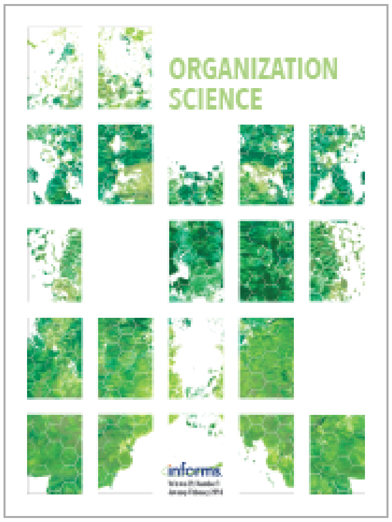

Organization Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Context and Aggregation: An Experimental Study of Bias and Discrimination in Organizational Decisions

Michael Christensen, Christian M. Dahl, Thorbjørn Knudsen, Massimo Warglien

informs
http://pubsonline.informs.org/journal/orsc

ORGANIZATION SCIENCE
Articles in Advance, pp. 1–19
ISSN 1047-7039 (print), ISSN 1526-5455 (online)

Context and Aggregation: An Experimental Study of Bias and Discrimination in Organizational Decisions

Michael Christensen,[a] Christian M. Dahl,[b] Thorbjørn Knudsen,[a,c,d] Massimo Warglien[e]

[a] Strategic Organization Design Unit, University of Southern Denmark, DK-5230 Odense M, Denmark; [b] Department of Economics, University of Southern Denmark, DK-5230 Odense M, Denmark; [c] Frankfurt School of Finance and Management; [d] Danish Institute for Advanced Study, University of Southern Denmark, DK-5230 Odense M, Denmark; [e] Center for Experimental Research in Management and Economics, Department of Management, Ca' Foscari University of Venice, 30123 Venice, Italy

Contact: chm@sam.sdu.dk, https://orcid.org/0000-0002-0299-1237 (MC); cmd@sam.sdu.dk, https://orcid.org/0000-0001-9880-2818 (CMD); tok@sam.sdu.dk, https://orcid.org/0000-0003-2798-7485 (TK); warglien@unive.it, https://orcid.org/0000-0002-0615-4938 (MW)

Received: September 15, 2019
Revised: June 30, 2020; January 15, 2021; April 30, 2021
Accepted: May 16, 2021
Published Online in Articles in Advance: October 21, 2021

https://doi.org/10.1287/orsc.2021.1502

Abstract. This paper addresses a notable gap at the intersection of organizational economics and organization science: how does organizational context influence aggregation of individual behavior in organizational decisions? Using basic centralized versus decentralized organizational structures as building blocks for our experimental design, we examine whether assignment of organizational positions, incentive schemes, and structural configuration induce endogenous adaptation in the form of change in reservation levels (bias) or modified discrimination capability in subjects' behavior. We found that evaluators adapted their reservation and discrimination levels in centralized structures, whereas they did *not* generally adapt their reservation and discrimination levels when placed in decentralized structures. We identify mechanisms that explain these findings; explain how they influence aggregate, organizational behavior; and discuss implications for research and practice.

History: This paper has been accepted for the *Organization Science* Special Issue on Experiments in Organizational Theory.

Funding: M. Warglien acknowledges financial support from the European Union's Horizon 2020 research and innovation programme under [Grant Agreement 732942].

Supplemental Material: The online appendix is available at https://doi.org/10.1287/orsc.2021.1502.

Keywords: decision making • organizational behavior • organizational design • laboratory research • experimental designs • research design and methods

Introduction

Organizational economics and organization science offer competing as well as complementary logics and methods to understand the nature, design, and performance of organizations (Levinthal 2011, Gavetti et al. 2012, Gibbons and Roberts 2013). Although each has contributed enormous literatures, there has been limited cross-talk between the two "camps." Drawing on Cyert and March (1963), a common source of inspiration for both camps, this paper addresses a notable gap at their intersection: understanding the nature and sources of predictable (statistical) bias and imperfect discrimination in organizations.

In this regard, bias is a technical concept, defined in terms of a deviation from the best estimate of a threshold above which an organization's net gains are positive. There is a subtle relation between individual managers' bias and aggregate, organizational-level bias: if managers are *unbiased*, most aggregation procedures generate *biased* outcomes (Sah and Stiglitz 1986, Christensen and Knudsen 2010). Consider a procedure to screen suppliers on a composite score on technical adequacy, safety policies, and legal compliance. If a single manager decides on the bid for contract, a supplier is more likely to win approval than if three such managers must unanimously approve. The unanimity rule generates an aggregate bias because it increases the threshold for approval relative to each manager's threshold. To ensure the aggregate decision is unbiased, the committee members must therefore lower their individual approval thresholds. More generally, most aggregation procedures introduce bias unless decision makers adjust behavior conditional on organizational context. Although this problem has been noted as critical for designing organizations with desired properties (Sah and Stiglitz 1986, Reitzig and Maciejovsky 2015, Puranam and Maciejovsky 2020), the literature is just beginning to catch up (Piezunka and Schilke 2021). The equally important problem of endogenous change in discrimination has received even less attention.[1]

In this paper, we examine the effect of organizational context on aggregate bias and discrimination. A vigorous stream of literature in organization science has, over the last two decades, made notable advances in understanding how aggregation of individual choices influences organizational behavior and performance (Christensen and Knudsen 2002, 2010, 2013, 2020; Knudsen and Levinthal 2007; Csaszar 2012; Csaszar and Eggers 2013; Reitzig and Maciejovsky 2015). The basic premise in this line of research is that the organization has a fundamental role in aggregating choice functions, which characterize the organizational members' ability to pass judgment. A choice function, also known as a screening function, maps an individual's assessment of the consequences of a choice onto an appropriate action, for example, screening applicants for credit (Christensen and Knudsen 2020), screening stocks for mutual fund investment (Csaszar 2012), or screening crowdsourced ideas for improving a firm's operations (Reitzig and Maciejovsky 2015). In these examples, screening involves (statistical) bias and imperfect (noisy) discrimination. Given the agents' inherent bias and ability to discriminate, organization design aims to maintain unbiased decisions (eliminate systemic error) and, at the same time, increase discrimination (reduce random error in choice).

Although prior research has made significant progress in understanding how organization design can improve organizational-level decisions, an important gap in our knowledge is whether individual behavior changes conditional on organizational context, that is, if agents in organizations adapt their bias and/or discrimination level. Csaszar (2012) provided the first large-scale empirical test of the fundamental dyadic decision structures presented in Sah and Stiglitz (1986), which showed that decentralized structures (polyarchies) relative to centralized structures (hierarchies) accept more projects, make fewer omission errors, and make more commission errors. Although this finding suggests that decision structures can be designed to predictably generate desired aggregate outcomes, it does not examine endogenous adaptation. However, recent papers have documented that organizational context may induce endogenous adaptation that alters aggregate outcomes in decision processes (Reitzig and Maciejovsky 2015, Christensen and Knudsen 2020, Piezunka and Schilke 2021). Taken together, these findings suggest that individuals may adjust behavior conditional on organizational context and thereby endogenously modify their screening functions. As a result, the aggregate organizational-level screening function may change in ways that are surprising and possibly undermines the intentions of the organization designer. Thus, an alternative approach, associated with organizational economics, emphasizes the importance of multiagent incentive problems and commonly invokes the assumption that organizational members deploy strategically sophisticated behaviors as they aim to further their individual interests (Gibbons and Roberts 2013). Although strategic sophistication can trigger shirking and free riding when motivations are misaligned, it may also improve agents' coordination when their interests are aligned. In this regard, Sah and Stiglitz (1986) suggested that managers may rationally adjust their behavior to counter the bias that an organizational structure may induce, for example, members of committees that require unanimous approval would lower their individual approval thresholds.

However, it is an open question if the organizational context—the configuration of organizational "pipes" or the provision of incentives—would induce changes in individual behavior that influences aggregate outcomes. What predictable screening biases will emerge as a result of varying the organizational structure? Do individuals rationally adapt their behavior to counter the bias induced by the structural configuration? Will different organizational structures prime different levels of sophistication in the behavior of organizational members? Does a manager's position in the decision-making process affect strategic behavior? Does provision of incentives influence the provision of effort? The purpose of this paper is to address these open questions on organizational bias and discrimination in organizational decisions

Using the basic Sah and Stiglitz (1986) structures—the hierarchy (H) and the polyarchy (P)—as building blocks for treatments in our experimental design, we examine whether assignment of organizational positions, incentive schemes, and the structural configuration of the organization induce endogenous adaptation in the form of bias or modified discrimination capability in subjects' behavior. These conditions jointly represent potential sources of adaptive behavior—for example, reduced effort or limited understanding of the organizational context the actor is operating in—that are commonly highlighted in the literature (e.g., Adner and Levinthal 2008, Gibbons and Roberts 2013). If present at the individual level, these conditions will likely induce bias and lower discrimination capability at the group and organizational level (Sah and Stiglitz 1988). To assess whether these effects are present in our experimental data, we contribute a statistical modeling framework that allows identification of bias and discrimination in organizational behavior at three levels: (1) aggregate organizational behavior; (2) endogenous behavior of pairs of individuals (groups), conditional on type of organization; and (3) endogenous behavior of individuals, conditional on the organizational position they occupy. The following sections detail the theoretical motivation for our study. We then present our experimental design, statistical model, results, and robustness tests. We conclude by discussing implications for theory and practice.

Theoretical Motivation

We draw our theoretical motivation from research on screening that, over the last two decades, has emerged as a major stream of research in organization science (Joseph and Gaba 2020). This literature originates in Christensen and Knudsen (2002, 2010) who extended the work of Sah and Stiglitz (1985, 1986, 1988) to analyze a broad array of organizational forms. The fundamental building blocks in this approach are two dyads: hierarchies and polyarchies, which are (respectively) stylized forms of centralized and decentralized organizational forms. From these two building blocks, any organizational form can be constructed and its performance can be analyzed using the approach in Christensen and Knudsen (2010). The motivating question in this line of research is as follows: "How do you structure decision rights in a group, such that they make fewer errors of omission and commission?" (Joseph and Gaba 2020, p. 273). This research question has motivated a growing body of studies (e.g., Csaszar 2012, 2013; Csaszar and Eggers 2013; Reitzig and Maciejovsky 2015); it is also motivating our approach, which we detail in the following.

The general aim in the design of decision-making structures is to construct an organizational form with the desired level of efficiency. The extent to which this aim can be achieved depends on (a) how capable the individual decision makers are, (b) how they apply these capabilities in an organizational context, and (c) how individual contributions are aggregated. In the present study, we limit our focus to Sah and Stiglitz's (1986) fundamental project selection model, which was based on two mutually exclusive assumptions. First, by assuming that the behavior of the individual evaluator does not change conditional on organizational context, they demonstrate that organizational structure has profound impact on organizational performance. Second, by assuming that agents are sensitive to organizational context via its impact on incentives, they demonstrate that endogenous adaptation of individual behavior may result in unbiased, optimal aggregate performance. Consistent with the first assumption, theoretical frameworks from the literature on organization science inform design of complex, efficient decision-making organizations (e.g., Gavetti et al. 2005, Gulati et al. 2005, Rivkin and Siggelkow 2007, Csaszar 2012). In contrast, organizational economics maps out optimal incentive schemes for simple (most often dyadic) structures under the second assumption (e.g., Gibbons and Roberts 2013). Here we develop a laboratory experiment to test both of these basic assumptions and their implied theoretical and practical consequences. To put our experiment into perspective, we briefly recap the original project selection model of Sah and Stiglitz (1986), focusing on the propositions that will inform our investigation.

Consider an organization that is screening a stream of projects to assess whether their net gain makes them worth undertaking, for example, screening applicants for credit, screening suppliers of raw materials, or screening proposals for investing in start-ups. Let x denote the quality of a project. The net benefit to the organization, here denoted as the value of the project $V(x)$, can be positive or negative depending on the quality. Let $\rho(x)$ denote the distribution of quality across possible projects. For each project, the organization must decide whether it will obtain an uncertain net value associated with committing to its realization or rather forgo the opportunity and realize a certain gain of zero. The decision is delegated via the organizational structure to the individual members of the organization, whose task it is to evaluate ("screen") the projects. Each evaluator will either assess the project value to be positive (a "good" project) or negative (a "bad" project), resulting in acceptance of the project (commitment), rejection of the project (omission), or passing the project on for further evaluation. Thus, communication among organizational members is constrained to delegation of decisions. If screenings were perfect, the evaluators would accept all projects of positive net value and reject all projects of negative net value. In that case, all evaluators would make identical decisions and these would be perfectly aligned with the final organizational-level decision. As it is human to err, let the "screening function" $p(x)$ be the probability that the individual evaluates project x to be good, that is, the evaluator believes the project is associated with a net gain. Intuitively, for very bad projects there is only a very small probability that the project will be evaluated as good, whereas for very good projects that probability is close to unity.[2] Two properties of this screening function, bias and discrimination, are central to the information processing perspective of organizational design. These terms have a particular and technical meaning in the project selection model. The bias is measured relative to the projects for which the individual is indifferent, which is conveniently represented by a reservation level on net value, that is, the perfect screener has a reservation level of zero, whereas the biased screener has a nonzero reservation level. The discrimination level captures sensitivity to changes in project value, which is represented by the slope of the screening function,[3] for example, the perfect screener approaches infinite discrimination ability. Any systematic or random errors associated with limited human ability, such as bias and discrimination, can, in principle, be eliminated at the organizational level by redesigning the way information is aggregated (Christensen and Knudsen 2010).

Figure 1 shows Sah and Stiglitz's (1986) archetypical organizational architectures, hierarchy and polyarchy. The hierarchy features centralized decision making and represents a bureaucracy-oriented organization, whereas the polyarchy features decentralized decision making and represents a market-like organization. Both organizing principles abound in structures and processes found in markets, institutions, and within firms (Csaszar and Eggers 2013, Reitzig and Maciejovsky 2015).

In the Sah and Stiglitz (1986) hierarchy, the evaluators have fixed positions and the projects always arrive at the same evaluator for a first screening. If this initial evaluator evaluates a project to have negative value, the project is rejected on behalf of the organization. Otherwise, the project is sent to the second evaluator for another screening, which ultimately decides the fate of the project. Thus, in a hierarchy, the evaluators operate according to a strict sequential order, which implies that both evaluators must accept a project before the organization commits to its realization. In many ways, the polyarchy operates in the exact opposite fashion. Each project arrives at a randomly appointed evaluator (50-50 chance), and this first evaluator has the authority to accept the project on behalf of the organization. Alternatively, if the evaluator finds that the project has negative value, it is sent to the second evaluator, who then decides the fate of the project. Thus, in Sah and Stiglitz's (1986) version of a polyarchy, the evaluators are equals and interchangeable and both evaluators must reject a project before the organization will pass on the opportunity to realize it (omission). The two structures are contrasted in Figure 1 together with the solitary agent (A) who is working unconstrained of organizational context.

The probability p^G that organization G accepts a project, that is, the organization's screening function, can be calculated from the individual screenings p_i and p_j of the two evaluators. Sah and Stiglitz's (1986) proposition 1 states that hierarchies will reject more projects than polyarchies. To make this intuitive, consider a solitary evaluator (A)–the number of projects to which he or she commits will be determined by his or her individual screening p_i. Adding a second layer of decision making in a hierarchical structure will imply that some projects that the first evaluator accepts will be rejected by the second evaluator (unless the screening capabilities of both agents are perfect). The opposite holds true for polyarchies. In both cases, the aggregation of individual contributions results in bias, that is, the organization's reservation level is shifted relative to the agent's reservation level. Figure 2 displays the effects of such organizational bias for both architectures—the intercept of the organization's curve shifts, which results in a change in errors of commission and omission relative to the individual agent.

Of course, evaluators may be aware of the potential bias induced by the decision-making structure and react to the implied loss of payoff. Indeed, the hierarchy may become too conservative in the sense that it omits

Figure 1. Project Flow of the Basic Sah and Stiglitz (1986) Hierarchies (H) and Polyarchies (P) Together with the Unitary Evaluator (A)

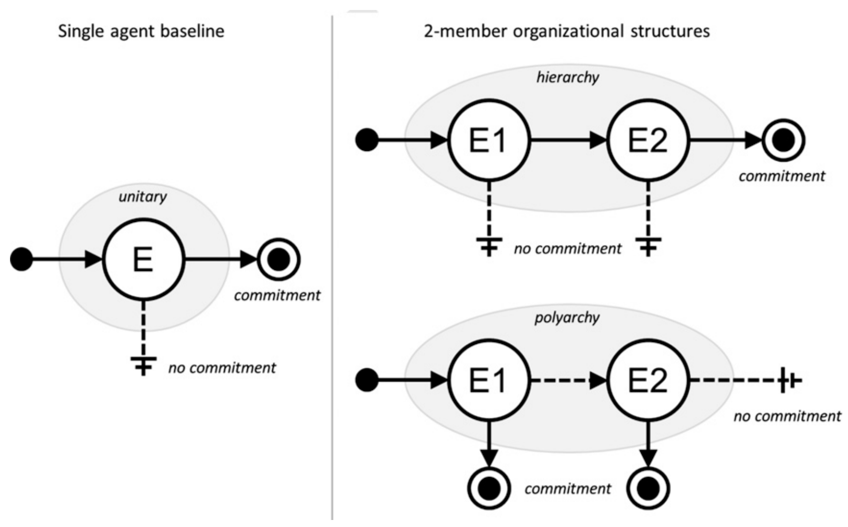

Figure 2. (Color online) Comparison of the Aggregate Screening Function of a Hierarchy and Polyarchy to that of the Unbiased Solitary Evaluator

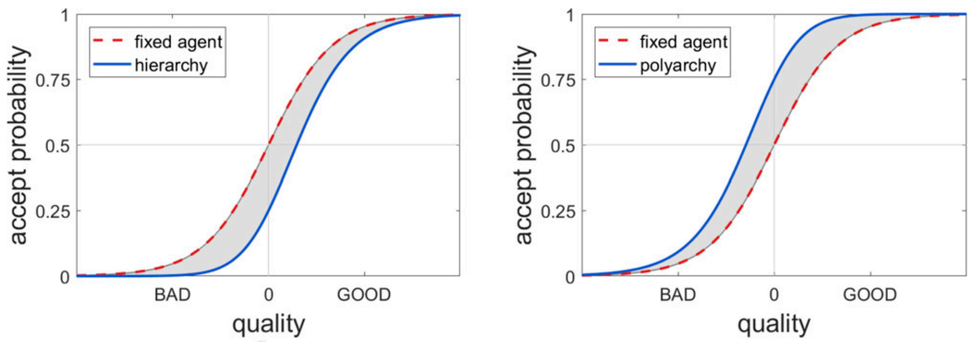

Notes. The example assumes that evaluators have immutable (fixed) screening functions. The grey area is the change in errors of commission and omission.

so many good projects that the evaluators' performance (and pay) will decline. In consequence, evaluators in a hierarchy will have an incentive to become less conservative, shifting their individual reservation levels, R, toward a lower value (to the left in Figure 2). Relative to the solitary agent operating on his or her own, this shift in reservation level induces a bias, that is, the individual becomes more permissive. The reverse argument goes for the polyarchy.

Sah and Stiglitz (1986) analyze such endogenous adaptation under several conditions. The hierarchy is constructed such that the evaluators operate under broad incentives, that is, the two evaluators share the value V upon realization of the project (commitment). By contrast, the polyarchy is constructed so that individuals operate under narrow incentives, that is, the evaluator accepting a project claims all value V, whereas the other evaluator receives a payoff of zero. In this sense, the polyarchy represents a stylized form of market-like organization. Because decentralized structures are also found inside bureaucratic organizations, Sah and Stiglitz (1986) introduced a third variant, the "coordinated polyarchy" (CP), which is simply a polyarchy with *shared* rather than individual payoffs. Sah and Stiglitz (1986) prove (in their proposition 4) that $R^{H_i} < R^{P_i} < R^{CP_i}$, that is, that evaluators, i, in hierarchies should always have the least conservative reservation level, whereas individuals in coordinated polyarchies should have the most conservative reservation level. Evaluators who are able to perform such endogenous adaptation of their reservation level will generate superior organizational performance relative to evaluators whose reservation levels are fixed once and for all (Figure A.1 in the online appendix illustrates).

Because organizational decision making is essentially a screening process or, in the words of Drucker (1967, p. 92), a risk-taking judgment, it is important to test the two core propositions of Sah and Stiglitz (1986) in an experimental setting. This is because a deeper understanding of the extent to which evaluators adapt their reservation levels and perhaps also the sharpness of their judgment (discrimination levels) will allow researchers and practitioners to develop superior organizational design. In this regard, Sah and Stiglitz's (1986) two core propositions capture a notable tension between two camps that have developed distinct approaches to organizational design. One camp, associated with research on organization science (Christensen and Knudsen 2010), team theory (Marschak and Radner 1972), and electronic engineering (Moore and Shannon 1956a, 1956b), has advanced models that provide insights on the way engineering of information flows in decision making can improve efficiency of decision structures. The assumption in these models is that incentives are aligned and that individual behavior is immutable. In contrast, models from the other camp, associated with organizational economics, have focused on how individual actors strategically adapt behavior in response to alternative mechanisms (e.g., allocation rules) and structural conditions (e.g., conflict of interest). Consistent with the idea that evaluators adapt behavior strategically, Sah and Stiglitz (1986) restrict their predictions of endogenous adaptive behavior to the symmetric Nash equilibrium. However, it is not given that evaluators behave consistent with this prediction, even in a simple laboratory setting. Advancing knowledge that can inform these broader issues further motivates the use of our laboratory experiment, which offers ideal conditions for identification of mechanisms that can account for human behavior.

Finally, Sah and Stiglitz's (1986) analysis is limited to investigating the potential effect of evaluators who adapt one parameter, the reservation level, R, which determines the *bias* of the screening function. The slope of the screening function (denoted β), which captures the evaluator's ability to *discriminate* quality, is considered as a given, immutable property of the evaluator.

However, the possibility that evaluators adapt their level of discrimination conditional on context is of obvious importance for the theory and practice of organization design. For example, difficulties relating to distributing feedback to organizational members may impede organizational learning or the provision of incentives may regulate the provision of cognitive effort. It is therefore an important open question whether the organizational structure, through its effect on decision processes, changes the evaluators' discrimination level (sharpness of screening).

In what follows, we provide an experimental design and a statistical modeling approach that provide answers to these open questions relating to endogenous adaptation of bias and discrimination of the evaluators' screening functions.

Experimental Design

We conducted an experiment with six different treatments. Each treatment corresponds to a particular decision-making structure, described by number of members, aggregation rule (consensus), incentives, and assignment of positions (fixed versus random). As shown in Table 1, the treatments include the original Sah and Stiglitz (1986) architectures H/P/CP and two additional hybrids, hierarchy with randomly varying positions (RH) and coordinated polyarchy with fixed positions (FCP), that were introduced to assess the effects of randomly varying the individual's position while keeping incentive structures constant. There is also an individual treatment A in which a subject makes a solo decision without being inserted into any decision-making structure. Table 1 illustrates the conceptual relations between the two-member organizations, and Table 2 summarizes the treatments. We present the approach and results of all treatments together. In the following, we detail the basic evaluation task, the layout of the experiment, and the different treatments.

Table 1. Overview of Two-Member Organizations Included in Our Experimental Design

		Assignment of individuals	
		Fixed positions	Random positions
Decision rule	*Unanimous accept, individual reject*	H	RH
	Unanimous reject, individual accept,	FCP	P[a], CP

Note. The solitary agent (A) is our baseline treatment. FCP, coordinated polyarchy, fixed positions; RH, hierarchy, random positions.

[a]Note that P has narrow incentives, whereas all other structures have broad incentives.

Evaluation Task

We designed a binary (good/bad) categorization task. To give all test subjects the same amount of ex ante experience with the task, we constructed a new task and allowed preliminary training instead of relying on common or preknown tasks. The task must allow a simple parametrization of projects and their values, and the task must also be relatively easy to learn. Inspired by the signal detection literature (Sorkin et al. 1998) we chose the task to screen two-dimensional geometric images whose "quality" can be determined by a linear function of two features, diameter and position. Each image has a fixed background consisting of a horizontal black line evenly dividing a yellow square surface. Onto this background, a red circle is placed with its center on the line. This red circle will randomly vary in size and horizontal position. The circle is generated by randomly drawing two independent and uniformly distributed real numbers $x_1, x_2 \in [-1, +1]$. The first variable, x_1, is mapped linearly onto the middle section of the line to represent the center (c) of the circle. The second variable, x_2, determining the diameter (d) of the circle, is mapped linearly to the interval [1/50,1/5] of the width of the background. More details on the mapping can be found in the online appendix. An example of the visual stimuli is presented in Figure 3.

Each subject evaluates a sequence of such randomly generated images, that is, an image serves the role of a proposal whose value must be assessed. Each image is assigned the quality $x = 4/9\, x_1 + 5/9\, x_2$, and the value function used for determining payoffs is simply $V(x) = sign(x)$. Thus, the goal for each subject is to accept whenever the combination of diameter and position generates a positive payoff. The functions determining x and $V(x)$ are not communicated to the test subjects, but they are informed that the value of each image does not depend on time or previous choices. For the two-member treatments, the organizational design aggregates individual choices into an organizational choice as will be described in the section below. In these treatments, acceptance/rejection by an individual organizational member signals whether this member is willing to commit to getting a nonzero payoff and thereby guides the organizational choice according to the aggregation rule of the organization. Of course, the organization may ultimately choose differently if the other organizational member chooses differently.

After each decision to accept/reject, the test subject is presented with a feedback screen displaying the payoff. Subjects are given feedback on the quality of the image (good/bad) if and only if the subject participated in the decision and received a nonzero payoff.

Table 2. Treatments and Effective Sample Sizes

Treatment	Samples	Description of organization
A	18	Solitary individual
H	16	Hierarchy of two members
P	15	Polyarchy of two members (individual payoffs)
CP	15	Coordinated polyarchy of two members (shared payoffs)
RH	15	Random-position hierarchy of two members
FCP	18	Fixed-position, coordinated polyarchy of two members

Note. The two structures below the line were added to the original Sah and Stiglitz (1986) structures to control for the effect of fixed versus randomly varying positions.

Treatments

Each treatment is divided into two stages, first a training stage and then a performance stage. Each stage is initiated with a set of instructions and a questionnaire to check the subjects' understanding of the instructions. The subjects perform the experiment separately in small cubicles, and they do not communicate except through observation of decision outcomes via the feedback screen. The different treatments refer to differences in organizing principles applied to the test subjects during the latter performance stage. The treatment labels, descriptions, and sample sizes are shown in Table 2.

Sample size was determined on the basis of an ex ante power analysis. We estimated the required sample size to achieve a power of 0.9 on proposition 1 of Sah and Stiglitz (1986), that is, that hierarchies reject more and that polyarchies accept more, respectively, than the solitary agent. The power analysis implied a sample size of 15. The details are reported in the online appendix.

The training stage is common to all treatments and consists of evaluating the value of 100 images (rounds). During the training stage, the subjects work in solitude and they are free to experiment and learn as the outcome of this stage does not contribute directly to the final payment. The number of images for the training stage was decided from prototype trials, indicating a sharp initial rise in hit rates during the first 15–50 decisions, followed by moderately increasing hit rates for the subsequent images.

Figure 3. (Color online) Screenshot of a Typical Image that Subjects Encountered

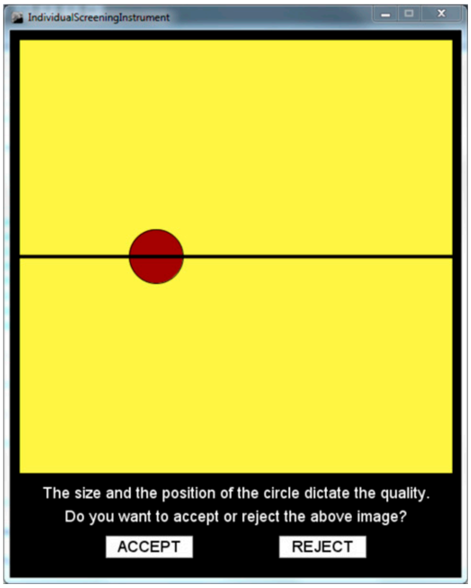

For the performance stage, each subject is assigned to a treatment, that is, to one of the six kinds of organization structures. The instructions for the performance stage of treatments A, H, and P are featured in the online appendix. Note that the test subjects are not informed about the position that they are assigned to within the organization (first/second evaluator) or whether the positions are fixed or randomly varying from image to image. This choice of opacity, apart from keeping instructions simpler, aligns the experiment with Sah and Stiglitz's (1986) focus on the symmetric equilibrium point. However, we cannot rule out that agents are able to infer their position during the experiment (e.g., by observing changes in the portfolio of proposals they receive or by comparing waiting times before they can act/after they receive feedback). This possibility will be checked in our statistical analysis.

The assignment of subjects to treatments is done randomly but within the constraints of a plan devised by the laboratory assistant prior to sign-in of subjects, always involving an optional instance of treatment A in case an odd number of test subjects would show up to the laboratory sessions. Similarly, gender balance is managed by admitting equal numbers of each gender into the experiment. For the two-member treatments (not A), the subjects are permanently and anonymously paired with a partner by the server software running the experiment.[4] These pairs (or individuals in treatment A) are the independent units of observation, henceforth simply denoted as pairs. The pairing is made at random. The performance stage consists of evaluating 200 proposals (images associated with a particular value), determined by balancing the need for the learning dynamics to reach a stable state against the need to limit session length to avoid that the subjects' concentration would falter as well as the more pragmatic concern of avoiding excessive costs (the subjects are paid a show-up fee that in conjunction with bonus payments equate local minimal hourly wages). During evaluation of these 200 proposals, test subjects accumulate payoffs in terms of a numerical score. The final score generates a bonus on top of the show-up fee (see the online appendix). Subjects are told that a higher score would yield a higher payment. A full session is scheduled to approximately an hour. Specifically, through our laboratory (Laboratory@SDU), we recruited 176 participants (50% female), who took part in sessions that averaged 65 minutes and who received an average of DKK 171.14 (about USD $25) for participation.

Statistical Model

In this section, we outline the statistical modelling framework and present the empirical estimation and hypothesis testing strategy. The approach we develop contributes to advance experimental work on organizations as it facilitates statistical testing across levels of analysis. Specifically, our approach allows a statistical test of differences between and within organizations: (1) differences in aggregate behavior of organizations; (2) differences in joint behavior of individuals conditional on the organization they are located in but ignoring differences in positions within the organization;[5] and (3) differences in behavior of individuals, conditional on the position they occupy when performing an evaluation in a given organization.

Assume that an evaluator observes a project and, on that basis, forms a subjective estimate of its value, y^*, given as

$$y^* = \alpha + \beta x + \sigma\epsilon,$$

where $\epsilon \sim i.i.d.$ with zero mean and unit variance, σ is the scale of the noise internal to the evaluator, and x denotes the actual quality of the project under evaluation. Projects are accepted if $y^* > R$ for some reservation level, R. If the evaluator accepts the project, we observe $y = +1$; if the project is not accepted, we observe $y = -1$. The probability that the evaluator accepts the project is then given as

$$P(y^* > R \mid x) = P(y = 1 \mid x) = P\left(\epsilon > \frac{R - \alpha}{\sigma} - \frac{\beta}{\sigma} x \middle| x\right).$$

By symmetry of the distribution of ϵ this simplifies to

$$P(y = 1 \mid x) = P\left(\epsilon < \frac{\alpha - R}{\sigma} + \frac{\beta}{\sigma} x \middle| x\right) = P\left(\epsilon < \tilde{R} + \tilde{\beta} x \mid x\right),$$

where $\tilde{R} = (\alpha - R)/\sigma$ and $\tilde{\beta} = \beta/\sigma$ uniquely determine the distribution of the evaluator's actions (choice behavior), even though we cannot fully estimate how the evaluator judged the projects (evaluation behavior) because of the indeterminate scale σ.

We apply a logit model to estimate all screening functions and we use a linear probability to check for robustness (see the online appendix). If it is assumed that the distribution of ϵ is independent and identically distributed for some symmetric noise distribution, the accept probability becomes

$$p(x) = P(y = 1 \mid x) = \Lambda\left(x; \tilde{\beta}, \tilde{R}\right) = \frac{1}{1 + \exp\left(-(\tilde{R} + \tilde{\beta} x)\right)}, \tag{1}$$

where the marginal effects of interest are given as

$$P_{\tilde{R}}(y = 1 \mid x) = \Lambda\left(x; \tilde{\beta}, \tilde{R}\right)\left(1 - \Lambda\left(x; \tilde{\beta}, \tilde{R}\right)\right),$$
$$P_x(y = 1 \mid x) = \Lambda\left(x; \tilde{\beta}, \tilde{R}\right)\left(1 - \Lambda\left(x; \tilde{\beta}, \tilde{R}\right)\right)\tilde{\beta}.$$

The coefficients of interest $(\tilde{\beta}, \tilde{R})$ are estimated jointly by maximum likelihood. The marginal effects $P_{\tilde{R}}(y = 1 \mid x)$ and $P_x(y = 1 \mid x)$ (slope of the screening function) can be computed as the average marginal effects (see, e.g., Greene 2018), and the associated standard errors are computed based on clustering by pairs.

Framework for Test of Individual Behavior

We use the following tests to study the effects of organizational structures on agents' evaluation behavior. Assume that we have two types of evaluators. They could (but need not) be a first evaluator, denoted by f, and a second evaluator, denoted by s. The evaluations for evaluator $i = f, s$ can then be summarized as

$$y^{*,i} = \alpha^i + \beta^i x + \sigma^i \epsilon^i,$$

with choices

$$y^i = +1 \text{ if } y^{*,i} > R^i \text{ and } y^i = -1 \text{ otherwise.}$$

Assume we want to test if the evaluators have the same evaluation behavior. This can be done by stacking $y = (y^f, y^s)$ and $x = (x^f, x^s)$ and by writing the estimation equation as

$$P(y = 1 \mid x) = \Lambda\left(x, I(s); \tilde{\beta}^f, \tilde{\beta}^s, \tilde{R}^f, \tilde{R}^s\right)$$
$$= \left(1 + \exp\left(-(\tilde{R}^f + \left(\tilde{R}^s - \tilde{R}^f\right) I(s) + \tilde{\beta}^f x + \left(\tilde{\beta}^s - \tilde{\beta}^f\right) I(s) x)\right)\right)^{-1},$$

where $I(s) = 1$ implies that the second type of evaluator has a task to perform and $I(s) = 0$ implies that second evaluator is not active. By defining $\Delta\tilde{R} = \tilde{R}^s - \tilde{R}^f$ and $\Delta\tilde{\beta} = \tilde{\beta}^s - \tilde{\beta}^f$, we can test the composite hypothesis

$$H_0: \Delta\tilde{R} = 0 \wedge \Delta\tilde{\beta} = 0,$$
$$H_1: \textit{not } H_0,$$

from the logit regression model

$$P(y=1\,|\,x) = \left(1+\exp\left(-(\tilde{R}^f + \Delta\tilde{R}I(s) + \tilde{\beta}^f x + \Delta\tilde{\beta}I(s)x)\right)\right)^{-1}.$$

It is important to stress that because it is not possible to identify the deep behavioral parameters of the model, that is, $(\alpha^f, \alpha^s, \beta^f, \beta^s, \sigma^f, \sigma^s, R^f, R^s)$, we cannot reject that the evaluators might have different evaluation behavior even when H_0 cannot be rejected, that is, the evaluators could change in ways that cancel out in the joint estimation. However, if H_0 is rejected, we can conclude that the evaluators do not have identical evaluation behavior. To control for potential heterogeneity/clustering, we include random effects, defined over pairs, in the statistical modelling framework when we perform the estimation. Importantly, this does not alter the testing framework (we are providing the details in the online appendix).

Hypotheses and Results

Our hypotheses, tests, and results encompass three levels: aggregate behavior of organizations; adjustment of behavior in pairs conditional on structure, assuming that both agents adjust in the same way (group adjustment); and individual organizational members' adjustment of behavior conditional on structure (individual adjustment). We address each level in turn.

Aggregate Behavior of Organizations

Per our theoretical framework, we test if organizational behavior in the aggregate is consistent with the predictions of Sah and Stiglitz (1986). Referring to the aggregate screening function, the prediction is that hierarchies are more conservative than the solitary agent (A) who is more conservative than polyarchies. This implies that, relative to the solitary agent (A), a hierarchy, in the aggregate, increases its reservation level, whereas a polyarchy, in the aggregate, reduces its reservation level, that is, theory predicts the following: $R^H > R^A$ and $R^A > R^P$. Because our statistical model identifies the quantity $\tilde{R} = (\alpha - R)/\sigma$, the sign in the estimated value of $\tilde{R}$ is *opposite* to the sign in the reservation level, R. In other words, if $R^H > R^A$, then $\tilde{R}^H < \tilde{R}^A$. Thus,

Hypothesis 1. *Hierarchies (H– denotes either H or RH) are more conservative than the solitary agent (A) who is more conservative than polyarchies (P– denotes either P, CP, FCP). This implies the following:*

Hypothesis 1.1. *Hierarchies (H, RH) have a more conservative reservation level than the solitary agent (A),* $\tilde{R}^{H-} < \tilde{R}^A$.

Hypothesis 1.2. *The solitary agent (A) has a more conservative reservation level than polyarchies (P, CP, FCP), i.e.,* $\tilde{R}^A < \tilde{R}^{P-}$.

Turning to the results, Table 3 shows descriptive statistics for the last 50 periods of the performance stage. Based on acceptance of 47% of the proposals and rejection of 53% of the proposals, the solitary agent (A) has an accuracy of 0.871, that is, the solitary agent makes a correct decision in about 87% of the cases. The way the solitary agent succeeds and fails is described by the fraction of bad proposals that are correctly rejected (recall) and the fraction of good proposals that are wrongly rejected (fallout), which translates into a fraction of type II (commission) and type I (omission) errors out of the total number of decisions.[6] In comparison with the solitary agent (A), the hierarchies (polyarchies) accept less (more) proposals and have higher (lower) recall and fallout. Even so, the *sum* of type I and II errors does not differ much between hierarchies and polyarchies, as reflected in the measure of accuracy.

Table 4 shows results based on the logit specification for each of the six treatments (see the Statistical Model section). In this, and all other analyses, we consistently use data from the last 50 periods of the performance stage of our experiment (out of a total of 200 time periods). This is because our experiment involved a learning task. During a session, test subjects improve performance (reducing type I and type II errors) with a decreasing rate toward a stable performance level.[7] Because we are testing theory with

Table 3. Descriptive Statistics at the Organizational Level Obtained Over the Last 50 Periods of the Performance Stage

	Accept rate		Recall (TPR)		Fallout (FPR)		Type I error rate		Type II error rate		Accuracy	
Treatment	est.	s.e.	est.	s.e.	est.	s.e.	est.	s.e.	est.	s.e.	est.	s.e.
A	0.473	0.022	0.886	0.022	0.150	0.021	0.073	0.010	0.056	0.010	0.871	0.014
H	0.372	0.023	0.927	0.015	0.294	0.039	0.139	0.018	0.038	0.008	0.824	0.020
P	0.603	0.020	0.756	0.030	0.040	0.009	0.020	0.005	0.123	0.017	0.857	0.015
CP	0.607	0.038	0.723	0.063	0.046	0.013	0.024	0.007	0.140	0.034	0.836	0.033
RH	0.392	0.026	0.924	0.016	0.262	0.041	0.127	0.022	0.040	0.009	0.833	0.022
FCP	0.636	0.021	0.692	0.037	0.063	0.014	0.034	0.008	0.151	0.020	0.814	0.017

Notes. TPR and FPR denote the True Positive Rate and False Positive Rate. In the column heads, the terms estimator and standard error are abbreviated as est. and s.e., respectively.

Table 4. Test of Aggregate Effects

	(1)			(2)			(3)			(4)			(5)			(6)		
$\tilde{R}$	est.	s.e.	p	est.	s.e.	p	est.	s.e.	p	est.	s.e.	p	est.	s.e.	p	est.	s.e.	p
A	−0.24	0.16	0.132*	−0.27	0.19	0.154*	−0.27	0.33	0.414*	−0.25	0.18	0.169*	−0.25	0.21	0.243*	−0.25	0.24	0.294*
H	−0.85	0.24	<0.001													−0.89	0.35	0.006
P				1.32	0.28	<0.001										1.22	0.36	<0.001
CP							1.63	0.50	<0.001							1.46	0.36	<0.001
RH										−0.66	0.27	<0.001				−0.65	0.36	0.036
FCP													1.55	0.31	<0.001	1.58	0.34	<0.001
$\tilde{\beta}$	7.32	0.36	<0.001*	8.75	0.45	<0.001*	8.45	0.44	<0.001*	7.87	0.40	<0.001*	7.52	0.37	<0.001*	7.68	0.23	<0.001*
Random effects																		
σ^2	3.29			3.29			3.29			3.29			3.29			3.29		
τ_{00}	0.26			0.41			1.74			0.40			0.61			0.84		
ICC	0.07			0.11			0.35			0.11			0.16			0.20		
N	34			33			33			33			36			97		
Observations	1,700			1,650			1,650			1,650			1,800			4,850		
Conditional R^2	0.722			0.779			0.711			0.740			0.730			0.730		
AIC	1,165.511			1,015.441			1,062.474			1,100.530			1,220.088			3,225.331		
log-likelihood	−578.755			−503.720			−527.237			−546.265			−606.044			−1,604.666		

Notes. Mixed logit estimation of organizational-level changes in reservation level (bias), $\tilde{R}$, and slope of screening function (discrimination), $\tilde{\beta}$. Included are random effects terms by individuals in treatment A and by pairs under treatments H, P, CP, RH, and FCP. The parameters σ^2 and τ_{00} denote the variances of the idiosyncratic error term and the fixed effects term, respectively. In the column heads, the terms estimator, standard error, and p-value are abbreviated as est., s.e., and p, respectively. ICC denotes the fraction of the variance that is "explained" by the random effects, and N denotes the number of pairs/random effects terms. Data are from the last 50 time periods of the performance stage. AIC, Akaike's information criteria.

*The *p*-values with an asterisk are two-sided, whereas all other *p*-values are computed under the assumption of the directional Sah and Stiglitz (1986) hypothesis, that is, one-sided test.

unambiguous predictions of the direction of change in reservation levels, all *p*-values for reservation levels are computed under the assumption of the directional hypothesis provided by Sah and Stiglitz (1986), that is, a one-sided test.

The results in Table 4 are based on estimation of each treatment pairs (columns (1)–(5)) as well as simultaneous estimation of all treatments (column (6)) under the assumption that the slope of the organizational-level screening function does not differ across treatments, that is, $\tilde{\beta}$ is constrained.[8] The results support Hypothesis 1.1. The hierarchy significantly increases the reservation level relative to the solitary agent (A). That is, the hierarchy only accepts a project if it has a higher positive value than the solitary agent (A) requires, that is, $\tilde{R}^H < \tilde{R}^A$. Once again, recall that the estimated parameter, $\tilde{R}$, has the opposite sign relative to the actual reservation level, R. In addition to testing the basic Sah and Stiglitz (1986) hierarchy relative to A, Table 4 shows results for RH, the hierarchy with randomly varying positions. Hypothesis 1.1 is supported also for this variant of the hierarchy.

The results in Table 4 also support Hypothesis 1.2. The basic Sah and Stiglitz (1986) polyarchy significantly decreases the reservation level relative to the solitary agent (A) and therefore accepts projects with lower positive value than the solitary agent, that is, $\tilde{R}^A < \tilde{R}^P$. The test of variants of the polyarchy with shared pay (CP) and fixed positions (FCP) also support Hypothesis 1.2. In summary, our test of aggregate organizational behavior shows that, relative to the solitary agent, all variants of hierarchies that we tested have a more conservative reservation level, whereas all the variants of the polyarchy that we tested have a more permissive reservation level. This result provides strong evidence in support of Hypothesis 1. A cross-comparison of *differences* in organizational bias ($\Delta\tilde{R}$) across organizational structures provides a clear pattern that lends additional confirming evidence for this claim (see the online appendix, Table A.5). These results show that the difference in organizational bias ($\Delta\tilde{R}$) between hierarchies (H, RH) and all versions of the polyarchy (P, CP, FCP) is significant ($p < 0.001$). In contrast, there is *no significant difference* of $\Delta\tilde{R}$ between H and RH ($p = 0.622$), P and CP ($p = 0.766$), P and FCP ($p = 0.835$), or CP and FCP ($p = 0.87$).

At this point, it may be useful to strengthen the intuition of the relation between our estimated screening function and its fit to the empirical data. As Figure 4 shows, the S-shaped screening function for the solitary agent (A) obtained from the logit model fits the empirical data well. The obtained goodness of fit supports the visual impression: $R^2 = 0.73$ with $\tilde{\beta}$ constrained in simultaneous estimation of all treatments (see Table 4) and $R^2 = 0.778$ obtained with $\tilde{\beta}$ as free parameter (see the online appendix,

Figure 4. Fit of Estimated Screening Function (Mixed Model Logit Fit) to Empirical Data of Treatment A

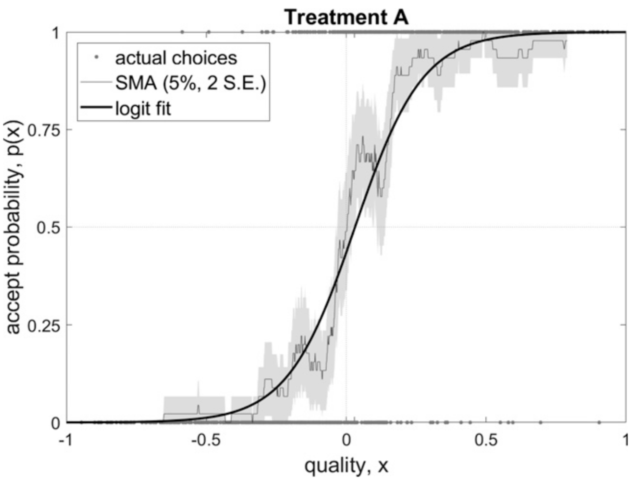

Notes. The empirical data are obtained from all 18 subjects over the last 50 time periods of the performance stage. Five percent of these 900 observations are binned with respect to the distribution of quality, *x*, and a simple moving average (SMA) computed, here shown with 2 standard errors (S.E.).

Table A.3).[9] For comparison, we have also estimated a mixed linear probability model (see the online appendix, Tables A.7 and A.9). The results are in general robust to the choice of functional form (see the section on robustness). To further strengthen the intuition relating to the magnitude of the effects shown in Table 4, we refer to results on the accept probability of the marginal project (online appendix, Table A.3) as well as average marginal effects (online appendix, Tables A.12, A.13).

Endogenous Behavior at the Group Level

According to Sah and Stiglitz (1988, p. 465), "the individuals' approval errors are, in general, endogenous. Among the features on which these errors may depend are the organization's structure and the nature of individuals' information processing. In addition, these errors may also depend on whether individuals consider others' errors in choosing their own decision rules; that is, whether or not they attempt to accept or reject projects strategically to offset others' errors." We first test this conjecture at the group level, which assumes that both individuals within an organization adapt their reservation levels in the same way (in a distribution sense). More precisely, at the group level we analyze individual decisions under the assumption that individuals endogenously change reservation levels (or slopes/discrimination level) conditional on organizational type (*not position* within the organization). We refer to such change in behavior as pairwise adjustment. Because the stylized organizations we study have two members, we examine whether there is pairwise adjustment of reservation levels in any of the treatments.

Although different individual endogenous adjustment cannot be excluded even if we fail to observe pairwise adjustment of reservation levels, it is highly unlikely because it would require that the two agents adapt their reservation levels exactly the same amount in opposite directions. Of course, a test of the behavior of each separate individual member within an organization provides a definitive answer. We provide this extension after examining what we refer to as group-level adjustment.

If the agents jointly adapt to the organizational context, aggregate organizational behavior will in all likelihood deviate from the predictions of Hypothesis 1. Thus, we examine whether group-level adjustment occurs. Sah and Stiglitz (1986, 1988) propose that individuals endogenously adjust the reservation levels of their screening functions such that the organization (hierarchy or polyarchy) achieves optimal performance when incentives are aligned. Per our theoretical framework, we therefore test if pairs of individuals in hierarchies and polyarchies rationally adapt the reservation levels of their screening functions so that optimal performance is achieved in the aggregate:

Hypothesis 2. *Pairs of individuals, i, rationally adapt the reservation level ($\tilde{R}^{G_i}$) of their screening functions to achieve optimal aggregate performance for the structure (G) they are placed in (polyarchies or hierarchies). This implies the following:*

Hypothesis 2.1. *Relative to the solitary agent (A), pairs of individuals reduce their reservation levels ($\tilde{R}^{H-_i}$) toward the value ($\tilde{R}^{*H-}$), which is optimal for a given hierarchy (H, RH), $\tilde{R}^{H-_i} > \tilde{R}^{A}$.*

Hypothesis 2.2. *Relative to the solitary agent (A), pairs of individuals increase their reservation levels ($\tilde{R}^{P-_i}$) toward the value ($\tilde{R}^{*P-}$), which is optimal for a given polyarchy (P, CP, FCP), $\tilde{R}^{P-_i} < \tilde{R}^{A}$.*

Table 5 shows the results on group-level behavior based on the logit specification. The results support Hypothesis 2.1 as $\tilde{R}^{H_i} > \tilde{R}^{A}$. The pair of individuals within the hierarchy has significantly reduced their joint reservation level relative to that of the solitary agent (A). As Sah and Stiglitz (1986) predicted, these two agents have adjusted reservation levels to become more permissive. Table 5 shows that Hypothesis 2.1 is also supported for the variant of the hierarchy with randomly varying positions. In summary, agents in hierarchies (H, RH) jointly adapt their reservation levels in a direction that is consistent with rational adaptation, even if it is not clear whether they approach optimality ($\tilde{R}^{*H-}$). Below, we test if this is the case.

Although we observe group-level adaptation in the hierarchy, this is not the case for the polyarchy. In the

Table 5. Test of Endogenous Adaptation at the Group Level

	(1)			(2)			(3)			(4)			(5)			(6)		
$\tilde{R}$	est.	s.e.	p	est.	s.e.	p	est.	s.e.	p	est.	s.e.	p	est.	s.e.	p	est.	s.e.	p
A	−0.21	0.15	0.161*	−0.24	0.16	0.126*	−0.23	0.20	0.253*	−0.23	0.15	0.133*	−0.24	0.18	0.195*	−0.22	0.17	0.204*
H	0.40	0.21	0.028													0.41	0.25	0.048
P				0.08	0.22	0.642										0.04	0.25	0.563
CP							0.18	0.30	0.726							0.14	0.25	0.712
RH										0.46	0.22	0.018				0.45	0.25	0.036
FCP													0.52	0.26	0.977	0.43	0.24	0.963
$\tilde{\beta}$	5.96	0.26	<0.001*	7.37	0.33	<0.001*	6.83	0.31	<0.001*	6.89	0.31	<0.001*	7.21	0.32	<0.001*	6.22	0.15	<0.001*
Random effects																		
σ^2	3.29			3.29			3.29			3.29			3.29			3.29		
τ_{00}	0.26			0.25			0.57			0.25			0.43			0.37		
ICC	0.07			0.07			0.15			0.07			0.12			0.10		
N	34			33			33			33			36			97		
Observations	2,116			2,040			2,044			2,036			2,222			6,858		
Conditional R^2	0.629			0.716			0.659			0.688			0.708			0.646		
AIC	1,691.012			1,424.690			1,521.358			1,495.400			1,588.315			5,342.504		
log-likelihood	−841.506			−708.345			−756.679			−743.700			−790.157			−2,663.252		

Notes. Mixed logit estimation of joint change in the subjects' reservation level, $\tilde{R}$, and slope of screening function, $\tilde{\beta}$ (discrimination). In the column heads, the terms estimator, standard error, and p-value are abbreviated as est., s.e., and p, respectively. Included are random effects terms by individuals in treatment A and by pairs under treatments H, P, CP, RH, and FCP. The parameters σ^2 and τ_{00} denote the variances of the idiosyncratic error term and the fixed effects term respectively, whereas ICC denotes the fraction of the variance that is "explained" by the random effects. N denotes the number of pairs/random effects terms. Data from last 50 time-periods of the performance stage. AIC, Akaike's information criteria.

*The *p*-values with asterisk are two-sided, whereas all other *p*-values are computed under the assumption of the directional Sah and Stiglitz (1986) hypothesis, that is, one-sided test.

basic polyarchy (P) as well as the other variants of the polyarchy (CP, FCP), we must therefore reject Hypothesis 2.2. That is, we conclude that pairs of agents in polyarchies do *not* (rationally) adapt their reservation levels. We now turn to our results at the individual level.

Endogenous Behavior at the Individual Level

Individual-level behavior captures the possibility that the two individuals within an organization adapt their reservation levels in different ways. It is also possible that the individuals adapt their discrimination levels (Sah and Stiglitz 1986, pp. 719–720). Thus, we test both possibilities.

Hypothesis 3. *Individuals adapt their behavior in a direction consistent with rational adaptation to the structure they are placed in (polyarchy or hierarchy).*

This implies a test of adaptive behavior for the first (f) and second (s) agent in each organizational structure:

Hypothesis 3.1. *Individuals reduce their reservation levels* ($\tilde{R}^{H_i}$) *toward the value* ($\tilde{R}^{*H}$)*, which is optimal for performance of the basic hierarchy,* $\tilde{R}^{H_i} > \tilde{R}^{A}$ *for* $i = f,s$.

Hypothesis 3.2. *Individuals reduce their reservation levels* ($\tilde{R}^{RH_i}$) *toward the value* ($\tilde{R}^{*RH}$)*, which is optimal for performance of the hierarchy with randomly varying positions,* $\tilde{R}^{RH_i} > \tilde{R}^{A}$ *for* $i = f,s$.

Hypothesis 3.3. *Individuals increase their reservation levels* ($\tilde{R}^{P_i}$) *toward the value* ($\tilde{R}^{*P}$)*, which is optimal for performance of the basic polyarchy,* $\tilde{R}^{P_i} < \tilde{R}^{A}$ *for* $i = f,s$.

Hypothesis 3.4. *Individuals increase their reservation levels* ($\tilde{R}^{CP_i}$) *toward the value* ($\tilde{R}^{*CP}$)*, which is optimal for performance of the coordinated polyarchy,* $\tilde{R}^{CP_i} < \tilde{R}^{A}$ *for* $i = f,s$.

Hypothesis 3.5. *Individuals increase their reservation levels* ($\tilde{R}^{FCP_i}$) *toward the value* ($\tilde{R}^{*FCP}$)*, which is optimal for performance of the coordinated polyarchy with fixed positions* $\tilde{R}^{FCP_i} < \tilde{R}^{A}$ *for* $i = f,s$.

As Figures 2 and 3 showed, screening functions are two-dimensional objects. Testing adaptation of the individual's reservation level captures the screening function's first dimension. Assuming that agents hold their discrimination levels constant, this was the possibility that Sah and Stiglitz (1986) examined in detail. However, individuals may, independent of adjusting their reservation levels, reduce the slope, or discrimination level, $\tilde{\beta}$, which captures the screening function's second dimension (Sah and Stiglitz 1986). There are at least two reasons why individuals would reduce the slope of their screening functions: (1) when they face an uncertain situation (see, e.g., Luce 1959, McKelvey and Palfrey 1995) and (2) when they free ride on others' effort (see, e.g., Gibbons and Roberts 2013). For these reasons, we also test if individuals will reduce the slope of their screening functions, that is, $\tilde{\beta}^{Organisation} < \tilde{\beta}^{A}$ for all individuals in all organizations (H, RH, P, CP, FCP). Rather than separately enumerating the hypotheses for these tests of changes in the screening function's slope, we shall clearly separate the results for $\tilde{\beta}$ and $\tilde{R}$ in our exposition.

Individual Adaptation of Individuals in Hierarchies

Table 6 shows results on individual-level behavior based on the logit specification. These tests are based on an extension of the stacking approach that we have developed (see model specification). Specifically, we test whether there is a difference in reservation level, $\Delta\tilde{R}$, between the first agent, f, in a structure and the solitary agent (A), that is, $\Delta\tilde{R} = \tilde{R}^{f} - \tilde{R}^{A}$. We then perform the same test for the structure's second agent, s, that is, $\Delta\tilde{R} = \tilde{R}^{s} - \tilde{R}^{A}$. This framework is also used for testing adaptation of the individuals' discrimination level, $\Delta\tilde{\beta}$.

Turning to the results on *reservation levels*, $\Delta\tilde{R}$, shown in Table 6, individuals adapt their reservation level when placed in a hierarchy. But in what direction? The coefficient of $\Delta\tilde{R}$ is positive, which implies that $\tilde{R}^{f} > \tilde{R}^{A}$. Recalling that $\tilde{R}$ has the opposite sign of R, we observe that individuals in the hierarchy *lower* their reservation levels relative to A. This is adaptation in the direction of optimality, which Sah and Stiglitz (1986) predicted. The effect of the agent's adaptation is, in this instance, to improve the overall aggregate behavior of the hierarchy, possibly without inducing bias in the hierarchy's aggregate reservation level. If this condition is achieved, the agents' adaptation is optimal, and the corresponding optimal reservation level, $\tilde{R}^{*}$, is displayed in the last column of Table 6. In summary, when placed in a hierarchy, individuals adapt their reservation levels in a direction that is consistent with rational adaptation as Sah and Stiglitz (1986) predicted.

Figure 5. Estimated Screening Function (Mixed Model Logit Fit) of First (A1) and Second Evaluator (A2) in the Basic Hierarchy Compared with the Screening Function of the Solitary Agent A

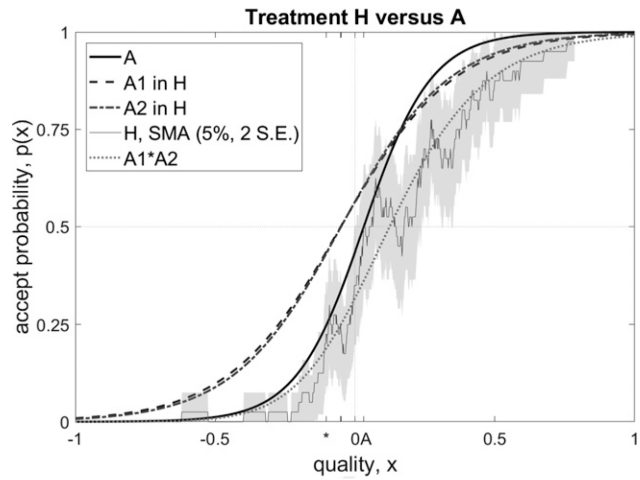

Notes. The basis for the estimation of screening functions is empirical data from the last 50 time periods of the performance stage. Five percent of these observations are binned with respect to the distribution of quality, x, and a simple moving average (SMA) computed, here shown with 2 standard errors (S.E.).

Figure 5, which is based on our experimental data, shows how the first (A1) and second (A2) agent in the basic hierarchy have adapted their reservation and discrimination levels. The figure nicely illustrates that both individuals in H have adapted their reservation levels in the direction that is optimal—this is marked

Table 6. Test of Endogenous Adaptation at the Individual Level, $\tilde{\beta}$ Unconstrained

	$\Delta\tilde{R}$			$\Delta\tilde{\beta}$			Model			F-test		$\tilde{R}$*	
Coefficient	est.	s.e.	p	est.	s.e.	p*	obs.	cond., R^2	L.R.	est.	p*	Value	p*
First agent in H	0.512	0.281	0.034	−3.337	0.660	0.000	1,700	0.686	−637.8	14,442	0.000	0.852	0.226
Second agent in H	0.518	0.296	0.040	−3.106	0.812	0.000	1,316	0.719	−458.9	7,829	0.000	0.852	0.259
First agent in RH	0.495	0.223	0.013	−2.014	0.703	0.004	1,650	0.720	−582.7	6,570	0.001	0.852	0.109
Second agent in RH	0.604	0.278	0.015	−2.842	0.874	0.001	1,286	0.724	−446.9	6,387	0.002	0.852	0.372
First agent in P	0.014	0.228	0.524	−1.239	0.742	0.095	1,650	0.735	−564.6	1,399	0.247	−0.326	0.137
Second agent in P	0.150	0.228	0.744	−1.730	0.942	0.066	1,290	0.732	−432.8	2,146	0.117	−0.326	0.037
First agent in CP	0.139	0.262	0.702	−2.714	0.689	0.000	1,650	0.702	−606.7	7,827	0.000	−0.852	0.000
Second agent in CP	0.252	0.392	0.739	−2.213	0.948	0.020	1,294	0.704	−438.1	3,315	0.036	−0.852	0.005
First agent in FCP	0.526	0.263	0.977	−0.335	0.818	0.682	1,800	0.777	−561.8	2,069	0.126	−0.852	0.000
Second agent in FCP	−0.208	0.311	0.251	−5.169	0.769	0.000	1,322	0.694	−488.6	23,048	0.000	−0.852	0.039

Notes. Mixed model logit (with random effects) estimation of change (compared with treatment A) in each agent's reservation level, $\tilde{R}$ slope of screening function, $\tilde{\beta}$ (discrimination), and F-test of H_0: $\Delta\tilde{R} = \Delta\tilde{\beta} = 0$. The column marked R^* shows the optimal reservation assuming no change to discrimination. In the column headings, the terms estimator, standard error, p-value, observations, conditional R^2, and Likelihood Ratio are abbreviated as est., s.e., p, obs., cond. R^2, and L.R., respectively. Data from last 50 time periods of the performance stage.

The p-values in the p columns are two-sided, whereas all other p-values are computed under the assumption of the directional Sah and Stiglitz (1986) hypothesis, that is, one-sided test.

by the small ticks on the x-axis; the optimal reservation level, $\tilde{R}^*$, is marked with a star. In Table 6, the reservation of the evaluator, $\tilde{R} = \Delta\tilde{R} + \tilde{R}^A$, is tested against the optimal $\tilde{R}^*$. The optimal endogenous shift $\tilde{R}^*$ is derived and calculated as demonstrated in the online appendix, using the unconstrained estimate for $\tilde{\beta}$ for the single agent (Table A.3 in the online appendix).[10] Figure 5 also illustrates that both individuals in H appear to have reduced their discrimination levels, that is, the slope of their screening functions is reduced. Finally, it is remarkable that the two agents in H have adapted their screening functions in the same way, that is, there is little difference between the group and individual level. The online appendix features similar figures for the remaining four structures (RH, C, CP, FCP). Notably, these figures illustrate that the two agents in FCP have adapted in different ways—one adapts the reservation level, the other the discrimination level.

According to the results on *discrimination levels* shown in Table 6, all subjects in hierarchies (H, RH) reduced the slope of their screening function, $\tilde{\beta}$ ($p < 0.05$, two-sided test). Although the individuals in the hierarchies adapt their reservation levels in a direction that improves aggregate performance, the endogenous adaptation of their discrimination levels has the opposite effect as it both increases type I and type II errors.

Table 6 also provides a joint test of reservation and discrimination levels for H and RH, that is, H_0: $\Delta\tilde{R} = \Delta\tilde{\beta} = 0$. In summary, the hypothesis that individuals do not adapt reservation *and* discrimination levels is therefore rejected for H and RH. As predicted by Sah and Stiglitz (1986), the reservation levels are adapted in a direction that is consistent with endogenous rational adaptation. The discrimination levels are adapted such that the slope of the screening function is reduced (more noisy screening).

Individual Adaptation of Individuals in Polyarchies

According to the results for adaptation of *reservation levels*, $\Delta\tilde{R}$, featured in Table 6, we can reject Hypotheses 3.3–3.5 for all agents in all variants of the polyarchy. When placed in a polyarchy, individuals do not adapt reservation levels in the direction associated with rational adaptation; for this reason, they, obviously, do not approach the optimal level (last column in Table 6). Relative to the basic polyarchy, it does not matter if the individuals have shared pay (Hypothesis 3.4). This is remarkable because it shows that, in contrast to Sah and Stiglitz's (1986) prediction, differences in incentive structure (P versus CP) do not affect bias in polyarchies. Turning to FCP, polyarchies with fixed position (Hypothesis 3.5), we find that evaluators do *not* adapt rationally when placed in FCP but rather produce the opposite effect. That is, subjects who go first in FCP change their reservation levels in a counterintuitive way. Comparing this result with the behavior of individuals placed in hierarchies provides remarkable evidence that different organizational structures prime different levels of sophistication in the behavior of organizational members. Considering the results for adaptation of *discrimination levels*, $\Delta\tilde{\beta}$, shown in Table 6, both agents in CP and the second agent in FCP reduced the slope of their screening function, $\tilde{\beta}$ ($p < 0.05$, two-sided test). At $p < 0.10$ (two-sided test), the picture is clearer; all agents in polyarchies apart from the first agent in FCP reduced the slope of their screening function, $\tilde{\beta}$. We elaborate on these results in the discussion section. In summary, individuals do not adapt their reservation and discrimination levels when placed in polyarchies (P, CP) unless positions are fixed (FCP).

Robustness Tests

For robustness, we conducted a complete analysis of all hypotheses with a mixed linear probability model, including random effects (see the online appendix, Tables A.7–A.11). The results do not deviate much. At the aggregate level, we obtained identical results for the hypothesis tests. At the next lower level, joint behavior at the group level, we also obtained identical results for the hypothesis tests. Finally, at the lowest level, individual behavior of each member within an organization, we obtained the same overall pattern as in the logit model. In summary, our results are generally robust with respect to the choice of functional form specification: logit model versus linear probability model.

Discussion

We began by noting the need for research on the effect of organizational context on managerial and organizational behavior. To advance research on this topic, we investigated whether context influences bias and discrimination in individual and aggregate organizational decisions. Our experimental findings showed that evaluators adapt their reservation and discrimination levels in centralized structures (hierarchies), whereas they did *not* generally adapt their reservation and discrimination levels when placed in decentralized structures (polyarchies). For this reason, decentralized structures exhibit greater *aggregate* bias than do centralized structures. We here provide plausible explanations for these findings and highlight implications for theory and practice.

Explaining Bias

As we previously suggested, the Sah and Stiglitz (1986) model offers two different perspectives on organizational decision making that, with a slight metaphorical abuse, can be labeled as "the information engineer's perspective" and "the game theorist's perspective." In Sah and Stiglitz (1986), proposition 1 represents the "information engineer's perspective." It states that a polyarchy selects a larger proportion of the available projects than does a hierarchy. It is similar to an engineering perspective because this proposition is derived by considering the performance of different decision-making circuits employing human processors whose individual screening functions are assumed to be *independent* from the architecture in which they are embedded. As in prior work (Csaszar 2012), the prediction of Sah and Stiglitz's (1986) proposition 1 receives robust support in our experiments. Relative to the solitary agent (A), all variants of hierarchies (H, RH) that we tested have a more conservative reservation level, whereas all variants of the polyarchy that we tested (P, CP, FCP) have a more permissive reservation level. This result provides strong evidence that organizational structure, in the aggregate, induces a predictable (statistical) bias relative to the behavior of the solitary agent.

Considering the next lower level, behavior at the group level (in our experiment, pairs of agents) captures the possibility that the individuals within an organization jointly adapt their reservation levels in response to features of the decision-making architecture (the "game theorist's perspective," as per proposition 4 in Sah and Stiglitz 1986). We found that subjects in hierarchies (H, RH) adapt their own reservation levels, $\tilde{R}$, although not enough to offset the bias induced by the organizational structure. Nevertheless, *subjects in hierarchies do adapt* their reservation levels in the direction of optimality, that is, the pair of subjects in a hierarchy have a more permissive screening function than does the solitary agent. This is not the case for the polyarchies. In the basic polyarchy as well as the coordinated polyarchy, subjects do *not* adapt in the direction of optimality— they do not jointly increase their reservation levels in comparison with the solitary agent. Finally, in the coordinated polyarchy where agents have fixed positions (FCP), we observed group-level adaptation in the direction *opposite* of rational behavior. Taken together, these findings are notable because they highlight that the subjects jointly differentiate their responses to the generic type of organizational structure they operate in, independent of the individual organizational positions (roles) they are assigned to (fixed versus random) and independent of incentive schemes (shared versus individual pay).

We are thus left with puzzling evidence, as agents seem to display differentiated levels of endogenous responses in different types of structures. What are the key differences between the two types of architecture—hierarchies and the polyarchies—that might explain this result? A closer look at the experimental task suggests an important asymmetry between *accept* and *reject* decisions that has different implications for behavior in hierarchies and polyarchies. A decision to accept a project is associated with a risky payoff and uncertain feedback on project quality. In contrast, a decision to reject a project is equivalent to null payoff and no feedback. Note that this asymmetry is not an artifact of our experimental setting; it is intrinsic to the task in the Sah and Stiglitz (1986) model as well as many sequential screening tasks in organizations (Csaszar 2012, Reitzig and Maciejovsky 2015, Christensen and Knudsen 2020). In learning to screen the alternatives with uncertain value, accepting an alternative is more focal and salient than rejecting it (and has more informational value).

In a Sah-Stiglitz hierarchy (H, RH), accepting is necessarily the outcome of decisions made by both agents that work in it. In hierarchies, therefore, the salience of accepting alternatives brings awareness of their mutual dependence to the attention of both agents. In all polyarchies (P, CP, FCP), by contrast, each of the two agents can individually accept a proposal, which implies that the agents' interdependence is *less salient* than it is for the hierarchy. This differential awareness of interdependence with others can explain why we observe systematic differences in the way subjects in hierarchies and polyarchies adapt their reservation levels. Subjects in hierarchies act consistent with mutual awareness of their interdependence in the decision-making process and respond correctly to this condition. The case is different for polyarchies (C, CP): lower awareness of strategic interdependence implies that subjects do not counter the bias induced by the organizational structure.

Our explanation of why FCP deviates from P and CP (by adapting reservation levels, but on the wrong direction) is not based on salience of acceptance. In contrast, it is premised on differences in the positions (roles) the individual agents occupy. In FCP, the subjects' positions are fixed—the same agent always makes the first choice. The first agent in FCP has a peculiar organizational role that strongly empowers this agent relative to the second one. Most decisions of *accept* are in his or her hands—the second agent will only receive the cases the first agent discarded, and these are mostly "bad cases." We therefore suggest that the strong power the first agent in FCP holds with respect to accepting proposals induces a mixture of overconfidence[11] and unwillingness to delegate decisions of acceptance to the second agent. This explains why the first agent in FCP adapts its reservation levels to become more permissive. In contrast, the second agent in FCP will likely act more at random because this agent has few, mostly bad, alternatives to learn from. Our findings relating to endogenous adaptation of discriminating capability (see below) offer additional support for this conjecture.

Explaining Reduced Discrimination

Bias—the systematic deviation of the reservation level—is one aspect of the agent's screening function (Sah and Stiglitz 1986). Another aspect of interest is the agent's capability to discriminate between good and bad alternatives. Discriminating capability is defined as the screening function's slope (in our model, a higher value of $\tilde{\beta}$ implies higher screening capability). In the literature on decision making, low discriminating capability in screening is commonly associated with "noise" (Luce 1959, McFadden 1976, McKelvey and Palfrey 1995). Our experimental findings show that subjects who are placed in organizational structures are "noisier"—they display lower discriminating capability—than are solitary agents. Subjects who go first in FCP are an exception to this finding, as they do not reduce their discriminating capability relative to the solitary agent; their case will be discussed separately.[12]

To explain why subjects in hierarchies or polyarchies display lower discriminating capability than solitary agents, we must understand how agents learn from experience in such structures. From a learning point of view, more experience is likely to generate steeper screening functions, that is, to higher values of the $\tilde{\beta}$ parameter in our model (e.g., see McKelvey and Palfrey 1995). However, compared with solitary agents, pairs in hierarchies have to solve what is known as the "credit assignment" problem in learning (Samuel 1959, Holland 1975, Axelrod and Cohen 2001, Fang 2012): when outcomes are the result of a sequence of decisions, how are the merits for the right (or demerits for the wrong) choice best allocated? This problem is amplified by the fact that the first agent in a hierarchy does not receive direct feedback for its choices, whereas the second necessarily learns from a biased sample, that is, the first agent has already removed most of the bad alternatives (Christensen and Knudsen 2010). As a result of these concurring difficulties, learning in hierarchies becomes noisier and the slope of the subjects' screening functions is reduced.

Subjects in polyarchies (except FCP) also lowered their discriminating capability relative to the solitary agent (at $p < 0.10$). In the case of P and CP, an agent only receives feedback when this agent is the first to accept a project. Relative to the solitary agent, the feedback each subject in P and CP receives is therefore effectively decreased. A further consequence of the structural constraints operating on agents in a polyarchy is that the first agent sends a high proportion of "bad projects" to the second agent, that is, projects that the first agent believes should be rejected. That is, at each decision-making round, the first agent in a polyarchy acts as a filter that separates good and bad projects—and the more the agents learn, the stronger is the separation of bad and good projects. Paradoxically, learning therefore increases the variance of the samples that the subjects in P and CP observe as they randomly shift between going first and second over multiple rounds of the experiment. Thus, reduced feedback and increased variance in the project portfolio jointly contribute to explain why subjects in polyarchies (C, CP) lowered their discrimination capability, that is, reduced $\tilde{\beta}$.

Even if our conjectures may have merit, we cannot exclude the complementary explanation that agents in polyarchies lowered their discrimination level because of a reduction in their supply of cognitive effort—a form of cognitive free riding when decisions are taken by pairs. This explanation is supported by the fact that we observe less lowering of discriminating capability ($\tilde{\beta}$) when incentives to free riding are absent—that is, in the case of P as opposed to CP. However, the very same fact that agents still reduce discrimination also in P rules out cognitive free riding as the main explanation of our findings.

Turning to the coordinated polyarchy with fixed positions, our finding—that subjects who go second in FCP lower their discriminating capability relative to the solitary agent, whereas subjects who go first do not—stands out as a special case. In FCP, the first agent effectively receives experiential feedback very similar to that of the solitary agent—and, unsurprisingly, the discriminating capability ($\tilde{\beta}$) of the first agent in FCP is undistinguishable from the discriminating capability of the solitary agent (at $p < 0.10$). On the other hand, the second agent in FCP receives very little useful feedback—it receives a high proportion of bad projects from the first agent. Because bad projects are most of the time rejected and because subjects only receive feedback on accepted projects, the subjects who go second in FCP are deprived of feedback, which is reflected in a very low discriminating capability, $\tilde{\beta}$.

Relations to Previous Literature: Epistemic Interdependence

Our explanation of findings relating to bias are premised on the subjects' differential awareness of interdependence relating to accepting and rejecting alternatives. From a theoretical point of view, our results relating to bias, therefore, provide an interesting example of epistemic interdependence in organizations (Puranam et al. 2012) and the failures to recognize it. Epistemic interdependence arises when one agent must have predictive knowledge about the behavior of another agent to reach a common goal. Epistemic interdependence requires two conditions: (1) at least one agent faces broad (interdependent) incentives and (2) the same agent is scheduled to act before knowing the behavior of the other. As noted by Puranam et al. (2012, p. 427), "[i]n a dyad there is, by definition, epistemic interdependence if one agent's optimal action depends on a prediction of the action of another." In our experimental setting, agents are located in dyads. So, the agents are epistemically interdependent when (1) they are jointly rewarded for performance and (2) at least one agent must correctly predict how the effect of structure influences the other agent's behavior. The first condition does not hold for P. The other cases, however, do support the existence of epistemic interdependence and the need for agents to understand how organizational structure will affect the other agent's choices. Subjects in hierarchies act consistently with an understanding of epistemic interdependence, albeit not perfectly so. In contrast, subjects in CP and FCP fail to recognize epistemic interdependence. This failure points to salience and attentional factors as critical in understanding how epistemic interdependence affects organizational performance—with important implications for organization design that we discuss below.

Relations to Previous Literature: Depth of Strategic Thinking

The claim that the depth of strategic thinking by agents may be affected by the perceived structure of interaction is not entirely new, and our results resonate with former experimental observations. For example, it has been noted that agents are more prone to consider the behavior of others in games where moves are sequential rather than simultaneous, even when the games are strategically equivalent (Schotter et al. 1994). Camerer has suggested that different levels of strategic thinking can be induced by prompting beliefs about the other player (Camerer et al. 2004) and by inducing alternative representations of games (Camerer 2003). Research has also found that the complexity of the game structure (Devetag and Warglien 2003) as well as role ambiguity may inhibit strategic thinking (Agranova and Schotter 2012). Our experiment extends these findings by highlighting that organization design, in its own right, may shape the levels of strategic thinking that organizational members employ.

Methodological Contribution

In addition to our theoretical contribution, we offer a methodological contribution as we introduce tools to analyze nested levels of change in social organizations (i.e., organization, group, and individual level) that may be useful for future experimental work on organizations. In particular, we suggest that our approach may benefit research on organizations that aims to disentangle behavior at different levels of analysis. This would include studies of organizational search and learning and, more generally, efforts to identify how the microstructure of organizations affect systemic change at the macro level (Puranam 2018).

Limitations

While pointing to significant novel insights, we are aware that our experimental setup has several limitations. First, it only considers sequential decision processes. Of course, committee voting and many other organizational decision processes occur in parallel. As previously mentioned, experiments in behavioral game theory find that parallel decision processes make it harder to take the intentions of others into account and therefore limit the depth of strategic thinking (Schotter et al. 1994). At the same time, parallel processing increases the need for predictive knowledge to avoid an increase in coordination costs. This creates a potential design tension that deserves further enquiry. A second limitation is that our experiment only considers dyads. The two types of dyad we studied, hierarchies and polyarchies, can be thought of as modules from which more complex assemblies can be constructed (Christensen and Knudsen 2010). Little is known about how such assemblies will modify the behavior of lower-level modules—a question over which laboratory experiments may provide new important insights. Thirdly, our experiment assumes a stationary task environment. How would different architectures affect the capability of organizations to respond to changes in the environment? And how would learning and strategic adaptation interact during change? These are core questions about which we know very little (see Csaszar and Eggers (2013) for a computational exploration). As our experimental setup and statistical approach are very well suited to address such questions, we welcome future research that employs them.

Conclusion

Our results on bias and discrimination in organizational decisions provide novel insights that are useful both for research and for the practice of organization design. The literature on screening theory, that over the last two decades has emerged as a major stream of research in organization science (Joseph and Gaba 2020), can trace its roots to information economics (Marschak and Radner 1972; Sah and Stiglitz 1985, 1986) and information engineering (Moore and Shannon 1956a, b). Screening theory (e.g., Sah and Stiglitz 1986, Christensen and Knudsen 2010) considers group decision making without conflict of interest, which implies that strategic response to design features that influence alignment of interests and incentives are assumed away. Although this approach facilitates powerful analysis of complex decision structures, it also imposes limitations for practical applications that aim to reduce problems related to politics, conflict, and opportunistic behavior in organizations. In contrast, organizational economics traditionally has its main focus on designing organizational mechanisms that reduce the negative consequences of strategic behavior of agents (Gibbons and Roberts 2013). Our findings indicate that there would be notable gains from increasing the level of cross-talk between the two camps. In the intellectual tradition of Herbert A. Simon, organizations are solutions to the problems created by the agents' bounded rationality. Our contribution to theory broadens this perspective by suggesting how organization design may *shape* levels of rationality of organizational members. Thus, we show that the configuration of pipes through which information flows among organizational members may regulate the agents' (mutual) awareness of their interdependence.

Our study also has implications for practice. Perhaps the most important is the novel finding that organization designers can pick organizational configurations to regulate the level of the employees' awareness of interdependencies. Even though most real-world organizations are more complicated than the simple dyads we studied, the latter are building blocks, which are found in most organizations, for example, in decision processes relating to approval of firm policies and selection of job applicants. Thus, organization designers can regulate the level of the employees' awareness of (mutual) interdependencies and thereby tune the employees' strategic sophistication. As we have shown, applying structures that are equivalent to hierarchies would increase the employees' awareness of interdependencies and their strategic sophistication. In contrast, polyarchies have the opposite effect—with the possible exception of the case in which most decision power is assigned to one agent. Thus, organization designers can assess trade-offs regarding gains and losses from tuning the employees' mutual awareness and strategic sophistication. For example, reducing awareness of interdependence in conflictual settings might dampen the negative effects of the underlying divergent interests. A further related practical implication can be drawn from our finding that incentives have no effect on the subjects' performance in polyarchies. If this conjecture holds more generally, the implication is that monetary incentives are effective in structures that increase the employees' awareness of interdependencies. Of course, we should add the caveat that the outlined implications for practice are limited because they have been extracted from a laboratory study. Thus, we encourage others to explore this exciting new path at the intersection of organization science and organizational economics.

Endnotes

[1] Discrimination is here a technical term that refers to uncertainty about the actual value of a proposal (how noisy an evaluation is). The less uncertainty about a proposal, the higher a manager's ability to discriminate.

[2] Technically, we require the value function to be (piecewise) continuous.

[3] Assuming continuity, discrimination can be defined either as the steepest slope on the screening or as the slope at the reservation level. Although in principle they must be treated slightly differently, in practice, as well as in our treatment, they coincide.

[4] Server and client were coded in Processing 2/3 using packages controlP5 for the graphical user interface and oscP5 for communication.

[5] This level of analysis is aligned with the symmetry assumption of Sah and Stiglitz (1986), for example, that both individuals in a hierarchy do not shift reservation levels or that both individuals shift reservation levels to the symmetric Nash equilibrium point.

[6] Type I error rate is calculated as the ratio of false positives (rejected good) to total images, and the type II error rate is calculated as the ratio of false negatives (accepted bad) to total images. Accuracy, the fraction of decisions that are correct, captures both bias and noisy discrimination. It can be calculated as 1-type I error rate - type II error rate.

[7] Evidence in support of this claim is shown in the online appendix.

[8] For robustness, we estimated each treatment with $\tilde{\beta}$ as a free parameter. These results, shown in the online appendix (Table A.3), do not change the results in a qualitative way, neither do they change the conclusions reported here.

[9] Throughout the paper we are using the conditional R^2 for mixed models as defined by Nakagawa and Schielzeth (2013).

[10] The calculation of the optimal shift was performed 100 times by drawing normally distributed $\tilde{\beta}$ values with the mean and standard error of Table A.3 in the online appendix, then calculating the average and standard error of this sample. The significance test includes the prediction error, which is roughly 0.002.

[11] The association of overconfidence and excess of "positive" decisions (to buy, to enter in a market) has been observed both in the field and in the laboratory in many behavioral studies (e.g., Camerer and Lovallo 1999, Barber and Odean 2001)

[12] At $p < 0.10$ (two-sided test), except for the first agent in FCP, all agents in hierarchies and polyarchies lowered their discrimination capability. At $p < 0.05$ (two-sided test), both agents in CP and the second agent in FCP lowered their discrimination capability.

References

Adner R, Levinthal DA (2008) Doing vs. seeing: Acts of exploitation and perceptions of exploration. *Strategic Entrepreneurship J.* 2(1): 43–52.

Agranov M, Schotter A (2012) Ignorance is bliss: An experimental study of the use of ambiguity and vagueness in the coordination games with asymmetric payoffs. *Amer. Econom. J. Microeconom.* 4(2):77–103.

Axelrod R, Cohen MD (2001) *Harnessing Complexity* (Basic Books, New York).

Barber BM, Odean T (2001) Boys will be boys: Gender, overconfidence, and common stock investment. *Quart. J. Econom.* 116(1): 261–292.

Camerer CF (2003) Behavioural studies of strategic thinking in games. *Trends Cognitive Sci.* 7(5):225–231.

Camerer C, Lovallo D (1999) Overconfidence and excess entry: An experimental approach. *Amer. Econom. Rev.* 89(1):306–318.

Camerer CF, Ho TH, Chong JK (2004) A cognitive hierarchy model of games. *Quart. J. Econom.* 119(3):861–898.

Christensen M, Knudsen T (2002) *The Architecture of Economic Organization: Toward a General Framework* (University of Southern Denmark, Odense, Denmark).

Christensen M, Knudsen T (2010) Design of decision-making organizations. *Management Sci.* 56(1):71–89.

Christensen M, Knudsen T (2013) How decisions can be organized and why it matters. *J. Organ. Design* 2(3):41–50.

Christensen M, Knudsen T (2020) Division of roles and endogenous specialization. *Indust. Corporate Change* 29(1):105–124.

Csaszar FA (2012) Organizational structure as a determinant of performance: Evidence from mutual funds. *Strategic Management J.* 33(6):611–632.

Csaszar FA (2013) An efficient frontier in organization design: Organizational structure as a determinant of exploration and exploitation. *Organ. Sci.* 24(4):1083–1101.

Csaszar FA, Eggers JP (2013) Organizational decision making: An information aggregation view. *Management Sci.* 59(10): 2257–2277.

Cyert RM, March JG (1963) *A Behavioral Theory of The Firm* (Prentice-Hall, Inc., Englewood Cliffs, NJ).

Devetag G, Warglien M (2003) Games and phone numbers: Do short-term memory bounds affect strategic behavior? *J. Econom. Psych.* 24(2):189–202.

Drucker PF (1967) The effective decision. *Harvard Bus. Rev.* 45: 92–104.

Fang C (2012) Organizational learning as credit assignment: A model and two experiments. *Organ. Sci.* 23(6):1717–1732.

Gavetti G, Levinthal DA, Rivkin JW (2005) Strategy making in novel and complex worlds: The power of analogy. *Strategic Management J.* 26(8):691–712.

Gavetti G, Greve HR, Levinthal DA, Ocasio W (2012) The behavioral theory of the firm: Assessment and prospects. *Acad. Management Ann.* 6(1):1–40.

Gibbons R, Roberts J (2013) *The Handbook of Organizational Economics* (Princeton University Press, Princeton, NJ).

Greene WH (2018) *Econometric Analysis*, 8th ed. (Pearson, London).

Gulati R, Lawrence PR, Puranam P (2005) Adaptation in vertical relationships: Beyond incentive conflict. *Strategic Management J.* 26(5):415–440.

Holland JH (1975) *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (University of Michigan Press, Ann Arbor).

Joseph J, Gaba V (2020) Organizational structure, information processing, and decision-making: A retrospective and road map for research. *Acad. Management Ann.* 14(1):267–302.

Knudsen T, Levinthal DA (2007) Two faces of search: Alternative generation and alternative evaluation. *Organ. Sci.* 18(1): 39–54.

Levinthal DA (2011) A behavioral approach to strategy–What's the alternative? *Strategic Management J.* 32(13):1517–1523.

Luce D (1959) *Individual Choice Behavior* (Wesley, New York).

Marschak J, Radner R (1972) *Economic Theory of Teams* (Yale University Press, New Haven, CT).

McFadden DL (1976) Quantal choice analysis: A survey. *Ann. Econom. Soc. Measurement* 5(4):363–390.

McKelvey RD, Palfrey TR (1995) Quantal response equilibria for normal form games. *Games Econom. Behav.* 10(1):6–38.

Moore EF, Shannon CE (1956a) Reliable circuits using less reliable relays, part I. *J. Franklin Inst.* 262(September):191–208.

Moore EF, Shannon CE (1956b) Reliable circuits using less reliable relays, part II. *J. Franklin Inst.* 262(October): 281–297.

Nakagawa S, Schielzeth H (2013) A general and simple method for obtaining R^2 from generalized linear mixed-effects models. *Methods Ecology Evolution* 4(2):133–142.

Piezunka H, Schilke O (2021). The effect of organizational aggregation structures on individuals' voting behavior: An experimental investigation. Working paper, INSEAD, Fontainebleau, France.

Puranam P (2018) *The Microstructure of Organizations* (Oxford University Press, Oxford, UK).

Puranam P, Maciejovsky B (2020) Organizational structure and organizational learning. Argote L, Levine JM, eds. *The Oxford Handbook of Group and Organizational Learning* (Oxford University Press, Oxford, UK), 521–533.

Puranam P, Raveendran M, Knudsen T (2012) Organization design: The epistemic interdependence perspective. *Acad. Management Rev.* 37(3):419–440.

Reitzig M, Maciejovsky B (2015) Corporate hierarchy and vertical information flow inside the firm—A behavioral view. *Strategic Management J.* 36(13):1979–1999.

Rivkin JW, Siggelkow N (2007) Patterned interactions in complex systems: Implications for exploration. *Management Sci.* 53(7): 1068–1085.

Sah RK, Stiglitz JE (1985) Human fallibility and economic organization. *Amer. Econom. Rev.* 75(2):292–297.

Sah RK, Stiglitz JE (1986) The architecture of economic systems: Hierarchies and polyarchies. *Amer. Econom. Rev.* 76(4):716–727.

Sah RK, Stiglitz JE (1988) Committees, hierarchies and polyarchies. *Econom. J. (London)* 98(391):451–470.

Samuel AL (1959) Some studies in machine learning using the game of checkers. *IBM J. Res. Development* 3(3):210–229.

Schotter A, Weigelt KW, Wilson C (1994) A laboratory investigation of multiperson rationality and presentation effects. *Games Econom. Behav.* 6(3):445–468.

Sorkin RD, West R, Robinson DE (1998) Group performance depends on majority rule. *Psych. Sci.* 9(6):456–463.

Michael Christensen is associate professor at the Strategic Organization Design unit at the University of Southern Denmark (SDU). He received his PhD in physics from SDU. His research interests include learning and decision making in organizations, competition and industry dynamics, and emergent properties of complex systems.

Christian M. Dahl is professor of economics at the Department of Business and Economics, University of Southern Denmark. He received his PhD from Aarhus University in 2000. His research interests are broadly defined within the area of theoretical and applied econometrics with particular focus on causal identification and estimation, time series analysis, and business data science.

Thorbjørn Knudsen is professor of strategic organization at the University of Southern Denmark and the Frankfurt School of Finance and Management—and a chair of social sciences at the Danish Institute for Advanced Study. He received his PhD from the University of Southern Denmark in 1999. His research interests center around evolutionary and adaptive processes in organizations and the way organization design can shape these processes.

Massimo Warglien is a professor of management at the Department of Management, Universita` Ca' Foscari Venezia and member of the Bliss research center. His research and teaching are in the areas of organization theory, language and interaction, behavioral decision making, digital humanities, and multispecies studies.