Role ambiguity and endogenous adaptation

Michael Christensen, Thorbjørn Knudsen & Massimo Warglien

Abstract

In a laboratory experiment, we used a screening task in which Sah-Stiglitz-dyads screen if a proposal's quality justifies its acceptance (or rejection). We find that individual actors reduce supply of effort when they clearly understand how the organizational arrangement can be gamed. In contrast, actors experiencing role ambiguity may not clearly understand the organizational arrangement and how it can be gamed. This insight explains our finding that individuals adapt behavior when their role is fixed while their behavior is unaltered when exposed to role ambiguity. The implication of this finding is that, paradoxically, role ambiguity makes individual behavior predictable. For this reason, it is possible to design organizations that have predictable desired effects, such as canceling individual cognitive biases and higher performance. While our results are only suggestive in this regard, the implication is that organizations can be designed to remove biases in financial decision making, hiring decisions, and perhaps even overcome deep cultural biases. Our findings also add a new perspective to the idea that individuals respond to incentives in the sense that behavior can roughly be described as expected value maximization conditional on limitations of human actors. While we do not dispute this idea, we offer a complementary perspective that removing the basis for gaming the organization also removes the need for inventing incentives that may prevent such gaming.

Keywords: organization design; organizational structure; decision making; role ambiguity; endogenous adaptation

Role ambiguity and endogenous adaptation

Introduction

Extensive prior research has examined how structuring decision processes facilitates desired organizational level performance. Recently there has been a confluence of diverse research streams that encompasses economic decision theory (Marschak & Radner 1972; Sah & Stiglitz 1986; 1988; Sah 1991; Stiglitz 2002, 2004), behavioral theory of organizational decision processes (Simon 1947; March and Simon 1958; Cohen, March and Olsen 1972), voting theory (Condorcet 1785; Fishburn 1977: Young 1988), committee decision making (Gerling, Grüner, Kiel & Schulte 2005), and wisdom-of-crowds (Galton 1907; Surowiecki 2004; Page 2007). This confluence has advanced a vigorous stream of theories of information aggregation in organization science (Christensen & Knudsen 2002; 2010, 2013, 2019; Knudsen & Levinthal 2007; 2013, 2018; Csaszar 2013; Csaszar & Eggers, 2013).

Central to these theories is that organizations have a fundamental role in aggregating choice functions of individual agents. A choice function, also known as a decision rule, is a function which maps observations about alternatives onto an appropriate action, e.g. whether job candidates should be hired or investment projects declined. Organizational decisions are reached by aggregating individual choice functions, which implies that different aggregation procedures will produce different organizational level choice functions. More precisely, an aggregation procedure is a composition of a set of choice functions. For example, the well-known procedure of Condorcet voting (e.g. Fishburn 1977), where a decision is reached by majority vote, is a composition of individual preferences about alternative candidates that result in selection of one of them. More generally, aggregation is a fundamental process that joins individual contributions and thereby determines properties and outcomes at the organizational level. There are many aggregation procedures and they usually lead to very different organizational properties and outcomes (Christensen & Knudsen 2010; Csaszar 2013). For example, changing voting procedures may change outcomes in committee meetings or political elections, e.g. a different candidate may be selected if voting rules are changed in political elections

(Arrington & Brenner 1980). For another example, the design of decision procedures in a bank's credit unit will, in the aggregate, influence net revenue from the loan business (Christensen & Knudsen 2010, 2019).

Research on organizational decision making has identified aggregation procedures that result in superior organizational level decision functions that maximize expected performance or minimize risk relating to the task environment within which they operate (Christensen & Knudsen 2002, 2010, 2019; Csaszar 2013; Csaszar & Eggers 2013). While prior research has made significant progress in understanding organizational level decisions, an important gap in our knowledge is the possibility that individual behavior is endogenous to the organizational structure. Csaszar (2012) provided the first large scale empirical test of the fundamental dyadic decision structures presented in Sah & 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. This finding suggests that design of decision structures can be used to generate organizational level outcomes that predictably improves the behavior of the individuals involved in the process. However, this claim is based on the premise that individuals do not modify their choice functions conditional on the structural configuration they are located in. Drawing on field data and model-guided experimental studies, Reitzig & Maciejovsky (2014) showed that this premise is not generally valid. They found that mid-level managers in the field pass on fewer ideas the steeper the hierarchy that surrounds them, suggesting that individual behavior is endogenous to organizational structure. In a study of credit assessment in a bank, Christensen & Knudsen (2019) found that division of roles in a decision structure generates endogenous specialization.

Taken together, these findings suggest that individuals adjust behavior, and thereby modify their choice functions as they are placed in an organizational setting (Sah & Stiglitz 1986). As a result, the organizational level choice function may change in ways that are surprising, and possibly undermines intended effects, e.g. by inducing freeriding in sequential decision structures (Swank & Visser 2008) or inspiring tactical voting in a Condorcet setting (Gerling et al. 2005). The purpose of this paper is to examine if, or how, individual actors who are tasked with making decisions adapt to the organizational structure they are located in. Prior research that extended Sah & Stiglitz (1986) to a general theory of designing organizational decision structures assumes that individual human actors are endowed with fixed choice functions. These choice functions, commonly referred to as screening functions, capture an agent's ability to pass judgment (e.g. Christensen & Knudsen 20109; Csaszar 2013; Csaszar & Eggers 2013). While this simplifying assumption preserves mathematical tractability when deriving organizational level choice functions, it could be seriously misleading if individual behavior is adaptive (Sah & Stiglitz 1986). Sah & Stiglitz (1986) suggested that actors may endogenize to optimally counter a bias caused by the effect of organization design, e.g. countering the bias in the direction of commission errors that decentralized structures (polyarchies) produce. However, there is also the mentioned possibility that individual actors may game the decision structure to extract rents by freeriding. Of course, there is yet a further possibility that agents, in some situations, do *not* endogenously adapt their screening function in response to the structure they are located in. While the available empirical evidence supports the conjecture that individual behavior can be adaptive to organizational structure (Reitzig & Maciejovsky 2014; Christensen & Knudsen 2019), it does not systematically identify the conditions under which such endogenous adaptation occurs and what form it takes. The implications of settling this gap in our knowledge are far reaching. If agents endogenously adapt to the way decision structures are designed, unforeseen effects will interfere and possibly undermine the intended aims of the design. It would therefore be important to gauge the overall net effect of organization design. If, on the other hand, agents do *not* adapt to the structural arrangement, organization design has predictable effects, which would broaden its scope of application.

To address the highlighted research gap, we designed a laboratory experiment in which agents located in Sah-Stiglitz-dyads perform a screening task. Using the standard sequential decision procedures introduced by Sah & Stiglitz (1986), we examine how binary choice functions are influenced by aggregation procedures that capture two fundamentally different organizational settings: the decentralized structure known as a 2-member *polyarchy*, and the centralized structure known as the 2-member *hierarchy*. These dyadic structures, which were proposed by Sah & Stiglitz (1986), are the fundamental building blocks in theories of information aggregation in organization science (Christensen & Knudsen, 2010; Csaszar, 2013 & Eggers, 2013). Our aim is to address the open questions implied by the research gap we identified. Do individuals adjust behavior when located in either of these organizational settings? If so, will the organizational setting inspire differences in the way individuals adjust behavior? To investigate whether an actor's comprehension of the mechanics of the structural configuration promotes or prevents adjustment of her/ his behavior, we introduce role ambiguity – induced by randomly switching the individuals' position in each of the two structures. This allow us to address the question: does fixed roles versus fluid roles influence if, or how, individuals adjust behavior? We design a laboratory experiment that allows us to address these questions.

Model

As we are interested in examining how the organizational setting possibly influences individual behavior, we design an experiment where: 1) individual choice functions can reliably be extracted, 2) the organizational setting is the simplest possible, and 3) the organizational setting captures aggregation rules that are fundamentally different. This is achieved by developing binary choice tasks involving evaluation of 2D visual stimuli. Binary choice tasks, binary choice functions, and aggregation procedures are defined consistent with prior work in decision theory (Sah & Stiglitz, 1986, 1988, 1991; Christensen & Knudsen 2010) and signal detection theory as applied to analysis of group performance (Sorkin, Hayes & West 2001; Sorkin, West, & Robinson 1998; Klein & Epley 2015).

Individual choice: In our experimental set-up, individual choice is restricted to the binary option of accepting or rejecting proposed commitments. It captures a situation where an individual must choose to accept an alternative based on noisy and incomplete observations. For example, observations from an interview with a job-candidate may, in conjunction with secondary observations regarding work-history and education, be decisive for a hiring decision. More formally, an alternative is represented as a list of variables, x, that describes the available observations on the alternative, which is generally conceived of as an investment project. The quality of such a project, $y = Y(x)$, is a function of the observable variables and it is independent of past and future choices. Quality indicates the value of projects to the organization, and projects are labeled *good* or *bad* depending on the sign of the quality. The available projects are independent and identically distributed, with density $g(x)$, inducing a distribution over quality.

Individual choice functions: Binary choice functions, which are commonly referred to as psychometric functions or screening functions, are completely characterized by the probability for acceptance (or similarly, the error rate) for each option. The screening function, $p_a(x)$, is the probability that agent *a* will accept the project x. Agents make conditionally independent choices for each separate project, i.e. any pair of decisions by two agents is independent for a given x.

Aggregation procedure: The organizational design consists of an aggregation procedure, which can be thought of as an authority structure that delegates decision rights, as well as an incentive structure assigning payoffs to agents according to their choices. The structures considered here (as in Sah & Stiglitz 1986; Christensen & Knudsen 2010; Csaszar, 2013 & Eggers 2013), are the 2-member *hierarchy* (H) and the 2-member *polyarchy* (P). While the hierarchy is a stylized representation of centralized firm-like bureaucracies, the polyarchy represents decentralized entities competing in a market setting (Sah & Stiglitz 1986). These structures are illustrated in **Fejl! Henvisningskilde ikke fundet.**. Likewise, two different incentive structures are considered here, a collaborative and a competitive one. As control group, we use the single agent acting alone (A). For a single agent acting alone, the aggregation procedure is trivial; the individual screening function simply becomes the probability for organizational acceptance.

– Figure 1 here –

In the hierarchy, both agents must promote a project for the organization to accept and thereby commit to the project. This is analogous to the logical *and*-operator: Organizational acceptance is conditioned on the first *and* the second evaluator's acceptance. This decision rule delegates different rights to the agents. The first evaluator considering the project has the right to single-handedly (independent of the other agent) reject the project on behalf of the organization but is not empowered to force an acceptance. In contrast, the second evaluator has the right to make an organizational choice of both acceptance and rejection. The project is delegated from the first to the second evaluator via an acceptance by the first evaluator. Agents often have fixed roles in a firm, i.e. one agent always "goes first." An incentive structure that inspires collaboration is chosen for the hierarchy by letting the agents share a common payoff function.

In the polyarchy, both of the individual agents must reject a project for the organization to reject and thereby not commit to the project. This is analogous to the logical *or*-operator: Organizational acceptance is conditioned on acceptance by the first evaluator *or* on the acceptance by the second evaluator. This decision rule allocates decision rights in a way that differs from the hierarchy. In the polyarchy, the first evaluator considering the project has the right to single-handedly accept the project on behalf of the organization but is not empowered to force a rejection. The second evaluator in the sequence, by contrast, is empowered to force a final rejection. The project is delegated from the first to the second evaluator via a rejection by the first evaluator. Agents often have randomly assigned roles in a market, i.e. they race to be the first evaluator. An incentive structure that mimics competition can be chosen for the polyarchy by letting the first promoting agent receive payoff while giving nothing to the other. Alternatively, the polyarchy may be thought of as a collaborative structure, in which agents are incentivized by equal shares of payoffs. This, configuration is referred to as a *coordinated polyarchy* (CP).

The two basic aggregation procedures – polyarchy versus hierarchy – have very different properties when viewed as functions over individual screenings. For this reason, the polyarchy and the hierarchy will select different subsets of a pool of projects, implying significant differences in organizational level performance. Of course, the predicted behavior of the two structures is derived under the assumption of fixed agent behavior, i.e. agents do not endogenously adapt their screening functions. In the symmetric situation where the agents in a given structure have identical, fixed screening functions, the probability of organizational acceptance for the hierarchy is $f_H = {p_H}^2$. Similarly, the probability of organizational acceptance for the polyarchy is $f_P = p_P\,(2 - p_P)$.

Prediction of performance levels are complicated by the possibility of interdependency between aggregation procedures and individual choice functions. In the above expressions for organizational acceptance probabilities, the agent screenings are labeled with indices that identify the structure in which the agent resides, i.e. polyarchy versus hierarchy. This indicates that an agent's choice may not only depend on the project itself, but also on the agent's perception of the structural configuration, which defines that agent's decision rights.

At a coarse-grained macro-perspective, where organizations are viewed as unitary actors represented by organizational level choice functions, problems relating to the inner workings of the organization are bracketed. A more fine-grained micro-perspective, where each organization is comprised of multiple actors, introduces well-known considerations relating to the organization's incentive structure. Each agent may observe several performance indicators: the overall organizational performance, the individual share of the organization's payoff, and several additional metrics such as error rates and accuracy measures. The rational agent will choose to optimize individual payoffs, but even when the interests of the individual agent and the organization are aligned, the optimal choice function is not the same for the individual and the organization since the aggregation procedure is generally not an identity mapping. For example, in a very conservative organization, like a hierarchy, the agent might compensate for the organization's tendency to reject too many projects by accepting a larger number of projects. This would be achieved by giving projects of marginal quality the benefit of doubt. In any case, considering aggregation of contributions from multiple agents involves game theoretical considerations in terms of *level-k thinking* (Nagel 1995; Camerer, Ho & Chong 2004) regarding how two or more agents collaborate. In contrast, the boundedly rational agent is likely to evade advanced rationalizations and revert, deliberately or unwittingly, to optimization under simplifying assumptions, e.g. that it is useless to game the organizational structure because it is opaque. This contrast between characterizing agents as fully or boundedly rational captures behavioral assumptions that are difficult to justify from an empirical perspective. In Sah & Stiglitz (1986), this contrast, or tension, is represented by two mutually exclusive behavioral assumptions:

- Alternative a) Agent screening functions are fixed.
- Alternative b) Agent screening functions are strategically adapted.

The first alternative (fixed screening) is consistent with the situation where agents are unwilling or unable to adapt their behavior. It is also consistent with the situation where boundedly rational agents try to optimize organizational performance while completely ignoring the effect of aggregation, i.e. they pick choice functions the would be optimal for a single agent acting alone. The second alternative (adaptive screening) is consistent with rational agents who condition their behavior relative to the impact of the aggregation procedure. Figure 2 illustrates the difference between fixed and adaptive screening functions for the polyarchy and the hierarchy.

– Figure 2 here –

The assumption of fixed versus adaptive screening lead to testable predictions as stated in propositions 1 and 4 in Sah & Stiglitz (1986). The first concerns organizational bias relative to the individual agents reservation level (Sah & Stiglitz, 1986, p. 722): “a project is accepted if its [desired] profit is above the reservation level, R, and it is rejected otherwise.” Thus, an unbiased agent has symmetrical Type I and Type II error around its reservation level. Under the assumption that agents have fixed ability to screen proposals, both hierarchies and polyarchies will be biased relative to the individual agent. In this sense, Sah & Stiglitz’s Proposition 1 addresses organizational bias: Given identical agents, a polyarchy selects a larger proportion of the available projects (more commission errors) while a hierarchy selects a smaller proportion of the available projects (more omission errors). Their Proposition 4 addresses endogenous screening: Based on the organizational context, the agents adjust their reservation level so that in a polyarchy they are more conservative than in a hierarchy, and in a cooperative polyarchy they are more conservative than in a competitive polyarchy. Because the validity of the behavioral assumptions that justify Proposition 4 are largely unknown, we devised experiments to probe if individuals adjust behavior when located in either of the two organizational settings (polyarchy versus hierarchy) and, if so, what form such endogenous adaptation would take.

Experiments

We conducted two experiments. The first tested the three classical decision structures: hierarchy (centralized), polyarchy (decentralized), and the coordinated polyarchy (incentive-aligned version of polyarchy). It generated two surprising findings. First, differences in incentive schemes matter very little for the behavior of individuals in the coordinated and uncoordinated polyarchy. Second, and more puzzling, we found that behavior may either be fixed or adaptive contingent on the type of organization. To uncover whether role assignment or delegation was driving the response, we conducted a second experiment.

Experiment 1. Inspired by the Signal Detection literature (Sorkin et al. 1998) we submitted subjects to a sequence of trials in which their task was to categorize a 2D geometric image. Each image has a fixed background consisting of a horizontal black line evenly dividing a yellow square surface (500 pixels wide on identical 1024x768 resolution screens). Onto this background is put a red circle whose center is placed on the line. This red circle will randomly vary in size and 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 [1/10,9/10] 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], the horizontal measure of the surface containing this image. Each image is assigned the quality $Y = 4/9\, x_1 + 5/9\, x_2$, meaning that images with positive and negative values are equally likely. For each image that served as input, exemplified in Figure 3, the task of the subject was to "accept" or "reject" that input. The organizational design aggregates individual choices into an organizational choice function, which determines whether an image is "accepted" or "rejected." Organizational acceptance implies commitment to receive feedback (good vs bad) and to receive payoff (+1 vs –1) according to the quality of the project (positive or negative). Organizational rejection implies absence of commitment, and therefore zero payoff and no feedback. The subjects were not informed about the quality mapping, as they were supposed to learn from feedback.[1] With this specification of visual stimuli and the definition of a mapping that captures the "quality" of the images, the (organizational) task can roughly be describe as accepting circles that are either large or located to the right, and otherwise reject. The subjects performed the experiments separately in small cubicles, and they did not communicate except through choices in joint decision making.

The experiment was divided into two stages, a training stage and a performance stage. Each stage was initiated with a set of instructions and a questionnaire to check whether the subjects had internalized these instructions. During the 100 trials of the training stage, the subjects worked in

[1] For robustness a number of variants of the experiment were prototyped for the control group setting (A), using different visualizations and different dimensionalities of the problem. The symmetric feedback condition was also checked, i.e. providing feedback on project quality both in case of rejection and acceptance. It sped up learning but yielded very similar individual screening functions.

solitude and they were free to experiment and learn as the outcome of this stage did not generate pay offs. For the performance stage, subjects were allocated to each the 4 following organization types: control group which is the single agent acting alone (A), the hierarchy (H), the (competitive) polyarchy (P), or the (cooperative) coordinated polyarchy (CP). The subjects were then permanently assigned, randomly and anonymously, to one of these organizations. Their task was to work for 200 trials in an attempt to achieve a high bonus payment. The layout of the experiment is sketched in

Figure **3**.

– Figure 3 here –

Although the categorization task was the same in both stages, the introduction of an organizational context added significant complexity compared to the training stage, particularly regarding the ordering/timing of the evaluations and the division of feedback and payoff between agents.

Each subject evaluates a sequence of randomly generated images. For the hierarchy, the same subject was always tasked with evaluating first, whereas for the polyarchy, the roles as first and second evaluator was, for each new image, assigned at random. Although the subjects were instructed and quizzed on the details of the aggregation procedure and the incentive structure of their particular organization, the trial-by-trial information of the evaluations was not explicitly available.

– Figure 4 here –

A test subject receives feedback under the form of payoff only if (s)he observed the input and the final decision was to accept. If the input was finally rejected, no feedback is given. The flow, which was repeated many times during each of the two stages is illustrated in

Figure **4** for case of the polyarchy. For each visual stimulus, the two agents were randomly assigned to a role as either first (A) or second (B) evaluator. The first evaluator received the input (indicated by small white dot) and was allowed to decide, while the other evaluator was put on wait without seeing the input (indicated by the black dot). If a second opinion was not required after the first agent had made a choice, a report screen was shown to both agents (omitting feedback for the inactive party), and both were prompted to continue to the next input. On the other hand, if a second opinion was required, the first evaluator was put on wait while the case was shown to the second evaluator who then had to decide, leading both to a report screen (containing feedback for both parties) and a prompt to continue to the next input.

We decided on the number of subjects on the basis of prototype experiments, which indicated that a sample size of 15 subjects in each treatment is necessary to obtain statistical significance in difference tests. Subjects were incentivized by a show-up fee and an additional bonus obtained during the performance stage. For each trial during the two stages (training, performance), data were recorded on the input (center, diameter, quality) and on the decision process (sequence of individual decisions, payoff, response times).

Experiment 2. The hierarchy and the (coordinated) polyarchy differ in two main aspects. First, the hierarchy allocates fixed roles to the evaluators, e.g. the same agent is always the first to evaluate a case, whereas the polyarchy has random, case-by-case role assignment. Second, the hierarchy delegates decision rights for the cases which the first evaluator assigns positive value to (logical AND), i.e. if the first evaluator assigns positive value an image, a second opinion is requested. In contrast, the polyarchy delegates decision rights for the cases which the first evaluator assigns negative value to (logical OR). Thus, we are unable to resolve any underlying differences in reaction patterns observed in experiment 1. For this reason, we defined two additional treatments that allowed us to identify to which extent role-assignment (fixed versus random) or logical function of the decision structure (AND versus OR gate) were driving the results. The two new treatments were: random-position hierarchy (RH) and the fixed-position (coordinated) polyarchy (FCP). In RH the roles as first and second evaluator are assigned randomly on a case-by-case basis while organizational acceptance still requires that both evaluators accept (logical AND). In FCP the roles as first and second evaluator are fixed once and for all while organizational rejection still requires that both evaluators reject (logical OR). According to *Experiment 1* incentives mattered little for the polyarchy (random roles) or the coordinated polyarchy (fixed roles). Therefore, we focused on shared payoffs when implementing RH and FCP in *Experiment 2*. All other conditions and parameters of *Experiment 2* were identical to that of *Experiment 1*.

Results

The data were analyzed for consistency, effort and learning, and for properties of the individual and organizational decision processes.

Experiment 1. A total of 107 subjects participated in the first experiment, including the control group (15 subjects). The gender balance of subjects was approximately equal. The distribution of generated image parameters was uniform (according to our post hoc analysis), and no individual subject (or pair hereof) was presented with a degenerate set of images. We did not identify any outliers based on analysis of response times (see plot in Figure 6).

To measure that the subjects were able to perform the task, the receiver operating characteristics (ROC) were compared to random behavior, i.e. the diagonal in the ROC diagram. With the null hypothesis "*the project is good*", the recall (TPR) is the fraction of bad projects that get (rightly) rejected, whereas the fall-out (FPR) is the fraction of good projects that get (wrongly) rejected. For example, fall-out is 0.160 and recall is 0.822 on average in control group A. Observations confirm that the fraction of good projects (images) rejected is significantly (95% confidence level) different from the fraction of bad images rejected, except for one subject (in CP2 treatment) who was therefore removed from the sample. The resulting distribution of subjects on treatments is shown in Table 1. Lacking any prior knowledge of the experiment, we assume that all subjects learned to distinguish between good and bad projects (images) during the experiment.[2]

– Table 1 here –

Experiment 2. A total of 69 subjects participated in the second experiment. Again, the gender balance of subjects was approximately equal. Two subjects (one in each of the RH2 and FCP2 treatments) were removed from the data set because their actions do not deviate significantly from random choice (95% CI). The effective sample sizes are shown in Table 1.

Analysis

[2] The evidence suggests subjects learned from experience. The number of correct decisions across subjects and treatments deviates significantly (5 std) from the expectation obtained from a coin-flipper. The time from which the cumulative success-count consistently deviates more than 5 std from a random walk serves as a conservative upper bound on the point in time subjects have learned from experience. During our prototype experiments the mean of this upper bound was estimated to 59 trials. This estimate served to define the length of the subjects' training session. Post hoc analysis of the conducted experiments do not call this estimate into question.

Testing the fixed versus adaptive agent hypotheses requires that we extract screening functions from individuals. Assuming individual screening is a logit

$$p(x) = \frac{1}{1 + \exp\left(-\frac{y - y_0}{\Delta y}\right)}$$

over quality y, the two parameters y_0 (denoting the reservation level) and Δy (denoting the discriminating ability) can be estimated with the maximum likelihood method. Results are shown in Table 2 both at the organizational and individual level together with the optimal individual reservation level (optimal y_0). The optimal y_0 is calculated via optimization by integration under the assumption that the discrimination ability of the control group A is the best available, while the reservation level is a parameter that can be picked freely. Table 2 shows that the choice of organization (polyarchy versus hierarchy) significantly changes the reservation level compared to the control group (A). The polyarchy produces a significant negative shift in reservation level (more commission errors) while the hierarchy produces a significant positive shift in reservation level (more omission errors). At the individual level, we observed no significant shift in the bias or discrimination of the subjects' screening functions.

– Table 2 here –

Tables 3 and 4 examine the fixed versus adaptive agent hypotheses using acceptance rate, recall, fall-out, and accuracy. For the fixed agent predictions, we use the screening of the individual acting alone as our benchmark, i.e. control group A shown in Table 2. For the adaptive agent predictions, the screening of the individual is assumed to have the discriminating ability of the average subject in control group A while the reservation level is set to the optimal value according to organizational form. The comparisons are performed at both the organizational and the individual level. Estimates for individuals acting in organizations are obtained by averaging over the 2 agents in the organizations. The averages are estimated over the entire performance stage *t*=1-200.[3]

– Table 3 here –

[3] A similar analysis was performed on the latter half of the performance stage *t*=101-200 to exclude initial adaptation to the organizational context taking place within the first 100 trials. The conclusions are the same as those reported above. The shifts in rates are slightly larger but no less significant because of the reduction of sample size.

The results from *Experiment 1* featured in Table 2 and 3 jointly document significant organizational biases. This finding on organizational level shift in bias of screening functions is consistent with prior large-scale empirical research (Csaszar 2012). Figure 5 graphically illustrates the estimated organizational biases (y_0) shown in Table 2. These biases give rise to the observed organizational acceptance rates found in Table 3: polyarchies shift the aggregate, *organizational* screening function to the left (*higher* acceptance rates than the control group, A) while hierarchies shift the aggregate screening function to the right (*lower* acceptance rates than the control group, A). However, the accept rates for each organization or the individuals in those organizations do not differ significantly from the fixed agent predictions. In other words, the agents do *not* significantly adapt the bias in their screening functions. The plots featured in Figure 5 show the average fitted logit at organizational and individual levels for treatments H2 and CP2 (*Experiment 1*) against the two predictions (fixed versus adaptive screening), and against the two treatments RH2 and FCP2 from *Experiment 2*. Even with aligned incentives, individuals do not adapt to remove the biases that these simple organizations introduce. This result implies that the organizational forms in our study can be designed such that they, in the aggregate, eliminate inherent individual biases. Why? Because individuals do not endogenously change bias, it can be eliminated by design.

– Figure 5 here –

While agents do *not* change bias (adapt reservation level), our results from *Experiment 1* indicate that individuals do adapt their discrimination level in some of the treatments. In hierarchies (H2), the recall drops dramatically at the individual level, leading to acceptance of more bad projects (images), an effect that is echoed in a slightly less pronounced drop in organizational recall and accuracy. In contrast, the polyarchies (P2) do not deviate from the fixed agent predictions, neither at the organizational nor the individual level. Why? As previously mentioned, two possible factors may explain why individuals adapt in the hierarchy (H2) and *not* in the polyarchy (P2): role-assignment (fixed versus random) or logical function of the decision structure (AND versus OR gate). As explained below, *Experiment 2* allowed us to determine, which it is.

Further considering the results from *Experiment 1*, note that the hierarchy (H2) also deviates significantly from the adaptive agent predictions. So, while individuals arranged in a hierarchy (H2) respond by adapting, this adaptation does not involve shifting the reservation level to the rationally optimal reservation level. Thus, the idea of rational adaptive agent responses to the aggregation procedure that Sah & Stiglitz (1986) proposed is refuted. In fact, the lack of significant change in acceptance rates at the level of the individual is an indication that the agent's discriminating ability rather than the bias is changed. How is the discriminating ability changed? Figure 5 and Table 2 indicate that the individuals organized in hierarchies (H2) have relaxed their screening, i.e. diminished the slope of their screening function. This suggests that agents in the hierarchy (H2) engage in free-loading.

As do hierarchies (H2), polyarchies (P2) deviate significantly from the adaptive agent predictions. That is, individuals in polyarchies (P2) do not adapt. Interestingly, there is no significant difference between the two polyarchy treatments (P2, CP2), so the subjects appear indifferent to being incentivized for individual or joint performance.

In summary, the findings of *Experiment 1* show that all 2-member organizations (H2, P2, CP2) introduce significant organizational level decision biases. These biases are *not* countered by individual level adaptation of reservation levels, i.e. individual bias does not change when agents are placed in these dyadic structures. This implies that individuals fail to adapt optimally in the sense of maximizing payoffs. This would entail that agents maintain their discrimination level while they, at the same time, optimally shift the reservation level (as formally derived in Sah & Stiglitz 1986). Finally, we showed that individuals in hierarchies (H2) adapt by reducing discrimination while individuals in polyarchies (P2) neither adapt reservation levels (bias) nor discrimination (relaxing the screen). Thus, adaptation that is endogenous to the organization appears to be conditional on the nature of the structure that agents are located in.

– Table 4 here –

Experiment 2 was designed to settle what property of the organization that causes the observed differential response, i.e. reduced discrimination in the hierarchy (H2) versus no discrimination in the polyarchy (P2). The results featured in Table 3 and 4 show that neither RH2 nor FCP2 adapt in an optimal way as Sah & Stiglitz (1986) predicted. Comparing RH2 and FCP2 with the "fixed agent" also show that these two structures differ significantly (Table 3). The adaptive response is strongest in the fixed-position polyarchy (FCP2) and it mirrors the agents' response when located in the hierarchy (H2), suggesting that fixed roles are triggering adaptive responses for individuals in organizations. Reversely, role ambiguity, induced by random role assignment, appears to prevent adaptation to the organizational structure. Since the agents' adaptation in the two fixed-position treatments (H2, FCP2) relax their observed level of discrimination – a response consistent with free-loading – organizational performance is reduced. The implication is that designing organizations with a non-trivial level of role ambiguity prevents gaming the structure, and, more generally, make individual screening performance predictable. In different words, role-ambiguity prevents free-loading because it makes individuals indifferent to working alone or in organizations.

Disturbing this clear picture is the fact that the random-position hierarchy (RH2) also displays a deviation from the fixed agent prediction: the recall drops, albeit less than in the fixed-position hierarchy (H2). The symmetric response is not found in the coordinated polyarchy (CP2), indicating that subjects do perceive the delegation rules of the hierarchy and polyarchy differently. This conjecture is supported by observed differences in response times as shown in Figure 6, i.e. P2 is associated with faster response than RH2. A possible explanation for the difference between P2 and RH2 is as follows. While RH2 has random assignment of positions, the delegation rule is the AND gate, which captures logical conjunction. Many situations in life that imply logical conjunction invite free-loading. If both agents are assumed to take responsibility for accepting a candidate, each will be tempted to free-load. In this case, both are assumed to opt in, so it may intuitively appear easy to "weasel" out. By contrast, the delegation rule in P2 is the OR gate, which captures logical disjunction. Both agents are assumed to opt out, so opting in and supplying effort to free-load may, frankly, appear counterintuitive.

– Figure 6 here –

In summary, the findings of *Experiment 2* show that role-clarity triggers a strong but non-rational adaptative response, which is symmetric across the modes of delegation, i.e. the decision rules. In the presence of role ambiguity, however, only the hierarchical decision rule (AND gate: all must accept) triggers a weaker, non-rational adaptational response (in RH2). The explanation for this could be that the delegation rule of the hierarchy (RH2) appears more intuitive than the delegation rule of the polyarchy. Thus, some subjects may, with difficulty, understand how RH2 can be gamed, i.e. let the other agent do the job independent of your own position in the structure. Given the less intuitive decision rule of the classical polyarchy (P2), which has random assignment of roles, the effective level of role-ambiguity is so high that it effectively prevents attempts to game the structure. The evidence on response times shown in Figure 6 are consistent with this conjecture in the sense that subjects make relatively fast decisions in the hierarchy (H2), which has fixed roles, while they spend significantly more time in the polyarchy even when it also has fixed roles (FCP2). Finally, our results are easily related to two testable predictions proposed by Sah & Stiglitz (1986). We find strong support for their proposition 1 on organizational bias: A polyarchy does select a larger proportion of the available projects (more commission errors) than does a hierarchy (more omission errors). However, we find no support for their proposition 4 on endogenous screening: Individuals do not adapt their reservation levels rationally with respect to expected payoff.

Discussion

Our experiments examined how alternative ways of structuring decision processes affect the quality of organized decisions. We created a screening task in which simple dyads screen a series of proposals and, for each proposal, judge whether it has a quality that justifies acceptance or rejection. This task represents familiar examples of screening that occur when organizations evaluate job candidates, innovation projects, and loan applications. The basic measures of decision making quality in such screening tasks summarize errors made when "good" projects are rejected (Type I error) and "bad" ones accepted (Type II error). Since both types of error have the same cost in our experiments, the sum of such errors is a valid measure of the overall quality of the decisions that get made.

We examined the simplest possible decision making structures that can be created with two agents. These correspond to what Sah & Stiglitz (1986) labeled "hierarchies" and "polyarchies." In hierarchies, the agents decide, in a fixed sequence, whether to reject a proposal; rejection by a single organizational member is irreversible. In polyarchies, agents decide in random order whether to accept a project. Acceptance by any agent is irreversible. In their classical paper, Sah & Stiglitz (1986) derive some important testable propositions. Two of these are especially relevant to our study. The first concerns the bias induced by different decision making structures. Theoretically, hierarchies and polyarchies produce symmetric biases in decision making relative to the reservation levels of the agents located in the structure. An unbiased agent has symmetrical Type I and Type II error around its reservation level, which is the minimal desired level of profits obtained from accepting a project. Under the assumption that agents have fixed ability to screen proposals, both hierarchies and polyarchies will be biased relative to the individual agent. That is, hierarchies reject too many good proposals (Type I error) and polyarchies accept too many good proposals (Type II error) relative to an unbiased individual agent.

This raises the question if theoretical predictions based on fixed ability to screen proposals are borne out on practice. Do individuals for some reason endogenously adjust behavior when located in either of these organizational settings, and if so, will the role they occupy inspire differences in the way they adjust behavior? Our experiments were designed to address this question. We examined if agents endogenously adjust behavior conditional on the structure (hierarchy versus polyarchy) and the role they occupied in that structure (fixed versus random order of evaluation). If agents are aware of the way decision processes are structured, they can rationally (e.g. via Bayes rule) adjust their screening function to modify the bias induced by the organization. Thus, we should observe that agents "correct" their reservation level such that it is higher in hierarchies than in polyarchies – in other words decision makers should be more "conservative" in polyarchies than in hierarchies. Another open question concerns incentives. Sah & Stiglitz (1986) define two different incentive systems for polyarchies. In one system – the basic 2-member polyarchy P2 – each agent is rewarded for individual performance. In the second system – the coordinated polyarchy CP2 – agents are incentivized to maximize joint outcomes. The prediction is that Bayesian agents should be more conservative (higher reservation level) when they have an incentive to maximize the quality of their joint output (Sah & Stiglitz 1986).

Consistent with the findings in Csaszar (2012), the results of our experiments provide strong support for the hypothesis that hierarchies and polyarchies generate organizational level bias. The way the decision process is structured significantly affects the quality of the output generated at the organizational level. As predicted, the effect is remarkably symmetrical. As expected, hierarchies err on the conservative side, denying acceptance of too many good projects, while polyarchies are looser and accept too many bad projects. Thus, the decision structure matters and its effects are predictable.

At the same time, the hypotheses about rational individual adaptation of bias find very little support in our experiments. As a benchmark for individual adaptation we use treatments of individuals acting alone, i.e. uncontaminated from participating in any decision making structure. In this condition, we show that individuals are unbiased. When individuals are embedded in decision making structures, we show that they do not adapt the bias of their screening function relative to the benchmark. That is, the observed bias of decision making structures results from the (on average) unbiased behavior of individuals. No signs of adaptation are found in polyarchies. However, some adaptation is seen in hierarchies where agents exhibit a level of recall that is significantly lower relative to the benchmark, i.e. reduced discrimination. It is further notable that this effect of reduced individual discrimination fails to show up in the aggregate behavior of hierarchies. A plausible explanation is that agents in the hierarchy "relax" their screening function symmetrically around their reservation level. This effect is consistent with free-riding, which is made possible by the fixed role-structure of hierarchies. Finally, we show that incentivizing subjects at the individual or organizational level in the polyarchy does not alter their behavior (as measured by the screening function). In different words, the experimental subjects appear to be substantially indifferent to the incentive structure.

Study 2 addresses some of the questions about endogenous adaptation to organizational structure that were left open by study 1. Why should there be an asymmetry between polyarchy (no individual adaptation) and hierarchy (moderate individual adaptation)? One possible explanation is that the decision process in the polyarchy is cognitively more complex than in the hierarchy. For example, the "mental model" (Johnson-Laird 1983, 2001) representing the decision process in a polyarchy requires more features than the representation of a hierarchy does. Thus, it may be easier for subjects embedded in hierarchies to "read the structure" and take it into consideration, e.g. acting to improve it or perhaps rather game it to reduce supply of effort. An alternative explanation is that the agents condition their behavior on the role they occupy, but they can only do so if the role is fixed. In our study 1, subjects in hierarchies have a fixed position in the decision process, while subjects in polyarchies are assigned a random position. Recent experiments (Agranova & Schotter 2012) show that role ambiguity can affect the capacity of individuals to respond strategically to a given interactive situation. To examine whether this explanation is consistent with our findings, we reversed the role structure used in study 1. That is, we ran experimental treatments in which the position of each agent in a *hierarchy* was randomly assigned, while it is fixed in the case of the *polyarchy*.

Our results support the explanation that fixed roles allow agents to condition behavior relative to the role they occupy. When the role in a polyarchy is fixed (FCP2), agents endogenously adapt their screening behavior (fall-out is significantly lower than in the benchmark provided by an individual agent acting alone). As in the hierarchy with fixed roles (H2), the agents "relax" their screening function symmetrically around their reservation level, an effect that is consistent with free-riding. In contrast, in the treatment where the role in the hierarchy is randomly assigned (RH2), this response is damped. Why does the classical polyarchy (P2), in which roles are randomly assigned, completely eliminate endogenous adaptation while there is some endogenous adaption in the hierarchy with randomly assigned roles (RH2)? Our conjecture, while only suggestive, is that some agents game the hierarchy even if roles are randomly assigned because the delegation rule of the hierarchy appears more intuitive than the delegation rule of the polyarchy. Given the less intuitive decision rule of the classical polyarchy (P2), the effective level of role-ambiguity may be so high that it prevents attempts to game the structure. Even if the evidence is not conclusive, our finding opens the exciting perspective that role ambiguity may be a design variable that can be administered to remove free-loading. Equally important, role ambiguity, paradoxically, makes individual behavior predictable. Our conjecture is that this effect occurred because individual agents who experience role ambiguity desist adapting to the structure and rather continue to behave as individuals acting alone.

Thus, study 1 and 2 jointly imply that: (1) the decision structure does *not* influence behavior of agents who experience a high level of role-ambiguity, (2) agents adapt behavior if they are located in a structure with fixed roles, in which case (3) they endogenously change behavior in a myopic way – to reduce supply of effort – rather than optimally correcting the bias introduced by the structure. Our results have implications for the study of organization design as they indicate organizations can be designed to reduce, and possibly remove, cognitive biases in decisions. This exciting perspective complements prior research that has almost exclusively focused on identifying, explaining, and possibly countering biases of individual human actors (see e.g. Gärling, Kirchler, Lewis, van Raaij

2009). Of course, it is important to understand how pervasive cognitive biases may undermine markets, cause financial crises (Gärling et al. 2009), and influence hiring decisions. While diversity initiatives, process audits, testing and increased self-awareness can help (e.g. Bendick & Nunes 2012), we offer a complementary perspective that organizations can be designed to reduce biases. In this regard, it is very encouraging that our results show that individuals may *not* change bias conditional on the structure they are embedded in. Equally important, with high levels of role ambiguity, individual agents appear to give up adapting to the structure and rather continue to behave as individuals acting alone. The implication is that, given sufficient high levels of role ambiguity, structures can be designed to remove bias and increase tolerance. This effect will be present at the organizational level. A hiring committee that would process applicants according to a polyarchy rule effectively becomes tolerant. Of course, the cost is an increase in commission error. Thus, an important topic for future research is to calibrate the gains of reduced bias in decision processes that can be obtained from organization design against the cost of applying such design.

In this regard, an important limitation of the present study is that our results were obtained from experiments on the simplest possible dyadic decision structures, the polyarchy and the hierarchy. It is important to point out that the results from the present study cannot be directly extrapolated to larger structures. For example, Christensen & Knudsen (2010) have shown that there are more complex configurations of the screening process that can be structurally bias-free. As these structures are more complex than the simple dyadic structures we have explored, an interesting question arises concerning the design trade-offs between the cost of more complex structures and the cost of biased choices in simpler ones. Another related question concerns the effect of role ambiguity in more complex structures. Surprisingly, the random allocation of roles in the 2-member polyarchy in conjunction with the somewhat opaque delegation rule in that structure (OR gate) induced a level of role ambiguity that completely removed endogenous adaptation to the structure. A further important question to be settled in future research is to determine the level of complexity of decision structures – in terms of size, delegation rule, and role assignment – that removes individual attempts to free-ride or otherwise game the organization. For the theorist, this suggests that realistic models of organizations should assume that levels of strategic response of agents to organizational mechanisms are limited

(Nagel 1995; Camerer, Ho & Chong 2004; Costa-Gomes & Crawford 2006). For the practitioner, a benevolent and behaviorally aware designer should take into account such limits to maximize collective performance.

Conclusion

This study advances prior empirical work that pointed to the existence and complications of endogenous adaptation in decision structures (Reitzig & Maciejovsky 2014; Christensen & Knudsen 2019). Our contribution is to systematically identify conditions that predict the presence/ absence of endogenous adaption to the structure individuals are located in. Thus, our results from multiple lab experiments show that role ambiguity reduces endogenous adaptation. As role ambiguity is increased, adaptation ceases and the subjects behave as individuals acting alone even when they are located in an organization. In contrast, role clarity leads to endogenous adaptation in the sense that the subjects behaved as if they reduced supply of effort.

Our research addresses an important gap in the literature on information aggregation in organization science (Christensen & Knudsen 2002, 2010, 2013, 2019; Knudsen & Levinthal 2007; 2013, 2018; Csaszar 2013; Csaszar & Eggers, 2013). Central to these theories is that organizations have a fundamental role in aggregating choice functions of individual agents. From a theoretical as well as a practical perspective, it is critical to understand if, how, and when individuals endogenize behavior conditional on the structure of the organization they are located in. From a theoretical perspective, such endogenous adaptation complicates derivation of clear predictions about the systematic behavior of the organization (Sah & Stiglitz 1986). The related practical implication is that organized decision structures will behave in ways that are ill understood and potentially counter-productive. In this perspective, our finding that role ambiguity may dampen or even remove endogenous adaptation is very encouraging. When agents behave as individuals acting alone independent of the structure they are located in, organization design can shape organizational decision structures in ways that potentially generate robust and easily predictable behavior.

More broadly, we have developed an experimental design that can bridge the divide between the micro-perspective of interacting decision makers and the macro-perspective of the resulting organizational behavior and performance. Theories of organizational decision making explicitly embraces both the micro- and the macro perspective and provide a powerful mathematical approach to derive aggregate decision functions that capture the way individuals interact (Sah & Stiglitz 1986; Christensen & Knudsen 2010; Csaszar & Eggers 2013). Of course, this implies that the specification of individual-level assumptions will influence predictions regarding organizational-level behavior and performance. Our approach allowed us to test the behavioral assumptions of organizational decision theory (Christensen & Knudsen) and identify individual-level processes that caused endogenous adaptation to the structural configuration at the organization level. A critical topic in this regard is to determine the effect of role ambiguity in more complex structures. When role ambiguity removes individual-level endogenous adaptation, the derivation of organization-level predictions from processes becomes much more reliable. Thus, we suggest that researchers determine how behavioral assumptions are contingent on the context within which individuals interact, and then specify aggregation processes that are sensitive to those assumptions. We are enthusiastic about the outlined *neo-contingency approach* and hope it will inspire further progress in the study of organizational decision processes.

References

Agranova M, Schotter A (2012) Ignorance is bliss: an experimental study of the use of ambiguity and vagueness in the coordination games with asymmetric payoffs. *American Economic Journal: Microeconomics*. 4(2): 77-103.

Arrington TS, Brenner S (1980) The Advantages of a Plurality Election of the President. *Presidential Studies Quarterly*. 10: 476-482.

Bendick Jr. M, Nunes AP (2012) Developing the Research Basis for Controlling Bias in Hiring. *Journal of Social Issues*. 68(2): 238-262.

Camerer CF, Ho TH, Chong JK (2004) A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*. 119(3): 861–898,

Christensen M, Knudsen T (2002) *The architecture of economic organization: Toward a general framework*. (Mimeo, University of Southern Denmark: Odense, Denmark).

Christensen M, Knudsen T (2010) Design of Decision-Making Organizations. *Management Science*. 56(1): 71-89.

Christensen M, Knudsen T (2013) How Decisions Can Be Organized and Why it Matters. Journal of Organization Design. 2(3): 41-50.

Christensen M, Knudsen T (2019) Division of Roles and Endogenous Specialization. *Industrial and Corporate Change*. (Special Issue in the Honor of James G. March, Forthcoming).

Cohen MD, March JG, Olsen JP (1972) A garbage can model of organizational choice. *Administrative Science Quarterly*. 17(1):1-25

de Condorcet J (1785) Essai sur l' application de l' analyse à la probabilité des décisions: rendues à la pluralité des voix. (Imprimerie Royale, Paris).

Costa-Gomes MA, Crawford VP (2006) Cognition and Behavior in Two-Person Guessing Games: An Experimental Study. *American Economic Review*. 96(5): 1737- 1768

Csaszar FA (2012) Organizational structure as a determinant of performance: evidence from mutual funds. *Strategic Management Journal*. 3(6): 611–632.

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

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

Fishburn PC (1977) Condorcet Social Choice Functions. SIAM Journal on Applied Mathematics 33: 469-89.

Galton F (1907) ”Vox populi” *Nature*. 75 (1949): 450–451.

Gerling K, Grüner HP, Kiel A, Schulte E (2005) Information Acquisition and Decision Making in Committees: A Survey. *European Journal of Political Economy*. XXI:563–597.

Gärling T, Kirchler E, Lewis A, van Raaij F (2009) Psychology, Financial Decision Making, and Financial Crises. *Psychological Science in the Public Interest*. 10(1): 1-47.

House RJ, Rizzo JR (1972) Role conflict and ambiguity as critical variables in a model of organizational behavior. *Organizational Behavior and Human Performance*. 7: 467-505.

Jackson SE, Schuler RS (1985) A Meta-analysis and Conceptual Critique of Research on Role Ambiguity and Role Conflict in Work Settings. *Organizational Behavior and Human Decision Processes*. 36: 16-78.

Johnson-Laird PN (1983) *Mental models: Towards a cognitive science of language, inference, and consciousness* (Harvard University Press, Cambridge MA).

Johnson-Laird PN (2001) Mental models and deduction. *Trends in Cognitive Sciences*. 5(10): 434-442.

Klein N, Epley N (2015) Group discussion improves lie detection. *Proceedings of the National Academy of Sciences of the United States of America*. 112(24): 7460-7465.

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

March JG, Simon H (1958). *Organizations* (Wiley, New-York).

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

Nagel R (1995) Unraveling in Guessing Games: An Experimental Study. *American Economic Review*. 85(5): 1313-1326.

Page S (2007) *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies* (Princeton University Press, Princeton, NJ).

Sah RK, Stiglitz JE (1986) The architecture of economic systems: Hierarchies and polyarchies. *The American Economic Review*. 76(4): 716-727.

Sah RK, Stiglitz JE (1988) Committees, hierarchies and polyarchies. *The Economic Journal*. 98(391): 451-470.

Sah RK (1991) Fallibility in human organizations and political systems. *J. Econom. Perspect*. 5(2): 67–88.

Simon HA (1947) *Administrative behavior: A Study of Decision-Making Processes in Administrative Organizations* (Free Press, New York, NY).

Sorkin RD, Hays CJ, West R (2001). Signal-detection analysis of group decision making. *Psychological Review*. 108(1): 183-203.

Sorkin RD, West R, Robinson DE (1998) Group performance depends on majority rule. *Psychological Science*. 9(6): 456-463.

Stiglitz JE (2002) Information and the Change in the Paradigm in Economics. *The American Economic Review*. 92(3) : 460-501.

Stiglitz JE (2004) Information and the Change in the Paradigm in Economics, Part 2. *The American Economist.* 48(1) : 17-49.

Surowiecki J (2004) *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations* (Doubleday & Co., New York, NY).

Swank O, Visser B (2008) The consequences of endogenizing information for the performance of a sequential decision procedure. *Journal of Economic Behavior & Organization*. 65: 667–681.

Young HP (1988) Condorcet's theory of voting. *American Political Science Review*. 82: 1231-1244

Tables

	Treatment	Samples	Description of organization
Control	**A**	18	Single agent
Exp 1	**H2**	16	Hierarchy of 2 members
	P2	15	Polyarchy (individual payoffs) of 2 members
	CP2	14	Coordinated polyarchy (shared payoffs) of 2 members
Exp 2	**RH2**	14	Random-position hierarchy of 2 members
	FCP2	17	Fixed-position, coordinated polyarchy of 2 members

Table 1: The effective sample sizes by treatment for experiment 1 and 2.

Level	**Organizational**		**Individual**		
Treatment	y_0	Δy	y_0	Δy	**Optimal y_0**
A	–0.0209	0.1476	–0.0209	0.1476	+0.0000
H2	+0.1486 ***	0.1465	–0.0501	0.1847	–0.1244
P2	–0.1619 **	0.1297	–0.0004	0.1478	+0.0480
CP2	–0.1956 **	0.1387	–0.0288	0.1519	+0.1244
RH2	+0.1431 ***	0.1622	–0.0306	0.1681	–0.1244
FCP2	–0.2213 ***	0.1598	+0.0424	0.2869	+0.1244

Table 2: Average screening parameters (reservation level y_0 and discrimination Δy) of organizations and individuals by treatment. The variation of screening capabilities is large across individuals, and the hierarchical treatments show significantly larger variation than the control group. Consequently, the differences to the control group are only significant at the organizational level († p<0.1, * p<0.05, ** p<0.01, *** p<0.005). The optimal, individual reservation is calculated by numerical integration assuming fixed Δy=0.1476.

level	treat-ment	observed average				fixed agent prediction			
		acc. rate	recall	fall-out	accuracy	acc. rate	recall	fall-out	accuracy
organization	A	50.4%	82.2%	16.0%	83.1%	51.6%	81.8%	15.0%	83.4%
	H2	38.0%	92.3% *	30.0% †	81.4% *	40.1%	94.1%	25.8%	84.1%
	P2	62.7%	70.0%	5.4%	82.7%	63.2%	69.4%	4.1%	83.4%
	CP2	62.7%	69.2%	3.6%	82.4%	63.2%	69.4%	4.1%	83.4%
	RH2	38.9%	92.2% *	28.9%	81.6% †	40.1%	94.1%	25.8%	83.4%
	FCP2	65.0%	64.0% †	6.0% *	78.9% **	63.2%	69.4%	4.1%	83.4%
individual	A	50.4%	82.2%	16.0%	83.1%	51.6%	81.8%	15.0%	83.4%
	H2	63.0%	70.0% ***	16.2%	80.1% ***	60.5%	81.8%	15.0%	83.4%
	P2	42.0%	83.8%	18.6% †	82.9%	42.6%	81.8%	15.0%	83.4%
	CP2	42.0%	82.9%	16.4%	83.3%	42.6%	81.8%	15.0%	83.4%
	RH2	60.1%	76.4% *	15.4%	81.3% †	60.5%	81.8%	15.0%	83.4%
	FCP2	39.5%	79.9%	28.6% ***	79.2% ***	42.6%	81.8%	15.0%	83.4%

Table 3: Properties of average decisions at the organizational and individual levels by treatment compared to the predicted rates when assuming a fixed agent screening equal to that of the control group A of $y_0 = -0.0209$ and $\Delta y = 0.1476$ († p<0.1, * p<0.05, ** p<0.01, *** p<0.005).

level	treat-ment	observed average				adaptive agent prediction			
		acc. rate	recall	fall-out	accuracy	acc. rate	recall	fall-out	accuracy
organizational	A	50.4%	82.2%	16.0%	83.1%	50.0%	83.5%	16.5%	83.5%
	H2	38.0% ***	92.3% ***	30.0% ***	81.4% ***	48.4%	87.3%	15.8%	85.7%
	P2	62.7% *	70.0% *	5.4% †	82.7% †	57.8%	77.1%	7.4%	84.9%
	CP2	62.7% ***	69.2% ***	3.6% ***	82.4% †	51.6%	84.2%	12.7%	85.7%
	RH2	38.9% ***	92.2% ***	28.9% ***	81.6% *	48.4%	87.3%	15.8%	85.7%
	FCP2	65.0% ***	64.0% ***	6.0% ***	78.9% ***	51.6%	84.2%	12.7%	85.7%
individual	A	50.4%	82.2%	16.0%	83.1%	50.0%	83.5%	16.5%	83.5%
	H2	63.0%	70.0%	16.2% ***	80.1% †	59.7%	72.0%	8.7%	81.7%
	P2	42.0% *	83.8% †	18.6%	82.9%	46.2%	86.9%	20.6%	83.1%
	CP2	42.0%	82.9% ***	16.4% ***	83.3%	40.3%	91.3%	28.0%	81.7%
	RH2	60.1%	76.4% *	15.4% ***	81.3%	59.7%	72.0%	8.7%	81.7%
	FCP2	39.5%	79.9% ***	28.6%	79.2% *	40.3%	91.3%	28.0%	81.7%

Table 4: Properties of average decision at the organizational and individual levels by treatment compared to the predicted rates when assuming an adaptive agent screening equal to that of the control group A with respect to discrimination, $\Delta y = 0.1476$, and of optimal y_0 († p<0.1, * p<0.05, ** p<0.01, *** p<0.005).

Figures

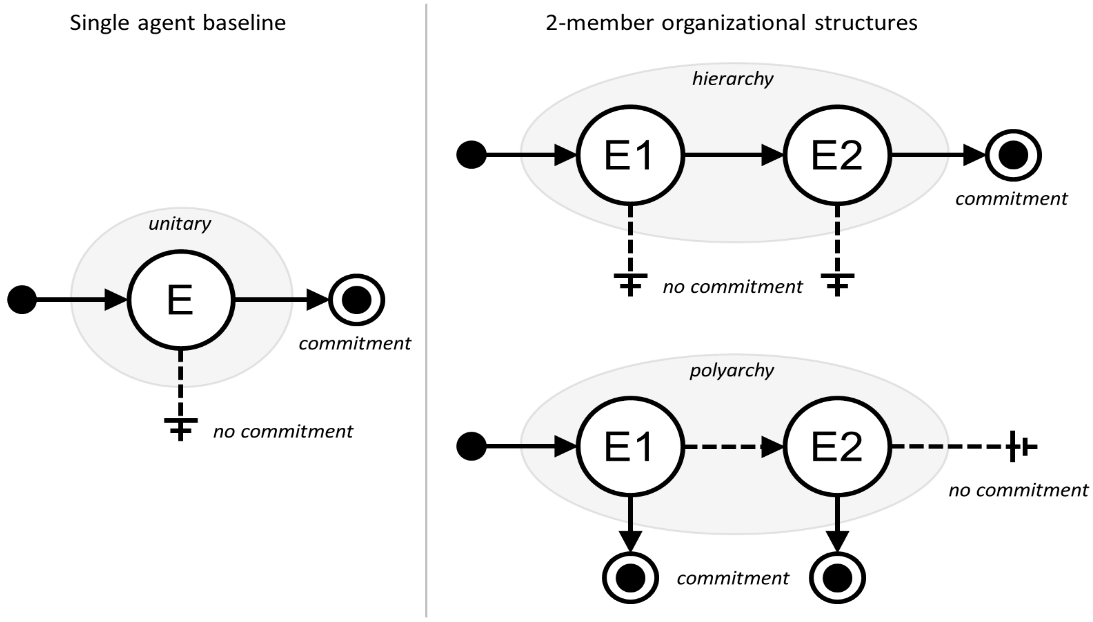

Figure 1: In the bureaucratic hierarchy (a) for a proposal to be accepted it must be evaluated and promoted at each department in the fixed chain of command. In the market-like polyarchy (b) the firms race to first implement promising business opportunities. These two stylized economic systems differ in three ways: 1) the assignment of roles as first (E1) and second (E2) evaluator, 2) the structure for delegating decision rights, and 3) the incentive structure for awarding feedback and payoff.

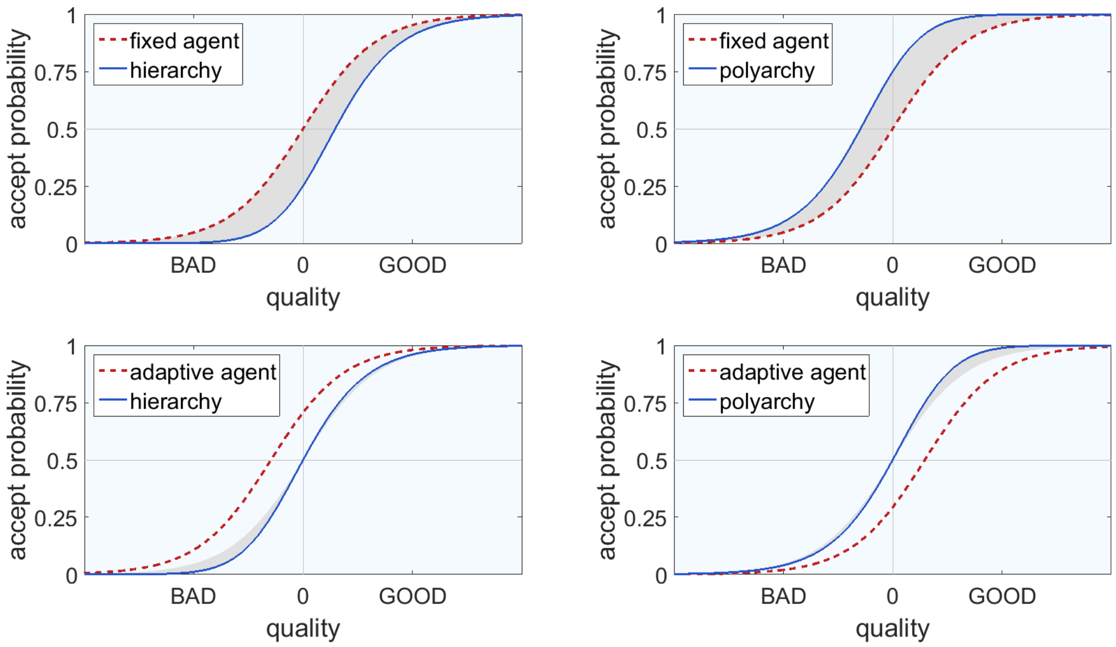

Figure 2: Assuming agents do not change when organized, the two upper graphs compare the ability of the hierarchy and polyarchy, respectively, to that of the unbiased agent. The greyed area is the change in errors of commission and omission. Assuming that agents adapt strategically by picking a suitable reservation level, the two lower graphs illustrate it leads to superior performance (measured by error reduction) for both organizational forms.

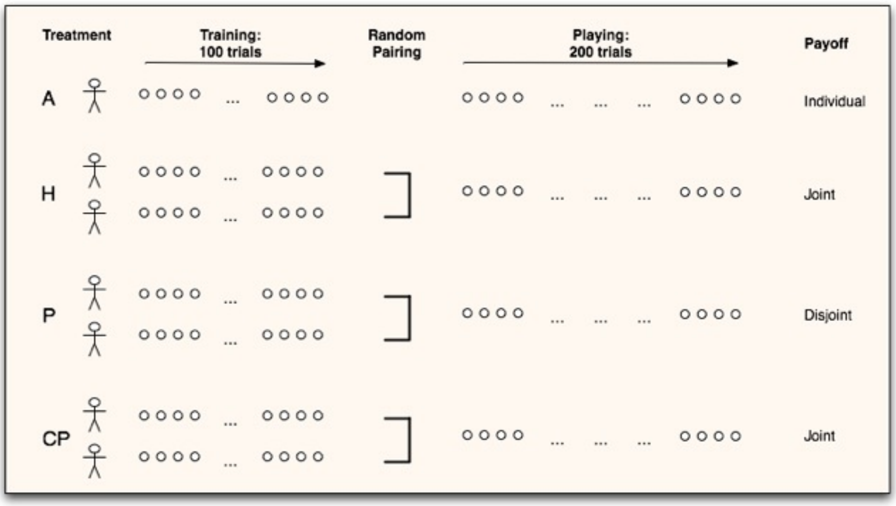

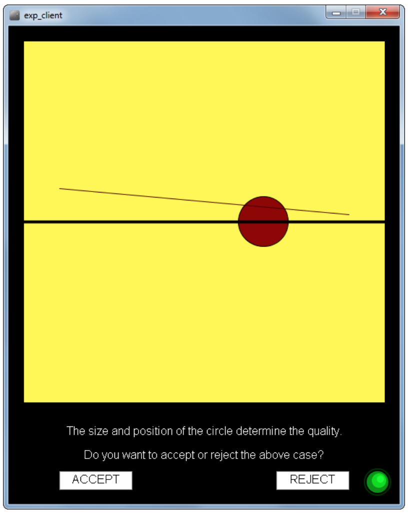

Figure 3: The organizational screening experiment 1 includes an initial individual training run and a subsequent organizational performance run. Between the two stages the subjects were randomly selected and paired into the different structural treatments: single agent acting alone (A), hierarchy (H), polyarchy (P), and coordinated polyarchy (CP). The screenshot to the right shows an example of the typical image that the subjects encountered, except for the slanted line (image is good if the red circle touches it) which was not present in the experiment.

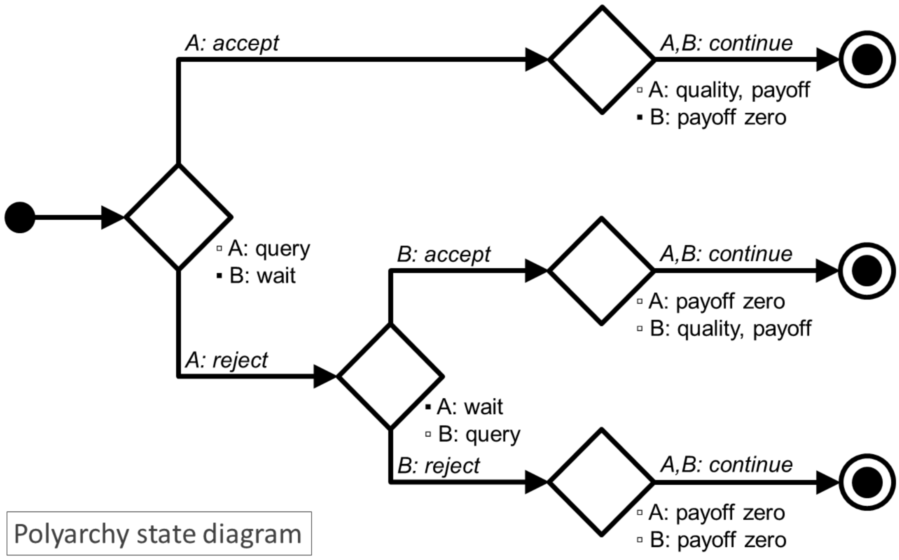

Figure 4: The state diagram of the decision flow for a single trial in the polyarchy (P) shows the choices and consequence of the first and second evaluator, here labeled A and B respectively. In contrast, the hierarchy (H) would terminate without requiring B to decide in the case where the first evaluator rejects, and the only route to non-zero shared payoff is via two acceptances.

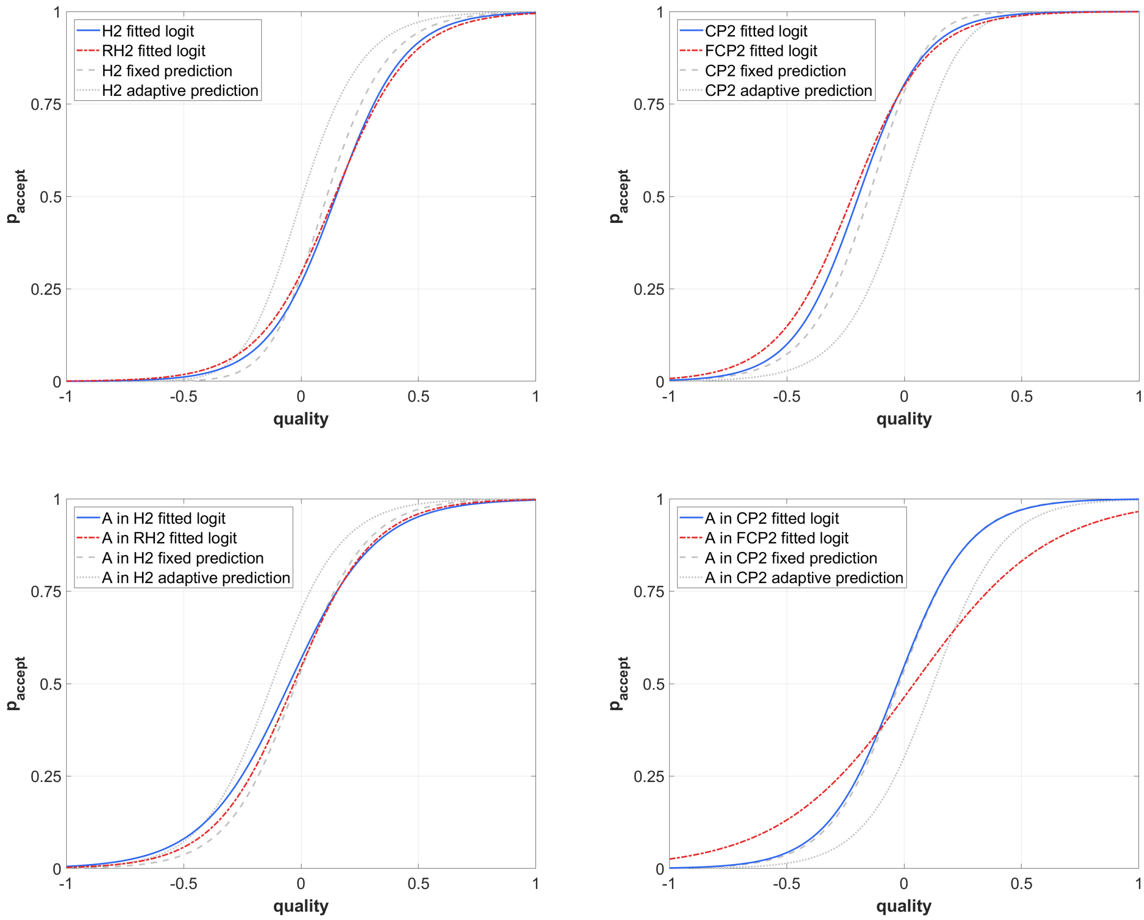

Figure 5: The average organizational (top) and individual (bottom) screening functions for the hierarchy (left) and the polyarchy (right) compared to the fixed and adaptive predictions (greyed out). Solid lines represent experiment 1 organizational forms while dashed lines represent experiment 2 alternative forms.

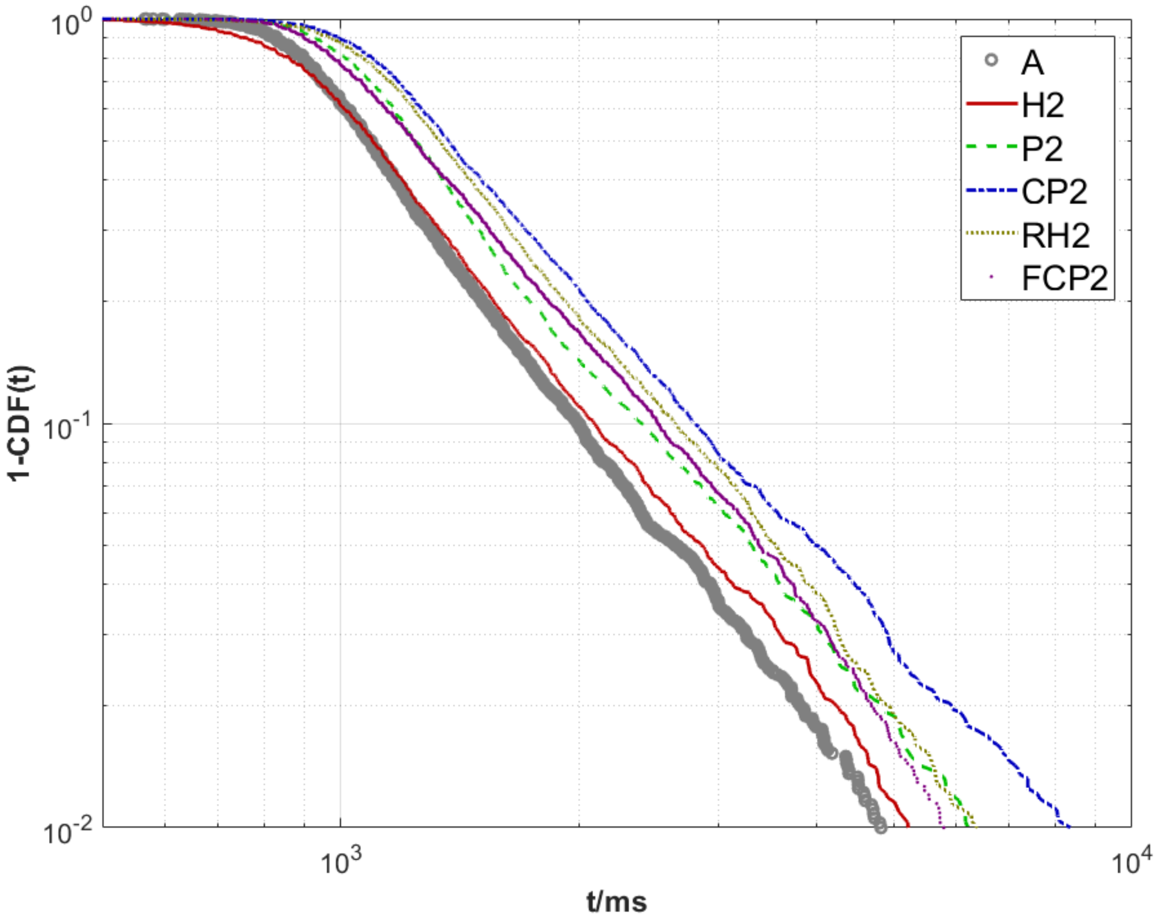

Figure 6: The distribution of response times varies across the treatments, showing significant differences in the time subjects took to deliberate their choices in different organizational designs.