

An Agent-Based Model of Delegation Relationships With Hidden-Action: On the Effects of Heterogeneous Memory on Performance

Patrick Reinwald

Department of Management Control
and Strategic Management
University of Klagenfurt
Klagenfurt, Austria
Email: patrick.reinwald@aau.at
ORCID: 0000-0002-2907-7939

Stephan Leitner

Department of Management Control
and Strategic Management
University of Klagenfurt
Klagenfurt, Austria
Email: stephan.leitner@aau.at
ORCID: 0000-0001-6790-4651

Friederike Wall

Department of Management Control
and Strategic Management
University of Klagenfurt
Klagenfurt, Austria
Email: friederike.wall@aau.at
ORCID: 0000-0001-8001-8558

Abstract—We introduce an agent-based model of delegation relationships between a principal and an agent, which is based on the standard-hidden action model introduced by Holmström and, by doing so, provide a model which can be used to further explore theoretical topics in managerial economics, such as the efficiency of incentive mechanisms. We employ the concept of agentization, i.e., we systematically transform the standard hidden-action model into an agent-based model. Our modeling approach allows for a relaxation of some of the rather "heroic" assumptions included in the standard hidden-action model, whereby we particularly focus on assumptions related to the (i) availability of information about the environment and the (ii) principal's and agent's cognitive capabilities (with a particular focus on their learning capabilities and their memory). Our analysis focuses on how close and how fast the incentive scheme, which endogenously emerges from the agent-based model, converges to the solution proposed by the standard hidden-action model. Also, we investigate whether a stable solution can emerge from the agent-based model variant. The results show that in stable environments the emergent result can nearly reach the solution proposed by the standard hidden-action model. Surprisingly, the results indicate that turbulence in the environment leads to stability in earlier time periods.

Keywords—Agent-based modeling and simulation; Management control system; Information asymmetry; Complexity economics; Agentization.

I. INTRODUCTION

The standard hidden-action model introduced by Holmström [1] describes, in general, a delegation relation between a principal and an agent. It covers a situation in which the principal delegates a task to an agent. The agent selects an effort level in order to carry out this task, which is not observable by the principal (it is *hidden*). The agent's effort together with some environmental impact produces outcome which is to be shared between the principal and the agent. The principal's objective is to maximize her share of the outcome associated with the task while the agent strives for maximizing his share of the outcome minus his disutility from making effort to carry out the task. The standard hidden-action model proposes an incentive scheme (i.e., a rule to share the outcome) which aligns the agent's and the principal's objective so that the principal's utility finds its maximum. As only the outcome (but not the effort level) is observable for the principal, the sharing rule is based on outcome only.

In order to derive the optimal sharing rule, agency theory makes rather "heroic" assumptions about the capabilities of both parties with respect to (i) information processing capacity, (ii) availability of information, and (iii) capability to find the optimal solution immediately. Axtell [2] argues that the most "heroic" assumptions are agent homogeneity, non-interactiveness, and the existence of equilibrium solutions and refers to them as "neoclassical sweetspot". The result of these rather "heroic" assumptions is that the explanatory and the predictive power regarding real world problems substantially decreases [3]. Critics often refer to principal-agent models as "toy problems" and argue that solutions derived from such "toy problems", where authors can "assume complicating things away", are of limited use and only merely help to solve real world problems [4]-[6].

We take up on this critique and put the "heroic" assumptions included in the standard hidden-action model [1] in the focus of this paper. In the vein of Leitner and Wall [7], we transfer the standard hidden-action model into an agent-based model following a procedure introduced by Guerrero and Axtell [8] and Leitner and Behrens [9], which allows us to relax some of the included assumptions. In particular, we relax the assumptions related to the principal's and the agent's (i) information processing capacity and (ii) the availability of information. We model situations in which the principal and the agent no longer have full information about the environment but have to learn this information over time. In consequence, they can no longer find the optimal (second-best) solution immediately, like it is the case in the standard hidden-action model. We, therefore, endow the principal and the agent with the capability to search for the best possible incentive scheme over time by employing a hillclimbing algorithm. Please notice that the optimal solution, which is derived from the standard hidden-action model for cases in which the the agent's effort is *not* observable, is referred to as *second-best solution*. The first-best solution, on the contrary, assumes that the agent's effort is observable [10], which, in consequence, means that no incentive problem in the above outlined sense arises.

The remainder of this paper is organized as follows: Section II introduces two variants of the hidden-action model. First the main features of the standard hidden-action model are

presented. We, then introduce the agent-based model variant which relaxes some of the assumptions included in the standard hidden-action model. In Section III, we elaborate on the simulation setup, and introduce and discuss the results. Section IV concludes the paper and gives an outlook on future work.

II. THE HIDDEN-ACTION MODEL

This section summarizes the main features of the standard hidden-action model introduced in [1] and proposes an agent-based representation of the hidden-action problem in which the principal and the agent are endowed with limited and heterogeneous memory.

A. The standard hidden-action model

The standard hidden-action model, which describes a delegation relation between a principal and an agent, was first described by Holmström [1]. This model covers a situation in which a principal offers an agent a contract upon a task to be carried out and a sharing rule over the generated outcome, among other things. If the agent agrees on the condition stated in the contract, he exerts effort to complete the specified task. Together with an exogenous factor, his effort generates outcome, but this also leads to disutility for him. For the principal, however, the effort the agent had carried out is unobservable, which results in a situations where only the outcome can be used as a basis for the sharing rule. Both the principal and the agent are individual utility maximizers [10], [5]. Furthermore, it is assumed that the principal is risk neutral and characterized by her utility function

$$U_P(x, s) = x - s(x) , \quad (1)$$

whereby x represents the generated outcome and $s = s(x)$ is the function for the sharing rule. As mentioned before, the outcome

$$x = f(a, \theta) , \quad (2)$$

is a function of the agent's effort a and the exogenous factor θ . The agent, who is assumed to be risk averse, is characterized by the utility function

$$U_A(s, a) = V(s) - G(a) , \quad (3)$$

where $V(s)$ represents the generated utility from the compensation and $G(a)$ denotes to disutility from the exerted effort. For the contract to be effective, two additional constraints have to be fulfilled. First, the participation constraint

$$E(U_a(s, a)) \geq \bar{U} , \quad (4)$$

which assures that the agent gets at least the utility he would get from the best outside option \bar{U} . Second, the incentive compatibility constraint

$$a \in \arg \max_a E\{U_A(s(x), a')\} , \quad (5)$$

which aligns the agent's goal (maximizing his utility) with the goal of the principal. For an extensive review and the formal solution to this problem, the reader is, for example, referred to Holmström [1] and Lambert [10].




Figure 1. Flow diagram

B. Agent-based model variant

We transfer the standard hidden-action model [1] into an agent-based model following the agentization procedure introduced by Guerrero and Axtell [8] and Leitner and Behrens [9]. We limit the principal's and the agent's availability of information about the environment and endow them with the capability to learn about the environment over time. As a consequence of this change, the principal and the agent can no longer find the optimal (second-best) contract immediately but search for it over time. This stepwise search for the optimal contract requires us to switch from a one-periodic to a multi-periodic model in which the principal and the agent agree upon a contract in every timestep. Due to the employed learning mechanism (which is introduced below), the principal's and the agent's states of information are changing over time. We do not give the agent the possibility to allocate effort over multiple periods. A condensed overview of the sequence of the events is provided in Figure 1.

We indicate time steps by $t = 1, \dots, T$. After the adjustment to a multi-periodic model, the principal's utility function introduced in (1) takes the form of

$$U_P(x_t, s(x_t)) = x_t - s(x_t) , \quad (6)$$

where x_t denotes the outcome and $s(x_t)$ is the agent's compensation in time step t . The production function introduced in (2) is adapted, so that

$$x_t = a_t + \theta_t , \quad (7)$$

where a_t stands for the effort-level selected by the agent from a set of all feasible actions A_t in t . A_t is a subset of A with the range from zero to the value a_t , whereby the latter refers to the effort-level which is required to achieve the second-best solution according to the standard hidden-action model. Setting the boundaries in this way is in line with [1] and assures the feasibility of the solution. We denote the effort level for which the principal designs the contract by \tilde{a}_t and refer to it as incited effort. The variable θ_t denotes the realized exogenous factor in t which follows a Normal distribution. Given these adaptations, the agent's compensation function $s(x_t)$ takes the form of

$$s(x_t) = x_t * p_t , \quad (8)$$

where $p_t \in [0, 1]$ is the premium parameter in t . After the adaption to a multi-periodic model, the agent's utility function

takes the form of

$$U_A(s(x_t), a_t) = \frac{\overbrace{1 - e^{-\eta * s(x_t)}}^{V(s(x_t))}}{\eta} - \frac{\overbrace{a_t^2}}{2}, \quad (9)$$

where η represents the agent's Arrow-Pratt measure of risk-aversion [11]. The agent strives for maximizing this utility function but, due to the adaptations outlined below, replaces the environmental factor θ_t by his expectation thereof. For the computation of the agent's expectation about the environmental factor see (11) and (13).

The most central change during the process of agentization is the relaxation of the availability of information about the environment for the principal and the agent. In the standard hidden-action model [1], it is assumed that both parties have knowledge about the distribution of the exogenous factor and its parameterization, so that they can compute an expected value for the environmental variable. In contrast to this, in the agent-based model variant they have to learn about the environment over time. The assumptions related to the environment are relaxed in the following way:

- 1) The principal and the agent no longer have full information about the environmental factor at the beginning of the simulation.
- 2) The principal and the agent are able to individually learn about the environment over time.
- 3) The principal and the agent also have the cognitive ability to store the gathered information about the environment in their memory.

These changes are implemented by a simultaneous and sequential learning model (see also [7]). Part of this learning model is not only the process of gathering and storing the information but also the application of this information to compute the estimation of the exogenous factor. Both the principal and the agent are now able to individually learn about the environment and estimate the exogenous factor in the same way: They make their conclusions about the environment on the basis of the observable outcome x_t (see (7)), which is possible because the only unknown information in the equation is the realized environmental factor. Recall, both the principal and the agent know the realized outcome x_t , the agent knows the exerted effort a_t , and the principal knows the incited effort \tilde{a}_t . Using this information, the principal and the agent can infer the estimations of the environmental factor from the realized outcome, according to

$$\tilde{\theta}_t = x_t - \tilde{a}_t, \quad (10)$$

and

$$\theta_t = x_t - a_t, \quad (11)$$

respectively. Furthermore, they are able to privately store the estimations in their memory until their defined cognitive capacity (m_P for the principal and m_A for the agent) is reached. Once the cognitive capacity is reached, the oldest entries are replaced by the most recent observations. Their expectation about the environmental factor is computed by averaging all privately stored and retrievable estimations of

TABLE I. NOTATION FOR THE AGENT-BASED MODEL VARIANT

Description	Parameter
Timesteps	t
Principal's utility	U_P
Agent's utility	U_A
Agent's Arrow-Pratt measure of risk-aversion	η
Agent's share of outcome in t	$s(x_t) = x_t * p_t$
Outcome	$x_t = a_t + \theta_t$
Principal's expected outcome	\tilde{x}_{Pt}
Premium parameter in t	p_t
Exerted effort level in t	a_t
Induced effort level by the principal in t	\tilde{a}_t
Set of all feasible actions in t	A_t
Exogenous (environment) variable in t	θ_t
Principal's estimation of the realized exogenous factor in t	$\tilde{\theta}_t$
Mental capability of the principal	m_P
Mental capability of the agent	m_A
Averaged expected exogenous factor of the principal	$\hat{\theta}_{Pt}$
Averaged expected exogenous factor of the agent	$\hat{\theta}_{At}$

exogenous factors, so that

$$\hat{\theta}_{Pt} = \begin{cases} \frac{1}{t-1} \sum_{n=1}^{n=t-1} \tilde{\theta}_n & \text{if } m_P = \infty, \\ \frac{1}{m_P} \sum_{\substack{n=1 \\ \forall t \leq m_P: n=1 \\ \forall t > m_P: n=t-m_P}}^{n=t-1} \tilde{\theta}_n & \text{if } m_P < \infty, \end{cases} \quad (12)$$

for the principal and

$$\hat{\theta}_{At} = \begin{cases} \frac{1}{t-1} \sum_{n=1}^{n=t-1} \theta_n & \text{if } m_A = \infty, \\ \frac{1}{m_A} \sum_{\substack{n=1 \\ \forall t \leq m_A: n=1 \\ \forall t > m_A: n=t-m_A}}^{n=t-1} \theta_n & \text{if } m_A < \infty, \end{cases} \quad (13)$$

for the agent.

Next, the principal randomly discovers two alternative effort levels in the search space A_t which together with \tilde{a}_t serve as candidates for \tilde{a}_{t+1} . Using her expectation about the environment $\hat{\theta}_{Pt}$, the principal computes the expected outcome \tilde{x}_{Pt} (according to (7)) and her expected utility (according to (6)) associated with all three candidates. Notice that for all effort levels the associated premium parameter is computed according to

$$p_t = \max_{p \in [0,1]} U_p(\tilde{x}_{Pt}, s(\tilde{x}_{Pt})). \quad (14)$$

Finally, the principal selects the candidate with the highest expected utility as desired effort level \tilde{a}_{t+1} for period $t+1$ and communicates the associated premium parameter p_{t+1} to the agent. In period $t+1$, the agent starts over the procedure with selecting an effort level a_{t+1} using p_{t+1} (see Figure 1).

III. SIMULATIONS PARAMETERS AND RESULTS

We have conducted 8 scenarios with different underlying assumptions about the cognitive capacity (memory) of the principal and the agent and two different levels of environmental turbulence.

Recall that the environmental variable is modeled to follow a Normal distribution. The variations in environmental turbulence are operationalized by altering this distribution's standard deviation which is set relative to the optimal outcome x of the standard hidden-action model (second-best solution in [1]). For the case of a rather stable environment we set $\sigma = 0.05x$ and for a rather unstable environment we set $\sigma = 0.45x$. The distribution's mean is fixed at zero.

In the 8 investigated scenarios, we have, in general, two configurations related to the cognitive capacity of the principal and the agent. The first one assumes that the principal has advantage in cognitive capacity and the second assumes the opposite situation. Each of them consist of 4 scenarios where the one in advantage always has unlimited memory and the other one has either a memory with the length of one or five. The agent's Arrow-Pratt measure is always set to $\eta = 0.5$, which characterizes a risk-averse agent.

For every scenario, we perform 700 repetitions ($R = 700$), our analysis focuses on the first 20 time steps ($t = 1, \dots, 20$). In order to implement our simulation model we use MATLAB®.

We report the averaged normalized effort level carried out by the agent in every period t as performance measure. Therefore in every simulation run $r = 1, \dots, R$ and for every period $t = 1, \dots, T$ we track the effort level a_{tr} chosen by the agent, and normalize it by the optimal level of effort a^* .

$$\phi_t = \frac{1}{R} \sum_{r=1}^{r=R} \frac{a_{tr}}{a^*} \quad (15)$$

We report the averaged normalized effort level as a measure of performance, as provides fundamental insights into the functioning of the emergent incentive schemes without further perturbation caused by environmental turbulence.

A. Results 1: Advantage in information for the agent

First, we analyze the results in the situation with the advantage for the agent. In the two scenarios with a relatively stable environment ($\sigma = 0.05x$) we can see that performance is significantly higher (at the end of the observation period) for cases in which the principal's cognitive capacity is higher, too (in terms of an increase in memory length): As the principal's memory increases from $m_P = 1$ to $m_P = 5$, the observed final averaged normalized effort level increases from 0.9063 to 0.9474 (see the two top subplots in Figure 2). We can observe similar results for the scenarios with a rather unstable environment ($\sigma = 0.45x$). Here, the final performance measure increases from 0.8135 to 0.8785 as the principal's memory increases from $m_P = 1$ to $m_P = 5$ (see the two bottom subplots in Figure 2). Furthermore, in line with intuition, we can observe that an increase in environmental turbulence leads to a decrease in the effort exerted by the agent.

We further analyze the effect of the principal's memory on the average variance of the results (i.e., the average variance of the effort exerted by the agent throughout the entire observation period): For the scenarios in a rather unstable environment, the average standard deviation significantly decreases from 0.17405 (for $m_P = 1$) to 0.15719 (for $m_P = 5$). For the cases in a rather stable environment no significant differences can be observed. The significance at the 99%-level was confirmed using an F-test. Thus, increasing the principal's cognitive capacity in an unstable environment significantly reduces the variance of the effort induced by the incentive scheme (and exerted by the agent). In other words, if the principal manages to increase her memory in an turbulent environment, she not only increases the average normalized effort level but also significantly reduces the risk of extreme deviations from this value.

Finally, we take a look on the stability of the averaged normalized effort level. We regard a solution to be stable as soon as (i) the averaged normalized effort level at period t




Figure 2. Situations with advantage regarding cognitive capacity for the agent (A) and the cognitive capacity of the principal (P) with either $m_P = 1$ or $m_P = 5$. Scenarios a plotted for two different environmental situations, stable ($\sigma = 0.05x$) and unstable $\sigma = 0.45x$.

is not significantly different from the same measure in period $t-1$, and given that (ii) this condition does not change after this point in time. For this analysis, we perform a T-test and set $\alpha = 0.01$. First, we focus on rather stable environments: For cases with a low (high) cognitive capacity for the principal of $m_P = 1$ and $m_P = 5$, we observe a stable solution for periods $t = 7$ and $t = 9$ onwards, respectively. For unstable environments, we can observe that a stable solution emerges earlier in both scenarios. For $m_P = 1$ ($m_P = 5$) the solution becomes stable in period $t = 4$ ($t = 6$). This is a counter-intuitive result as one would expect that turbulence in the environment leads to instability in the emergent solution.

B. Result 2: Advantage in information for the principal

This section analyses the cases in which the principal has an advantage in information. In stable environments ($\sigma = 0.05x$), we cannot observe significant differences: For the agent's memory of $m_A = 1$ and $m_A = 5$, the observed final averaged normalized effort levels are 0.9784 and 0.9786, respectively (see the two top subplots in Figure 3). The same result can be observed for unstable environments ($\sigma = 0.45x$). Here, the final performances are 0.9507 and 0.9512 for $m_A = 1$ and $m_A = 5$, respectively (see the two bottom subplots in Figure 3). Additionally, in line with intuition, we can observe that an increase in environmental turbulence leads to a decrease in the effort exerted by the agent.

Further, we analyze the effects of the agent's memory on the average variance of the results (i.e., the average variance of the effort exerted by the agent throughout the entire observation period). For unstable environments ($\sigma = 0.45x$), the observed results are similar to the ones presented in the previous section: The variances of the exerted effort levels are significantly different at the 99%-level for $m_A = 1$ and $m_A = 5$, for the former (latter) we observe a standard deviation of 0.1953 (0.15781). For stable environments, the results indicate that increasing the agent's memory significantly decreases the standard deviation of the exerted effort at the 95%-level:




Figure 3. Scenarios with advantage regarding cognitive capacity for the principal (P) and the cognitive capacity of the agent (A) with either $m_A = 1$ or $m_A = 5$. Scenarios are plotted for two different environmental situations, stable ($\sigma = 0.05x$) and unstable $\sigma = 0.45x$.

We observe 0.14089 and 0.13821 for $m_A = 1$ and $m_A = 5$, respectively. Thus, increasing the agent's cognitive capacity significantly reduces the variance of the effort induced by the incentive scheme (and exerted by the him). In other words, if the agent manages to increase his memory, he significantly reduces the risk of extreme deviations from the performance value.

Finally, we investigate the stability of the averaged normalized effort level. First, we focus on the scenarios with rather stable environments: For the cognitive capacity of the agent of $m = 1$ and $m = 5$ we reach the stable point at period $t = 14$ and $t = 13$, respectively. For unstable environments we can observe in both scenarios that a stable solution emerges earlier. For $m = 1$ ($m = 5$) the performance reaches a stable point at period $t = 9$ ($t = 7$). This leads to the same counter-intuitive result as in the section before, as one would expect that turbulence in the environment leads to instability in the emergent solution.

IV. CONCLUSION AND FUTURE WORK

The results presented in this paper deliver some insights about the effects of heterogeneous agents in a hidden-action setting:

- Our results suggest, that gathering information about the environment is a good strategy for the principal to increase his utility especially situations in which the environment is rather turbulent.
- In turbulent environments, increasing the memory of both the principal and the agent always has a positive effect on the variance of the results. This means that increasing the memory significantly reduces the risk of extreme deviations from the performance measures reported above.
- In stable environments, the results, on the contrary, suggest that only an increase of the agent's memory leads to a significant decrease in the exerted effort's variance.

- Surprisingly, the presented results indicate, that turbulence has a positive effect on stability, so that a stable solution emerges earlier in turbulent environments.

Future work might want to deeper investigate the effects of heterogeneous memory in the hidden-action setting (e.g. more memory length) and also include cognitive biases when characterizing the principal's and the agent's cognitive capabilities (such as the recency or the primacy effect [12]). Another potentially fruitful option avenue for future research might be to limit the principal's knowledge about the characteristics of the agent (such as the utility function).

ACKNOWLEDGMENT

This work was supported by funds of the Oesterreichische Nationalbank (Austrian Central Bank, Anniversary Fund, project number: 17930).

REFERENCES

- [1] B. Holmström, "Moral hazard and observability," *The Bell Journal of Economics*, vol. 10, no. 1, 1979, p. 74.
- [2] R. L. Axtell, "What economic agents do: How cognition and interaction lead to emergence and complexity," *The Review of Austrian Economics*, vol. 20, no. 2-3, 2007, pp. 105–122.
- [3] C. Schneeweiss, "On a formalisation of the process of quantitative model building," *European Journal of Operational Research*, vol. 29, no. 1, 1987, pp. 24–41.
- [4] M. C. Jensen, "Organization Theory and Methodology," *The Accounting Review*, vol. 58, no. 2, 1983, pp. 319–339.
- [5] K. M. Eisenhardt, "Agency theory: An assessment and review," *The Academy of Management Review*, vol. 14, no. 1, 1989, p. 57.
- [6] S. P. Shapiro, "Agency Theory," *Annual Review of Sociology*, vol. 31, no. 1, 2005, pp. 263–284.
- [7] S. Leitner and F. Wall, "Decision-facilitating information in hidden-action setups: an agent-based approach," *Journal of Economic Interaction and Coordination*, 2020, pp. 1–36.
- [8] O. A. Guerrero and R. L. Axtell, "Using Agentization for Exploring Firm and Labor Dynamics," in *Emergent Results of Artificial Economics*, S. Osinga, G. J. Hofstede, and T. Verwaart, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 139–150.
- [9] S. Leitner and D. A. Behrens, "On the efficiency of hurdle rate-based coordination mechanisms," *Mathematical and Computer Modelling of Dynamical Systems*, vol. 21, no. 5, 2015, pp. 413–431.
- [10] R. A. Lambert, "Contracting theory and accounting," *Journal of Accounting and Economics*, vol. 32, no. 1-3, 2001, pp. 3–87.
- [11] K. J. Arrow, "The Role of Securities in the Optimal Allocation of Risk-bearing," *The Review of Economic Studies*, vol. 31, no. 2, 1964, p. 91.
- [12] A. Mantonakis, P. Rodero, I. Lesschaeve, and R. Hastie, "Order in choice: effects of serial position on preferences," *Psychological science*, vol. 20, no. 11, 2009, pp. 1309–1312.