ONLINE APPENDIX 1:

DETAILED DESCRIPTION OF MODEL COMPONENTS & PARAMETERS

This appendix presents the equivalent of section 3 of the main paper (Model Structure), but provides a more detailed description of parameters and their implementation. Its objective is to facilitate further developments of the model by other authors.

1. Task environment

Task environments capture a mapping between the possible courses of action available to agents, and the likelihood for each course of action to attain an agent's desired goal. In our model, the task environment is characterized by the pattern of *interdependence* between agents in the organization. This is referred to as the "***task structure***" throughout the paper and denoted by ***T***.

We assume that T arises through a process whereby an organization designer uses formal authority to divide the organization's goals into a set of tasks, and allocate bundles of such tasks ("formal roles") to specific agents. This is the process of division of labor. In a world where the activities of organizations are imperfectly decomposable into independent clusters of subtasks and boundedly rational designers are unlikely to identify such clusters even if they existed, it is inevitable that organizations often allocate different agents with tasks which are interdependent, i.e. whose values affect one another (Heath and Staudenmayer 2000, Simon 1962, Sosa et al. 2004, Thompson 1967). For instance, an agent's task may constitute the input to another agent's task, or two agents may work jointly on the same input. Whenever two agents are allocated interdependent tasks, a need for interaction arises: agents must exchange information and/or material (Thompson 1967).

We thus represent an m-member organization's task environment T as an m x m matrix. Whenever agent i depends on agent j in the task structure, the cell T_{ij} takes a value of 1, and 0 otherwise. The number of dependences between agents in the task environment is tuned by the parameter $\gamma_T \in [0,1]$ which constitutes the density of T, i.e. the proportion of off-diagonal cells filled with ones in the matrix T.

High values of γ_T indicate that the task environment is characterized by a high level of interdependence between agents. In the basic version of this model, we consider a case where the task structure is symmetric and moderately dense: any pair *(i,j)* of agents has a 0.5 probability of being interdependent, and all interdependences are reciprocal so that $T_{ij} = T_{ji} = 1$.

The realized ***social interaction structure***, ***R***, represents the pattern of realized interactions between agents. We represent this structure as a symmetric m x m matrix R where cells filled with 0 indicate that two agents do not interact, while cells filled with 1 indicate that two agents do interact. Each realized link in the m x m matrix R arises from the double correspondence of interaction requests between two agents, i.e. when two agents are both willing to interact. The matrix R represents the intra-organizational social structure typically studied by network scholars, often referred to as “task-related” (e.g. Casciaro and Lobo forthcoming) or “advice” network (e.g. Sparrowe and Liden 2005). Within this network, links enable information and resources to flow between agents (Tsai and Ghoshal 1998: 467), so that two agents who have a social interaction are better able to collaborate in managing their dependencies.

The symmetry of R (i.e. $R_{ij} = R_{ji} = 1$) reflects the idea that social relationships only exist when both sides acknowledge them (Homans 1958, Levine and Prietula 2012: 1750, Salancik 1995: 346-347). This is also a realistic representation of information flows between agents: while a specific piece of information may flow from one agent to another in a directed manner (Gargiulo et al. 2009), two agents who consider each other as bound by a social relationship will typically exchange information in a bi-directional manner over a period of time (Morelli et al. 1995). Finally, this is also in accordance with previous empirical results showing that inward and outward flows of information within a pair of agents tend to be highly correlated (e.g. Cross et al. 2001: 229). The reciprocity inherent in social interactions has prompted network scholars to consider social capital as a “jointly owned asset” (Burt 1992: 9).

The ***performance of agents***—and, in aggregate, the performance of the organization—results from the correspondence between the task structure T and the social interaction structure R. It is thus a measure of correctness of R given T. This is because interdependence creates a requirement for interaction between agents (Thompson 1967, Tushman and Nadler 1978), so that mismatches between the matrices *T* and *R* create a potential for coordination and/or cooperation failures (Puranam et al. 2012: 425).

More specifically, agents can suffer from two types of errors in forming interactions with each other (viewed purely from the perspective of organizational performance). An ***error of omission*** happens when an agent fails to interact with another agent whom it depends on (i.e. $T_{ij} = 1$ and $R_{ij} = 0$). This type of error damages performance by preventing agents from discovering and acting on interdependence between them. Errors of omission have received the most attention from management scholars: both research and practitioner writings have been concerned with how to prevent the damages caused by lack of communication between agents who should exchange knowledge, especially across organizational boundaries (Carlile 2004, Cross et al. 2002, Doz et al. 2001: 115-138). An ***error of commission*** happens when an agent interacts with another agent whom it does not depend on (i.e. $T_{ij} = 0$ and $R_{ij} = 1$). While this type of error has received less scholarly attention, it can also be a significant burden for organizations: as Hansen (2009) put it, collaboration is likely to undermine performance if the projected return from collaboration does not exceed the sum of opportunity costs and collaboration costs. Such costs include the time and attention put into maintaining interactions with actors whom the focal agent does not need to interact with in order to fulfil its tasks, and which could have been allocated to more valuable activities (Levine and Prietula 2012: 1749).

In any period *t*, each agent *i* obtains a performance score P_{it} between zero and one. P_{it} is obtained as a ratio of the number of agents with whom the focal agent coordinates properly (i.e. $R_{ij} = 0$ when $T_{ij} = 0$, or $R_{ij} = 1$ when $T_{ij} = 1$) over the total number of other agents. We assume that the cost of omission errors is equal to the benefit which would have been obtained if an interaction had been present between two interdependent agents. Similarly, the cost of commission errors is assumed equal to the benefit which would have been obtained if an interaction had been absent between two independent agents. In the basic version of this model, errors of omission and commission are assumed to have equal cost, and the payoff distribution of correct interactions is assumed to follow a power law: while the average payoff of correct interactions (i.e. where $R_{ij} = 1$ and $T_{ij} = 1$) is 1, we assume that some interdependences are more pronounced than others (which seems more realistic, but does not affect our results). The performance of the organization is calculated as the average performance of all agents.

2. Representation of the task environment

As our model assumes limited rationality both for the designer and for the employees of an organization, neither side has *ex ante* a perfect understanding of the task environment. The designer may have some notions of T that have resulted from her allocations of tasks to agents with varying degrees of accuracy, and the agents may each have some limited understanding of which of the m-1 other agents they would benefit from interacting with.

We introduce a parameter $\boldsymbol{\Omega}$ which expresses the ***accuracy of the designer's architectural knowledge***—i.e. the designer's accuracy in representing T. For any task structure T, the designer perceives accurately a proportion Ω of all interdependences between agents. This parameter reflects the bounded rationality of organization designers: although the designer of an organization decides upon the division and allocation of tasks to agents and therefore causes some agents to depend on each other in the task structure, this does not mean that the designer perceives these interdependences accurately (Adler 1995, Eppinger 2001, Gokpinar et al. 2010, Heath and Staudenmayer 2000). Consider the manager of an IT company responsible for developing some ERP (Enterprise Resource Planning) software, who may have identified the development of marketing and human resource management applications as two separable tasks and allocated them to two different agents. It does not necessarily follow that the manager perceives the need for coordination between these two agents—for instance, a manager lacking technical knowledge of the task at hand may not perceive that integration difficulties would arise if the two agents fail to coordinate and develop their applications in two different programming languages. This type of coordination neglect is often observed in organizations (Heath and Staudenmayer 2000).

The parameter Ω may be understood as the designer's intrinsic ability to form an understanding of interdependences between agents, so that high values of Ω reflect designers with a higher ability to perceive interdependences between agents. Another possible interpretation would be to consider Ω as reflecting the difficulty of forming an understanding of task interdependences in the organization's environment; high values of Ω would then be representative of settings where there is low ambiguity in determining whether two agents are interdependent. Our investigation focuses on the case where the designer's knowledge is no better than chance (i.e. $\Omega = 0.5$), with additional analysis exploring the effects of increasing accuracy. We also do not consider situations where the designer is just systematically wrong (so that $\Omega < 0.5$), as that leads rather obviously to mal-adaptation.

We assume that the formal structure is generated based on the designer's perception of the task structure, so as to create the structural conditions enabling collaboration between interdependent agents. Thus, the formal structure encourages social interactions between agents who are perceived by the designer to depend on each other and discourages interactions between those who aren't (Brass et al. 2004: 796). This "formal" structure encompasses all elements of the organizational structure that are mandated through formal authority, such as grouping decisions and linking mechanisms, reporting relationships, systems, incentives, allocation of design rights, locus of activities or even hiring decisions (Gulati and Puranam 2009, Nadler and Tushman 1997).

We provide a simple representation of the formal structure through a ***distance matrix D***. In our formulation, the distance between two agents reflects any formal structural element that affects their probability of interacting. We recognize that many structural elements may affect this probability: for instance, the designer may reduce the probability of interaction between two agents by placing them in different administrative units or physical locations (Doz et al. 2001, Hansen 1999, Hinds and Kiesler 2002), or by hiring agents whose personal traits are different enough to preclude homophily-based interaction (McPherson et al. 2001).

In this model, we abstract from the multiplicity of antecedents and variants to formal structure, and model structure through its consequences: formal structure is implemented in the model as a symmetric distance matrix D of dimensions m x m, with each cell D_{ij} containing a zero or one. This amounts to a simplification whereby, for any pair of agents, the formal structure either encourages or discourages interaction between the two agents[1]. The distance matrix is thus equivalent to the grouping structure in an organizational chart. To see this, consider that a clustering algorithm applied to the designer's distance matrix will produce clusters that then correspond to organizational units, where ties in clusters are within-unit interactions, and ties across clusters are committees or cross-functional teams. For any task structure T with m agents, the distance matrix D has $\Omega * m$ cells where D_{ij} is equal to T_{ij}—and thus $(1-\Omega) * m$ cells where D_{ij} equals $(1 - T_{ij})$.

The agents' representation of the task environment is captured via a ***social attraction*** matrix, ***S***. This is a time varying matrix of dimensions *m* x *m*, where each element S_{ijt} constitutes the probability for agent *i* to approach agent *j* with an interaction request (i.e. an expression of willingness to interact) in period *t*. Each element in *S* is a convex combination of the corresponding elements in:

a) The distance matrix *D*, with weight *w*
b) The realized social structure of the previous period, R_{t-1} (to be defined below under "Choice Process"), with weight *1-w*

Thus an agent's representation of partner desirability depends both on the formal structure (via *D*) and history (via R_{t-1}). Notice that since the model is computed over multiple periods, the realized network of each period reflects not only the influence of the last period's interactions but also the interactions of all previous periods. The ***enforcement of formal structure, w***, can be tuned to obtain representations where the formal structure is very influential (high *w*), as opposed to representations where the past history of social interaction is very influential (low *w*). Higher values of *w* have a natural interpretation as a "strongly enforced" or "well implemented" formal structure, as a higher weight is placed by agents on forming interactions that are encouraged by the formal structure when *w* is high.

[1] Continuous measures of distance could also easily be implemented in this model, to give a sense of priority; for instance, the designer may believe that two sets of agents must all speak with each other but more so within each set. Among the interactions to be formed, some interactions (e.g. hierarchical ones) could be stressed a bit more, others a little less. In the current formulation, the designer just has an opinion on "who needs to talk to whom", not necessarily how much. However, our results look qualitatively similar if we allow for some variation in the strength of the interactions being formed.

3. Choice process

In this model, we focus on adaptation by agents over multiple periods within a given formal structure. We thus assume a constant formal structure. Therefore, the designer's choice only arises in the initial period t_0, when the designer generates a formal structure[2] (i.e. the proximity matrix D) based on his understanding of the task structure *T*.

The agents' choice process, on the other hand, operates in every period *t*. This process involves each agent indicating willingness to interact ("interaction requests") to the *m*-1 other agents (or not), based on its *social attraction* vector, S_{it}, described in the previous section. In keeping with the principle of non-greedy action selection in models of adaptive rationality, the actual choice of whom to make interaction requests to is noisy: whenever there is conflict between the formal structure and past realized interactions regarding a certain partner (e.g. $D_{ij}= 1$ and $R_{ijt-1} = 0$), the corresponding cell in the partner desirability matrix (S_{ijt}) takes a value between zero and one which turns probabilistically into an ***interaction request*** $\boldsymbol{I_{ijt}}$: the new interaction request status I_{ijt} is equal to D_{ij} with probability *w*, or equal to R_{ijt-1} with probability (1-*w*). The probabilistic nature of interaction requests captures the noisy nature of attempts to form social connections: because of luck and propinquity, agents are influenced by formal structure and past realized interactions in a non-deterministic manner.

While the influences of formal structure and past interactions lead agents to form an initial representation of the interaction requests they wish to send, agents may alter this initial representation through exploration. Consistent with the principle of problemistic search, agents engage in exploration for new partners (Baum et al. 2005, Lavie and Rosenkopf 2006) only when their performance falls below an ***aspiration level***, $\boldsymbol{\alpha}$, for performance. While this process of trigger-based search is consistent with traditional treatments of performance-based adaptation (Cyert and March 1963, March and Simon 1958), it is worth noting that the meaning of "aspirations" can be interpreted more broadly when applied to individual agents—as in our model—than to organizations or organizational units. In our model, an

[2] One of our additional experiments incorporates a process whereby the designer chooses the distance matrix D out of many possible alternatives and adapts choices over time. However, much of the additional insight from such an exercise may be obtained by thought experiments within the current set-up (at least with reference to the key arguments in this paper).

agent's performance is measured as its contribution—through the satisfaction of work-related interaction needs—to the performance of the organization. Therefore, the aspiration level may be interpreted as a measure of goal alignment between the organization and its agents: each agent will seek to modify its interactions with other agents until these interactions match the task structure closely enough to yield a contribution to organizational performance which is higher or equal to the aspiration level, α. An aspiration level, α, taking a value between 0 and 1, is determined for all periods of the simulation.

As described above, *whether* an agent explores for new partners is determined through a comparison between performance and aspirations: the agent explores whenever $P_{it-1} < \alpha$. *How much* an agent explores is determined by the ***search rate β***: a realized set of interaction requests for each agent is a vector I_{it} of dimension m, whose elements are 1 or zero with probabilities dependent on S_{it}, but flipped with probability β if the agent performed below aspirations in the last period (i.e. I_{ijt} becomes $(1-I_{ijt})$ with probability β) [3].

As the concluding step of this process, a new matrix R_t arises from the double correspondence of interaction requests: whenever two agents both decide to interact with each other (i.e. $I_{ijt} = 1$ and $I_{jit} = 1$), then an interaction happens between them (i.e. $R_{ijt} = R_{jit} = 1$).

4. Transformation of representation

Finally, the agent's social attraction vector is constantly transformed over time as long as $0<w<1$, because of the dual weight on past realized interaction patterns R_{t-1} and on the proximity matrix D. Since the current version of the model assumes a constant formal structure for each project started by an organization, the designer's representation of the task structure is assumed constant for each project. The change in representation for the agents thus comes via past patterns of realized interaction structures (weighted by $1-w$) and their own exploration (via β).

[3] We also implemented versions of the model where the search rate was determined by the *extent* of the discrepancy between performance and aspirations. That is, search is distance-based: the further away agents are from their aspiration levels, the more they explore. Specifically, the extent of exploration is determined jointly by the ratio of performance to aspirations (P_{it-1} / α) and a search threshold β' which takes values between 0 and 1. If an agent's performance falls below the aspiration level, each cell of his interaction request vector is flipped with probability $\beta = [(\alpha - P_{it-1})/ \alpha] * \beta'$.

Note on computational intensity

It is worth noting the computational intensity involved in simulating the dynamics of our model. The model essentially represents an organization as multiple layers of networks: the social network structure of the organization R_t is influenced by its past occurence R_{t-1} as well as a formal structure represented as a distance network D, itself influenced by a task structure represented as a network T. For each period of a run, the simulation computes a new matrix R_t, which itself requires the computation of two other matrices S_t, I_t, and a new performance vector P_t. Each cell of these elements requires computations based on cells located in other matrices (e.g. each cell S_{ijt} is a probability computed based on corresponding cells D_{ij} and R_{ijt-1}).

These computations are repeated in each period (i.e. 300 periods) of each run (i.e. 2000 runs), for each combination of parameters. The computational intensity grows linearly with the number of periods (t) and the number of runs, and grows exponentially with the number of agents (m) and with the granularity of the parameter space sampled. On a high-performance machine, many of the analyses featured in this paper took several hours to be computed, and some several days.

ONLINE APPENDIX 2:

GRAPHICAL REPRESENTATIONS OF RESULTS FOR THE ADDITIONAL EXPERIMENTS

Note: As in the main paper, graphs feature a dotted line (baseline with exploration rate beta=0) and a full line (case of interest with beta>0). Results are based on 1,000 runs of the model (versus 2,000 for the main results reported in the paper).

Figure A1: Organizational performance for different levels of enforcement of a random formal structure (Ω = 0.5), with a dense task structure (density = 0.8)

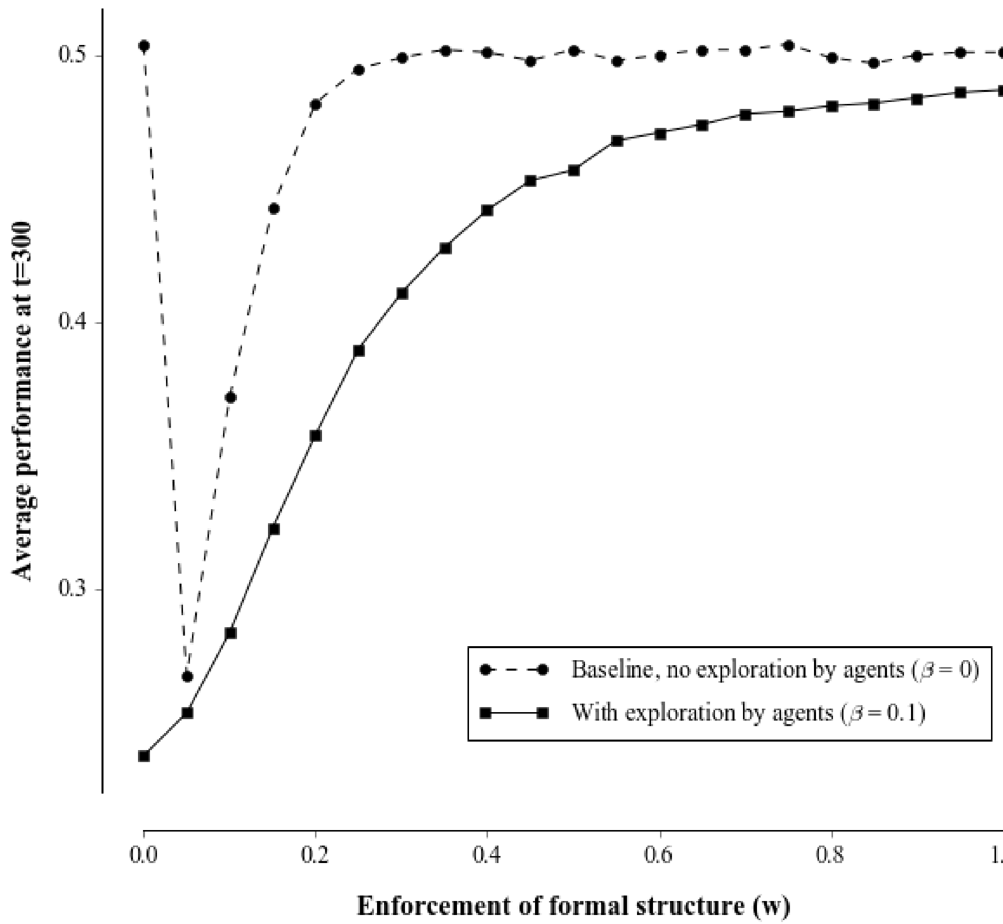

Note: the high level of performance at w=0 in the case without exploration ($\beta = 0$) is exclusively due to our assumptions about initial conditions: we assume that the initial social structure (R_0) has a density of 0.5, which is higher than the equilibrium level of density when agents are allowed to explore. At w=0 and without exploration, agents are fully inert so that density remains at 0.5 until the end of each run. Since this model features a high task density, a social structure with high density generates high performance. If we assumed an initial density of zero, then the performance at w=0 in the case without exploration would be the lowest level in the graph. Hence, as in our other models, it is much more meaningful to interpret results in the case where agents are allowed to explore (i.e. the full line with $\beta = 0.1$).

Figure A2: Organizational performance for different levels of enforcement of a random formal structure ($\Omega = 0.5$), with omission errors twice as costly as commission errors

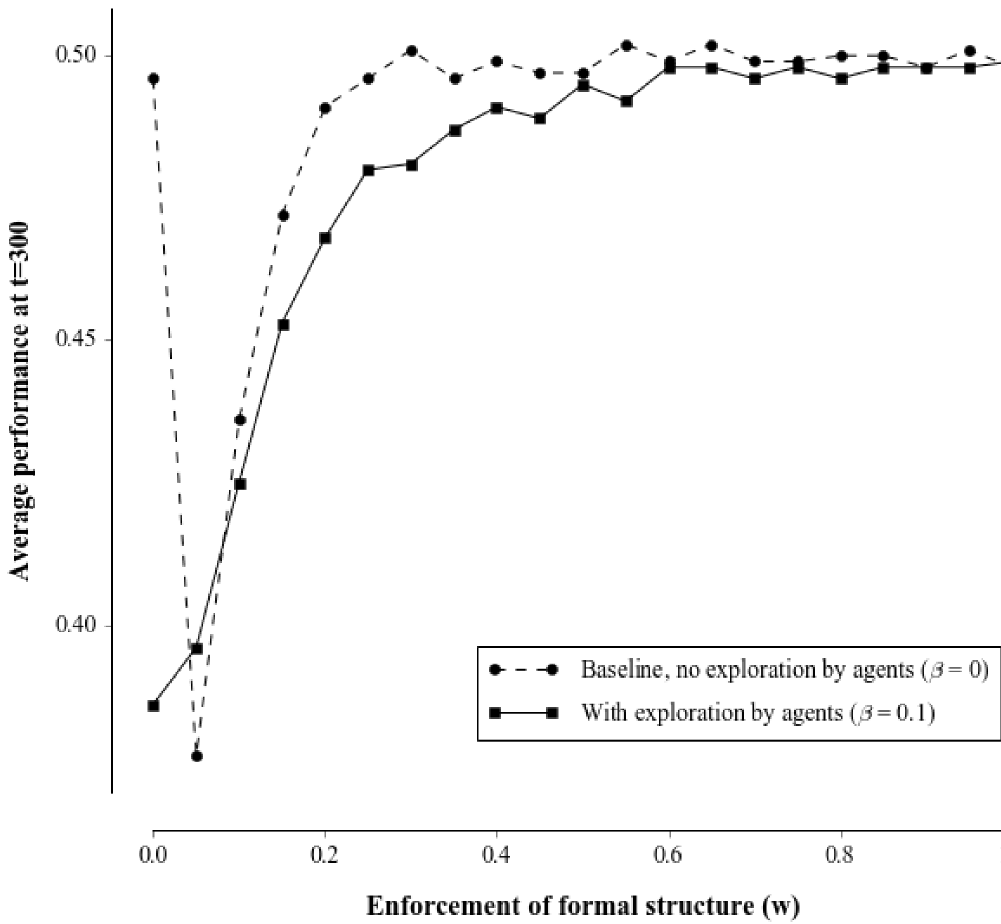

Note: the high level of performance at w=0 in the case without exploration ($\beta = 0$) can be explained in the same way as in Figure A1 (see note above). Hence, as in our other models, it is much more meaningful to interpret results in the case where agents are allowed to explore (i.e. the full line with $\beta = 0.1$).

Figure A3: Organizational performance for different levels of enforcement of a random formal structure ($\Omega = 0.5$), with formal structure as a constraint on agents' exploration ($w_2 = 1$)

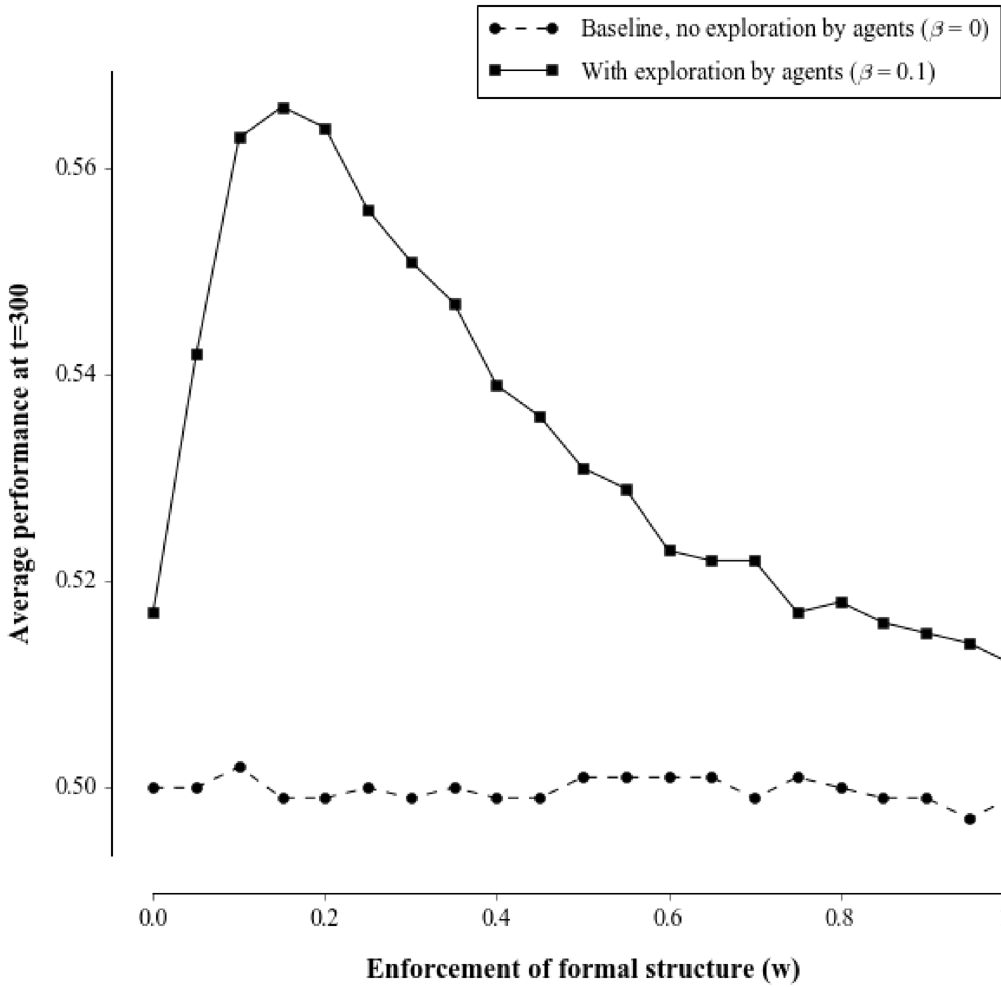

Note: this figure is to be compared with Figure 2 of the main paper (which shows the main results regarding the value of enforcing a random formal structure). The peak in performance at w=0.15 is higher when agents only explore locally within interactions encouraged by the formal structure ($w_2 = 1$).

Figure A4: Organizational performance for different levels of enforcement of a random formal structure (Ω = 0.5), with asymmetric interactions

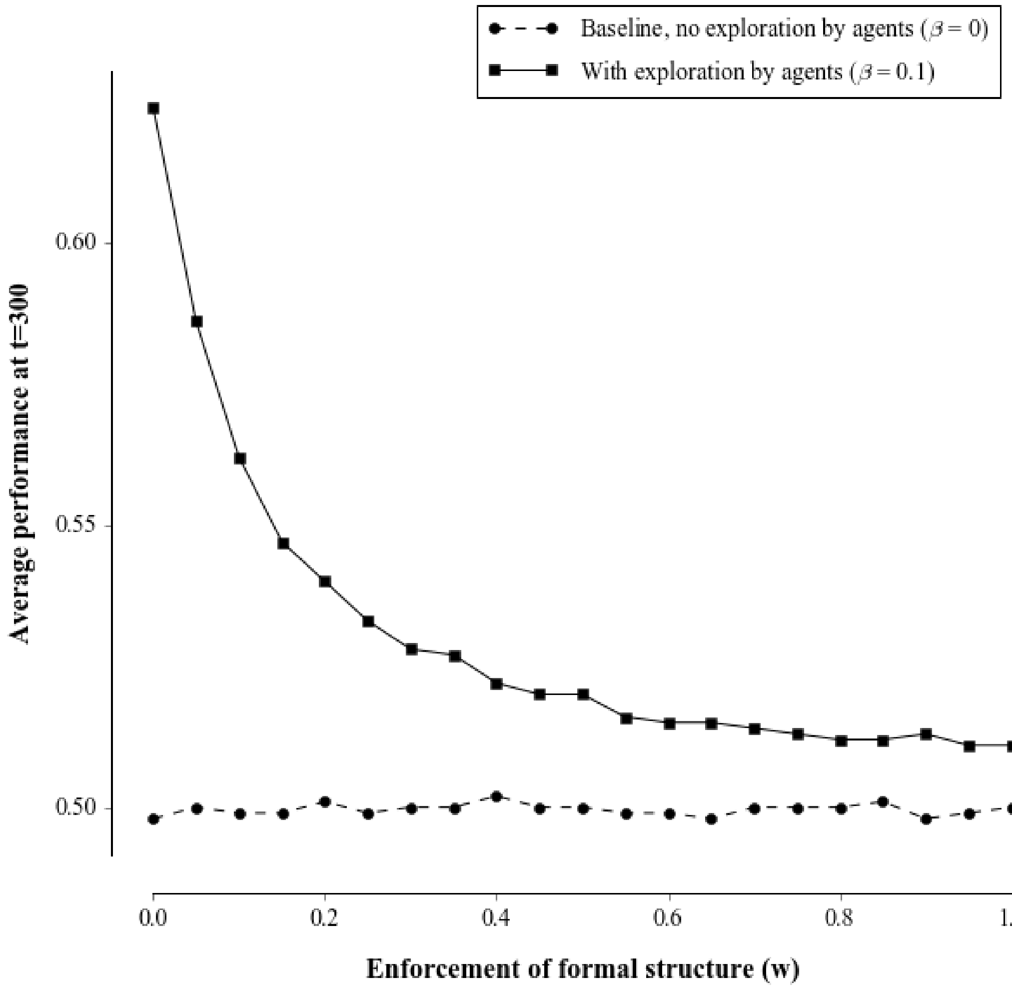

Note: this figure is to be compared with Figure 2 of the main paper (which shows the main results regarding the value of enforcing a random formal structure). When interactions are asymmetric, the enforcement of formal structure loses its value: this confirms that formal structure generates value by coordinating the decentralized efforts of agents in generating bi-lateral interactions.

Figure A5: Organizational performance for different levels of enforcement of a random formal structure (Ω = 0.5), with a culture of reciprocation (agents have a 50% probability of reciprocating interactions regardless of their perceived value)

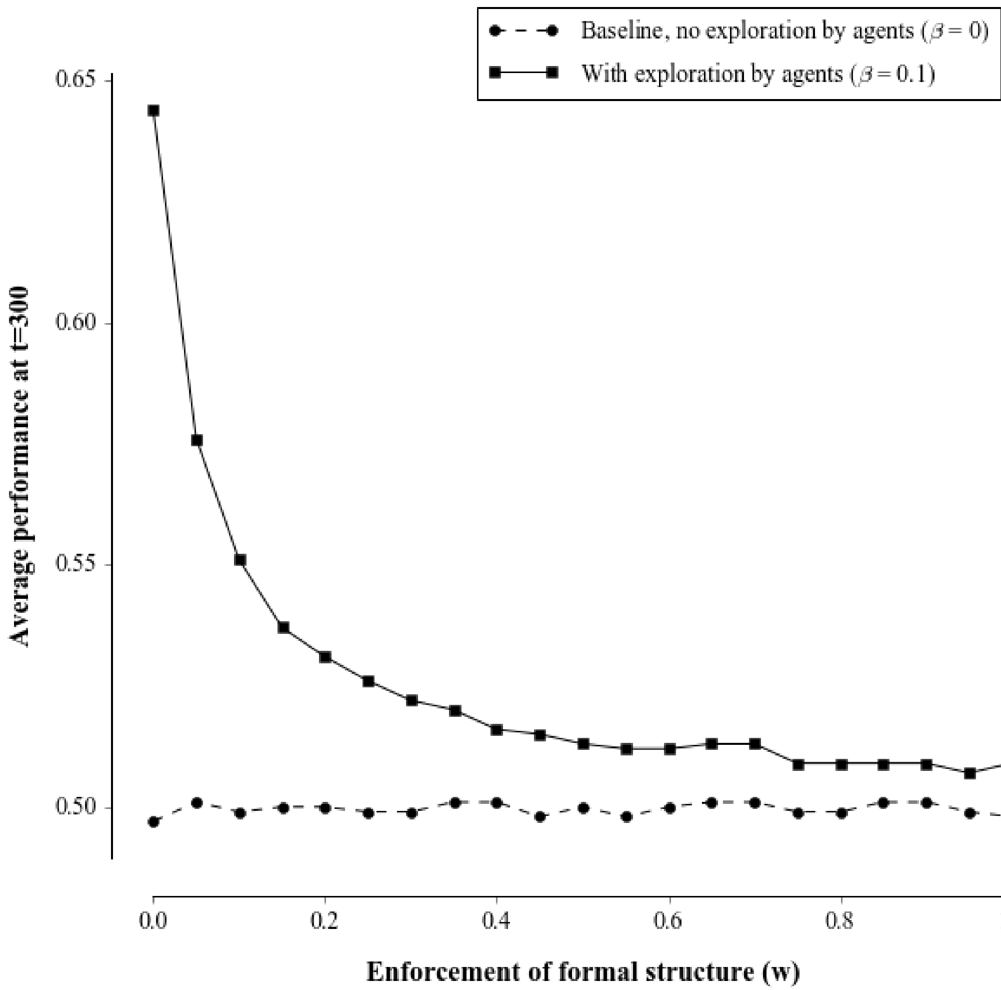

Note: this figure is to be compared with Figure 2 of the main paper (which shows the main results regarding the value of enforcing a random formal structure). When agents share a culture of reciprocity, the enforcement of formal structure loses its value (much like the previous figure, this provides confirming evidence of the mechanisms at play). This boundary condition to our result assumes moderate amounts of reciprocity. In the unrealistic situation of an organization where the culture of reciprocity is unlimited (i.e. 100% of interaction requests are reciprocated), the ability of agents to correct commission errors is much smaller, so that enforcing even random formal structures becomes valuable again. This is because it enables agents to experiment with deleting interactions (which they would otherwise always keep) since some of these interactions are now discouraged by the formal structure, and appropriately so in at least 50% of the cases.

Figure A6: Organizational performance for different levels of enforcement of a random formal structure (Ω = 0.5), with a capacity constraint on interactions (5 interactions per agent)

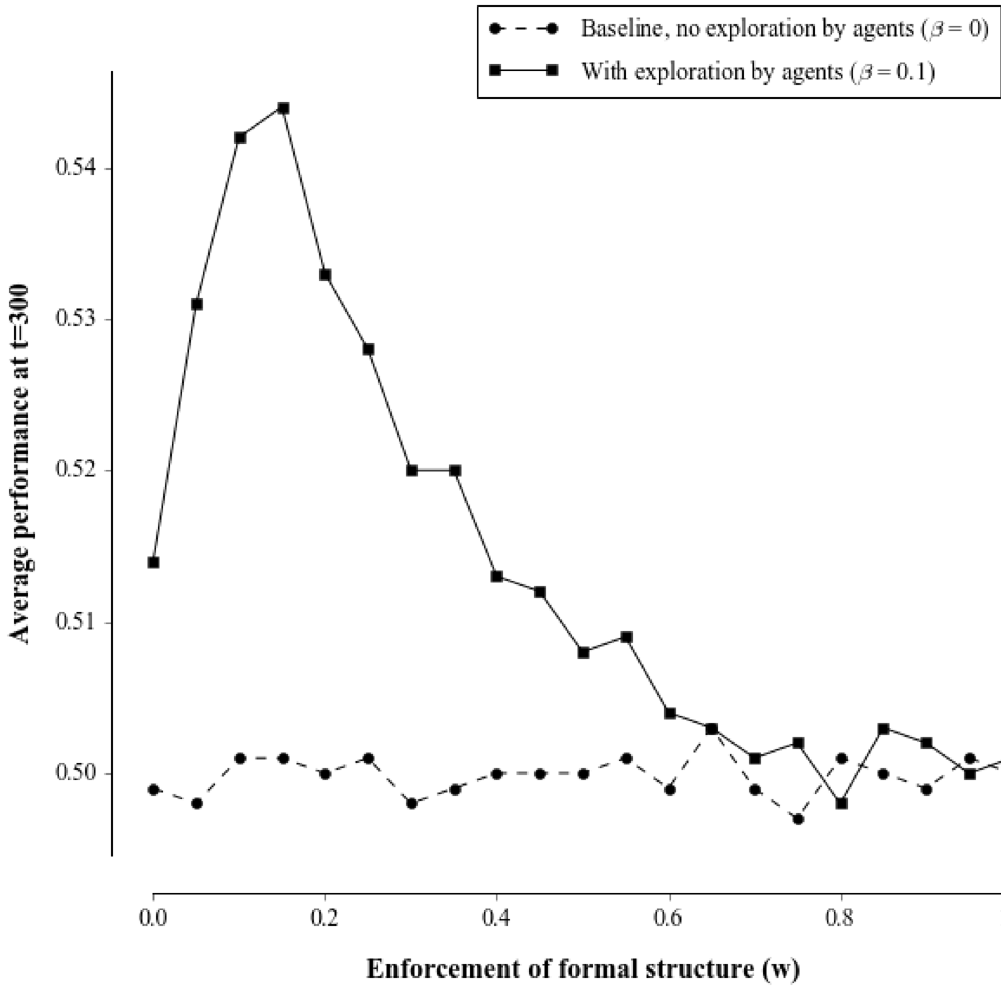

Note: this figure is to be compared with Figure 2 of the main paper (which shows the main results regarding the value of enforcing a random formal structure). Despite the constraint on the number of possible interactions per agent, the pattern of results remains similar to that of the main model.

Figure A7: Organizational performance for different levels of enforcement of a random formal structure (Ω = 0.5), with agents 50% less likely to explore by deleting correct ties

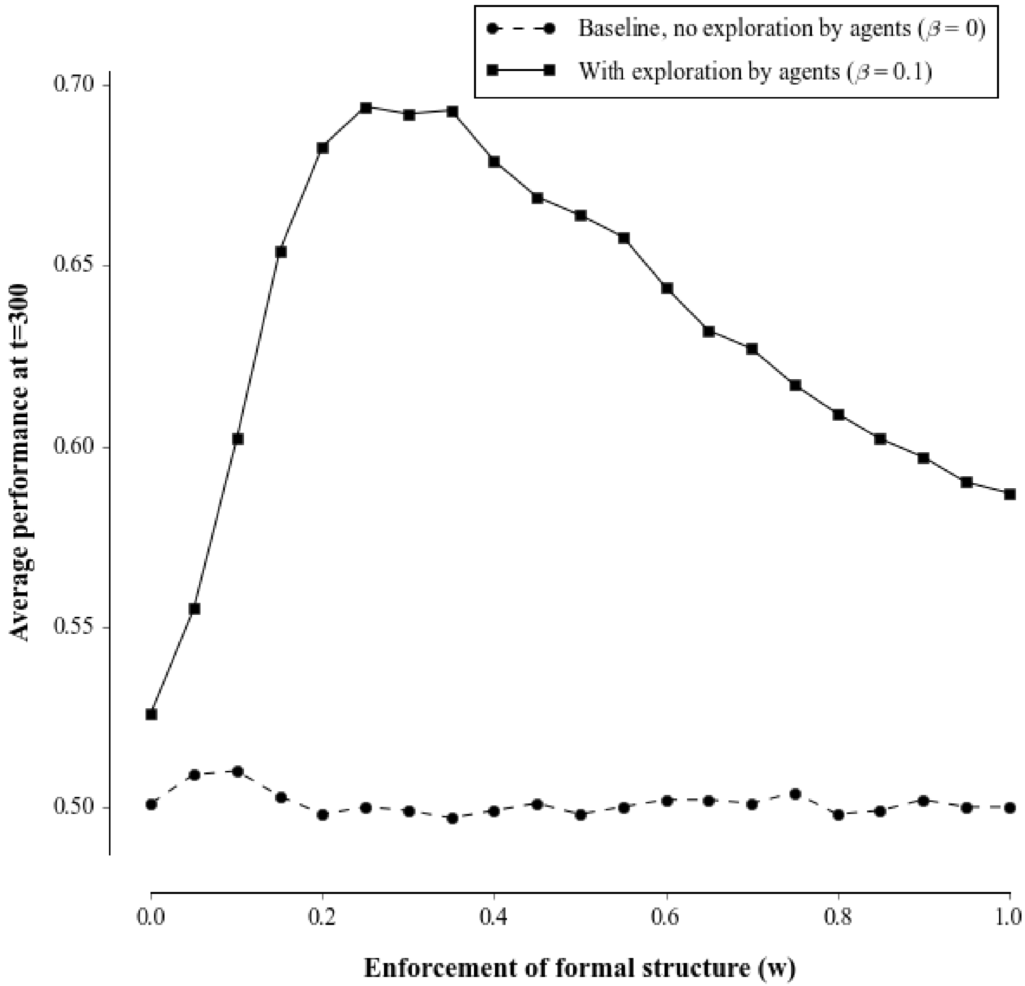

Note: this figure is to be compared with Figure 2 of the main paper (which shows the main results regarding the value of enforcing a random formal structure). The peak in performance is higher, and occurs at a higher enforcement of formal structure (w=0.3, versus w=0.15 in Figure 2).