here, puts a relatively higher weight on the performance of less able students, then we should observe a flatter effort-grade allocation schedule in place. This observation shows that changing teachers' incentives can have direct implications for student grading standards and, accordingly, effort levels demanded.

This finding is especially important in light of recent merit-pay programs introduced to improve teachers' incentives for motivating their students (see Lavy (2002); Atkinson et al. (2004); Lazear (2003)). Arguably, it is in the hands of social planners to affect the functional form of the teacher's utility function through pay-for-student-performance incentive schemes for teachers. Then, our model could be of help in designing optimal incentive schemes for teachers in order to align teachers' incentives more closely with those of social planners (or with socially desirable outcomes).

With respect to a job performance appraisal with ratings and its link to our model, it remains true that a manager wants to elicit more effort from her employees. Still, unlike in a teacher-student relationship,

[13]The above explanation is in sharp contrast to explanations from the psychological literature on subjective evaluation, a common example of which is given in Prendergast (1999, footnote 34): "An obvious reason for this [leniency bias] is that it is simply unpleasant for supervisors to offer poor ratings to workers, so they avoid this pain."

a manager's designed rating scheme is typically linked to the employees' pay, with a higher rating implying a higher pay, which, of course, does not fit the definition of "costless non-pecuniary rewards." However, having inquired into the inner workings of rating-pay schemes[14], in many instances we can still think of job performance appraisals in terms of a principal-agent model with costless rewards. Commonly, employees' performance is evaluated by their line managers, i.e., by low-rank managers. At the same time, managers' own incentive scheme may not fully internalize the payroll cost resulting from their evaluations of employees' performance, but it does depend on their aggregate performance (say, through bonuses at the end of the year). Therefore, we may obtain a situation where a line manager faces a soft budget constraint in designing rating schemes to elicit the highest effort expected from her employees. The possibility described, then, falls within our model. But, given its specific nature, it needs to be studied separately, and here it serves only as a potential alternative application of the model.

Mismatch between Grades and Abilities

Consider two classes of students, who come from two different populations of students, whose abilities are distributed on the same support $[\underline{\theta}, \overline{\theta}]$ according to twice differentiable distributions F_1 and F_2, respectively. Denote the student types from the two classes by θ_1 and θ_2, respectively, and let student type θ_2 be smaller than θ_1 in the likelihood ratio order, i.e.,

$$\frac{f_2(\theta)}{f_1(\theta)} \text{ decreases for all } \theta \text{ in } \left[\underline{\theta}, \overline{\theta}\right].$$

[14] I largely owe the following discussion to Ailko van der Veen.

The above stochastic dominance condition can be interpreted to mean students from the first class are held to be more able than those from the second class. Let $\{x_1(\theta), r_1(\theta)\}$ and $\{x_2(\theta), r_2(\theta)\}$ for θ in $[\underline{\theta}, \overline{\theta}]$ be the solutions to the optimization problem (1.1)–(1.4) for the student ability distributions F_1 and F_2, respectively. Then, the following holds.

Proposition 1.3 *If student type θ_2 is smaller than the student type θ_1 in the likelihood ratio order, then the grade allocations in the two classes satisfy $r_2(\theta) \geq r_1(\theta)$ for every student type θ in $[\underline{\theta}, \overline{\theta}]$.*

To put it in words, the lower the expectations the teacher holds about her students' abilities, the more lenient she should be when setting grading standards.

Proof. With reference to the pooling-equilibrium condition (1.11), define $G_i(\theta) = (1 - F_i(\theta))/(\theta f_i(\theta))$, and let $\theta_i^* = \min\{\theta : G_i(\theta) \leq 1, \theta \in [\underline{\theta}, \overline{\theta}]\}$, $i = 1, 2$. Since the likelihood ratio order implies the hazard rate order (see Shaked & Shanthikumar (1994)), which is $f_1(\theta)/(1 - F_1(\theta)) \leq f_2(\theta)(1 - F_2(\theta))$ for every θ, then it immediately follows that $G_1(\theta) \geq G_2(\theta)$, leading to $\theta_1^* \geq \theta_2^*$. Hence, we have $r_2(\theta) = 1 \geq r_1(\theta)$ for $\theta \in [\theta_2^*, \overline{\theta}]$. Next, consider θ in $[\underline{\theta}, \theta_2^*]$; denote the Lagrange multipliers from the two optimization problems, as defined in (1.14), by μ_i, $i = 1, 2$. Then, divide the equilibrium conditions for the effort levels x_1 and x_2 as in (1.13) to obtain for any θ in $[\underline{\theta}, \theta_2^*]$

$$\frac{V_x(x_1(\theta))g_x(x_2(\theta))}{V_x(x_2(\theta))g_x(x_1(\theta))} = \frac{\mu_1}{\mu_2}\frac{f_2(\theta)}{f_1(\theta)}.$$

Since the effort level $x_2(\theta_2^*)$ must be at least as large as $x_1(\theta_2^*)$, which stems from the second teacher's incentive to expand the pooling equilibrium even further, then at $\theta = \theta_2^*$ the left-hand side of the above equation is greater than or equal to 1, and so is the right-hand side. Due to the decreasing likelihood ratio f_2/f_1, the right-hand side stays greater than 1 for any θ in $[\underline{\theta}, \theta_2^*)$, and so does the left-hand side, implying that $x_2(\theta) \geq x_1(\theta)$ for every θ in $[\underline{\theta}, \theta_2^*)$, which subsequently leads to $r_2(\theta) \geq r_1(\theta)$. ■

Propositions 1.2 and 1.3 are illustrated in Figure 1. It depicts the optimal grading standards for two classes with the student type space $\Theta = [0.5, 1.5]$, where in the first class student type θ_1 is distributed according to the distribution with pdf $f_1(\theta) = \theta$, for every $\theta \in \Theta$, and in the second class student type θ_2 is distributed uniformly over Θ; and the functional forms are $V(x) = x$, $C(x, \theta) = x^2/(2\theta)$. The middle diagram of Figure 1 shows the optimal effort allocations (the dashed line for the first class and the dotted line for the second class), and the bottom diagram shows the optimal grade allocations for the two classes, respectively.

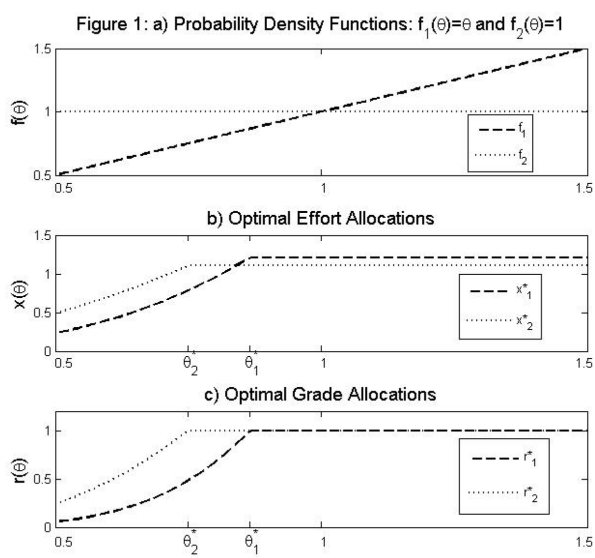

As we can see, both teachers pool some types of the most efficient student types. The teacher of the first class, however, offers the highest grade of 1 to fewer student types but against a higher cost (effort), while the teacher of the second class optimally chooses to be more lenient. This happens because the second teacher focuses on less able but more numerous students and attempts to extract more effort from them leaving more able students with high so-called information rent instead. The incentive compatibility constraint dictates *vice versa* that the teacher of the more able class of students offers a steeper grade-effort schedule in her attempt to extract more effort from more talented students.

In this light, it should not be surprising from the analysis perspective above that, given differences in student abilities among different classes, some teachers turn out to be more lenient than others. But, as we argue, this can be an outcome of the optimal design of grading standards, and not necessarily an outcome of some teachers' rent-seeking behavior, as sometimes is suggested (e.g., Johnson (1997)). In particular, our model predicts that high grades should be more easily attainable in classes with less able students, resulting in a mismatch and low correlation between students' grades and their abilities. For instance, if the population of mathematics students contains more talented people than, say, the population of economics students, then we should observe stricter grading standards applied in mathematics classes (as ample empirical evidence shows to be the case, which is explored in the following section). Furthermore, we could extend our argument to provide additional insight on grade inflation at universities (for an extensive study of this problem, see Johnson (2003)). With the number of students enrolled increasing (say, because of more accessible entry to universities) the overall distribution for students' abilities may become more skewed to the end of less able students. As a result, teachers may optimally respond by lowering grading standards (in order to maximize average effort).

5. Empirical Evidence

The theoretical predictions of the model, in particular that in Proposition 1.3, look to be empirically testable, because the data needed for this purpose, such as student grades and their ability proxy (like their performance on university entry exams or Scholastic Aptitude Test [SAT] scores), should be available at any university. Then, we would need, roughly speaking, to compare grading patterns for classes with different student ability distributions. However, and not surprisingly, there have been a number of empirical studies of the kind in the special literature of educational measurement (e.g., in academic periodicals such as the *Journal of Educational Measurement* or *Educational and Psychological Measurement*). Most importantly, those studies without exception report results that are fully in line with the model's predictions: fields with lower ability students studied as compared with those with higher ability students employ less stringent grading criteria. Even though many of those studies are fairly comprehensive in empirical matters, they lack any rigorous theoretical explanation for this phenomenon, their explanations hinging mainly on intuition or reference to similar phenomena from the adaptation-level theory in psychological literature. In what follows, we attempt to review in detail some of the empirical studies comparing grading standards over time and in different fields, and to show that our model proves helpful in explaining the empirical evidence obseved.

Aiken (1963) is one of the first empirical studies that suggest that grading behavior is dictated by the quality of students in the current class and not by some absolute invariant standards. Aiken (1963) presents time-series evidence from the Woman's College of the University of North Carolina that could imply that, with more able students in a class (as measured by their SAT scores and high-school rankings), teachers tend to apply more stringent grading standards. As for the theoretical explanation for this finding, the study only briefly mentions that it conforms with the adaptation-level theory or central tendency phenomenon, which basically concerns the tendency of supervisors to evaluate the performance of people supervised in relative terms rather than in absolute ones.

A much more comprehensive study Goldman &Widawski (1976) first notes the weaknesses of previous studies on grading patterns because of their using the total grade point average (GPA) as the criterion of grading standards. As they rightly argue, GPAs are not perfectly comparable either over time or among individual students because of the possibly different composition of courses included to compute grade averages. To remedy that, Goldman &Widawski (1976) employ a between-subjects design aimed at making grade comparisons more effective. They compute an index of grading standards using pairwise comparisons of grades in 17 major fields at the University of California, Riverside, from a random sample of 475 students. In particular, they perform the comparison of grading standards in one class (say, psychology) against those in another class (say, biology) by computing the difference in average grades of only those students who took both classes. After obtaining differentials in grading standards between any two classes (from the 17 classes available in their study), they construct an index of grading standards for each class, which is an average of all the differentials between that particular class and the rest of the classes. Finally, they correlate the computed indices of grading standards with the average scores on the verbal and mathematical portions of the SAT test and high-school GPAs (*i.e.*, student ability proxies) of all the students majoring in those 17 classes. The main empirical finding in Goldman &Widawski (1976) is that the constructed index of grading standards correlates highly in a negative direction with student ability proxies. In other words, they conclude that professors in a field containing more able students tend to grade more stringently than do professors in fields with lower ability students. As a result, they find that the past performance and abilities of students account for only slightly more than 50 percent of the variance in grades, and suggest introducing some grade adjustment mechanism to make grades more informative of students' true abilities. Again, in giving an explanation for the empirical results obtained, they restrict their argument simply by making a reference to the adaptation-level theory that people are judged in comparison to their peers.

A similar study Goldman & Hewitt (1975), which along with presenting the empirical results (which draw the same conclusions about grading behavior as in the studies mentioned above), also provides a more elaborate theoretical explanation for the results obtained. The authors think that the antecedents (e.g., student ability levels, work habits, etc.) and consequences (grading standards) of college grading are inextricably tied together by a personal characteristic of college instructors. This characteristic is the phenomenon of adaptation level, and it is so pervasive among college instructors and perhaps people in general, Goldman & Hewitt (1975) continue, as to be considered an almost inevitable factor in the college grading process. Consequently, through that personal characteristic link, grading standards would be partly determined by the ability level of the student population. However, along the lines of our model developed above, this personal characteristic, as envisaged by Goldman & Hewitt (1975), is not some intrinsic feature of human behavior but rather the outcome of optimal behavior.

A decade later, Strenta & Elliott (1987) replicated the study of Goldman &Widawski (1976) using data from a different institution, Dartmouth College, just to find that the differential grading standards exist in the same magnitude and in roughly the same order. Therefore, Strenta & Elliott (1987) argue that it remains the case that students with higher SAT scores tend to major in fields with more rigorous grading standards, and that factors attracting more talented students result in their being graded harder. (However, we would argue for the reverse direction of causation: since some fields attract more talented students, professors in those fields will grade their students more stringently, which is optimal in order to extract more effort.) As in previous studies, Strenta & Elliott (1987) argue that these differential grading standards serve to attenuate the correlation between the GPAs and SAT scores of the students, and they also show that the correlation increases sizably if GPAs are adjusted by accounting for differences in departmental grading standards. Finally, a similar study conducted at Duke University ((Johnson (2003)) confirmed the conclusions about systematic differences in grading standards from the previous studies.

Concerning the normative side of the differential grading standards discussed, there have been a number of papers proposing grade adjustment mechanisms (see, e.g., Johnson (1997)) in order to make grades more informative of students' actual abilities. Without going into the details of this literature, it is worth noting that, typically, those papers tend to assume that the true reason for differential grading standards lies with some personal features of the instructor (e.g., the adaptation level, unwillingness to spend office hours on dealing with students' complaints about low grades, etc.). Therefore, the proposed grade adjustment mechanisms would attempt to correct for presumed instructor-specific factors failing to recognize the possible endogeneity of those factors, which could lead the mechanism astray from the projected goals.

6. Conclusion

In this chapter, we solve for the optimal contract in an agency problem featuring costless (non-pecuniary) rewards, and apply the results obtained to provide alternative explanations for the compression of ratings and mismatch between students' abilities and grades. We argue that in equilibrium the variation in assigned rewards can be coarser than the underlying distribution for abilities. In particular, with the principal holding reasonable expectations about the overall distribution of agent abilities, setting uniform incentives for some most efficient agents can constitute an optimal contract. Specifically for student grading standards, if the teacher's goal is to induce her students to study as hard as possible (subject to the incentive-compatibility constraint), we should observe higher grades in classes with fewer able students. Significantly, the existing empirical evidence strongly supports the predictions of the model presented in this chapter, lending validity to the modeling technique chosen.

Therefore, the proposed framework could be used potentially as the "microfoundations" of student grading or job performance appraisal to analyze other related problems. For instance, one could explore how to design the performance evaluation process in order to reduce the inefficiencies observed (such as coarse grading/rating outcomes) or, with a dynamic version of the model, one could look into the phenomenon of grade inflation over time. Using the framework of this study, further research could also be done on studying the implications on student effort-grade allocations after the introduction of incentives for teachers or on developing grade-adjustment mechanisms to make the intercomparison of grades between various classes, departments, or schools feasible.

Bibliography

Aiken, Jr., Lewis R. 1963. The Grading Behavior of a College Faculty. *Educational and Psychological Measurement*, 23(2), 319–322.

Akerlof, George A. 1982. Labor Contracts as Partial Gift Exchange. *Quarterly Journal of Economics*, 97(4), 543–569.

Arcidiacono, Peter. 2004. Ability sorting and the returns to college major. *Journal of Econometrics*, 121(1-2), 343–375.

Atkinson, Adele, Burgess, Simon, Croxson, Bronwyn, Gregg, Paul, Propper, Carol, Slater, Helen, & Wilson, Deborah. 2004. *Evaluating the Impact of Performance-related Pay for Teachers in England*. CMPO working paper No. 04/113.

Berg, Joyce, Dickhaut, John, & McCabe, Kevin. 1995. Trust, Reciprocity, and Social History. *Games and Economic Behavior*, 10(1), 122–142.

Besley, Timothy, & Ghatak, Maitreesh. 2008. Status Incentives. *American Economic Review*, 98(2), 206–211.

Bolton, Patrick, & Dewatripont, Mathias. 2004. *Contract Theory*. MIT Press.

Brennan, Geoffrey, & Pettit, Philip. 2004. *The Economy of Esteem: An Essay on Civil and Political Society*. Oxford University Press.

Bénabou, Roland, & Tirole, Jean. 2003. Intrinsic and Extrinsic Motivation. *Review of Economic Studies*, 70(3), 489–520.

Dubey, Pradeep, & Geanakoplos, John. 2005. *Grading in Games of Status: Marking Exams and Setting Wages*. Cowles Foundation Discussion Paper No. 1544.

Falk, Armin, & Kosfeld, Michael. 2006. The Hidden Costs of Control. *American Economic Review*, 96(5), 1611–1630.

Frey, Bruno S., & Osterloh, Margit (eds). 2002. *Successful Management by Motivation: Balancing Intrinsic and Extrinsic Incentives*. Springer.

Fudenberg, Drew, & Tirole, Jean. 1991. *Game Theory*. MIT Press.

Goldman, Roy D., & Hewitt, Barbara Newlin. 1975. Adaptation-Level as an Explanation for Differential Standards in College Grading. *Journal of Educational Measurement*, 12(3), 149–161.

Goldman, Roy D., & Widawski, Mel H. 1976. A Within-Subjects Technique for Comparing College Grading Standards: Implications in the Validity of the Evaluation of College Achievement. *Educational and Psychological Measurement*, 36(2), 355–359.

Guesnerie, Roger, & Laffont, Jean-Jacques. 1984. A complete solution to a class of principal-agent problems with an application to the control of a self-managed firm. *Journal of Public Economics*, 25(3), 329–369.

Johnson, Valen E. 1997. An Alternative to Traditional GPA for Evaluating Student Performance. *Statistical Science*, 12(4), 251–269.

Johnson, Valen E. 2003. *Grade Inflation: A Crisis in College Education*. Springer.

Lavy, Victor. 2002. Evaluating the Effect of Teachers'Group Performance Incentives on Pupil Achievement. *Journal of Political Economy*, 110(6), 1286–1317.

Lazear, Edward. 2003. Teacher Incentives. *Swedish Economic Policy Review*, 10, 179–203.

Levin, Jonathan. 2003. Relational Incentive Contracts. *American Economic Review*, 93(3), 835–857.

MacLeod, W. Bentley. 2003. Optimal Contracting with Subjective Evaluation. *American Economic Review*, 93(1), 216–240.

Mirrlees, James A. 1971. An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies*, 38(2), 175–208.

Moldovanu, Benny, Sela, Aner, & Shi, Xianwen. 2007. Contests for Status. *Journal of Political Economy*, 115(2), 338–363.

Murphy, Kevin R., & Cleveland, Jeanette. 1995. *Understanding Performance Appraisal: Social, Organizational, and Goal-based Perspectives*. SAGE.

Ostrovsky, Michael, & Schwarz, Michael A. 2003. *Equilibrium Information Disclosure: Grade Inflation and Unraveling*. Harvard Institute of Economic Research Working Paper No. 1996.

Prendergast, Canice. 1999. The Provision of Incentives in Firms. *Journal of Economic Literature*, 37(1), 7–63.

Shaked, Moshe, & Shanthikumar, J. George. 1994. *Stochastic Orders and Their Applications*. Acedemic Press, Inc.

Sliwka, Dirk. 2007. Trust as a Signal of a Social Norm and the Hidden Costs of Incentive Schemes. *American Economic Review*, 97(3), 999–1012.

Strenta, A. Christopher, & Elliott, Rogers. 1987. Differential Grading Standards Revisited. *Journal of Educational Measurement*, 24(4), 281–291.

Chapter 2

Manager Accountability and Employee Wage Schedules*

Abstract

The goal of this study is to relate and jointly explain the empirical regularities of the compression of ratings and firm-size effects on wages. We develop a three-tier model of a firm's economic organization, which is centered on the empirical observation that managers face a soft budget constraint when evaluating their employees' performance. The model assumes that in small firms managers are held more accountable for payroll expense incurred than they are in large firms (because of lesser informational asymmetries). Incorporating these features into an otherwise standard agency model, this study predicts differences in pay schedules between small and large firms of the kind that is empirically observed. In particular, we argue that the large-firm wage premium and inverse relationship between wage dispersion and firm size can be the optimal outcomes of the agency problem with unaccountable managers that is studied here. The model also shows that the compression of ratings in job performance appraisals can be an equilibrium outcome.

1. Introduction

This chapter deals with the empirical regularities of the compression of ratings in job performance appraisals (Murphy & Cleveland (1995); Prendergast (1999)) and firm-size effects on employee wage schedules:

*I would like to thank Ailko van der Veen and Karl Wärneryd. Financial support from the Jan Wallander and Tom Hedelius Foundation is gratefully acknowledged.

the large-firm wage premium and inverse relationship between wage dispersion and firm size (Brown & Medoff (1989); Oi & Idson (1999)). We attempt to relate and jointly explain these phenomena with the help of a three-tier agency model of a firm's economic organization, where the tiers are the owner of a firm, the managers, and the employees. The model is developed around the observation from manager and employee compensation practices that managers do not fully internalize the payroll expenses incurred in evaluating their employees' performance (see, e.g., Longenecker et al. (1987)). There is also evidence indicating that in a small firm managers are held more accountable for their payroll expenses than they are in a large firm. This is because of lesser informational asymmetries between the owner of a smaller firm and her managers related to, in the words of Alchian & Demsetz (1972), "metering input productivity and metering rewards."

In the current chapter, we show that incorporating these features of vertical managerial relationships into an otherwise standard agency model produces theoretical predictions that offer a good match with the empirical stylized facts on wage patterns. Therefore, we argue that manager accountability (or rather the lack of it) can be a cause of the compression of ratings and firm-size effects on wage schedules. Our model is also consistent with the empirical evidence from the financial literature on small firms' higher stock returns and, supposedly, their higher profitability, see Banz (1981) and Fama & French (1992). We use the latter evidence to distinguish our explanation for the wage patterns observed in empirical studies from other alternative explanations.

Differently from Alchian & Demsetz (1972), here we study the problem of metering employees' inputs and rewards from the perspective of an owner-manager relationship. The idea is that the interests of an owner and a manager with respect to employee compensation may actually diverge. In practice, the monitoring and appraisal of employees' individual effort levels are done by managers—low- and middle-ranking managers in the corporate world—who are not residual claimants; nor can their pay be perfectly related to the firm's profits.[1] As an alternative to profit-sharing rules, the owner of a firm offers her managers a compensation scheme, which depends on their accomplishing individual objectives (so-called management by objectives), or on their performance evaluation adjusted for the firm's overall profitability (see Bruns & McKinnon (1992); Milkovich & Wigdor (1991)). In addition, the owner imposes objectives on managers that are to be achieved within a certain framework—performance appraisal standards on how to reward (or monitor) the performance of their employees. This is done in order to prevent managers from incurring great payroll expense when maximizing their compensation.

But, as is suggested from the incidence of the compression of ratings and other evidence of managers' lack of accountability, which will be discussed later, the existing practice of managerial compensation seems to have inefficiencies. By directly rewarding for managers' accomplishments and controlling for the costs they incur, the owner of a firm may fail to perfectly align managers' incentives with the firm's profit maximization. When designing a compensation scheme for her managers, the owner draws on her own knowledge about the workings of the manager's job—

[1] According to surveys by the US Bureau of Labor Statistics, in 1999 only 1.4 percent of US business establishments granted stock options to their nonexecutive employees. It is suggested that the reason for this is the limited incentive effects associated with stock options, see Besanko et al. (2007, p. 499). Moreover, among those firms that do offer stock options to all their employees, an incentive-based explanation for it is rejected, see Oyer & Schaefer (2005).

its contribution to the firm's profits and share of total costs—which may nonetheless be inaccurate. It allows managers to bargain for a compensation scheme more advantageous to them than to the firm.[2] As a result, in pursuing their goals, managers are likely to enjoy some leeway with respect to the payroll costs resulting from their evaluations of employees' performance.

According to surveys of business organizations (for a review, see Murphy & Cleveland (1995, p. 4)), most public and private companies in the US—more precisely, between 74% and 89% of those surveyed, with large companies somewhat more prevalent—practice a formal job performance appraisal system. Performance appraisals are mainly used for employee salary administration purposes (including managers). The usual way performance appraisals work is that a supervisor (manager) is asked to rate various aspects of his or her subordinates' performance on a prespecified scale, with their pay made proportional to the supervisor's rating.

However, the practice of performance appraisals has fallen short of the expectations about their utility. The distribution of ratings given typically exhibits a shallow differentiation of good from bad performance, arguably, leading to inefficient performance outcomes in the end. In the psychological literature, this has been labeled the "compression of ratings" phenomenon with, it is suggested, a "leniency bias" behind it (for comprehensive reviews, see Landy & Farr (1983) and Murphy & Cleveland (1995); for a case study, see Murphy (1992)). Economists see this phenomenon as one of the causes of the dominance of fixed wages in company payrolls (Prendergast (1999)), and, accordingly, raise the ques-

[2]See Milkovich & Wigdor (1991) for more on managerial compensation practices and managers' bargaining advantages.

tion of why the underlying economic incentives behind job performance appraisal do not work as planned (for a comprehensive discussion, see Bruns (1992)).

Another systematic pattern of employee pay practice is the large-firm wage premium, i.e., that large firms on average pay higher wages, *ceteris paribus*. It remains unclear (especially, in times of a slack labor market) why large firms pay higher wages when employees are ready to work for lower pay. This has been widely observed across different countries and industries: It seems that firm size matters. A number of explanations have been offered, some of which are discussed in the next section, but more research on this question seems called for (see Brown & Medoff (1989) and Oi & Idson (1999) for reviews). At the same time, despite paying on average lower wages, small firms reward their employees' abilities and acquired skills, such as experience, at a greater rate than do large firms (see Garen (1985); Evans & Leighton (1989)). Generally, an inverse relationship has been observed between wage dispersion and firm size (Stigler (1962)). All this hints at the possibility that economic incentives for employees are possibly better designed in small firms.

We approach these issues through a three-tier agency model of a firm's economic organization. In the firm, output is produced only by employees, who have different but privately observed ability levels. Managers are hired to supervise the employees' performance, and the owner supervises the managers. It is a two-stage agency problem. In the second stage, a manager provides incentives for and evaluates the performance (effort levels) of his employees in order to maximize his own compensation. In the first stage, the owner designs a compensation scheme for her managers. It rewards them for employee efforts extracted and controls payroll costs incurred through performance appraisal standards, which are assumed to have inefficiencies proportional to the size of the firm. Performance appraisal standards are modeled to take the form of an upper bound on employee pay that is at a manager's disposal, and inefficiencies take the form of a soft budget constraint that a manager faces when evaluating his employees' performance. The crucial feature of our model is that managers do not fully internalize the payroll cost resulting from their ratings, and it is the owner who bears the unaccounted part of this cost.

It will be shown that the softer a budget constraint is (i.e., the less accountable managers are), the more the owner limits her managers' discretion by lowering an upper bound on employee rewards. In response, a manager designs a pay-for-effort allocation schedule for his employees that pools some of the most efficient employee types subject to the highest reward. As a result, it leads to a coarse distribution of rewards (or ratings), which, we argue, could be behind the compression of ratings phenomenon. Furthermore, we show that the length of the pooling-equilibrium interval of employee types varies inversely with the manager's degree of accountability. Therefore, with the problem of a soft budget constraint more aggravated in larger firms, the model predicts that the average wage will increase but wage dispersion will decrease with firm size. At the same time, however, we find that small firms are more profitable than large ones.

The remainder of this chapter is organized as follows. Section 2 discusses related literature and provides motivation for the model, which is developed and solved in Section 3. Section 4 discusses the equilibrium properties obtained with respect to firm size, and Section 5 relates them to the existing literature. The last section concludes the study.

2. Background and Motivation

Industrial and organizational psychologists have traditionally viewed job performance appraisal and its consequences—the compression of ratings, in particular—as a measurement problem. They distinguish the three most frequently encountered measurement biases: the "halo effect", a tendency to rate the same on all dimensions, "centrality bias", an over-reliance on the middle of the rating scale, and "leniency bias", a tendency to give extreme ratings (which is the main focus of this study). Psychologists found no evidence that personal characteristics of raters or ratees have any explanatory power for the systematic patterns observed in performance appraisal, see Landy & Farr (1980). Instead, psychologists have come to think that performance appraisal cannot be adequately understood outside its organizational context, which is a major determinant of a rater's goal-oriented rating behavior, see Murphy & Cleveland (1995). In economic terms, this implies that job performance appraisal is an agency problem but, possibly, with some intrinsic inefficiencies.

The related psychological literature offers strong support for the modeling assumptions that are made in this work. As regards the assumption that managers may be facing a soft budget constraint, there is enough evidence to argue that managers do enjoy leeway when evaluating their subordinates' performance. For instance, Longenecker et al. (1987) show that managers manipulate the whole appraisal process to their own advantage; or they are not held accountable for their lenient appraisals, see Mero & Motowidlo (1995).

Furthermore, it has also been observed that employee performance appraisal standards vary greatly across different organizations, and one of the factors behind those differences is organization size. Landy & Farr (1983, p. 104–105) describe how many smaller organizations hold supervisor conferences to evaluate and, accordingly, reward the performance of each employee in turn, which is not feasible in large organizations. Murphy & Cleveland (1995, p. 355) see decentralization as a way to increase the efficiency of performance appraisal practice in organizations, because it would allow performance appraisal standards for every functional unit to be tailored more accurately. There is also experimental evidence showing that the degree of task interdependence among group members inversely affects the differentiation of good from bad performance, see Liden & Mitchell (1983).

The related literature in economics is too broad to be comprehensively discussed here in any greater detail; instead, the discussion is restricted to a few selected papers only. When it comes to explaining the compression of ratings phenomenon, this study is most closely related to principal-agent models with subjective evaluation, see MacLeod (2003) or Levin (2003). The distinctive feature of these models is that effort levels are non-contractible and are rewarded according to the principal's subjective evaluation. Under the threat of a conflict, the principal may find it futile to differentiate rewards solely on her subjective performance evaluations, when there is a great likelihood that the agent will think differently of his own performance. Unlike this strand of literature, the current study allows for verifiable and contractible employee effort levels, and our results hinge on asymmetric information and contract incompleteness between the owner and managers. In our model, the compression of ratings is the outcome of the optimal (effort-maximizing) incentive scheme offered by a manager to his employees.

As for the firm size effects—the large-firm wage premium and higher wage dispersion in smaller firms—there are numerous empirical and theoretical papers on the issue (for reviews, see Brown & Medoff (1989) and Oi & Idson (1999)). For instance, Stigler (1962) shows that in large companies the dispersion of wages is lower than that observed in small companies. He attributes this finding to the fact that the owner of a small company can better judge the quality of her employees' performance. Along the lines of "Stigler's conjecture", Garen (1985) develops a model based on the assumption that employees' monitoring and evaluation costs rise with firm size because of larger imperfections in acquiring information. He provides empirical evidence supporting his model's prediction that larger firms pay a smaller return to measured ability, but have a larger intercept in their wage equations, which also found support in Evans & Leighton (1989). (In our model, as will be shown, the same differences in pay schedules arise from the fact that the owner of a smaller firm can more accurately relate her managers' pay to the firm's profits, and *vice versa*.)

Regarding the large-firm wage premium, which means larger firms pay on average higher wages, everything else equal, one of the most frequently accepted explanations for this phenomenon is given in Idson & Oi (1999).[3] They argue that the shape of wage-size relation depends on worker preferences, working conditions, and, most importantly, technology. The idea is that large firms, exploiting their returns to scale, can invest in more productive labor tools. Idson & Oi (1999) argue that

[3]See also Bulow & Summers (1986) and Weiss & Landau (1984) for alternative explanations.

the systematic differences observed in wage schedules can arise because in larger firms employees, being better equipped, are more productive, as measured by output per hour, and, therefore, they command higher wages.[4] But this explanation fails to explain why there is a lower wage dispersion in larger firms or why large firms are less profitable (especially if it is argued they are more productive), as the financial empirical evidence indicates to be the case (Banz (1981); Fama & French (1992)). Here, we present a different interpretation of the empirical findings of Idson & Oi (1999). We argue that in a larger firm employees exert on average higher effort levels (and get paid more) because of more lenient incentive schemes set by their less accountable managers, which is, nonetheless, not in the best interest of the firm.

3. Model

Framework

Consider a profit-maximizing firm, where production is split among different production divisions. Every division consists of one employee and one manager, and it produces an input to the firm's final product using only the employee's labor services. A division manager's job is to extract effort from the employee, which the manager does by designing and implementing a pay-for-effort incentive scheme for his employee. It requires that the manager possesses specific knowledge about the division's production process which is not available to the owner of the firm to the full extent. At the same time, the owner designs a compensation scheme for her division managers that would maximize the firm's profits subject

[4]See also Hamermesh (1980) for a related argument.

to her informational constraints as defined later. We assume that the degree of asymmetry in information between managers and the owner varies inversely to the size of the firm (as measured, say, by the number of production divisions in the firm).

Next, consider a representative division of the firm, the workings and contribution to the firm's profits of which are similar to those of other divisions. An effort e, exerted by the division employee, results in the production of the division's output $V(e)$, where V is a production function with the properties $V_e > 0$ and $V_{ee} \leq 0$. It costs the employee a disutility of $C(e, \theta)$, where C is an effort cost function, and the parameter θ is the employee's privately known productivity level distributed on the finite support $[\underline{\theta}, \overline{\theta}]$ according to a twice differentiable common prior distribution F with the probability density function f ($f > 0$) satisfying the non-decreasing monotone hazard rate condition. For analytical convenience, assume that the effort cost function is separable in effort e and ability θ with its functional form $C(e, \theta) = g(e)/\theta$, where g is a strictly convex twice differentiable function. If offered a pay r in return for an effort e, the employee of ability θ can enjoy a net utility of

$$U^A(r, e, \theta) = r - C(e, \theta), \tag{2.1}$$

which needs to be at least non-negative for the employee to accept the offer (r, e).

The division manager designs pay-for-effort allocations for the employee to choose from, which the manager does trying to maximize his own reward coming from the compensation scheme offered by the owner. For the reasons previously explained, the owner offers the manager a compensation scheme that directly rewards the manager for his accomplishments by granting a fraction $\alpha \in (0,1)$ of the output V and that controls his payroll cost r. Because of asymmetric information, the owner can make the manager internalize only an $\alpha\lambda$ fraction of the payroll cost r, where the parameter $\lambda \in (0,1]$ is inversely related to the size of the firm, which is exogenously given. Furthermore, to alleviate the problem of the manager's having a soft budget constraint, the owner can impose an upper bound $\bar{r}$ on employee rewards that the manager can offer to his employee.[5]

Hence, the pay-for-effort allocation (r,e) implemented by the manager results in his reward of

$$U^M(r,e) = \alpha V(e) - \alpha\lambda r, \tag{2.2}$$

and the net profit accrued to the firm is equal to

$$\pi(r,e) = (1-\alpha)V(e) - (1-\alpha\lambda)r, \tag{2.3}$$

which is the output V less the employee payroll cost r and less the manager compensation.

The exogenously given parameter λ, observed by all the parties, can be interpreted as a characteristic of firm size, with a larger firm having a lower value of λ. The parameter λ is assumed to capture all the differences in information between the manager and the owner. A smaller value of λ implies a larger degree of asymmetric information, which translates into a softer budget constraint for the manager.

[5]An upper bound on employee rewards comes naturally from managers' practice of evaluating their employees' performance on a finite rating scale, which, together with ratings' monetary values, is set from above. In the model, setting an upper bound is all that the owner does when designing employee performance appraisal standards, other aspects of which are ignored for simplicity.

Game

Suppose that every division of the firm functions as the following two-stage game between the owner, manager, and employee, who are all rational utility maximizers. Given the framework described above, in the first stage the owner sets a compensation scheme for the manager. In the second stage, the manager, upon observing his own compensation scheme, designs a set of pay-for-effort allocations for his employee to choose from. The employee chooses the allocation that maximizes his utility, and after its implementation payoffs to all the parties follow.

More specifically, assuming that the manager's reward fraction α is exogenously determined, say, by the outside labor market for managers, the owner's action concerning the manager's compensation scheme is to set an upper bound $\bar{r} \in \mathbb{R}_+$ on employee rewards. The manager designs direct pay-for-effort allocations $x(\theta) = \{r(\theta), e(\theta)\}$ for every employee ability type $\theta \in [\underline{\theta}, \overline{\theta}]$, where the reward and effort allocations are functions defined, respectively, as $r : [\underline{\theta}, \overline{\theta}] \to [0, \bar{r}]$ and $e : [\underline{\theta}, \overline{\theta}] \to \mathbb{R}_+$. The employee of ability θ announces a type $\hat{\theta}$ from the type space $[\underline{\theta}, \overline{\theta}]$, which leads to the implementation of the allocation $(e(\hat{\theta}), r(\hat{\theta}))$. The resultant utility levels follow from (2.1) for the employee, from (2.2) for the manager, and, respectively, from (2.3) for the owner. All the utility levels are assumed to satisfy the von Neumann-Morgenstern axioms. Finally, to solve the game, we use the concept of Bayesian-Nash Equilibrium, which, in our setting, is a strategy profile $\{\bar{r}^*, x^*, \hat{\theta}^*(\theta)\}$ such that each type of every player plays her best reply given the strategies of the others.

Next, we solve the model by backward induction. Then, we discuss the properties of the solution obtained with respect to the parameter λ

(firm size). It will be shown that the smaller the value the parameter λ takes, the more the owner limits the manager's discretion by imposing a lower upper bound on employee rewards. It eventually leads to the manager's designing a flatter employee pay schedule with the ensuing compression of rewards (ratings) and firm-size effects of the type that are documented in the empirical literature.

The manager's problem, Stage 2

The manager faces a hidden information problem since the employee ability type θ is privately known. Given his own compensation scheme $(\alpha, \bar{r})$, with respect to direct pay-for-effort allocations $\{r(\theta), e(\theta)\}_{\theta \in [\underline{\theta}, \overline{\theta}]}$ the manager maximizes his expected utility

$$\int_{\underline{\theta}}^{\overline{\theta}} \alpha(V(e(\theta) - \lambda r(\theta))dF(\theta) \tag{2.4}$$

subject to

$$r(\theta) - C(e(\theta), \theta) \geq 0, \tag{2.5}$$

$$r(\theta) - C(e(\theta), \theta) \geq r(\hat{\theta}) - C(e(\hat{\theta}), \theta), \text{ and} \tag{2.6}$$

$$0 \leq r(\theta) \leq \bar{r}, \text{ for all } \theta \text{ and } \hat{\theta} \text{ in } [\underline{\theta}, \overline{\theta}]. \tag{2.7}$$

The first two constraints are the employee's participation and incentive compatibility constraints, respectively; and the last one is a constraint on employee rewards imposed by the owner in the first stage.

The solution to the manager's utility maximization problem without the upper-bound constraint, eq. (2.4)–(2.6), can be found by the well-established methods, following Mirrlees (1971). It is characterized by the functional equation

$$V_e(e(\theta)) - \lambda[C_e(e(\theta), \theta) - \frac{1 - F(\theta)}{f(\theta)} C_{e\theta}(e(\theta), \theta)] = 0. \qquad (2.8)$$

Let the effort function $e^u : [\underline{\theta}, \overline{\theta}] \to \mathbb{R}_+$ solve the above equation; then, the corresponding pay levels $r^u(\theta)$ are found from

$$r^u(\theta) = C(e^u(\theta), \theta) - \int_{\underline{\theta}}^{\theta} C_\theta(e^u(\tilde{\theta}), \tilde{\theta}) d\tilde{\theta}, \quad \text{for } \theta \in [\underline{\theta}, \overline{\theta}]. \qquad (2.9)$$

The assumed non-decreasing monotone hazard rate condition ensures that the effort schedule $e^u(\theta)$ is increasing in ability type θ and that the "no distortion at the top" property holds. The solution to the reduced problem $x^u(\theta) = \{r^u(\theta), e^u(\theta)\}_{\theta \in [\underline{\theta}, \overline{\theta}]}$ also constitutes the solution to the full problem if the left-out constraint is not binding, i.e., if $r^u(\overline{\theta}) \leq \bar{r}$.

If constraint (2.7) is binding, in order to solve the manager's problem we need to modify the solution method, which we do in the Appendix. But then, as our solution to the full problem will show, the "no distortion at the top" property is no longer preserved for the optimal pay-for-effort allocations. In particular, provided that the manager does not find it optimal to exclude some of the least efficient employee types—which is assumed to be the case throughout this chapter, implying that the mass of inefficient types is large enough—we show that the manager should offer a uniform pay-for-effort allocation to some of the most efficient types. The pooling of employee types takes place because the manager cannot design distinct incentive-compatible allocations for all the employee types if constrained in rewards. More precisely, since the manager cannot elicit the first-best effort level from the most efficient type due to the pay cap imposed, he has to revert to an effort level that is lower than the first-best one and make it available to a pool of employee types.

With a reference to the Appendix for the details of solving the full problem (2.4)-(2.7), its solution $x^*(\theta) = \{r^*(\theta), e^*(\theta)\}$ for $\theta \in [\underline{\theta}, \overline{\theta}]$ is given in Proposition 2.1 below.

Proposition 2.1 *Let $x^u(\theta) = \{r^u(\theta), e^u(\theta)\}$ for $\theta \in [\underline{\theta}, \overline{\theta}]$ be defined as in eq. (2.8) and (2.9). The solution $x^*(\theta) = \{r^*(\theta), e^*(\theta)\}_{\theta \in [\underline{\theta}, \overline{\theta}]}$ to the manager's problem (2.4)–(2.7) is as follows*

- *if $r^u(\overline{\theta}) \leq \bar{r}$, where $\bar{r}$ is the owner's imposed upper bound reward, then $x^*(\theta) = x^u(\theta)$;*

- *otherwise, for employee ability types θ in $[\underline{\theta}, \theta^p)$ the optimal pay-for-effort allocations are $\{r^*(\theta), e^*(\theta)\}$ and for types θ in $[\theta^p, \overline{\theta}]$ — $\{\bar{r}, e^*(\theta^p)\}$, where the starting point θ^p of the pooling interval $[\theta^p, \overline{\theta}]$ and the effort levels $e^*(\theta)$ for $\theta \in [\underline{\theta}, \theta^p]$ are jointly determined by*

$$\frac{1 - F(\theta^p)}{f(\theta^p)} = \frac{[V_e(e^*(\theta^p)) - \lambda C_e(e^*(\theta^p), \theta^p)]C_e(e^*(\theta^p))}{V_e(e^*(\theta^p))(-C_{e\theta}(e^*(\theta^p), \theta^p))} \tag{2.10}$$

$$C(e^*(\theta^p), \theta^p) - \int_{\underline{\theta}}^{\theta^p} C_\theta(e^*(\theta), \theta) d\theta = \bar{r}, \tag{2.11}$$

and

$$\begin{aligned} &[V_e(e^*(\theta)) - \lambda C_e(e^*(\theta), \theta)] + \lambda \frac{(1 - F(\theta))}{f(\theta)} C_{e\theta}(e^*(\theta), \theta) + \\ &+ \frac{(1 - F(\theta^p)}{f(\theta)} \frac{V_e(e^*(\theta^p)) - \lambda C_e(e^*(\theta^p), \theta^p)}{C_e(e^*(\theta^p), \theta^p)} C_{e\theta}(e^*(\theta), \theta) = 0. \end{aligned} \tag{2.12}$$

The pay levels $r^*(\theta)$ for $\theta \in [\underline{\theta}, \theta^p)$ are equal to

$$r^*(\theta) = C(e^*(\theta), \theta) - \int_{\underline{\theta}}^{\theta} C_\theta(e^*(\tilde{\theta}), \tilde{\theta}) d\tilde{\theta}. \tag{2.13}$$

Proof. See the Appendix. ■

The owner's problem, Stage 1

The owner's expected residual profit resulting from the manager's designed incentive scheme $x = \{r, e\}$ is given by

$$\Pi(x) = \int_{\underline{\theta}}^{\bar{\theta}} (1-\alpha)V(e(\theta)) - (1-\alpha\lambda)r(\theta)dF(\theta). \tag{2.14}$$

The owner's problem is to maximize (2.14) when designing a compensation package for her manager, i.e., when imposing an upper bound $\bar{r}$ on employee rewards. Since the rational owner can discern for herself the optimal employee incentive scheme x^*, designed by the manager in the second stage for a given upper bound $\bar{r}$, the owner's expected profit can be expressed solely as a function of her action $\bar{r}$.

Denote the expected profit function by $\tilde{\Pi}$, which is a mapping of an upper bound $\bar{r} \in \mathbb{R}_+$ into the profit $\Pi(x^*)$ as in (2.14), where x^* is the optimal pay-for-effort allocation schedule from Proposition 2.1 for a given $\bar{r}$. The function $\tilde{\Pi}$ is then defined by

$$\begin{aligned}\tilde{\Pi}(\bar{r}) = & \int_{\underline{\theta}}^{\tilde{\theta}^p(\bar{r})} (1-\alpha)V(e^*(\theta)) - (1-\alpha\lambda)r^*(\theta)dF(\theta) + \\ & + (1 - F(\tilde{\theta}^p(\bar{r}))[(1-\alpha)V(e(\tilde{\theta}^p(\bar{r}))) - (1-\alpha\lambda)\bar{r}],\end{aligned} \tag{2.15}$$

where $\tilde{\theta}^p$ is a mapping of an upper bound $\bar{r}$ into the starting point θ^p of the pooling interval $[\theta^p, \overline{\theta}]$, as defined in Proposition 2.1; $e^*(\theta)$ and $r^*(\theta)$ for $\theta \in [\underline{\theta}, \theta^p]$ are the optimal pay-for-effort allocations from Proposition 2.1.

The owner finds the optimal upper bound $\bar{r}$ maximizing (2.15) from the first-order condition of (2.15) with respect to $\bar{r}$, which is

$$V_e(e^*(\tilde{\theta}^p(\bar{r})))e^*_\theta(\tilde{\theta}^p(\bar{r}))\tilde{\theta}^p_r(\bar{r}) = \frac{1-\alpha\lambda}{1-\alpha}. \tag{2.16}$$

Differentiating (2.11) in Proposition 2.1 with respect to $\bar{r}$ gives

$$e^*_\theta(\tilde{\theta}^p(\bar{r}))\tilde{\theta}^p_r(\bar{r}) = \frac{1}{C_e(e^*(\tilde{\theta}^p(\bar{r}),\tilde{\theta}^p(\bar{r}))},$$

and plugging it into (2.16) renders the optimality condition for an upper bound $\bar{r}$:

$$\frac{V_e(e^*(\tilde{\theta}^p(\bar{r})))}{C_e(e^*(\tilde{\theta}^p(\bar{r})),\tilde{\theta}^p(\bar{r}))} = \frac{1-\alpha\lambda}{1-\alpha}. \tag{2.17}$$

Without going into any detail, the second-order condition is assumed to be satisfied (although ensuring this may require adding some additional mild assumptions on the functional forms of the production function V and effort cost function C or, alternatively, restricting parameter values).

Condition (2.17) has a natural interpretation. It requires setting an upper bound $\bar{r}$ so that in the optimum it equates the owner's marginal revenue $(1-\alpha)V_e(e^*(\tilde{\theta}^p(\bar{r})))$ from the highest effort level $e^*(\tilde{\theta}^p(\bar{r}))$ contracted by the manager with the corresponding marginal cost of $(1-\alpha\lambda)C_e(e^*(\tilde{\theta}^p(\bar{r})),\tilde{\theta}^p(\bar{r}))$. When $\lambda < 1$ (i.e., when the manager does not internalize his payroll expense incurred in full), the right-hand side of (2.17) is greater than one, implying that it is not in the owner's interest to have any first-best (socially optimal) effort level implemented (where the first-best level is determined from $V_e(e^{FB}(\theta)) = C_e(e^{FB}(\theta),\theta)$ for any θ). Therefore, if $\lambda < 1$, the owner imposes a binding upper-bound reward $\bar{r}$ on the manager in order to reduce the employee efforts he elicits below the socially optimal levels.

Equilibrium

Having established the conditions of the manager's and the owner's optimal play—Proposition 2.1 and eq. (2.17), respectively—we can solve for the Bayesian-Nash equilibrium of the game. In our derivations below, we make use of the assumption that the employee's effort cost function $C(e, \theta)$ is separable in effort and ability, i.e., $C(e, \theta) = g(e)/\theta$, which, of course, has no qualitative impact on the properties of the equilibrium obtained.

Plugging (2.17) into (2.10) from Proposition 2.1 renders the condition for the starting point θ^p of the pooling interval $\left[\theta^p, \overline{\theta}\right]$:

$$\frac{1-F(\theta^p)}{f(\theta^p)} = \theta^p \frac{1-\lambda}{1-\alpha\lambda}. \tag{2.18}$$

Since there may be no θ from $\left[\underline{\theta}, \overline{\theta}\right]$ satisfying the above condition, then the starting point θ^p of the pooling interval is more generally defined by

$$\theta^p = \min(\theta : \frac{1-F(\theta)}{f(\theta)} - \theta\frac{1-\lambda}{1-\alpha\lambda} \leq 0,\ \underline{\theta} \leq \theta \leq \overline{\theta}). \tag{2.19}$$

Similarly, plugging (2.17) into (2.12) from Proposition 2.1 renders the condition for the optimal effort levels $e^*(\theta)$ for ability types θ in $[\underline{\theta}, \theta^p]$:

$$\begin{aligned} &[V_e(e^*(\theta)) - \lambda C_e(e^*(\theta), \theta)] + C_{e\theta}(e^*(\theta), \theta)\times \\ &\times[\lambda\frac{(1-F(\theta))}{f(\theta)} + \frac{(1-F(\theta^p)}{f(\theta)}\frac{1-\lambda}{1-\alpha}] = 0. \end{aligned} \tag{2.20}$$

It is straightforward to see that the effort function e^* is continuous in employee type θ. The optimal pay schedule $r^*(\theta)$ for θ in $[\underline{\theta}, \theta^p)$ is given by (2.13), and it is also continuous in θ. Finally, the owner determines the optimal upper bound $\bar{r}^*$, ensuring condition (2.17) holds, from

$$\bar{r}^* = C(e^*(\theta^p), \theta^p) - \int_{\underline{\theta}}^{\theta^p} C_\theta(e^*(\theta), \theta) d\theta. \tag{2.21}$$

Proposition 2.2 below summarizes the above results and characterizes the equilibrium of the game studied above.

Proposition 2.2 *The Bayesian-Nash equilibrium of the game in question is the strategy profile* $\{\bar{r}^*, x^*, \hat{\theta}^*(\theta)\}$ *for* $\theta \in [\underline{\theta}, \overline{\theta}]$, *where*

- *the manager's optimal strategy* $x^* = (r^*, e^*)$ *is defined by:*
  - *for employee ability types* θ *in* $[\underline{\theta}, \theta^p)$, *with* θ^p *as in (2.19), the optimal allocation is* $x^*(\theta) = (r^*(\theta), e^*(\theta))$, *where the optimal effort and reward levels* $e^*(\theta)$ *and* $r^*(\theta)$ *are defined by (2.20) and (2.13), respectively;*
  - *for ability types* θ *in* $[\theta^p, \overline{\theta}]$, $x^*(\theta) = (\bar{r}^*, e^*(\theta^p))$, *where the effort* $e^*(\theta^p)$ *and reward* $\bar{r}^*$ *are found from (2.20) and (2.21), respectively;*
- *the owner's optimal strategy* $\bar{r}^*$ *is defined by (2.21);*
- *the employee of ability* θ *in* $[\underline{\theta}, \theta^p)$ *announces* $\hat{\theta}^*(\theta) = \theta$, *and of ability* θ *in* $[\theta^p, \overline{\theta}]$ — $\hat{\theta}^*(\theta) = \theta^p$.

4. Equilibrium properties

Below, we discuss the properties of the equilibrium obtained in their relationship to the parameter λ (firm size).

Pooling at the top

As it follows from Proposition 2.2 and the derivations preceding it, for the parameter λ values less than 1, the incentive scheme offered by the manager features a uniform pay-for-effort allocation for employee types θ from the non-empty interval $[\theta^p, \overline{\theta}]$ (if $\lambda < 1$, then $\theta^p < \overline{\theta}$ from (2.18)). The underlying reason for the existence of the pooling equilibrium is the misalignment of the owner's and manager's interests. When the manager is not fully accountable for the payroll costs incurred, the owner, who then bears a disproportionately larger share of costs, attempts to limit the manager's discretion by imposing a binding upper bound on employee rewards. Consequently, in response to the upper bound constraint imposed the manager optimally pools employee types and makes them subject to the highest available reward.

Moreover, the lower the value the parameter λ takes, the more the manager extends the pooling-equilibrium interval. Supposing that the starting point θ^p from (2.19) is in $(\underline{\theta}, \overline{\theta})$, it follows from (2.18) that the internal derivative $d\theta^p/d\lambda$ is positive:

$$\frac{d\theta^p}{d\lambda} = -\frac{\theta^p\left(\frac{1-\alpha}{(1-\alpha\lambda)^2}\right)}{\frac{d}{d\theta^p}\left(\frac{1-F(\theta^p)}{f(\theta^p)}\right) - \left(\frac{1-\lambda}{1-\alpha\lambda}\right)} > 0, \tag{2.22}$$

where in the denominator the derivative of the inverse hazard rate is negative (due to the assumption).

Proposition 2.3 summarizes the above findings.

Proposition 2.3 *With $\lambda < 1$, the employee types θ in $[\theta^p, \overline{\theta}]$, where $\theta^p < \overline{\theta}$ due to (2.19), are subject to the uniform pay-for-effort allocation $(\overline{r}^*, e^*(\theta^p))$, defined in Proposition 2.2. The length of the pooling-equilibrium interval $[\theta^p, \overline{\theta}]$ decreases in λ.*

With this result in mind, we argue later that the lenient job performance appraisal practice with the ensuing compression of ratings can, in fact, be an equilibrium outcome.

Wage dispersion

In this subsection, we argue that in equilibrium the range of rewards $[r^*(\underline{\theta}), \bar{r}^*]$ increases in the parameter λ, i.e., the smaller a firm is, the higher wage (pay) dispersion in the firm is. To make this argument, we take the internal derivatives $de^*(\theta^p)/d\lambda$ and $de^*(\theta)/d\lambda$ of equilibrium conditions (2.17) and (2.20), respectively, and show that the first one is positive and provide conditions when the second one is negative, from which the postulated result follows (in particular, we then have $d\bar{r}^*/d\lambda > 0$ and $dr^*(\underline{\theta})/d\lambda < 0$).

The owner's optimality condition (2.17) shows that with the parameter λ decreasing (which makes the right-hand side of (2.17) increase), the owner wants the manager's highest effort level contracted $e^*(\theta^p)$ to be lower. Formally, the internal derivative $de^*(\theta^p)/d\lambda$ of (2.17) is positive:

$$\frac{de^*(\theta^p)}{d\lambda} = -\frac{\alpha C_e(e^*(\theta^p), \theta^p)}{(1-\alpha)V_{ee}(e^*(\theta^p) - (1-\alpha\lambda)C_{ee}(e^*(\theta^p), \theta^p))} > 0.$$

Therefore, with smaller values of λ, in order to attain a lower effort level in equilibrium $e^*(\theta^p)$, the owner has to impose a lower upper bound on employee rewards, implying that $d\bar{r}^*/d\lambda$ is positive. To put it in words, the more unaccountable her managers are, the more the owner constrains their discretion about employee compensation.

Next, we take the internal derivative $de^*(\theta)/d\lambda$ of (2.20) to obtain that for every θ in $[\underline{\theta}, \theta^p)$:

$$\frac{de^*(\theta)}{d\lambda} = -\frac{-\frac{g_e}{\theta^2}[\theta + \frac{1-F(\theta)}{f(\theta)} - \frac{d\theta^p}{d\lambda}\frac{f(\theta^p)}{f(\theta)}\frac{1-\lambda}{1-\alpha} - \frac{1-F(\theta^p)}{f(\theta)(1-\alpha)}]}{V_{ee} - \lambda\frac{g_{ee}}{\theta} - \frac{g_{ee}}{\theta^2}[\lambda\frac{(1-F(\theta))}{f(\theta)} + \frac{(1-F(\theta^p)}{f(\theta)}\frac{1-\lambda}{1-\alpha}]},$$

where the arguments of functions V and C are dropped for more clarity, and we also use $C(e,\theta) = g(e)/\theta$. Since $V_{ee} \leq 0$ and $g_{ee} > 0$, the denominator of the above expression is negative. The numerator is also negative if the expression in the square brackets is positive, which, however, is dependent on parameter values. To have this expression positive, we make the following assumptions: the employee is cost-efficient enough, i.e., the parameter θ takes high enough values and/or the manager's share of output, the parameter α, is not too large. If these (well justifiable) assumptions are met, then $de^*(\theta)/d\lambda$ is negative for every θ, and so it is for $\theta = \underline{\theta}$, implying that $dr^*(\underline{\theta})/d\lambda < 0$ (as follows from (2.13)). Since the optimal effort and reward allocations are continuous in type θ, the dispersion of rewards increases in parameter λ.

Intuitively, this equilibrium property stipulates that with less accountable managers in her firm the owner tries to limit the payroll expenses they incur by lowering the upper bound on employee rewards. It eventually makes a reward-constrained manager distort the incentives of most efficient employee types even further by eliciting more effort from less able types.

Proposition 2.4 below summarizes the equilibrium property discussed, which is also illustrated in the numerical example of the next subsection.

Proposition 2.4 *The highest available employee reward $\bar{r}^*$ and lowest contracted reward $r^*(\underline{\theta})$, defined in Proposition 2.2, are, respectively, increasing and decreasing in parameter λ if the employee is cost efficient* *enough and the parameter α is not too large. Then, given the continuity of the equilibrium reward schedule r^*, the range of wage dispersion $[r^*(\underline{\theta}), \bar{r}^*]$ increases in parameter λ.*

Average wage

In this subsection, for different values of λ we estimate the employee's expected equilibrium wage, defined as $\int_{\underline{\theta}}^{\overline{\theta}} r^*(\theta)f(\theta)d\theta$, where r^* is the reward function defined in Proposition 2.2. In the previous subsection, we showed that the pattern of wage dispersion with respect to λ is parameter-dependent, and so is the pattern of the expected wage. Therefore, instead of deriving an analytical expression for the relationship between the expected wage and parameter λ, we present the argument with the help of a numerical example. We show how the large-firm wage premium can arise (i.e., that the expected equilibrium wage declines in λ) and that it is consistent with larger firms having lower profits.

For expositional convenience, consider the following specification of the model. The production function V is linear in effort, $V(e) = e$; the effort cost function takes the form $C(e, \theta) = e^2/(2\theta)$; the employee types are uniformly distributed on $[5, 10]$, i.e., $\underline{\theta} = 5, \overline{\theta} = 10$; the manager's output share $\alpha = 0.15$; and the parameter λ takes values from $[0.5, 1]$. Using this specification, we calculate the equilibrium results of Proposition 2.2, which are illustrated in Figure 1. Its diagrams a), b), and c) illustrate for different values of λ the employee's expected wage and the firm's expected profit, wage dispersion, and the pooling-equilibrium starting point θ^p, respectively.

Diagram a) of Figure 1 shows that the firm's expected profit monotonically increases in λ (see the dotted line; this result naturally follows from

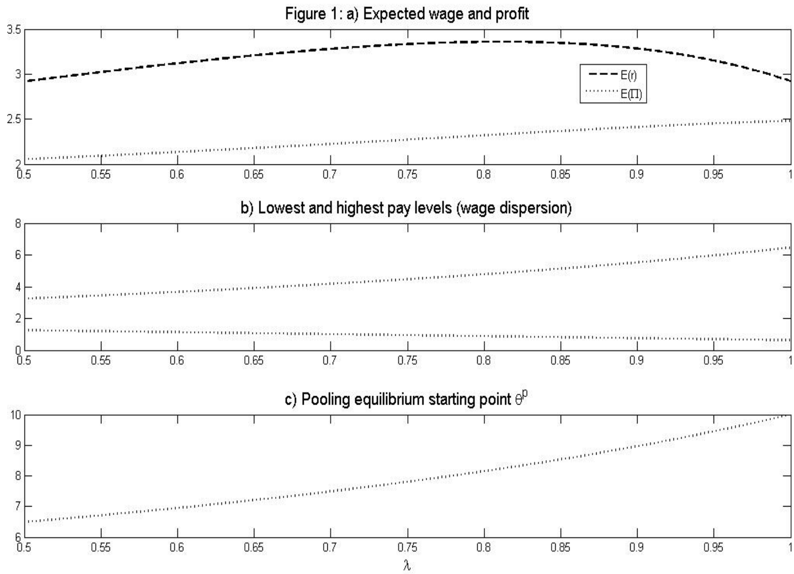

the model's setting). There is also an interval $[\bar{\lambda}, 1]$, with $\bar{\lambda} = 0.814$, at which the employee's expected wage declines in λ (see the dashed line; the employee's expected effort level, not shown in the diagram, also declines at this interval). To put it in words, a small firm's expected profit and payroll expense, everything else equal, can be respectively higher and lower than those of a larger firm (matching the empirical evidence of firm-size effects on average wages and profits). The reason for this, as we argue, is managers' lower degree of accountability in larger firms, which boosts employees' payroll expenses (and, correspondingly, their efforts exerted) above the firms' profit-maximizing levels, leading to profitability losses.

At the same time, as discussed in the previous subsections, we obtain the inverse relationship between firm size and wage dispersion: the gap between the highest and lowest pay increases in λ, see Diagram b). And

Diagram c) depicts the compression of ratings (rewards) phenomenon. For any λ less than one, different employee types at the high end of ability distribution are pooled for the same reward, and the pooling-equilibrium interval decreases in λ.

Furthermore, the patterns of wage variation in firms, demonstrated in Figure 1, are robust against other specifications of the model. For example, the same patterns—including the existence of a threshold value of λ, after which the expected wage declines in λ—are also observed for monotonically increasing, decreasing, or "bell-shaped" probability density functions f. Neither does the result change if the production function is taken to be strictly concave in effort e, e.g., $V(e) = e^{\beta}$ with $\beta < 1$, or to have returns to scale—$V(e) = e^{\beta}/\lambda^{s}$, where $s \in \mathbb{R}_{+}$ is a returns-to-scale parameter (then, there is a range of parameter s values, for which profitability still decreases with firm size).

Based on the above discussion, we make the following proposition.

Proposition 2.5 *Numerical tests of the model show that there is a restriction of parameter λ values, for which the employee's expected pay decreases in λ.*

5. Discussion

In the introduction, we raised the empirical stylized facts of the compression of ratings (rewards) and of the firm-size effects that we want to explain in this study. Below, we relate these facts with our theoretical results obtained. In addition, drawing on our findings, we provide different interpretations of some empirical evidence presented in the related literature.

Compression of ratings

It has been long observed that variation in rewards (ratings) is smaller than variation in the actual performance for which the rewards have been granted, see Murphy & Cleveland (1995). Relating this observation to our model, we argue that the compression of ratings can, in fact, be an outcome of managers' optimal performance evaluation strategy. If constrained in employee rewards, which he is only partially accountable for, a manager finds it optimal to extract more effort from low-ability employee types even at the expense of distorting the incentives of high-ability employee types. Given the results in Proposition 2.2 and 2.3, the manager differentiates only among those effort levels that are within the range $[e^*(\underline{\theta}), e^*(\theta^p)]$, and the width of this effort range decreases with firm size. So if an employee for one or another reason exerts an effort level above $e^*(\theta^p)$ the manager would still give her the same reward of $\bar{r}^*$.

Akerlof (1982) provides a specific example, where the incentives in place for cash posters at the Eastern Utilities Co. seemed to be suboptimal either from the employees' or the employer's perspective. In this example, employees were paid the same wage provided they recorded at least 300 postings per hour, and no bonuses or promotion promises were given for exceeding the limit. Some cash posters, however, did exceed the limit, but still were paid the same wage. It raised the question of why those "overworking" cash posters did not reduce their effort levels, or, on the other hand, why the employer did not provide additional incentives for them to extract even more effort.

In addition to the "gift-exchange" explanation by Akerlof (1982), our model can give another insight into the agency problem described. The fixed pay offered for at least 300 recorded postings could, in fact, constitute an optimal employee incentive scheme, where "optimal," from the manager's perspective, is to maximize the number of postings recorded. Technically, in our model, for low enough values of λ the pooling equilibrium may stretch out to comprise the whole employee type space (see condition (2.19)). To put it in words, if the manager is not held very accountable for the payroll expense he incurs, to set a uniform incentive scheme, just meeting the participation constraint of low-ability employees, can be optimal for the manager. However, why all the cash posters would not simply meet the prescribed limit is a question beyond the scope of our model.

Firm-size effects

The firm-size effects on wages take the form of a higher average wage and lower wage dispersion in larger firms (see Oi & Idson (1999); Garen (1985); Brown & Medoff (1989)). Given our assumption that a larger size means a larger asymmetry in information between the owner and managers, our model shows that the empirical regularities observed in practice can constitute an equilibrium outcome as well.

With regard to wage dispersion, we argue that the smaller a firm is (or the more accountable its managers are), the more efficient economic incentives for the firm's employees are put in place, and *vice versa*. It accordingly leads to the inverse relationship between wage dispersion and firm size (see Proposition 2.4). The reason for this result is that a larger firm has a more aggravated soft-budget-constraint problem, which prompts its owner to curb her managers' discretion about employee pay in order to avoid excessive payroll expenses. Managers respond to that, as discussed in the preceding subsection about the compression of ratings, by setting coarser reward schemes leading to a shallower differentiation of good from bad performance levels. This result has strong empirical support. Stigler (1962, Table 5) reports wage dispersion to vary inversely with firm size; Garen (1985) and Evans & Leighton (1989) report returns to employee ability and skills (experience) to be higher in smaller firms.

As for the large-firm wage premium, our model also offers a different view of this phenomenon. In Proposition 2.5, we argue that it can be an equilibrium outcome of the agency problem studied here that the larger a firm is, the higher its average wage is. A higher average wage comes from a higher average effort exerted, which empirically can be interpreted as meaning that workers are more productive in larger firms (see Idson & Oi (1999)). But as our model shows, it may not necessarily be the case. In larger firms, for the reasons explained before, managers design employee incentive schemes that elicit more effort from low-ability employees (whose incentives, otherwise, would be distorted to elicit more effort from high-ability employees). As a result, one can observe that employees in larger firms exert on average more effort, which, however, does not mean that they are more productive *per se*. It could be the incentive schemes offered by their managers that make them exert more effort on average, but this may not be in the firm's best interest.

In fact, our argument is reinforced by the empirical findings from financial studies about smaller firms having higher stock returns and, supposedly, higher levels of profitability (see Banz (1981); Fama & French (1992)). Hence, if workers in smaller firms are less productive (as argued, for example, in Idson & Oi (1999)), then how does this match with the fact that smaller firms have higher levels of profitability? Nonetheless, in our model, we do obtain that small firms are more profitable, which immediately follows from the model's structure. The owner of a smaller firm can more accurately align her managers' compensation scheme with the firm's profit maximization. At the same time, our model predicts that the average effort level decreases with firm size, but this is optimal from the firm's profit maximization perspective.

6. Conclusion

Based on the observation that managers have a soft budget constraint when evaluating their employees' performance, this study argues that the documented empirical regularities of the compression of ratings and firm-size effects can be the equilibrium outcomes of the model presented here. Given the idea that the owner of a firm cannot perfectly align her managers' incentives with the firm's profit maximization, the owner attempts to restrain her managers' payroll spending by putting an upper bound on employee rewards. This, subsequently, leads to managers designing flatter pay-for-effort allocations for their employees, which can be behind the compression of ratings phenomenon. Assuming that in smaller firms managers are held more accountable for their actions—as empirical evidence indicates—the model makes predictions that are in line with the empirical evidence from the industrial psychology, labor, and finance literature on firm-size effects. All in all, manager accountability can be a cause of the systematic differences observed in employee wage schedules. A further research direction could be to empirically test various predictions of the model in order to distinguish them more clearly from other alternative theories.

Appendix

The manager's problem, eq. (2.4)–(2.7)

Here, we prove Proposition 2.1, which is to solve for the optimal pay-for-effort allocations when the upper-bound constraint (2.7) is binding. First, we approach the problem through its discrete version, and then take the limit of the results obtained to arrive at the general solution.

Discretization

We discretize the employee type space $\left[\underline{\theta}, \overline{\theta}\right]$ into n discrete types $(\theta_1, ..., \theta_i, ..., \theta_n)$, where an employee type $\theta_i = \underline{\theta} + (i-1)\partial\theta$, for $i = 1, ..., n$, and $\partial\theta = (\overline{\theta} - \underline{\theta})/n$. Then, we discretize the initial (continuous) distribution F for employee types by defining probability weights $p(\theta_i) = \int_{\theta_i}^{\theta_i+\partial\theta} f(\theta)d\theta$ for every θ_i, which is the probability mass of the employee types within the interval $[\theta_i, \theta_i + \partial\theta]$. (From this discretization, we later switch to the continuous case by taking the limit $n \to \infty$, or $\partial\theta \to 0$.)

The discrete version of the manager's optimization problem eq. (2.4)–(2.7) is as follows. With respect to pay-for-effort allocations $\{r(\theta_i), e(\theta_i)\}$, $i = 1, ..., n$, the manager maximizes his expected utility

$$\sum_{i=1}^{n} p(\theta_i)\alpha[V(e(\theta_i)) - \lambda r(\theta_i)]$$

subject to

$$r(\theta_i) - C(e(\theta_i), \theta_i) \geq 0, \tag{P_i}$$

$$r(\theta_i) - C(e(\theta_i), \theta_i) \geq r(\theta_j) - C(e(\theta_j), \theta_i), \tag{IC_i}$$

$$0 \leq r(\theta_i) \leq \bar{r}, \quad \text{for every } i = 1, ..., n \text{ and } j \neq i. \tag{2.23}$$

Let us assume that the solution to the manager's problem is unique (the assumptions of our model ensure that it is indeed so). We start solving the problem by making the following conjecture.

Conjecture 2.1 *For any partition of the employee type space, the solution to the manager's problem consists of distinct pay-for-effort allocations for every employee type.*

Setting up the Lagrangian

As a standard approach to principal-agent problems with hidden information, we start with reducing our problem by getting rid of redundant constraints. First, the only binding participation constraint is that of the least efficient agent type from those contracted upon. We impose it to be P_1 assuming that in the population there is a large enough mass of inefficient employee types. In the optimum, the adjacent IC constraints need to be downward binding:

$$r(\theta_i) - C(e(\theta_i), \theta_i) = r(\theta_{i-1}) - C(e(\theta_{i-1}), \theta_i), \qquad i = 2, ..., n. \quad (2.24)$$

If in the solution the effort schedule is monotonically increasing in the employee type (which has to be checked separately), then due to the Spence-Mirrlees property the rest of incentive compatibility constraints are met.

The constraints in (2.24) together with the binding P_1 constraint determine the pay levels $r(\theta_i)$ for every $i = 2, ..., n$ as

$$r(\theta_i) = \sum\nolimits_{j=1}^{i} C(e(\theta_j), \theta_j) - \sum\nolimits_{j=2}^{i} C(e(\theta_{j-1}), \theta_j). \quad (2.25)$$

Since the upper-bound constraint (2.23) is binding, it follows from Conjecture 2.1 and the monotonicity of effort levels that only the most efficient employee type is subject to the highest reward. Then, together with (2.25) for $r(\theta_n)$, the upper-bound constraint takes the form of

$$\bar{r} - \sum_{i=1}^{n} C(e(\theta_i), \theta_i) + \sum_{i=2}^{n} C(e(\theta_{i-1}), \theta_i) = 0. \tag{2.26}$$

Next, we set the Lagrangian of the reduced optimization problem, which is

$$L(\{e(\theta_i)\}_{i=1}^{n}, \mu) = p(\theta_1)\alpha[V(e(\theta_1)) - \lambda C(e(\theta_1), \theta_1)] +$$
$$+ \sum_{i=2}^{n-1} p(\theta_i)\alpha[V(e(\theta_i)) - \lambda(\sum\nolimits_{j=1}^{i} C(e(\theta_j), \theta_j) - \sum\nolimits_{j=2}^{i} C(e(\theta_{j-1}), \theta_j))] +$$
$$+ p(\theta_n)\alpha[V(e(\theta_n)) - \lambda\bar{r}] + \mu(\bar{r} - \sum_{i=1}^{n} C(e(\theta_i), \theta_i) + \sum_{i=2}^{n} C(e(\theta_{i-1}), \theta_i)),$$

where μ is a Lagrange multiplier on the upper-bound constraint (2.26). (Other constraints enter the Lagrangian through $r(\theta_i)$ replaced by (2.25).)

The first-order conditions with respect to the effort levels $e(\theta_i)$ for $i = 1, ..., n-1$ are

$$p(\theta_i)\alpha[V_e(e(\theta_i)) - \lambda C_e(e(\theta_i), \theta_i)] - [\alpha\lambda \sum\nolimits_{j=i+1}^{n-1} p(\theta_j) + \mu] \times \tag{2.27}$$
$$\times (C_e(e(\theta_i), \theta_i) - C_e(e(\theta_i), \theta_{i+1})) = 0,$$

and with respect to $e(\theta_n)$ it is

$$p\left(\theta_n\right) \alpha V_e(e(\theta_n)) - \mu C_e(e(\theta_n), \theta_n) = 0. \tag{2.28}$$

Solving these n first-order conditions together with constraint (2.26) should give us the optimal effort levels $e^*(\theta_i)$ for $i = 1, ..., n$, with the corresponding pay levels $r^*(\theta_i)$ following from IR_1 and (2.25). If at the limit $n \to \infty$, the pay-for-effort allocations obtained are distinct for every employee type with the effort schedule monotonically increasing, then it is the solution to the manager's problem (2.4)–(2.7).

But, as is shown below, for fine partitions of the employee type space the perfect screening of employee types cannot be optimal. The manager can do better by pooling some of the most efficient types.

Pooling at the top

Let $\tilde{e}(\theta_i)$, $i = 1, ..., n$, solve the above first-order conditions. It must be that the effort level $\tilde{e}(\theta_n)$ aimed at the most efficient employee type is less than the first-best effort level defined as $e^{fb}(\theta_n) = \{e(\theta_n) : V_e(e(\theta_n)) - \lambda C_e(e(\theta_n), \theta_n) = 0\}$.[6] It results in the efficiency loss of $V_e(\tilde{e}(\theta_n)) - \lambda C_e(\tilde{e}(\theta_n), \theta_n) > 0$ and implies $\mu > p(\theta_n)\alpha\lambda$.

Next, through the Lagrange multiplier μ we combine the adjacent first-order conditions for $\tilde{e}(\theta_n)$ and $\tilde{e}(\theta_{n-1})$ to get

$$\frac{p(\theta_n)}{p(\theta_{n-1})} = \frac{[V_e(\tilde{e}(\theta_{n-1})) - \lambda C_e(\tilde{e}(\theta_{n-1}), \theta_{n-1})]C(\tilde{e}(\theta_n))}{V(\tilde{e}(\theta_n))[C_e(\tilde{e}(\theta_{n-1}), \theta_{n-1}) - C_e(\tilde{e}(\theta_{n-1}), \theta_n)]}. \tag{2.29}$$

Multiplying both sides by $\partial\theta$ and taking the limit $\partial\theta \to 0$ (equivalent to taking the limit $n \to \infty$) render that the left-hand side of the above expression tends to zero (since the limit $\lim_{n\to\infty} p(\theta_n)/p(\theta_{n-1}) = 1$). At the same time, the right-hand side is equal to

$$\frac{[V_e(\tilde{e}(\overline{\theta})) - \lambda C_e(\tilde{e}(\overline{\theta}), \overline{\theta})]C_e(\tilde{e}(\overline{\theta}))}{V_e(\tilde{e}(\overline{\theta}))(-C_{e\theta}(\tilde{e}(\overline{\theta}), \overline{\theta}))},$$

which remains strictly positive because of $V_e(\tilde{e}(\overline{\theta})) - \lambda C_e(\tilde{e}(\overline{\theta}), \overline{\theta}) > 0$.

Hence, for the continuum of employee types (or for fine partitions of the employee type space) the derived optimality (first-order) condi-

[6] To see this, if $\tilde{e}(\theta_n) = e^{fb}(\theta_n)$, then $\mu = \lambda\alpha p(\theta_n)$, from which it follows that the effort levels $e(\theta_i)$ for all i are identical to the optimal effort levels from the problem without the upper-bound constraint. But it would inevitably violate some incentive-compatibility constraints of the full optimization problem (provided, of course, the manager does not exclude any low types, and that is ruled out).

tions cannot support the distinct pay-for-effort allocations conjectured—Conjecture 2.1 does not hold at the limit. For fine type space partitions, to meet the optimality conditions the manager has to pool some of the most efficient employee types by making them subject to the highest reward of $\bar{r}$.

Then, we continue with gradually increasing the probability mass of employee types subject to the highest reward and denote this mass by $P(\theta_m) = \sum_{j=m}^{n} p(\theta_j)$, where $m = n-1, n-2, \ldots$. We repeat the above solution algorithm for different m (with m replacing n in the above derivations) until we have the optimality conditions met. In particular, for a given m, the first-order condition equivalent to (2.28) is:

$$P(\theta_m)\,\alpha V_e(e(\theta_m)) - \mu C_e(e(\theta_m), \theta_m) = 0, \tag{2.30}$$

while the rest of the first-order conditions for $i = 1, \ldots, m-1$ remain intact (again conjecturing that the effort schedule is increasing in the employee type).

The equivalent expression to (2.29) is

$$\frac{P(\theta_m)}{p(\theta_{m-1})} = \frac{[V_e(e(\theta_{m-1})) - \lambda C_e(e(\theta_{m-1}), \theta_{m-1})]C(e(\theta_m))}{V(e(\theta_m))[C_e(e(\theta_{m-1}), \theta_{m-1}) - C_e(e(\theta_{m-1}), \theta_m)]}. \tag{2.31}$$

Multiplying both sides by $\partial\theta$ and taking the limit $\partial\theta \to 0$ on both sides render the optimal pooling condition:

$$\frac{1 - F(\theta^p)}{f(\theta^p)} = \frac{[V_e(e(\theta^p)) - \lambda C_e(e(\theta^p), \theta^p)]C_e(e(\theta^p))}{V_e(e(\theta^p))(-C_{e\theta}(e(\theta^p), \theta^p))}, \tag{2.32}$$

where θ^p is the employee type for which the above optimality condition holds (which is exactly (2.10) in Proposition 2.1). The ability type θ^p is the starting point of the pooling interval $[\theta^p, \overline{\theta}]$, for which the uniform allocation $(e(\theta^p), \bar{r})$ applies. The effort level $e(\theta^p)$ is pinned down by the remaining optimality conditions as defined below.

The optimal allocations $\{e^*(\theta), r^*(\theta)\}_{\theta \in [\underline{\theta}, \overline{\theta}]}$

Having established the pooling condition (2.32) and reverting to the continuous case henceforth, from (2.30) the Lagrange multiplier is equal to

$$\mu = (1 - F(\theta^p)) \frac{aV_e(e(\theta^p))}{C_e(e(\theta^p), \theta^p)}.$$

Plugging it into the remaining first-order conditions (2.27) and taking the continuous version of them render for any $\theta \leq \theta^p$

$$[V_e(e(\theta)) - \lambda C_e(e(\theta), \theta)] + \lambda \frac{(1 - F(\theta))}{f(\theta)} C_{e\theta}(e(\theta), \theta) + \tag{2.33}$$
$$+ \frac{(1 - F(\theta^p)}{f(\theta)} \frac{V_e(e(\theta^p)) - \lambda C_e(e(\theta^p), \theta^p)}{C_e(e(\theta^p), \theta^p)} C_{e\theta}(e(\theta), \theta) = 0,$$

which is (2.12) in Proposition 2.1. Finally, the last condition that needs to be met is constraint (2.26), the continuous version of which is

$$\bar{r} = C(e(\theta^p), \theta^p) - \int_{\underline{\theta}}^{\theta^p} C_\theta(e(\theta), \theta) d\theta, \tag{2.34}$$

which is (2.11) in Proposition 2.1.

All in all, conditions (2.32)–(2.34) together determine the optimal effort levels $e^*(\theta)$ for all θ in $[\underline{\theta}, \overline{\theta}]$. Given the modeling assumptions imposed, one can easily verify from (2.33) that the monotonicity constraint for e^* to be increasing that has been omitted holds. Finally, the optimal pay levels $r^*(\theta)$ for θ in $[\underline{\theta}, \theta^p)$ follow from the continuous version of (2.25), which is (2.13) in Proposition 2.1.

Bibliography

Akerlof, George A. 1982. Labor Contracts as Partial Gift Exchange. *Quarterly Journal of Economics*, 97(4), 543–569.

Alchian, Armen A., & Demsetz, Harold. 1972. Production, Information Costs, and Economic Organization. *American Economic Review*, 62(5), 777–795.

Banz, Rolf W. 1981. The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1), 3–18.

Besanko, David, Dranove, David, Shanley, Mark, & Schaefer, Scott. 2007. *Economics of strategy*. Hoboken, N.J. : John Wiley & Sons.

Brown, Charles, & Medoff, James. 1989. The Employer Size-Wage Effect. *Journal of Political Economy*, 97(5), 1027–59.

Bruns, Jr., William J. (ed). 1992. *Performance Measurement, Evaluation and Incentives*. Boston, MA: Harvard Business School Press.

Bruns, Jr., William J., & McKinnon, Sharon M. 1992. Performance evaluation and managers'descriptions of tasks and activities. *Pages 17–36 of*: William J. Bruns, Jr. (ed), *Performance Measurement, Evaluation and Incentives*. Boston, MA: Harvard Business School Press.

Bulow, Jeremy I., & Summers, Lawrence H. 1986. A Theory of Dual Labor Markets with Application to Industrial Policy, Discrimination, and Keynesian Unemployment. *Journal of Labor Economics*, 4(3), 376–414.

Evans, David S., & Leighton, Linda S. 1989. Why Do Smaller Firms Pay Less? *Journal of Human Resources*, 24(2), 299–318.

Fama, Eugene F., & French, Kenneth R. 1992. The Cross-Section of Expected Stock Returns. *Journal of Finance*, 47(2), 427–465.

Garen, John E. 1985. Worker Heterogeneity, Job Screening, and Firm Size. *Journal of Political Economy*, 93(4), 715–739.

Hamermesh, Daniel S. 1980. Commentary. *Pages 383–388 of*: Siegfried, John J. (ed), *The Economics of Firm Size, Market Structure, and Social Performance*. Washington, D.C.: Federal Trade Commission.

Idson, Todd L., & Oi, Walter Y. 1999. Workers Are More Productive in Large Firms. *American Economic Review*, 89(2), 104–108.

Landy, Frank J., & Farr, James L. 1980. Performance rating. *Psychological Bulletin*, 87(1), 72–107.

Landy, Frank J., & Farr, James L. 1983. *The measurement of work performance: Methods, theory, and applications.* New York: Academic Press.

Levin, Jonathan. 2003. Relational Incentive Contracts. *American Economic Review*, 93(3), 835–857.

Liden, Robert C., & Mitchell, Terence R. 1983. The Effects of Group Interdependence on Supervisor Performance Evaluations. *Personnel Psychology*, 36(2), 289–299.

Longenecker, Clinton O., Sims, Jr., Henry P., & Gioia, Dennis A. 1987. Behind the Mask: The Politics of Employee Appraisal. *Academy of Management Executive*, 1(3), 183–193.

MacLeod, W. Bentley. 2003. Optimal Contracting with Subjective Evaluation. *American Economic Review*, 93(1), 216–240.

Mero, Neal P., & Motowidlo, Stephan J. 1995. Effects of rater accountability on the accuracy and the favorability of performance ratings. *Journal of Applied Psychology*, 80(4), 517–524.

Milkovich, George T., & Wigdor, Alexandra K. 1991. *Pay for performance: Evaluating performance and appraisal merit pay.* Washington, D.C. : National Academy Press.

Mirrlees, James A. 1971. An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies*, 38(2), 175–208.

Murphy, Kevin J. 1992. Performance measurement and appraisal: motivating managers to identify and reward performance. *Pages 37–62 of*: William J. Bruns, Jr. (ed), *Performance Measurement, Evaluation and Incentives.* Boston, MA: Harvard Business School Press.

Murphy, Kevin R., & Cleveland, Jeanette. 1995. *Understanding Performance Appraisal: Social, Organizational, and Goal-based Perspectives.* SAGE.

Oi, Walter Y., & Idson, Todd L. 1999. Firm Size and Wages. *Pages 2165–2214 of*: Ashenfelter, O., & Card, D. (eds), *Handbook of Labor Economics*, vol. 3. Elsevier.

Oyer, Paul, & Schaefer, Scott. 2005. Why do some firms give stock options to all employees?: An empirical examination of alternative theories. *Journal of Financial Economics*, 76(1), 99–133.

Prendergast, Canice. 1999. The Provision of Incentives in Firms. *Journal of Economic Literature*, 37(1), 7–63.

Stigler, George J. 1962. Information in the Labor Market. *Journal of Political Economy*, 70(5), 94–105.

Weiss, Andrew, & Landau, Henry J. 1984. Wages, Hiring Standards and Firm Size. *Journal of Labor Economics*, 2(4), 477–99.

Chapter 3

How Exposure to Markets Can Favor Inequity-Averse Preferences*

Abstract

This study shows how non-individualistic preferences can be individual fitness maximizing in the presence of general equilibrium externalities. In the model, individuals share an endowment among themselves and use the proceeds either for immediate consumption or for the purchase of consumption goods from merchants on the external market if such exists. Assuming that increased consumption means increased individual fitness, inequity-averse behavior with respect to endowment distribution can be an optimal response to merchants' price discrimination and lead to the evolution of inequity-averse preferences. The finding that members of a market-integrated society are more pro-social as compared to members of an isolated society is supported by empirical evidence.

1. Introduction

There is vast empirical evidence on people sharing money that shows that people seem to care, alongside their own pecuniary interest, about the well-being of other parties affected, to a larger or smaller extent, (for a comprehensive review, see Fehr & Schmidt (2006)). As documented and

*I am indebted to Karl Wärneryd for valuable comments throughout the progress of this work. I also thank Kaushik Basu, Martin Dufwenberg, and Cheng-Zhong Qin. The financial support from the Jan Wallander and Tom Hedelius Foundation and Inga och Sixten Holmquists stipendiestiftelse is gratefully acknowledged.

tested in numerous experiments, people's behavior seems to exhibit certain regular patterns such as inequity aversion—which will be the main subject of this study—described as "people are willing to give up some material payoff to move in the direction of more equitable outcomes."[1] This led to the thought that pro-social preferences—with inequity aversion, in particular—can be more characteristic of human nature than selfish ones. Popular representations of inequity-averse preferences are given in Fehr & Schmidt (1999); Bolton & Ockenfels (2000); Charness & Rabin (2002)), which are, in crude terms, extensions of own-regarding preferences to include inequity-aversion terms.

Further research has also shown that the form of revealed preferences (and, supposedly, the degree of inequity aversion) varies from society to society suggesting that preferences can, actually, be context dependent and shaped by the environment people live in. (For empirical evidence, see Buchan et al. (2002); Henrich et al. (2001); Henrich et al. (2004); whereas Bowles (1998) offers a systematic review of related theoretical and empirical literature.) The most compelling evidence on between-group differences in people's behavior comes from a project conducted in 15 remote small-scale traditional societies scattered around the globe, which is documented in studies Henrich et al. (2001) and Henrich et al. (2004). In particular, with the help of experiments, the researchers have discovered certain regularities in people's revealed amount of sociality. One of the regularities is that members of a market-integrated society (as measured, primarily, by the society's exposure to market exchange) behave on average more pro-socially than do members of an isolated society. Henrich et al. (2004, p.50–51), however, leave open the question

[1]Fehr & Schmidt (1999, p. 819)

of what explains this empirical pattern, calling for more research on this important finding, and the current study attempts to contribute toward a better understanding of this.

This study offers an evolutionary argument for inequity-averse preferences and shows how situational factors can influence their appearance. Aiming to explain the findings of Henrich et al. (2004), we provide a general equilibrium framework, where people's preferences for money distribution are endogenously determined by their society's exposure to merchandise markets and the structure of those markets. We demonstrate how, in societies with market exchange, inequity aversion can be individual fitness maximizing and eventually be favored by natural or rather cultural selection. The essence of the argument presented herein is that we measure evolutionary fitness *not* in terms of monetary returns, which are the direct object of people's decision making, but rather in terms of the consumption that those monetary returns can afford. Since these different measures of fitness in general equilibrium are not necessarily equivalent, sharp differences in the results can be obtained depending on what measure is used.[2]

We present the argument in a model—an extended dictator game with consumption—where in equilibrium the dictator can be better off (in terms of expected own consumption) by sharing the monetary endowment with others. As a simple example, illustrating the main idea presented here, consider a two-player dictator game with consumption, where the dictator is randomly chosen from the two identical individuals

[2]In a similar fashion, Huck & Oechssler (1999) develop an evolutionary argument for revengeful behavior presuming that the individual subjective payoff and subsequent evolutionary fitness resulting from strategies employed are not equivalent. The general models of evolution of preferences (see Ely & Yilankaya (2001); Ok & Vega-Redondo (2001); Dekel et al. (2007)) also differentiate between people's subjective and objective preferences.

to split a monetary endowment given exogenously. Suppose that the endowment distribution resulting from a split is public information, but the players' individual endowment shares are their private information. Let an individual's utility from an endowment split be measured in the units consumed of the only available good that the individual's endowment share can afford. Finally, there is a monopolist producer, who produces the good at some constant marginal cost, and, after learning about the endowment distribution, charges the price for a unit of the good that maximizes her profits from following simultaneous trades. Within the setting described, we raise the question of what is the optimal sharing rule maximizing the dictator's consumption? Obviously, it is not optimal for the dictator to keep all the endowment for himself, because in that case the producer targets only the dictator by setting the price equal to the whole endowment, leaving the dictator with only one unit of the good consumed. Instead, the dictator could increase his consumption by giving away to the other individual a portion of the endowment large enough to make the rational producer set the price aimed at both individuals, which would leave the richer one—the dictator—with some consumer surplus (or rather information rent) and more units consumed. Hence, from a conventional utility function for consumption we obtain a non-monotonic indirect utility function of money, which can be interpreted as having underlying inequity-averse preferences for money distribution. The intuition behind this result is that by sharing with others one can acquire information rent and, consequently, increase the purchasing power of one's own, even diminished, share.

In this chapter, we develop the above idea into a formal model. We take an evolutionary perspective to argue that, because of general equilibrium externalities, inequity-averse preferences for money distribution can render a higher material payoff than that rendered by individualistic preferences and, correspondingly, be favored by natural or cultural selection. We use an evolutionary approach in order to relax the rationality assumption *à la* "homo economicus" and allow the players to maximize their subjective preferences rather than their objective preferences, which, nevertheless, determine the players' reproduction success. While we adopt the "indirect" evolutionary approach (see Güth & Yaari (1992); Ely & Yilankaya (2001)), when showing that the equilibrium play of the game in question is evolutionary stable, the standard approach (Weibull (1995)) would render the same results, too. In fact, in our setting the two approaches are interchangeable, allowing us to relate the findings obtained with the literature on both approaches.

This study also contributes to the evolutionary literature by providing a distinct and empirically supported argument on how non-individualistic preferences in the individual selection framework may survive evolutionary pressures. Typically, evolutionary models in favor of non-individualistic preferences have required either a group-selection argument in the standard approach (for a review, see Bergstrom (2002)) or certain informational assumptions about the observability of the players' preferences in the "indirect" approach (for a concrete example, see Bester & Güth (1998); for a more general argument, see Dekel et al. (2007)). This study, however, bypasses all of the above: the result primarily hinges on general equilibrium effects.[3] Therefore, this study instead falls into the "game of life" paradigm, arguing that people's behavior should be examined in a

[3]Certainly, the current study is not unique in showing how individual selection can favor pro-social preferences. For instance, Becker (1976) presents a model in which egoists take actions as though they had altruistic preferences in order to benefit from others' altruism.

wider social context (see Binmore (1994, 1998); or Güth & Napel (2006) for an example related to the evolution of inequity-averse preferences).

The remainder of this chapter is organized as follows. Section 2 expands the example given above into a more general model and solves it. Section 3 discusses the results obtained, links them to empirical studies, and offers possible extensions. The last section concludes the study.

2. Model

Framework

After the land rewards a group of farmers[4] with a publicly observed harvest surplus, henceforth, the endowment S, the farmers share it among themselves, and, if used for their own consumption, a share $x \in [0, S]$ renders the material payoff of $U^0(x)$, $U^0_x > 0$, $U^0_{xx} < 0$. In the event the group is exposed to external trades, endowment shares can also be used as a means of exchange, i.e., as money, to purchase goods from merchants. It is assumed that the endowment distribution within the group, ensuing after an endowment split, is public information, while individual shares are only privately known.

Suppose that merchants can offer one type of goods—"the good"—which, on the other hand, can be produced in various quality q greater than or equal to some $\underline{q} > 0$ (this condition is for modeling purposes, as is explained later) with the production function $C(q), C_q > 0, C_{qq} > 0$, and the returns to scale from producing a given variety are constant. Assum-

[4]The "farmers" are chosen in order to allude to the historical division of labor into farmers, nomads, and merchants, which could potentially serve as a "real life" example in the subsequent argument about the evolution of inequity-averse preferences for money distribution. In addition, "farmer" economy is intended to refer to the traditional societies in Henrich et al. (2004), from which comes the empirical support for the results of the model.

ing that every farmer has a demand for at most one unit of the good, and accounting for the income distribution observed within the group and market competition, described more precisely below, merchants offer the farmers a menu of price-quality (p, q) bundles of the good to choose from, where the price p is gauged in terms of the endowment. The consumption of a (p, q) unit and of the remainder of the endowment share x renders a farmer the material payoff of $U^G(x-p,q), U_q^G > 0, U_{qq}^G < 0, U_x^G > 0, U_{xx}^G < 0, U_{qx}^G > 0$. The farmer will consider purchasing a variety (p, q) only if it results in a non-negative net utility level U, defined as $U(q,x,p) \equiv U^G(x-p,q) - U^0(x)$, which needs to be greater than or equal to 0 for the trade to take place. Correspondingly, the properties of the net utility function U are $U_x > 0, U_{xx} \leq 0, U_q > 0, U_{qq} \leq 0$, and $U_{qx} > 0$.[5] For convenience, let the function U be of the quasi-linear form in the price p :

$$U(q,x,p) = V(q,x) - p. \tag{3.1}$$

Finally, we shall consider three different scenarios of the external market structure: 1) merchants are absent (the farmers' economy is autarkic), 2) monopoly (there is a monopolist merchant), and 3) perfect competition (there are many competing merchants).

Game and natural selection

Along the lines of the above framework, consider a large population of farmers randomly and repeatedly matched to form separate groups of two

[5]All the listed properties of the utility function U are related to consumer preferences for normal goods (as in, e.g., Mas-Colell et al. (1995)). In particular, the positive partial derivative U_x implies that a richer individual derives a higher utility from the consumption of the good (due to, say, smaller opportunity costs). Similarly, the positive cross derivative U_{qx} can be interpreted as meaning that a richer consumer values quality more, which can be motivated by the convexity of the Engel curves for high-quality goods.

farmers, and every group is endowed with the same-size endowment S. In a group, Nature randomly selects a farmer, henceforth, the dictator, to divide the endowment into shares $s \in [0, S]$ for herself and $(S - s)$ for the other farmer.[6] Suppose that farmers have subjective preferences over endowment distribution (or, to put it differently, preferences for money distribution), characterized by the subjective utility function U^S with preference parameter (type) $\bar{s} \in [0, S]$ such that the (subjective) utility from an endowment share x, accrued to a farmer with a preference type $\bar{s}$, is

$$U^S(s, \bar{s}) = -\left|s - \bar{s}\right|.$$

Therefore, in what follows, a farmer of preference type $\bar{s}$, when selected to share the endowment, always keeps $\bar{s}$ for herself, leaving $S - \bar{s}$ to the other farmer in the match.[7] Next, suppose that the population distribution of subjective preference types is given by some distribution F over $[0, S]$.

The objective payoffs from a split, or the evolutionary fitness, are measured by the resulting material payoffs U^0 and/or U^G, which, on the other hand, depend both on the own-endowment share and on the menu of consumption bundles offered on the external market. The farmers, however, cannot discern for themselves what material payoffs their actions result in. Instead, they can be thought of as living behind the "veil of ignorance" about external markets or about what "global game" they are part of, and, therefore, they just divide the endowment according to

[6] As for the endowment sharing rule, we adopt the dictator-game framework, which is done mainly for modeling convenience; the main results are also robust against other modeling frameworks, e.g., the ultimatum game. What matters in the end is the presence of general equilibrium effects.

[7] While we are following the "indirect" evolutionary approach (Güth & Yaari (1992) and Ely & Yilankaya (2001)), alternatively, we could think of the farmers as being pre-programmed to split the endowment according to their preference types as in standard evolutionary models, see Weibull (1995). Due to the specificity of the game studied, the two approaches would render identical results, which is not necessarily the case in general (e.g., Huck & Oechssler (1999)).

their subjective preferences only. Significantly, knowledge of other farmers' preferences or the population distribution of preferences will not play any role in this game, although that is not generally the case (see, e.g., Ok & Vega-Redondo (2001) or Dekel et al. (2007)). Next, for modeling convenience, merchants are assumed to design consumption bundles for every match separately, and these bundles are not available to the farmers from other matches. This sequential structure of the model with merchants assumed to act in their profit maximizing way allows us to "prune" the production and consumption stages and to consider the reduced game only, with the external market structure and merchants' optimal play embodied in the players' (i.e., farmers') material payoff function.

As already specified above, we distinguish three cases of merchants' market: autarky, the monopolist market, and the perfectly competitive market, which give rise to three distinct farmers' material payoff functions, and we analyze the three cases in three different games Γ^A, Γ^M, and Γ^C, respectively. On the whole, each of these games will be a two-player dictator game with modified payoffs, measured in affordable consumption. More precisely, in every game, there are two players (two farmers), their action space is to choose an endowment share $s \in [0, S]$ for himself or herself, and the payoffs to the players from their actions are their expected material payoffs (evolutionary fitness), as will be defined below. Finally, every player has subjective preferences over an endowment split, characterized by a parameter $\bar{s}$, according to which he or she divides the endowment. Then, for every game examined separately, we shall tackle the question of what subjective preferences yield the greatest material payoffs and, accordingly, will be favored by natural selection, with their population share increasing at the expense of other less successful preferences.

We adopt the "indirect" evolutionary approach with a static stability concept of equilibrium so that in the equilibrium no mutation can give a higher payoff than that of the incumbent types. Based on the results of Ely & Yilankaya (2001) and applying them to our setting, evolution will select those subjective preferences, or, equivalently, actions over an endowment split that constitute an equilibrium of the global games Γ^A, Γ^M, and Γ^C, respectively, as more accurately discussed below and we shall call those preferences evolutionary stable.[8]

Equilibrium play in a match

Case 1: Autarky

Consider game Γ^A, where the farmers are not exposed to any external trades, making it a standard two-player dictator game. The material payoff from an endowment share x is $U^0(x)$; given the optimal play of the players (with respect to their subjective preferences), the expected material payoff to a farmer of preference type $\bar{s}$ when matched with a farmer of preference type $\bar{s}'$ is

$$\pi(\bar{s}, \bar{s}') = 0.5U^0(\bar{s}) + 0.5U^0(S - \bar{s}').$$

More generally, in game Γ^A the payoff (evolutionary fitness) to a preference type $\bar{s}$, given a population distribution of subjective preferences F,

[8] Ely & Yilankaya (2001) studies finite games, while in our model the action space is allowed to be infinite: $s \in [0, S]$. However, since we design our games in such a way that the existence of equilibrium is not an issue, then the results of Ely & Yilankaya (2001) apply to our setting as well despite a continuous action space. Alternatively, we could make our games studied finite by simply discretizing the players' action and preference spaces, and then the results of Ely & Yilankaya (2001) would apply directly.

is the average material payoff $\Pi(\bar{s})$ defined as

$$\Pi(\bar{s}) = 0.5U^0(\bar{s}) + 0.5E_{\bar{s}' \in C(F)}U^0(S - \bar{s}'),$$

where $C(F)$ is the support of the distribution F, and E is the expectations operator.

Since the second term of the above fitness expression does not depend on the own-preference type, the farmers of type $\bar{s} = S$ attain the highest fitness because of $U^0_x > 0$. Hence, the equilibrium of Γ^A (the evolutionary stable strategy) is to keep the whole endowment, resulting in the endowment split $(S, 0)$, which implies that in autarky selfish types $(\bar{s} = S)$ would prevail.

Case 2: Monopoly

In game Γ^M, to specify the players' material payoffs, first, we need to solve for the optimal consumption bundles offered by the monopolist profit-maximizing merchant. From the merchant's perspective, it is a mechanism design problem with hidden information, for a potential customer's wealth, i.e., his endowment share, is his private information. Once the menu of bundles is set, it is not subject to change, by which we rule out the possibility of the merchant's updating her beliefs about prospective buyers' wealth distribution after some trade has taken place (alternatively, we could assume that at the last stage only one trade with a random farmer takes place). The exposition of the merchant's problem closely follows Mussa & Rosen (1978).

Merchant's problem

After the endowment in a match is divided between the two farmers, the merchant learns about the ensuing endowment shares $\tilde{s}_1$ and

$\tilde{s}_2$, $\tilde{s}_1 \leq \tilde{s}_2$, and, accordingly, maximizes her expected profit $0.5(p_1 - C(q_1)) + 0.5(p_2 - C(q_2))$ with respect to price-quality bundles of the good $\{(p_1, q_1), (p_2, q_2)\}$, where the second bundle is aimed at the richer farmer (and the assumption is that every farmer has a demand for at most one unit of the good). There is no need to state the problem and solve it formally, for the solution to this type of problem is well established in the contract theory literature once specific conditions are met (such as the single-crossing property, which in our model is ensured by the assumption of the positive cross derivative of net utility function U, $U_{qx} > 0$). Hence, below we immediately proceed with describing the results for various scenarios of the endowment split.

In the special case of the equal endowment split $\tilde{s}_1 = \tilde{s}_2 = S/2$, the merchant offers one price-quality allocation $(p_1, q_1) = (p_2, q_2) = (p, q)$ such that

$$C_q(q) = V_q(q, \tilde{s}_2), \tag{3.2}$$

$$p = V(q, \tilde{s}_2), \tag{3.3}$$

which coincides with the first-best allocation under symmetric information, where consumers are left with no consumer surplus.

For an uneven split, $\tilde{s}_1 < \tilde{s}_2$, two cases need to be distinguished, depending on what the merchant finds optimal: 1) to serve both farmers, and 2) to ignore the poorer farmer and serve only the richer farmer. When the merchant serves both farmers, the optimal price-quality bundles (p_1, q_1) and (p_2, q_2) are found from

$$C_q(q_2) = V_q(q_2, \tilde{s}_2), \tag{3.4}$$

$$C_q(q_1) = 2V_q(q_1, \tilde{s}_1) - V_q(q_1, \tilde{s}_2), \tag{3.5}$$

for the quantities q_1 and q_2, and the prices follow from

$$U(q_1, \tilde{s}_1, p_1) = 0,$$

$$U(q_2, \tilde{s}_2, p_2) = U(q_1, \tilde{s}_2, p_1),$$

where the last two expressions are the binding individual rationality constraint of the poorer type ($\tilde{s}_1$) and incentive-compatibility constraint of the richer type ($\tilde{s}_2$), respectively. In this case, the richer farmer enjoys the information rent of size $U(q_2, \tilde{s}_2, p_2)$, while the poorer one is left with none. The condition for both farmers to be served is that the quality q_1 from (3.5) needs to be greater than or equal to the lowest feasible quality level $\underline{q}$, or the poorer farmer's endowment share $\tilde{s}_1$ to be greater than or equal to the threshold share s^* defined as

$$s^* = \{\tilde{s}_1 : \tilde{q}_1(\tilde{s}_1) = \underline{q}\}, \tag{3.6}$$

where $\tilde{q}_1(\tilde{s}_1)$ is the quality mapping from share $\tilde{s}_1$ to quality q_1 as found from (3.5), where $\tilde{s}_2$ is replaced with $S - \tilde{s}_1$.[9]

Hence, in the event of an endowment split with $\tilde{s}_1 < s^*$, the merchant does not serve the poorer farmer, but offers the first-best allocation to the richer farmer as in (3.2) and (3.3), thus, leaving the latter with no information rent.

Evolutionary fitness

As before, in game Γ^M the players' payoffs, or their evolutionary fitness, are their expected material payoffs. Given a population distribution of subjective preferences F, the expected evolutionary fitness of a player

[9]Had we $q \in [0, \infty)$, then the threshold share s^*, at which the merchant starts offering two non-zero consumption bundles, would not be precisely determined; neither would the equilibrium of game Γ^M.

of preference type $\bar{s}$ is

$$\Pi(\bar{s}) = 0.5Y(\bar{s}) + 0.5E_{\bar{s}' \in C(F)} Y(S - \bar{s}'), \tag{3.7}$$

where $C(F)$ is the support of the distribution F, and E is the expectations operator, and Y is the indirect utility function, which maps a player's endowment share into the resultant material payoff accounting for the merchant's optimal play. In particular, a player's indirect utility function Y of an endowment share $\tilde{s}$ (with the other player's share being $S - \tilde{s}$) is defined as

$$Y(\tilde{s}) = \begin{cases} U^0(\tilde{s}) & \text{if} \quad \tilde{s} \leq S/2, \\ U^0(\tilde{s}) + U(q_2, \tilde{s}, p_2) & \text{if} \quad S/2 < \tilde{s} \leq S - s^*, \\ U^0(\tilde{s}) & \text{if} \quad S - s^* < \tilde{s} \leq S, \end{cases} \tag{3.8}$$

where (p_2, q_2) is the price-quality allocation aimed at the richer player (farmer) as defined above (which is itself a function of an endowment share $\tilde{s}$); and s^* is the threshold endowment share as in (3.6) that provides the condition when both farmers are served. Assuming that function Y is increasing in $\tilde{s} \in (S/2, S - s^*]$[10], it takes the form as depicted in Figure 1 with a discontinuity at the point $S - s^*$ (the left- and right-hand limits of Y at $S - s^*$ are not equal because $U(q_2, S - s^*, p_2)$ is strictly positive).

From the definition of the function (3.8), we see that there is an upward shift $U(q_2, \tilde{s}, p_2)$ (which is also discontinuous at the right) in the values of indirect utility function Y over $(S/2, S - s^*]$, which, otherwise, takes the form of reservation utility $U^0(\tilde{s})$ only. As discussed above,

[10]The assumed monotonicity of indirect utility function Y over $(S/2, S - s^*]$ is, in fact, dependent on its own functional form as well as the merchant's production cost function C; however, this assumption has no impact on the main argument that follows and is made for expositional clarity to simplify the determination of the maximizer of Y over $(S/2, S - s^*]$.

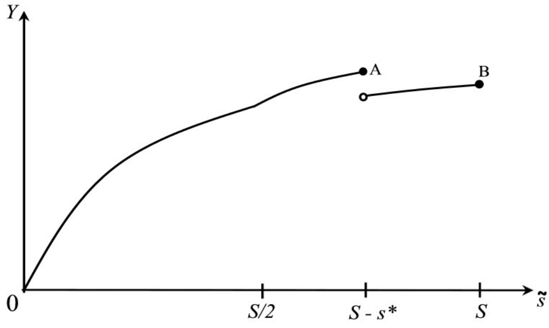

Figure 1. The indirect utility (material payoff) function Y

when inequality in endowment distribution is not too sharp and because of which the merchant finds it optimal to serve both farmers, the richer farmer enjoys the information rent of $U(q_2, \tilde{s}, p_2)$, which is, otherwise, fully extracted by the merchant. Since the value of function Y drops after $\tilde{s} = S - s^*$ (point A in Figure 1), which occurs when the merchant optimally shuts down on the poorer farmer, it is not straightforward to determine where function Y achieves its global maximum: at $\tilde{s} = S - s^*$ (point A) or at $\tilde{s} = S$ (point B). In other words, it is not obvious from the material payoff perspective whether the dictator should keep all the endowment for herself (and maximize her reservation utility U^0) or give away the share s^* to the other farmer (and enjoy some information rent). It depends on the size of information rent, which, on the other hand, is dependent on the form of the utility functions. Formally, if

$$U^0(S - s^*) + U(q_2, S - s^*, p_2) \geq U^0(S),$$

or

$$U^G(q_2, S - s^* - p_2) \geq U^0(S), \tag{3.9}$$

where (p_2, q_2) is the allocation aimed at the richer farmer at the endowment split $(S - s^*, s^*)$, then the dictator attains the highest material payoff when she shares the endowment with the other farmer (by giving the latter s^*). Intuitively, it is to require that farmers after a certain point become quickly satiated with the consumption of their own endowment (which is their land's produce) and value the outside good highly enough.

Returning to evolutionary fitness expression (3.7), we see that if condition (3.9) holds, the farmers with the preference type $\bar{s} = S - s^*$ acquire the highest expected material payoff. In other words, when the farmers are exposed to external trades run by a monopolist merchant, inequity-averse preferences may eventually be favored by natural selection, which will counter the merchant's monopoly power. All in all, the equilibrium split in game Γ^M is $(S - s^*, s^*)$ when condition (3.9) holds; otherwise, it is $(S, 0)$.

Case 3: Perfect competition

Consider game Γ^C, where there are many competing merchants on the external market. Given that all merchants' profits have to be equal to 0 in perfect competition, the competitive solution to a merchant's problem is easy to describe. The level of quality offered has to be as in the first-best case, while the price has to be equal to the total cost of producing that particular quality. Therefore, the price-quality allocation (p_j, q_j) aimed at a farmer with an endowment level $\tilde{s}_j$ is determined by (3.2) for the quality q_j, and the price $p_j = C(q_j)$, which, unlike in the monopoly case, is not a function of the endowment share $\tilde{s}_j$. Following the same logic as before, the selfish farmers with $\bar{s} = S$ attain the highest material payoff in game Γ^C, and this type of preference should eventually survive evolutionary pressures.

Summary

The proposition below summarizes the resultant evolutionary stable preferences for the environments analyzed.

Proposition 3.1 *In the two-player dictator game with consumption studied above, the evolutionary stable preference types* $\bar{s}_{es}$ *with respect to endowment distribution are*

- *in autarky (no external trades) —* $\bar{s}_{es} = S$*;*
- *with external trades run by a monopolist merchant — if condition (3.9) holds, then* $\bar{s}_{es} = S - s^*$*, where* s^* *is defined as in (3.6); otherwise,* $\bar{s}_{es} = S$*; and*
- *with external trades run by competitive merchants —* $\bar{s}_{es} = S$*.*

3. Discussion

Endogeneity of inequity aversion

The main result of the above proposition is that external factors, such as market exchange and market structure, can have an influence on people's behavior and the shape of their preferences. In particular, besides genuinely altruistic considerations for other people (frequently adhered to when explaining experimental evidence on people's behavior, e.g., Levine (1998)), people may also develop inequity aversion toward money distribution as an optimal response to merchants' monopolistic powers on the merchandise markets. In other words, people's preferences, as revealed by their behavior, may be endogenous to the environment they live in.

If we take the approach that behavioral traits are transmitted culturally rather than genetically (Cavalli-Sforza & Feldman (1981)), then people's revealed preferences may not prove conclusive of their true nature. To put it differently, the results obtained here can be interpreted as: *Even if* people are intrinsically selfish (with regard to their own consumption), they can still exhibit behavior as if they had inequity-averse preferences for money distribution (since it is own-consumption maximizing).

In the evolutionary literature, the group-selection argument is typically used to show how pro-social preferences can survive evolutionary pressures (see, e.g., Bergstrom (2002)). Unlike in standard models, our argument in favor of pro-social preferences does not hinge on the group-selection idea. Our explanation does not require that pro-social types of people, when matched with people like them, receive high enough payoffs to offset lower payoffs when matched with selfish types, thus, making their expected evolutionary fitness greater than that of the perfectly selfish types. Instead, our argument hinges on the idea that in order to subdue a third party's adverse impact people choose actions that, even if they seem to be to the overall benefit of their group, are, in fact, own-utility maximizing.

Needless to say, within the framework studied, other forms of evolutionary stable preferences could emerge depending on the degree of competition or monopoly on the market. A more general prediction would be that the more monopolist markets are, the less selfish people's preferences should be, offering an empirically testable hypothesis on the link between market structure and people's preferences, which is left for future research. At the same time, there is empirical evidence in favor of the model, which is described in the following subsection.

Empirical evidence

In empirical studies, it has been demonstrated that people's preferences, revealed by their exhibited behavior in sharing money, are not uniform across different societies and are rather shaped by socioeconomic and cultural factors (see, e.g., Henrich et al. (2001); Henrich et al. (2004); Buchan et al. (2002)). As already noted in the introduction, the most compelling evidence in support of our theoretical findings comes from the empirical project conducted in 15 small-scale traditional societies, which is documented in Henrich et al. (2001) and Henrich et al. (2004). The aim of this project was to look into the foundations of human sociality and its origins through studying small-scale societies, which could possibly shed light on the evolutionary transition of modern people's preferences (which are, actually, less diverse, see Roth et al. (1991)).

Henrich et al. (2001) found that people's preferences, revealed in playing the ultimatum, public good, and dictator games, differ across different groups, and that there are certain regularities in the documented differences. One of the regularities, relevant to our model, is that members of an isolated society behave less pro-socially than do members of a market-integrated society (as measured, primarily, by exposure to external market exchange). Henrich et al. (2001), however, provided no theoretical explanation for this important empirical finding, which is nonetheless fully in line with our theoretical predictions, presented in Proposition 3.1. Our explanation for this finding is that when people are exposed to external trades with merchants (who typically possess or collude to have some monopoly power), they are better off by sharing with others since it overcomes merchants' full-rent extraction.

Certainly, to make this explanation more credible one would need to more closely explore the relationship between market structure and preferences, for instance, as hypothesized in the previous subsection, which is left for further research.

Model extensions and research directions

Within a given society, the distribution of people's preferences is more diverse than just one type of preferences (see, e.g., Fehr & Schmidt (1999)). In our model—in particular, for game Γ^M with condition (3.9) being met—to achieve a non-trivial distribution of evolutionary stable preferences, we could elaborate by introducing a noisy signal that merchants receive about the income distribution resulting from an endowment split and depending on which they design consumption bundles. Then, due to the noisiness of a signal, there would be no single type of subjective preference that would be own-consumption maximizing for any realization of a signal. Instead, a different type of subjective preference would be evolutionary-fitness maximizing for a different realization of a signal, leading, eventually, to a more diverse distribution of evolutionary stable preferences. Similarly, we could subject the structure of the external market to different competitive shocks, which would also lead to a more diverse distribution of evolutionary stable preferences.

In a similar fashion, we can think of other mechanisms affecting the form of revealed (inequity-averse) preferences. For instance, within our model, consider the effect on people's (optimal) behavior after the introduction of a uniform sales tax on the outside good. If the public authority aims to maximize tax revenues, then the model would predict people responding to the tax by reducing inequality in wealth on the grounds similar to the case with a monopolist merchant. On the other hand, if the tax imposed by the public authority is negligible, then it would not have any impact on people's behavior. In other words, the importance of the government's role in the economy can also shape the appearance of people's preferences, with its more central role adding to inequity aversion among people.

An interesting extension would also be to consider an N–player dictator game, where the number of farmers, matched to play the dictator game, is larger than two (similarly to the framework in Ok & Vega-Redondo (2001)). Qualitatively, it should not change the results: in certain cases, inequity-averse preferences should still render the highest material payoff. Interestingly, in game Γ^M it may not be optimal (from the material payoff perspective) for the dictator to split the remaining endowment evenly among the rest of the players, provided he finds it optimal to give away some of the endowment. Instead, the dictator can do better by dividing the remaining endowment unevenly as it can be seen from the special case of $N = 3$ and $V(q, x) = qx$, which at the same time poses an interesting question of what is the optimal income distribution from the dictator's perspective in the game with more than two players.

4. Concluding remarks

We have argued that the inequity-averse preferences of the type documented in laboratory experiments may be a product of natural or cultural selection. It has been shown that inequity aversion to money distribution can be developed as an optimal response to the surrounding socioeconomic environment—monopolistic merchandise markets, for instance.

The findings presented herein can be thought of as an attempt to reconcile experimental evidence on people's (inequity-averse) preferences with the conventional modeling assumption about own-regarding preferences, which, as this study shows, can be consistent with each other. Nonetheless, the question of what preferences should be used in economic modeling remains open. In a partial equilibrium analysis, for instance, using the form of preference that is most characteristic of a given society or group of people (as, say, revealed with the help of experiments) would probably render more accurate predictions than would using the conventional assumption of own-regarding preferences. However, if a modeled policy change may have a substantial general equilibrium reach, then, along the lines of the model presented, it may also affect people's preferences through the social transmission of behavioral traits, complicating predictions of the modeled change in the longer run.

To make the results of this study more credible, more empirical research needs to be done on the interdependence between people's preferences and the environment they live in. In particular, one could examine the link between market concentration and people's preferences hypothesized here by regressing a market concentration index on a measure of inequity aversion across different countries.

Bibliography

Becker, Gary S. 1976. Altruism, Egoism, and Genetic Fitness: Economics and Sociobiology. *Journal of Economic Literature*, 14(3), 817–826.

Bergstrom, Theodore C. 2002. Evolution of Social Behavior: Individual and Group Selection. *Journal of Economic Perspectives*, 16(2), 67–88.

Bester, Helmut, & Güth, Werner. 1998. Is altruism evolutionary stable? *Journal of Economic Behavior & Organization*, 34(2), 193–209.

Binmore, Ken. 1994. *Game Theory and the Social Contract –Volume I: Playing Fair*. MIT Press.

Binmore, Ken. 1998. *Game Theory and the Social Contract –Volume II: Just Playing*. MIT Press.

Bolton, Gary E, & Ockenfels, Axel. 2000. ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review*, 90(1), 166–193.

Bowles, Samuel. 1998. Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions. *Journal of Economic Literature*, 36(1), 75–111.

Buchan, Nancy R., Croson, Rachel T. A., & Dawes, Robyn M. 2002. Swift Neighbors and Persistent Strangers: A Cross-Cultural Investigation of Trust and Reciprocity in Social Exchange. *American Journal of Sociology*, 108(1), 168–206.

Cavalli-Sforza, Luigi Luca, & Feldman, Marcus W. 1981. *Cultural Transmission and Evolution: A Quantitative Approach*. Princeton, N.J.: Princeton University Press.

Charness, Gary, & Rabin, Matthew. 2002. Understanding Social Preferences with Simple Tests. *Quarterly Journal of Economics*, 117(3), 817–869.

Dekel, Eddie, Ely, Jeffrey C., & Yilankaya, Okan. 2007. Evolution of Preferences. *Review of Economic Studies*, 74(3), 685–704.

Ely, Jeffrey C., & Yilankaya, Okan. 2001. Nash Equilibrium and the Evolution of Preferences. *Journal of Economic Theory*, 97(2), 255–272.

Fehr, Ernst, & Schmidt, Klaus M. 1999. A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics*, 114(3), 817–868.

Fehr, Ernst, & Schmidt, Klaus M. 2006. The Economics of Fairness, Reciprocity and Altruism—Experimental Evidence and New Theories. *Pages 615–691 of*: Kolm, Serge Christophe, & Ythier, Jean Mercier (eds), *Handbook on the Economics of Giving, Reciprocity and Altruism*, vol. 1. Elsevier.

Güth, Werner, & Napel, Stefan. 2006. Inequality Aversion in a Variety of Games—An Indirect Evolutionary Analysis. *Economic Journal*, 116(514), 1037–1056.

Güth, Werner, & Yaari, Menahem. 1992. An Evolutionary Approach to Explain Reciprocal Behavior in a Simple Strategic Game. *Pages 23–34 of*: Witt, Ulrich (ed), *Explaining Process and Change— Approaches to Evolutionary Economics*. The University of Michigan Press, Ann Arbor.

Henrich, Joseph, Boyd, Robert, Bowles, Samuel, Camerer, Colin, Fehr, Ernst, Gintis, Herbert, & McElreath, Richard. 2001. In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies. *American Economic Review*, 91(2), 73–78.

Henrich, Joseph, Boyd, Robert, Bowles, Samuel, Camerer, Colin, Fehr, Ernst, & Gintis, Herbert (eds). 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford University Press.

Huck, Steffen, & Oechssler, Jörg. 1999. The Indirect Evolutionary Approach to Explaining Fair Allocations. *Games and Economic Behavior*, 28(1), 13–24.

Levine, David K. 1998. Modeling Altruism and Spitefulness in Experiments. *Review of Economic Dynamics*, 1, 593–622.

Mas-Colell, Andreu, Whinston, Michael D., & Green, Jerry R. 1995. *Microeconomic theory*. Oxford University Press.

Mussa, Michael, & Rosen, Sherwin. 1978. Monopoly and product quality. *Journal of Economic Theory*, 18(2), 301–317.

Ok, Efe A., & Vega-Redondo, Fernando. 2001. On the Evolution of Individualistic Preferences: An Incomplete Information Scenario. *Journal of Economic Theory*, 97(2), 231–254.

Roth, Alvin E., Prasnikar, Vesna, Okuno-Fujiwara, Masahiro, & Zamir, Shmuel. 1991. Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study. *American Economic Review*, 81(5), 1068–1095.

Weibull, Jörgen. 1995. *Evolutionary Game Theory*. MIT Press.

EFI, The Economic Research Institute

Published in the language indicated by the title.

A complete publication list can be found at *www.hhs.se/efi*

Books and dissertations can be ordered from EFI via e-mail: EFI.Publications@hhs.se

Reports since 2006

2009

Books

Engwall, Lars. *Mercury meets Minerva: business studies and higher education: the Swedish case.*

Henriksson, Lars. *Marknad eller reglering?: vägval för europeisk telekommunikation.* Forskning i Fickformat.

Dissertations

Arbin, Katarina. *Individual information system acceptance behavior: an electronic ordering system case.*

Hellström, Katerina. *Financial accounting quality in a European transition economy: the case of the Czech republic.*

Hernant, Mikael. *Profitability performance of supermarkets: the effects of scale of operation, local market conditions, and conduct on the economic performance of supermarkets.*

Lakomaa, Erik. *The economic psychology of the welfare state.*

Melander, Ola. *Empirical essays on macro-financial linkages.*

Melén, Sara. *New insights on the internationalisation process of SMEs: a study of foreign market knowledge development.*

Törn, Fredrik. *Challenging consistency: effects of brand-incongruent communications.*

Wennberg, Karl. *Entrepreneurial Exit.*

Wetter, Erik. *Patterns of performance in new firms: estimating the effects of absorptive capacity.*

2008

Books

Breman, Anna. *Forskning om filantropi. Varför skänker vi bort pengar?* Forskning i Fickformat.

Einarsson, Torbjörn. *Medlemskapet i den svenska idrottsrörelsen: En studie av medlemmar i fyra idrottsföreningar.* EFI Civil Society Reports.

Helgesson, Claes-Fredrik and Hans Winberg (eds). *Detta borde vårddebatten handla om.*

Jennergren, Peter, Johnny Lind, Walter Schuster and Kenth Skogsvik (eds). *Redovisning i fokus.* EFI:s Årsbok 2008. EFI/Studentlitteratur.

Kraus, Kalle. *Sven eller pengarna? Styrningsdilemman i äldrevården.* Forskning i Fickformat.

Petrelius Karlberg, Pernilla. *Vd under press: om medialiseringen av näringslivets ledare.* Forskning i Fickformat.

Portnoff, Linda. *Musikbranschens styrningsproblematik.* Forskning i Fickformat.

Sjöstrand, Sven-Erik. *Management: från kontorsteknik till lednings- och organisationsteori: utvecklingen på Handelshögskolan under 100 år: 1909-2009.*

Östman, Lars. *Den finansiella styrningens realiteter och fiktioner: de finansiella styrformernas svenska historia, berättelser om Petersson och "Ericsson", finansiell styrning – en ansats till generell teori.*

Östman, Lars. *Mycket hände på vägen från Buchhaltung till Accounting: delar av Handelshögskolan under 100 år.*

Dissertations

Axelson, Mattias. *Enabling knowledge communication between companies: the role of integration mechanisms in product development collaborations.*

Benson, Ilinca. *Organisering av övergångar på arbetsmarknaden: en studie av omställningsprogram.*

Elhouar, Mikael. *Essays on interest rate theory.*

Farooqi Lind, Raana. *On capital structure and debt placement in Swedish companies.*

Granström, Ola. *Aid, drugs, and informality: essays in empirical economics.*

Hvenmark, Johan. *Reconsidering membership: a study of individual members' formal affiliation with democratically governed federations.*

Höglin, Erik. *Inequality in the labor market: insurance, unions, and discrimination.*

Johansson, Marjana. *Engaging resources for cultural events: a performative view.*

Kallenberg, Kristian. Business at risk. *Four studies on operational risk anagement.*

Kviselius, Niklas Z. *Trust-building and communication in SME internationalization: a study of Swedish-Japanese business relations.*

Landberg, Anders. *New venture creation: resistance, coping and energy.*

Pemer, Frida. *Framgång eller fiasko? En studie av hur konsultprojekt värderas i klientorganisationer.*

Rosengren, Sara. *Facing clutter: on message competition in marketing communication.*

Schilling, Annika. *Kan konsulter fusionera?: en studie av betydelsen av identitet vid en fusion mellan konsultföretag.*

Schriber, Svante. *Ledning av synergirealisering i fusioner och förvärv.*

Sjödin, Henrik. *Tensions of extensions: adverse effects of brand extension within consumer relationship.*

Strandqvist, Kristoffer. *Kritiska år: formativa moment för den svenska flygplansindustrin 1944–1951.*

Strömqvist, Maria. *Hedge funds and international capital flow.*

Söderström, Johan. *Empirical studies in market efficiency.*

Sölvell, Ingela. *Formalization in high-technology ventures.*

Thorsell, Håkan. *The pricing of corporate bonds and determinants of financial structure.*

Ulbrich, Frank. *The adoption of IT-enabled management ideas: insights from shared services in government agencies.*

Östling, Robert. *Bounded rationality and endogenous preferences.*

2007

Books

Andersson, Per, Ulf Essler and Bertil Thorngren (eds). *Beyond mobility.* EFI Yearbook 2007. EFI/Studentlitteratur.

Einarsson, Torbjörn and Filip Wijkström. *Analysmodell för sektorsöverskridande statistik: fallet vård och omsorg.* EFI Civil Society Reports.

Ericsson, Daniel. *Musikmysteriet: organiserade stämningar och motstämningar.*

Samuelson, Lennart (ed). *Bönder och bolsjeviker: den ryska landsbygdens historia 1902–1939.*

Dissertations

Ahlersten, Krister. *Empirical asset pricing and investment strategies.*

Alexius, Susanna. *Regelmotståndarna: om konsten att undkomma regler.*

Andersson, Magnus. Essays in empirical finance.

Berg, Bengt Åke. *Volatility, integration and grain bank: studies in harvests, rye prices and institutional development of the parish magasins in Sweden in the 18th and 19th centuries.*

Bianchi, Milo. *Of speculators, migrants and entrepreneurs: essays on the economics of trying your fortune.*

Brodin, Karolina. *Consuming the commercial break: an ethnographic study of the potential audiences for television advertising.*

Elger, Max. *Three essays on investment-specific technical change.*

Hagberg, Axel. *Bankkrishantering: aktörer, marknad och stat.*

Hinnerich, Mia. *Derivatives pricing and term structure modeling.*

Hjalmarson, Hanna. *En växande marknad: studie av nöjdheten med konsumtionsrelaterade livsområden bland unga konsumenter.*

Hjelström, Tomas. *The closed-end investment company premium puzzle: model development and empirical tests on Swedish and British data.*

Kraus, Kalle. *Sven, inter-organisational relationships and control: a case study of domestic care of the elderly.*

Lindqvist, Erik. *Essays on privatization, identity, and political polarization.*

Macquet, Monica. *Partnerskap för hållbar utveckling: systrar av Oikos och guvernanten som blev diplomat.*

Melian, Catharina. *Progressive open source.*

Nilsson, Daniel. *Transactions in cyberspace: the continued use of Internet banking.*

Petrelius Karlberg, Pernilla. *Den medialiserade direktören.*

Portnoff, Linda. *Control, cultural production and consumption: theoretical perspectives, empirical dilemmas, and Swedish music industry practices.*

Sköld, Martin. *Synergirealisering: realisering av produktsynergier efter företagssammanslagningar.*

Sonnerby, Per. *Contract-theoretic analyses of consultants and trade unions.*

Tyrefors, Björn. *Institutions, policy and quasi-experimental evidence.*

Valiente, Pablo. *Re-innovating the existing: a study of wireless IS capabilities to support mobile workforces.*

2006

Books

Lundeberg, Mats, Pär Mårtensson and Magnus Mähring (eds) *IT & business performance: a dynamic relationship.* EFI Yearbook 2006. EFI / Studentlitteratur.

Thodenius, Björn. *Organisering av kunskap: en studie av Wallenberg Consortium North.* EFI Civil Society Reports.

Wijkström, Filip and Torbjörn Einarsson. *Från nationalstat till näringsliv?: det civila samhällets organisationsliv i förändring.*

Wijkström, Filip, Stefan Einarsson and Ola Larsson. *Staten och det civila samhället: idétraditioner och tankemodeller i den statliga bidragsgivningen till ideella organisationer.*

Östman, Lars. *Lysande ögonblick och finansiella kriser: Dramaten under ett sekel.*

Dissertations

Argenton, Cedric. *Quality provision in duopoly.*

Beckerman, Carina. *The clinical eye: constructiong and computerizing an anesthesia patient record.*

Borglund, Tommy. *Aktievärden i fokus: internationell påverkan på intressentrelationer vid förvärv och fusion.*

Breman, Anna. *The Economics of altruism, paternalism and self-control.*

Edquist, Harald. *Technological breakthroughs and productivity growth.*

Eklund, Jana. *Essays on forecasting and bayesian model averaging.*

Frostenson, Magnus. *Legitimitetskontrollen: en studie av etiska värderingars roll i gränsöverskridande förvärv och fusioner.*

Gaspar, Raquel M. *Credit risk and forward price models.*

Gustafsson, Peter. *Essays on trade and technological change.*

Hopkins, Elisabeth. *Is a higher degree of local currency pricing associated with lower exchange rate pass-through?: a study of import pricing in 51 Swedish industries.*

Kling, Ragnar. *Developing product development in times of brutal change.*

Langenskiöld, Sophie. *Peer influence on smoking: causation or correlation?*

Lychnell, Lars-Olof. *"Och fungerar det inte, gör vi på något annat sätt": en klinisk fallstudie av IT-relaterat förändringsarbete i småföretag.*

Meitz, Mika. *Five contributions to econometric theory and the econometrics of ultra-high-frequency data.*

Mendicino, Caterina. *Financial market imperfections, business cycle fluctuations and economic growth.*

Ovanfors, Anna. *Essays on nonlinear time series analysis and health economics.*

Paltseva, Elena. *Essays on commitment and inefficiency in political economy.*

Rogberg, Martin. *Den modeföljande organisationen: om acceptansen av TQM och andra populära managementmodeller.*

Silvennoinen, Annastiina. *Essays on autoregressive conditional heteroskedasticity.*

Sjögren, Ebba. *Reasonable drugs: making decisions with ambiguous knowledge.*

Slinko, Irina. *Essays in option pricing and interest rate models.*

Wilander, Fredrik. *Essays on exchange rates and prices.*