The Dual Challenge of Search and Coordination for Organizational Adaptation:

How Structures of Influence Matter

Appendix

Contents:

§1. Effects of interdependence (χ) on payoffs

§2. Dominant influence structures at different levels of individual exploration (τ)

§3. Effects of symmetric and asymmetric influence in dyads

§4. Effects of acyclicity

§5. Varying *w* within flat and hierarchical teams

§6. Interdependence structures that are isomorphic to influence structure

§7. Generalization to other organizational structures and larger organizations

- **§7.1** Random graphs (communities) and branching hierarchies compared to crowds
- **§7.2** Larger organizations and search spaces

§8. References cited in Appendix

§1 The Effects of Interdependence on Payoffs

We consider a simplified task environment with S=2, with v_1 the global peak (with the value of 1) and v_2=0.75. Tables §1.1 and §1.2 show how the payoffs (π) an agent i choosing k receives depends on F_t (the fraction of other agents who pick the same choice k at time t), v_k (baseline value of choice k), and χ. When all other agents are on the global peak, an agent choosing the global peak gets the maximum payoff of 1, regardless of the value of χ. As the number of agents on the peak declines, the actual payoff to the agent of choosing the peak falls. At high levels of χ, this fall is steep. For instance, when χ is 0.8 (Table §1.2), the peak pays off only .6 when half the other agents are on the peak. Whereas when χ is 0.2 (Table §1.1), the payoff from picking the peak with half the other agents on it is .9.

Table §1.1: Payoffs when S=2, v_1=1 and v_2=0.75, χ=0.2

% on Peak	**0**	**.1**	**.2**	**.3**	**.4**	**.5**	**.6**	**.7**	**.8**	**.9**	**1.0**
Non-Peak Payoff	.75	0.735	.72	.705	.69	.675	.66	.645	.63	.615	0.6
Peak Payoff	.8	.82	.84	.86	.88	.9	.92	.94	.96	.98	1.0

Table §1.2: Payoffs when S=2, v_1=1 and v_2=0.75, χ=0.8

% on Peak	**0**	**.1**	**.2**	**.3**	**.4**	**.5**	**.6**	**.7**	**.8**	**.9**	**1.0**
Non-Peak Payoff	.75	.69	.63	.57	.51	.45	.39	.33	.27	.21	.15
Peak Payoff	.2	.28	.36	.44	.52	.6	.68	.76	.84	.92	1.0

Payoffs from picking the non-peak option are likewise tuned by χ. When the non-peak value is high (v_2 =0.75 in this example) and χ is high, this inferior state can be very attractive when even a moderate proportion of agents have picked it. For instance, in the example shown in Table 1.2, when the peak has 4 out of 10 agents on it, it is just marginally more attractive than the non-peak alternative. Thus, if a substantial number of agents are taking a common action, but not the action with the highest direct payoff (the v value), there is a tension between the returns to converging on a v value that is not the highest and the individual (and collective) search for a superior v value. This tension is heightened with χ. Put differently, both the penalty for non-convergence as well as the risk of entrapment to a non-peak are magnified when χ increases.

§2 Dominant Influence Structure at Different levels of τ

The graphs below show χ on the y-axis and show combinations of δ_0 and δ_1 on the x-axis. They combine the task environments shown in different graphs in Figure 2 in the text into a single graph, and also portray all levels of χ. We vary the individual exploration parameter (τ) across the rows of graphs below. In baseline analyses reported in the main body, τ=0.05. Here, we see that lower levels of τ give superiority to flat teams, which, through similar but diffuse beliefs, lead to exploratory behavior even with low τ. At higher levels of τ, flat teams lose their advantage because agents already have sufficient tendencies to explore.

Figure **§2.1**: Relative cumulative performance at T=50 [Key: Flat team (green), Hierarchical team (red) and Crowd (blue). Depth of shade shows magnitude of advantage over next best structure.

τ =0.01	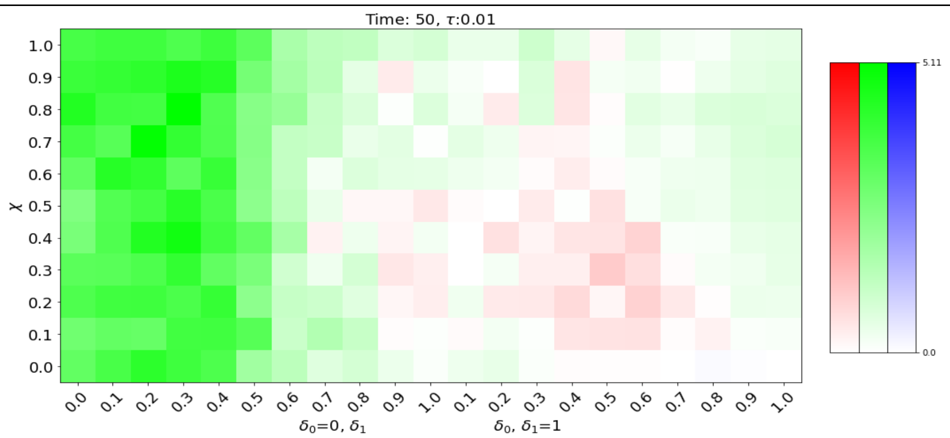
τ =0.05	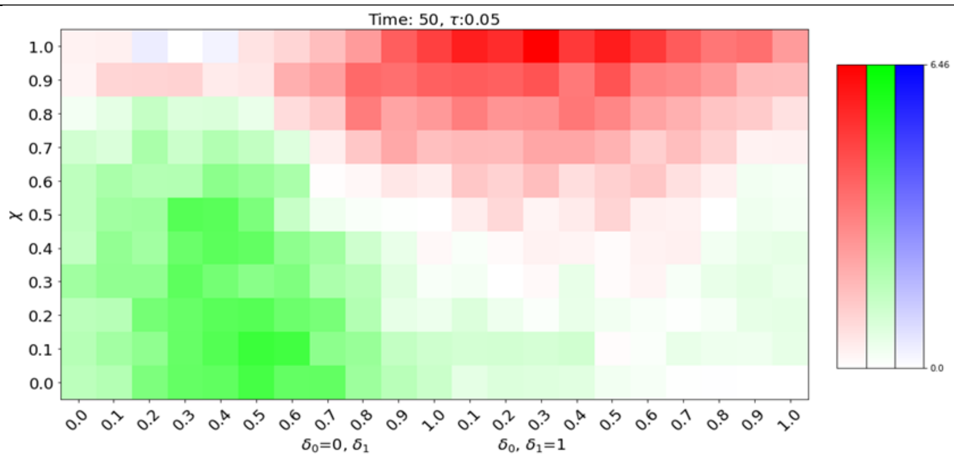

τ =0.1	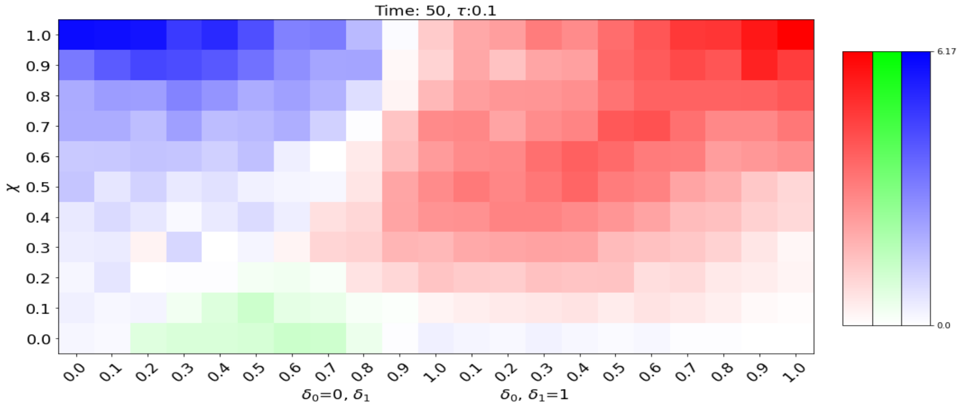Time: 50, τ:0.1 χ 1.0 0.9 0.8 0.7 0.6 0.5 0.4 0.3 0.2 0.1 0.0 0.0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1.0 0.1 0.2 0.3 0.4 0.5 0.6 0.7 0.8 0.9 1.0 δ0=0, δ1 δ0, δ1=1 6.17 0.0

Figure **§2.2**: Relative cumulative performance at T=500 [Key: Flat team (green), Hierarchical team (red) and Crowd (blue). Depth of shade shows magnitude of advantage over next best structure.

τ =0.01	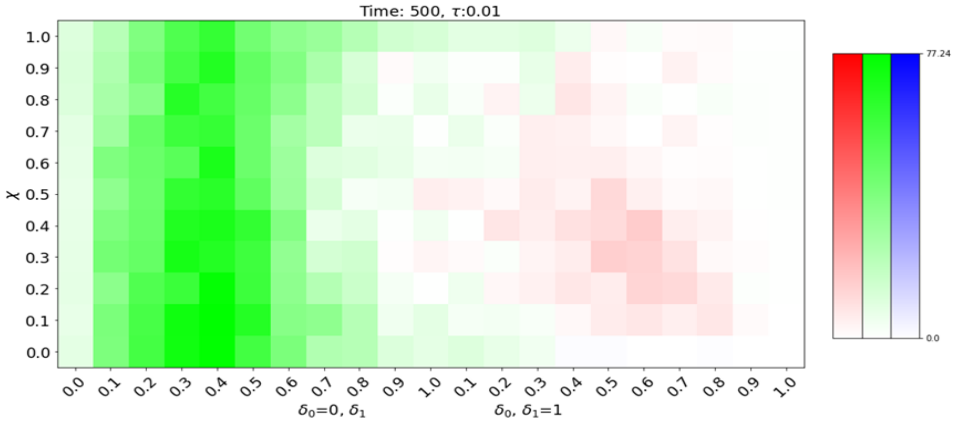
τ =0.05	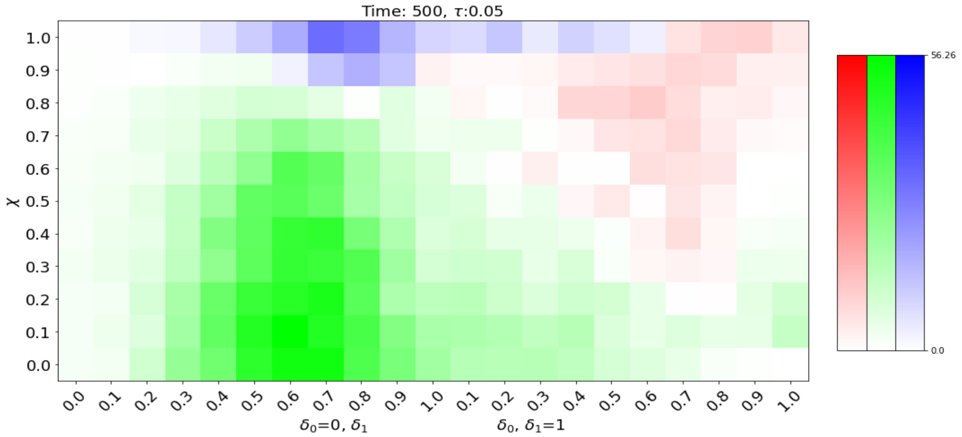
τ =0.1	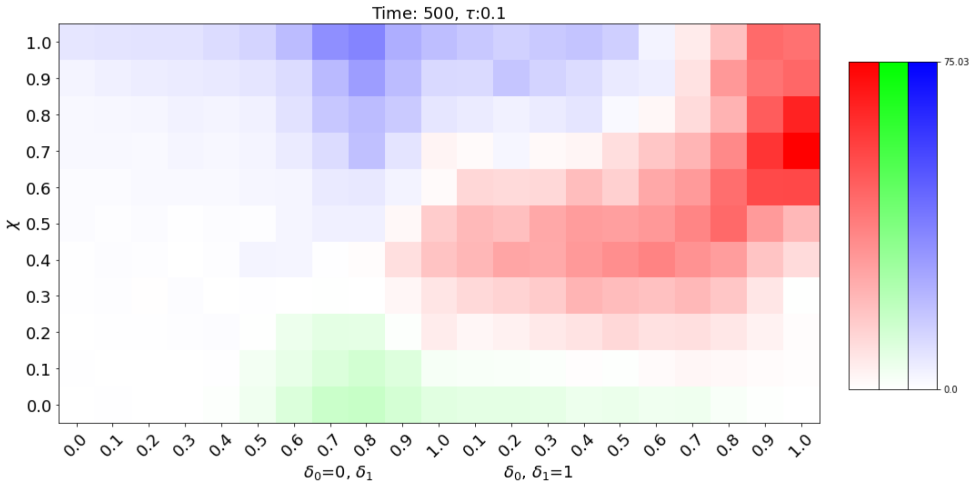

§3 Effects of symmetric and asymmetric influence in dyads

Figure **§3.1**. Dyadic elements of Organizational Forms (compare with Figure 1 in paper)

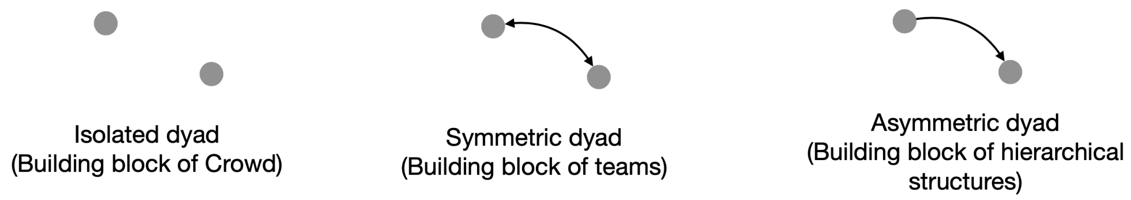

We present results from an analysis of a simple “off-line" version (i.e. with no learning from task environment- so that we can focus on influence alone) of our model over two periods. We modify equation 3 to allow for varying weights on own and other agent’s beliefs such that:

$$\mathbf{b}_{1,t} = w1 * \mathbf{b}_{1*,t} + (1-w1) * \mathbf{b}_{2,t}$$

$$\mathbf{b}_{2,t} = w2 * \mathbf{b}_{2*,t} + (1-w2) * \mathbf{b}_{1,t} \quad (4)$$

where w_1, w_2 in [0,1] and denote the weight on one’s own beliefs for each agent. At w=1, alter’s beliefs have no influence on ego’s beliefs, and at w=0, alter’s beliefs fully determine ego’s beliefs. For 0<w<1, influence produces a convex combination of one’s own beliefs with beliefs of the other as shown in equation 4 (DeGroot 1974; Friedkin and Johnsen 1990).

Let $\boldsymbol{x}$ and $\boldsymbol{y}$ be s dimensional vectors of beliefs at t=0 about the attractiveness of each of the s alternatives for agents 1 and 2 respectively. For simplicity of exposition, we assume that each vector contains a different permutation of the same realizations of the s values in {0,1}; that is, each agent shares the same set of attraction values but the assignment of a particular value to a given alternative may differ.[1] At t=0, the ***dissimilarity*** in beliefs between the two agents can be written simply as $|(\boldsymbol{y}-\boldsymbol{x})|$, with similarity simply being $1-|(\boldsymbol{y}-\boldsymbol{x})|$. For ease of exposition in the following discussion, we will study the degree of dissimilarity.

Post-influence at period $t+1$, the new beliefs for each agent are:

Agent 1: $w_1 \cdot \boldsymbol{x} + (1-w_1) \cdot \boldsymbol{y}$

Agent 2: $w_2 \cdot \boldsymbol{y} + (1-w_2) \cdot \boldsymbol{x}$

And the value of post-influence dissimilarity is:

$$| w_1 \cdot \boldsymbol{x} + (1-w_1) \cdot \boldsymbol{y} - w_2 \cdot \boldsymbol{y} - (1-w_2) \cdot \boldsymbol{x} |$$

$$= | (1 - w_1 - w_2) \cdot (\boldsymbol{y}-\boldsymbol{x}) | \quad (5)$$

Given w in [0,1], $| (1 - w_1 - w_2)|$ must always be <=1: *influence can never decrease similarity* (and if 0<w<1, it must necessarily increase it). This property produces the patterns observed in Figure

[1] This is a deterministic version of what will be true in expectation. For example, consider a situation where s= 4, the vectors x and y may take on the values such as x={0,0.1,0.3,0.8}, y={0.8,0.3,0.1,0}.

3 in the main text, where flat and hierarchical teams both attain greater similarity more rapidly than crowd.

Next, we parameterize the difference between the influence weights of the two actors by e such that $w_1 = u + e$ and $w_2 = u - e$ subject to the constraint that $e \leq u \leq 1-e$ (since influence weights must lie in [0,1]). This allows us to keep the average influence strength under symmetry ($u + u = 2u$)/2=u identical to that under asymmetry ($w_1 + w_2 = u + e + u - e = 2u$)/2=u.

Substituting these values into equation (4) yields the new expression for post-influence dissimilarity:

$$| (1 - u - e - u + e) \cdot (\boldsymbol{y} - \boldsymbol{x}) |$$

$$= | (1 - 2u) \cdot (\boldsymbol{y} - \boldsymbol{x}) | \quad (6)$$

Equation (6) shows that asymmetry of influence (e) does not affect post-influence dissimilarity of beliefs. Thus, the *asymmetry of influence has no direct impact on similarity once we keep average influence strength constant*.

Define "resolution" as the degree to which perceived value of one state is different than that of another state. Further, we define the operator $M\{\boldsymbol{b}\}$ as the index of the largest element in the belief vector $\boldsymbol{b}$ and, correspondingly, $m\{\boldsymbol{b}\}$ as the index of the smallest element. Thus, $M\{\}$ and $m\{\}$ refer to particular actions. Similarly, let $V\{\boldsymbol{b}\}$ represent the value of the largest element in the belief vector $\boldsymbol{b}$ and $v\{\boldsymbol{b}\}$ as the smallest value of this vector. The *resolution* of beliefs can then be defined as $V\{\}-v\{\}$. Given the random assignment of beliefs in the initial setting, the initial value of resolution is identical for agents. However, specific beliefs will diverge, implying that $M\{\boldsymbol{x}\}\neq M\{\boldsymbol{y}\}$ and $m\{\boldsymbol{x}\}\neq m\{\boldsymbol{y}\}$. Given this diversity in beliefs, it will be the case that $V\{\boldsymbol{x}\cdot w + \boldsymbol{y}\cdot(1-w)\}$ declines and $v\{\boldsymbol{x}\cdot w + \boldsymbol{y}\cdot(1-w)\}$ increases as w approaches ½. To see this, consider that if w=1 or 0, there is no change in average resolution across agents. It is only as comparable weight is placed on the other actor's beliefs ($w \rightarrow 1/2$), the beliefs will tend to converge to some intermediate value. Therefore, resolution, defined as $V\{\}-v\{\}$, must decline as $w \rightarrow$ ½. Simply put, *agents with different beliefs with the same initial resolution will have lower resolution beliefs after they influence each other symmetrically.* This produces the patterns observed in Figure 3, where flat teams systematically have lower belief resolution compared to the hierarchical team or crowd.

Asymmetry of influence also affects the relative resolution of beliefs across the agents. Recall that $V\{\boldsymbol{x}\cdot w + \boldsymbol{y}\cdot(1-w)\}$ declines as w decreases, and $v\{\boldsymbol{x}\cdot w + \boldsymbol{y}\cdot(1-w)\}$ increases as w decreases. If $w_1 > w_2$, then this implies that the beliefs of agent 1 will have higher resolution than those of agent 2. All else being equal, agents with higher resolution beliefs are more stable in their choice of actions, and more fuzzy beliefs lead to more switching in actions. This switching in behavior is a form of exploration and a lower resolution of belief is akin to a greater value of τ in the softmax action selection (equation 2) (cf., Posen and Levinthal 2012). This also helps to understand why influence structure and individual exploration parameter (τ) act as functional equivalents. However, because it operates asymmetrically within the dyad, in contrast to the parameter τ (the proclivity to explore) which is the same across agents, the differential degree of resolution makes it easier to find better combinations of actions when there is less uncoordinated simultaneous adjustment (Lounamaa and March 1987; Puranam and Swamy 2016).

To illustrate these results, Figure **§3.2** shows results on the three types of dyads (N=2, S=7) for just the first 20 periods**)** [2]

Figure §3.2. Dyadic models over time (χ=1.0, δ_0=0, δ_1=1, τ=0.1)

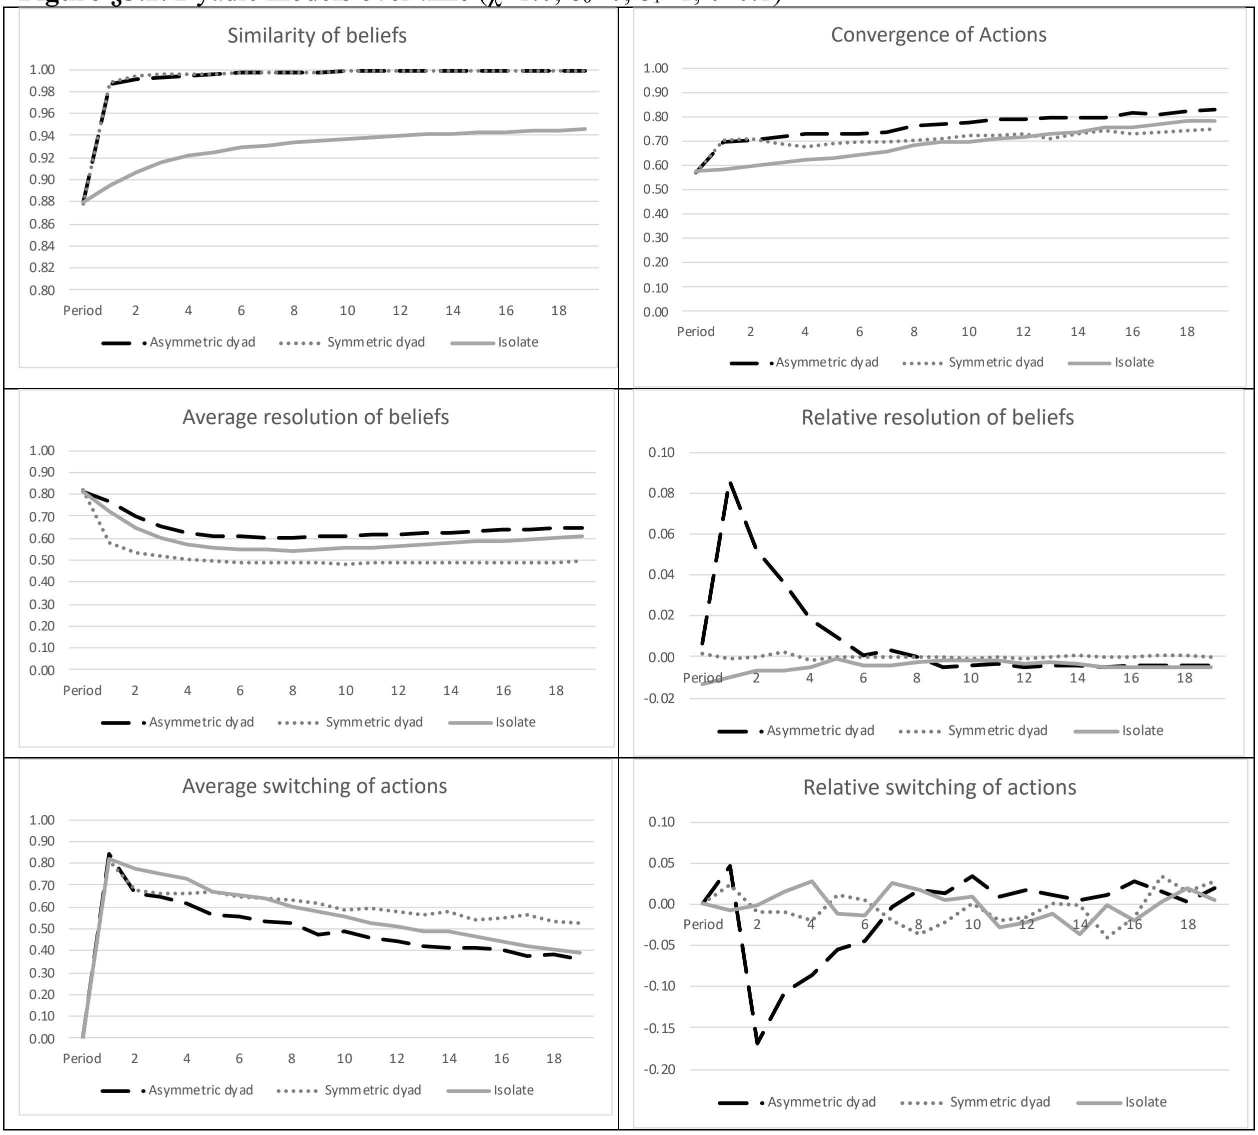

Isolated dyads (smallest units of crowds) have agents with $w_1=w_2=1$; symmetric dyads have agents with $w_1=w_2=0.5$; and asymmetric dyads have $w_1=1$, $w_2=0$. We see that both symmetric and asymmetric influence create higher similarity of beliefs between the two agents. However, asymmetric influence also creates asymmetry in resolution (difference in resolution between agents),

[2] Note given this is a dyadic relationship, the behavior converges to a steady-state much more quickly than the seven actor system we examine computationally.

as well as an asymmetry in the pattern of switching among actions. This analysis shows that the asymmetric influence inherent to hierarchical influence structures plays a crucial role in enhancing similarity while allowing differential switching, which aids coordinated search. This produces the rapid convergence to a good (though not necessarily best) organizational action that characterizes hierarchical structures.

§4. Effects of Acyclicity

Figure §4.1. Dynamics shown in Figure 3, compared to team with six randomly placed asymmetric ties (χ=1.0, δ_0=0, δ_1=1, τ=0.05)

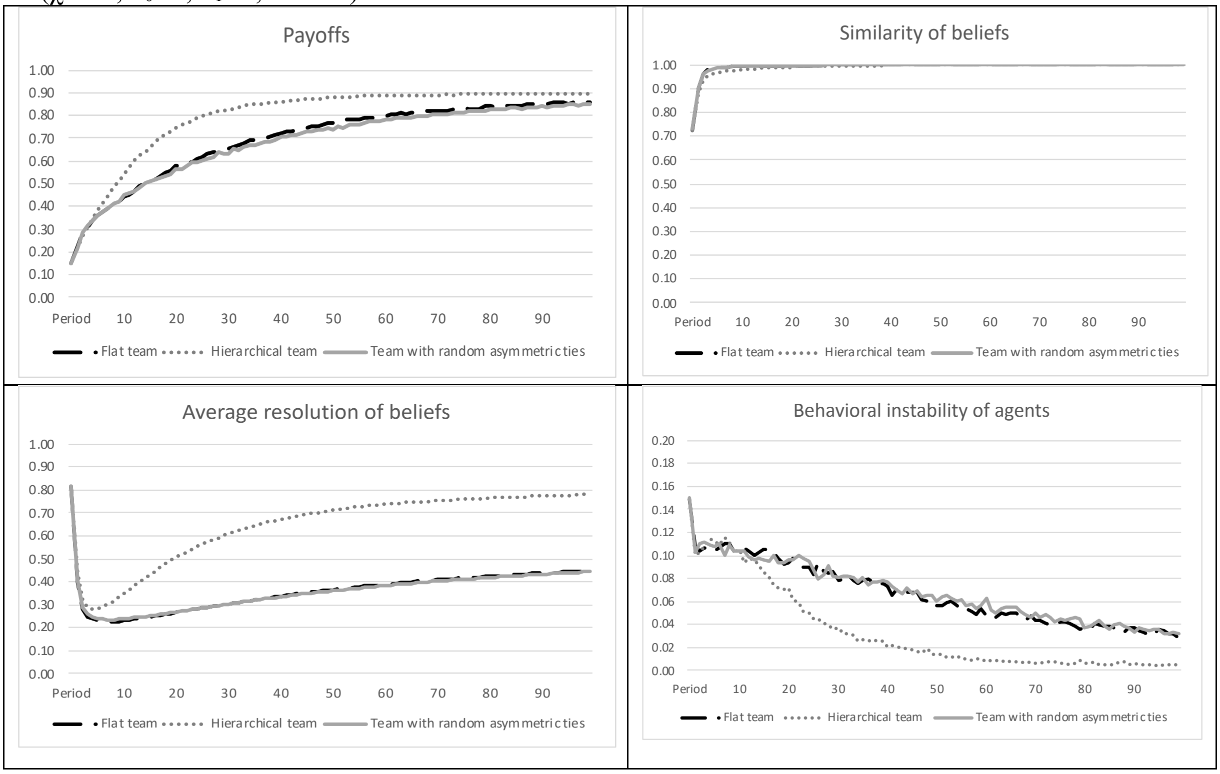

Figure §4.2. Making apex actor's ties symmetric in hierarchical team: Cumulative payoffs (N=S=7, δ_0 =0)

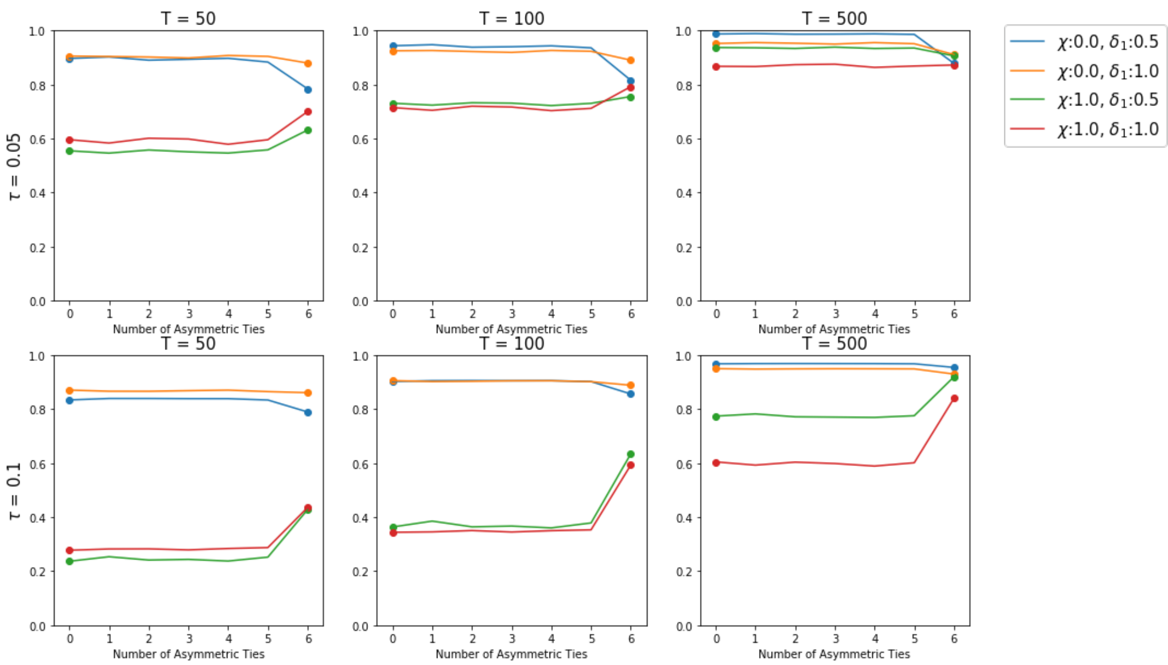

§5: Structural variations within teams

In Figure §5.1 we vary the influence weight w for all agents in a fully connected graph. When w=1, the structure is equivalent to the crowd, and when w=0, we get a form of pure diffusion (there is no learning from the task environment by the agents). In environments where only search matters, the crowd can be surpassed by the full density flat team; but when coordination also matters, given enough time the crowd dominates.

Figure §5.1 Cumulative performance across different levels of w for all agents in a fully connected graph (N=S=7, δ_0 =0)

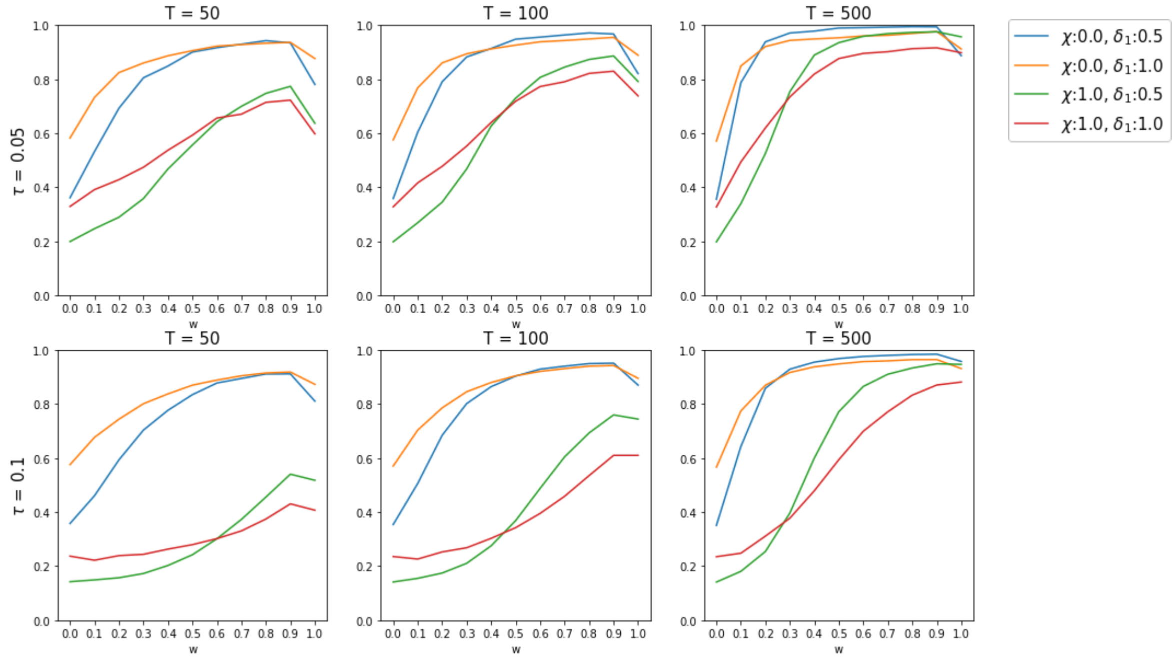

Figure §5.2 Cumulative performance across different levels of w for the apex agent (N=S=7, δ_0 =0)

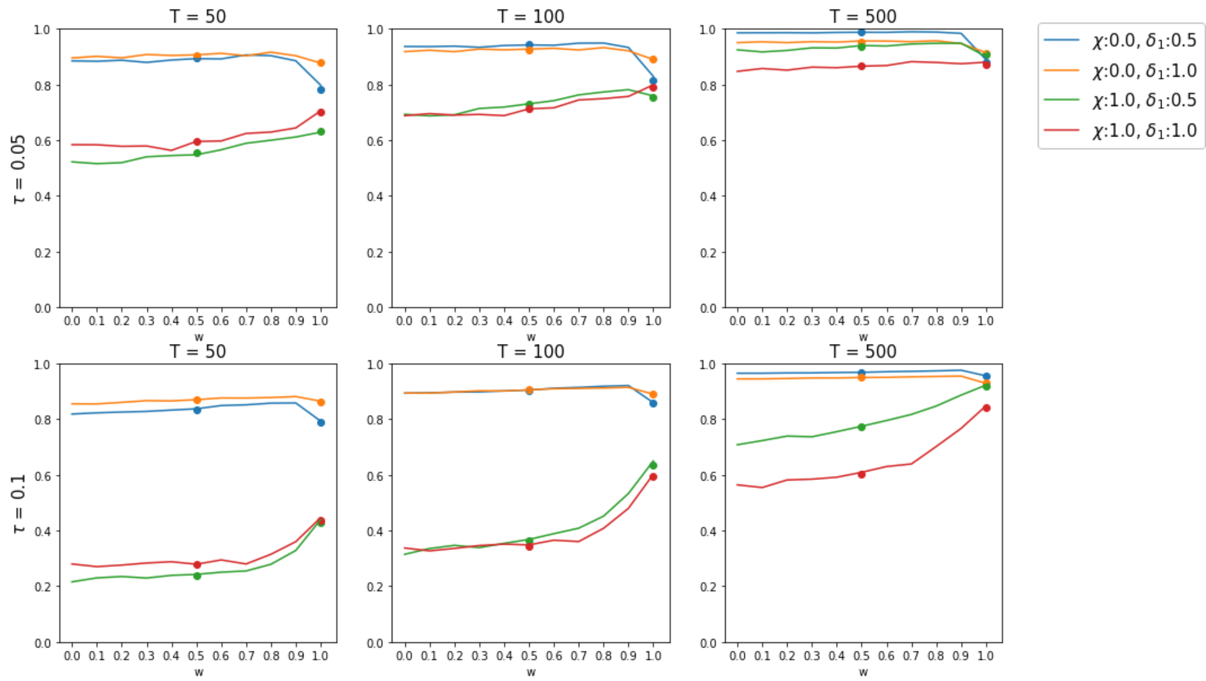

In Figure §5.2, we smoothly transition between flat team and hierarchical team via a single parameter – the influence weight on the apex actor. When this is ½ (as is the weight on all other actors), we get the baseline flat team structure. When w=1, we get the baseline hierarchical team. At w=0, apex agent's beliefs are influenced entirely by subordinates and not its own learning from the environment. The figure shows that in environments where only search matters, the leader is a liability, but when coordination also matters, the leader is valuable, even without any *ex ante* knowledge advantage.

Figure §5.3 Cumulative performance as w varies for all actors that are influenced (N=S=7, δ_0 =0, δ_1 =1 χ =1).

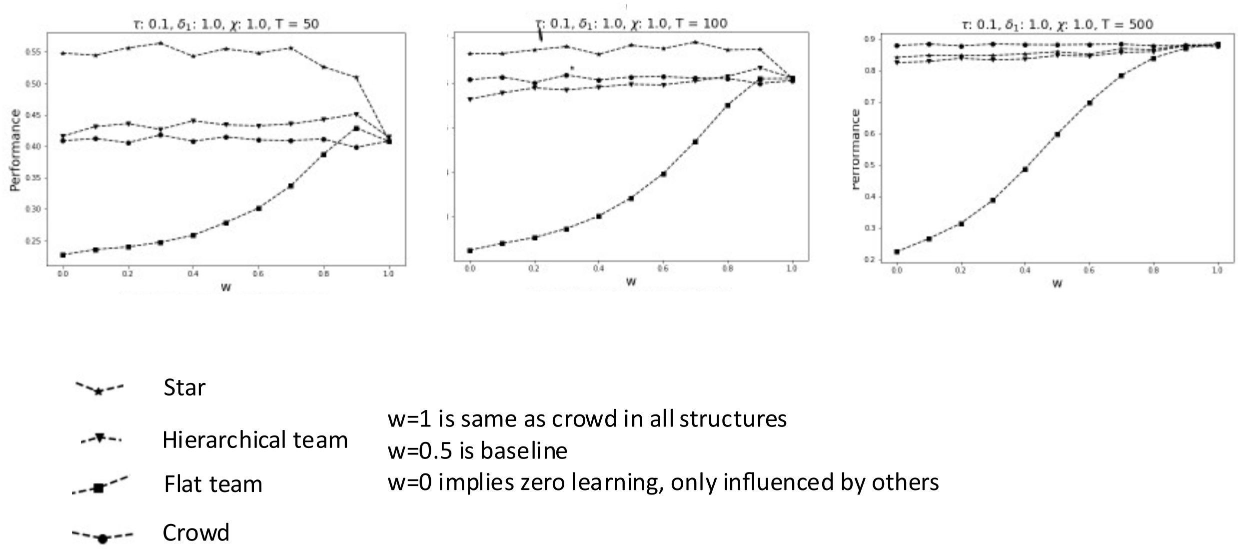

Figure §5.4 Cumulative performance as *w* varies for all actors that are influenced (N=S=7, δ_0 =0, δ_1 =1 χ =0).

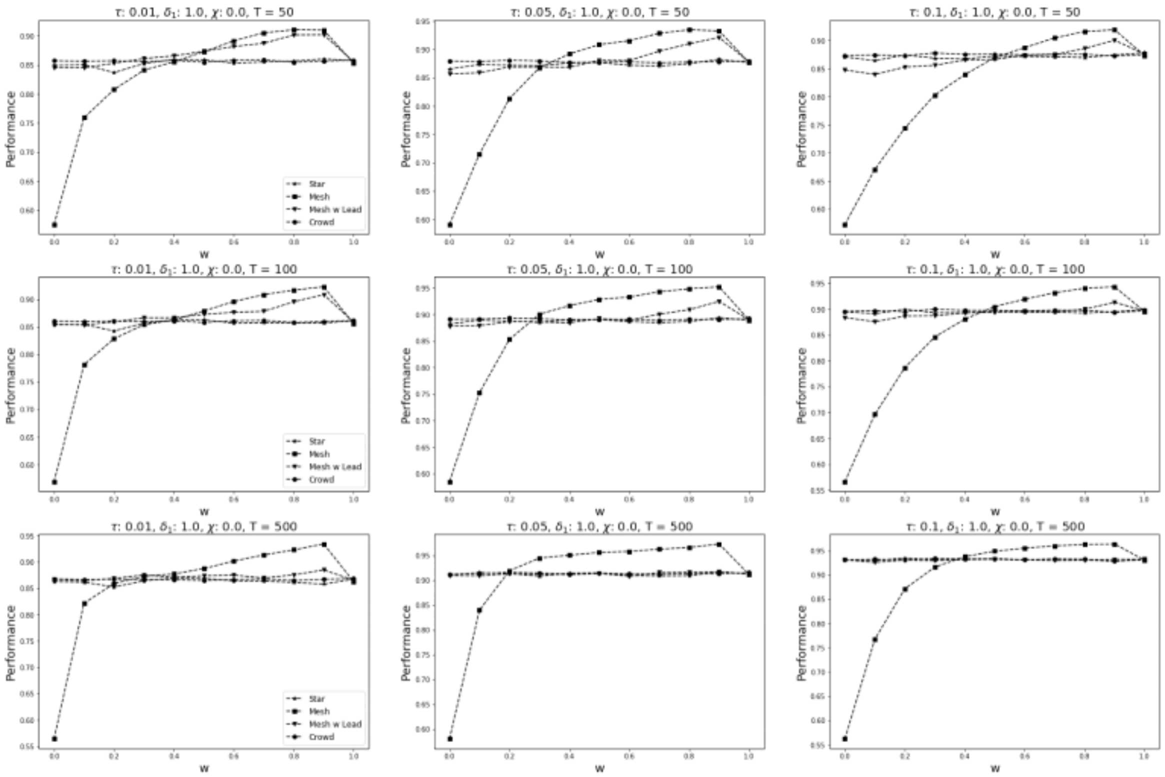

In Figures §5.3 and §5.4 we smoothly transition between flat teams, hierarchical teams and crowd by varying the (common) influence weight of all actors. We also add a "star" structure – with one central actor influencing all others. When w is ½ - we recover the baseline results for flat teams, hierarchical teams and crowd reported in the paper. When w→1 all structures converge to the crowd. When w→0, we get pure diffusion structures with zero learning from the environment.

§6 Interdependence structures that are isomorphic to influence structure (N=S=7, τ=0.05).

	χ=0	χ=0.5	χ=1.0
T=50	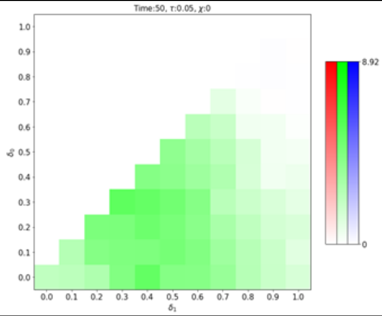	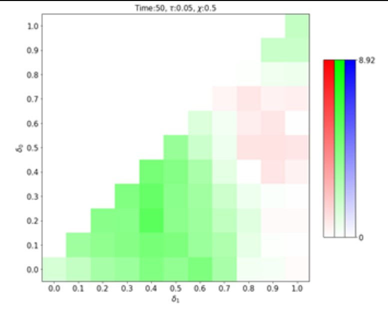	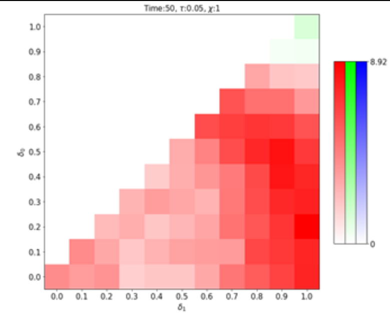
T=100	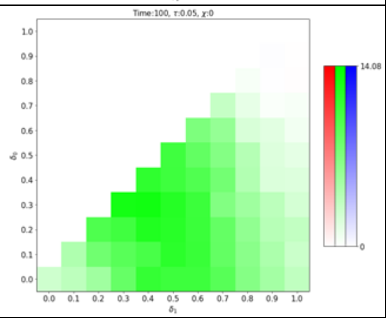	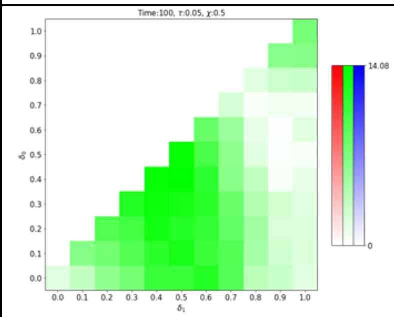	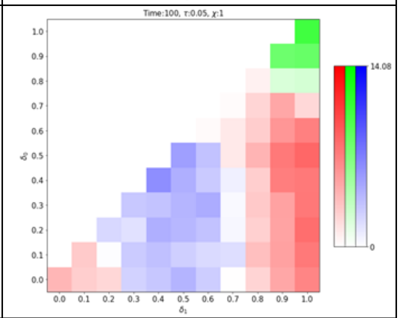
T=500	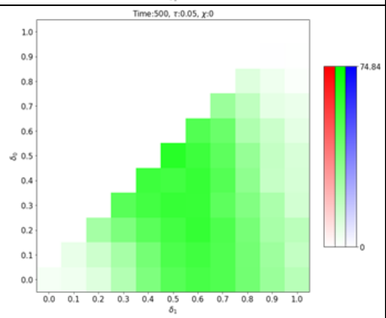	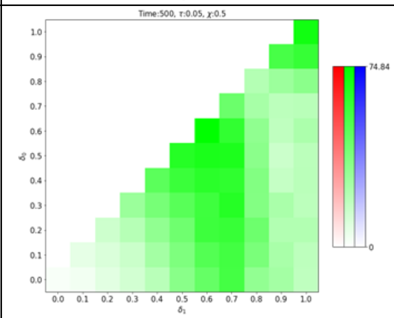	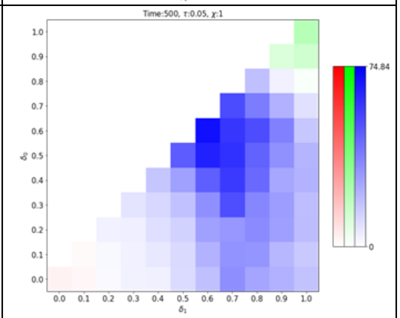

§7 Generalization to larger organizations

§7.1 Generalized learning in Random graphs & Branching Hierarchies vs. Crowds

Figure §7.1 Branching Hierarchy and Random Graphs vs. Crowds

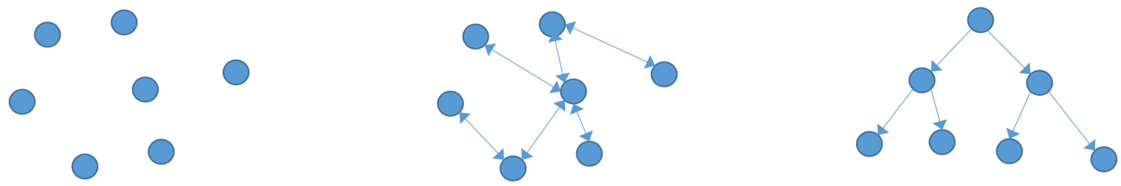

First, keeping the size of the system the same (N=S=7), we consider variants on flat (random graph-representing a community or "CO") and hierarchical teams (branching hierarchy or "BH"). The crowd ("CR") remains the comparison benchmark.

We also generalize the reinforcement learning process to incorporate the law of recency (Erev and Roth 1998). Each agent's payoff is used to update their assessment of attractiveness $b_{I,k,t}$, of the choice k they made at time t-1 for that agent, using the exponential recency weighted average rule:

$$b_{i,k,t}=b_{i,k,t-1} + \phi[\pi_{i,k,t}-b_{i,k,t-1}]$$

where $\phi \in [0,1]$ is called the step size parameter (Sutton and Barto 1998) and can be interpreted as a learning rate. The averaging rule used in the baseline analysis is equivalent to setting $\phi = 1/(n_k + 1)$ where n_k is the number of times the specific action has been selected before. Attractiveness of a choice increases when the current payoff exceeds past beliefs and vice versa.

In Figure §7.2 below (which can be viewed as a variant on Figure 3 in the paper), we can see that BH produces the greatest similarity in beliefs across agents, over all time scales. Resolution of beliefs first falls (through unlearning priors) then increases, for all organizational forms, with BH recovering the fastest. Beliefs eventually have the highest resolution in CR, suggesting that social influence, which makes beliefs similar in both CO and BH, also makes them lower resolution.[3] Greater resolution helps produce greater stability of choices from one period to the next. All three systems eventually reach fixed points, approaching stable behavior by T=200 periods or so, but at different speeds, with BH being the fastest. Mirroring our baseline results flat teams, the CO form attains steady state at a much slower pace than either BH or CR forms, because resolution in CO the lowest of the three forms.

[3] The kink in convergence and resolution effects in BH in the early periods reflects the beliefs of the apex agent being imposed on subordinates. This creates some convergence that creates good payoffs because χ>0. Subsequent learning by agents reduces the resolution (and convergence), which increases again as agents learn about their task environment.

Figure §7.2. Dynamics for alternative structures (N=7, S=7, δ_0=0, δ_1=1, χ=0.8, τ=0.1, ϕ=0.1)

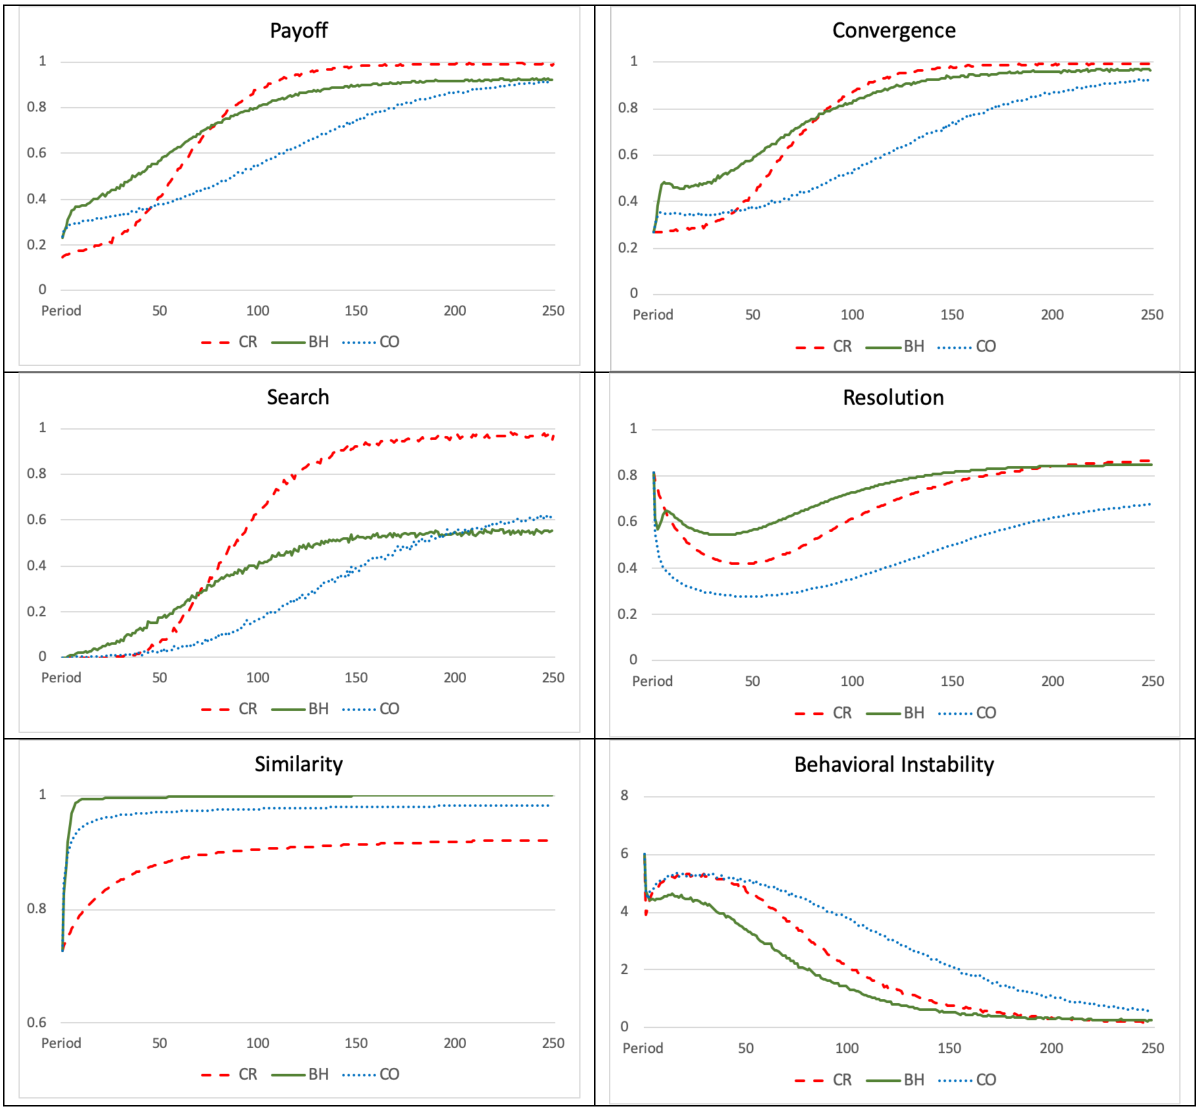

These results reiterate that the rapid convergence property of hierarchical influence (BH) is founded on its ability to create high similarity and high resolution of beliefs in a rapid manner.

Figure §7.3 Scaling up structure: Relative cumulative performance of BH, CR and CO by T=50 (δ_0=0, δ_1=1)

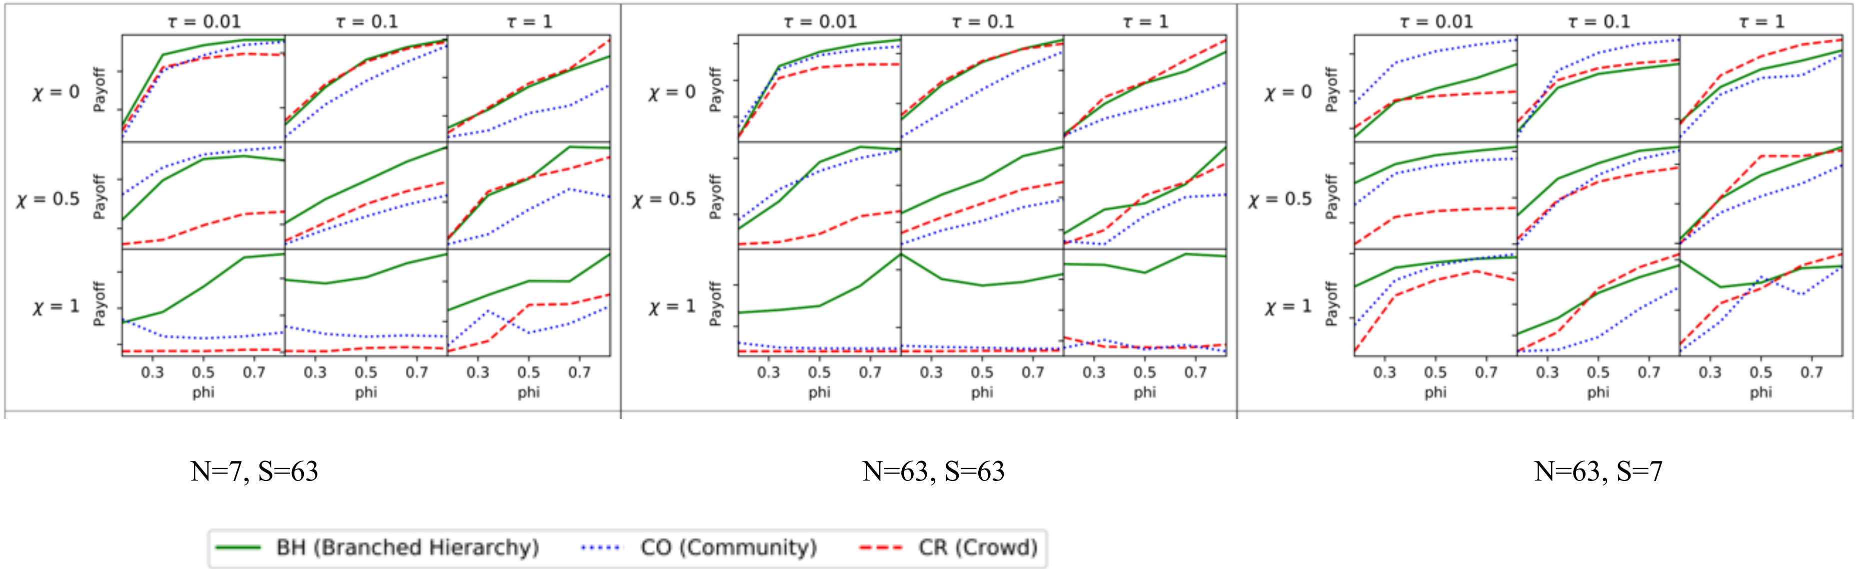

Increasing only the number of agents (N) heightens the challenge of coordination, as well as the risk of converging to the wrong organizational choice. A larger task environment (increasing S) makes convergence more challenging but also lowers the opportunity costs of converging to second or third best alternatives in that over a larger number of alternatives the second or third best will be more proximate to the largest value of v. Increasing both N and S shows that the increase in the search space drives most of the advantage for BH when scaling. These results confirm that while scaling the task environment (S) and organization (N) are alternate means to tune the challenges posed by variability and interdependence, the basic principle that hierarchic influence favors rapid convergence to good but not necessarily optimal organizational actions while symmetric influence promotes exploration, remains true.

§8. References

Friedkin NE, Johnsen EC (1990) Social influence and opinions. *Journal of Mathematical Sociology* 15(3-4): 193-206.

De Groot MH (1974) Reaching a consensus. *Journal of the American Statistical Association* 69(345): 118-121.

Holland PW, Leinhardt S (1976) Local structure in social networks. *Sociological Methodology* 7: 1-45.

Krackhardt D (1994) Graph theoretical dimensions of informal organizations. In KM Carley & MJ Prietula (Eds.), *Computational organization theory* (pp. 89-111). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

Lounamaa PH, March JG (1987) Adaptive coordination of a learning team. *Management Science* 33(1): 107-123.

Posen HE, Levinthal DA (2012) Chasing a moving target: Exploitation and exploration in dynamic environments. *Management Science* 58(3): 587-601.

Puranam, P, Swamy M (2016) How initial representations shape coupled learning processes. *Organization Science* 27(2): 323-335.