Agent-Based Analysis of Green Disclosure Policies and their Market-Wide Impact on Firm Behavior

Lingxiao Zhao
King's College London
London, UK
lingxiao.zhao@kcl.ac.uk

Maria Polukarov
King's College London
London, UK
maria.polukarov@kcl.ac.uk

Carmine Ventre
King's College London
London, UK
carmine.ventre@kcl.ac.uk

ABSTRACT

Green disclosure policies are designed to help firms communicate their environmentally friendly practices to investors. While most research has focused on the effects of these policies at the individual firm level, their influence within a system of multiple firms remains largely unexamined. To address this gap, we develop an agent-based model to simulate market dynamics among firms with varying levels of environmental performance and strategic responses. Using Empirical Game-Theoretic Analysis, we investigate how the costs associated with becoming greener and investors' valuation of these efforts shape equilibrium outcomes and the prevalence of green firms in the market. Our findings reveal that changes in the cost of green upgrades significantly influence firms' strategic choices and alter the equilibrium behavior of the other firms. Additionally, we analyze the effects of different green disclosure policies and find that under more relaxed policies, firms are more willing to incur into higher upgrade costs. Furthermore, we propose a two-stage disclosure policy that incentivizes all types of firms to improve their green practices, leading to broader adoption of sustainable energy solutions across the market.

KEYWORDS

Green Disclosure Policies, Agent-Based Simulation, Empirical Game-Theoretic Analysis

ACM Reference Format:
Lingxiao Zhao, Maria Polukarov, and Carmine Ventre. 2025. Agent-Based Analysis of Green Disclosure Policies and their Market-Wide Impact on Firm Behavior. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

In recent years, the global demand for environmental responsibility has grown, driven by climate change and the Sustainable Development Goals (SDGs). It has prompted companies to incorporate environmental factors into their operations. As a result, more firms and investors are adopting Environmental, Social, and Governance (ESG) frameworks and disclosing green initiatives to enhance credibility and transparency in society and the market.

Meanwhile, investors increasingly prefer firms that adopt green practices. The EU SEIP plans to mobilize €1 trillion for sustainability projects over the next decade. Similarly, the BlackRock Sustainable European Companies Fund also targets ESG-compliant firms.These ESG funds allocate investments based on firms' green disclosure states, further driving corporate transitions toward sustainability.

Green disclosure is the process by which companies share information about their eco-friendly practices with investors. This transparency facilitates efficient green capital allocation and encourages companies to adopt sustainable practices. Therefore, governments and regulators are increasingly mandating environmental disclosures. For example, the EIB promotes sustainable development by financing green disclosure projects as the EU's policy bank [5].

In general, green disclosure regulation requires firms to balance the costs of upgrading and transformation against the potential for higher valuations from investors. Meanwhile, the limited size of ESG funds intensifies market competition, requiring firms to attract investment and maintain a competitive edge against other firms.

To better analyze these multi-agent dynamics, we develop a green disclosure model for a large financial market following the dynamic disclosure model of [9] and employing the agent-based modelling (ABM) framework [13, 18]. Our approach innovates by extending the disclosure framework to a large-scale market, enabling a detailed analysis of corporate transformation costs, investor preferences, and market-wide interactions. This allows us to assess how different disclosure policies influence firms' green transitions.

Within this simulated financial market, limited ESG funds influence the strategic game among firms at various green levels. Using ABM simulations, we generate the game's payoff matrix and apply empirical game theory analysis (EGTA) method [20, 22] to identify equilibria, providing insights into agent strategies and system-wide green transition rates.

2 RELATED WORK

The primary function of green disclosure is to reduce the asymmetry of market information by creating a unified quality label, directing capital to green activities and encouraging corporate sustainability [12, 19]. Since its introduction, many studies have examined its impact on corporate operations, often focusing on specific disclosure policies [2, 14, 26]. For instance, Blanca et al. [4] found that companies operating in unregulated industries often disclose more green information to respond to stakeholder pressure. In contrast, highly internationalized firms face more complex trade-offs regarding disclosure transparency. Wu et al. [25] emphasized the critical role of ESG market transparency in enhancing market efficiency. They pointed out the phenomenon of "greenwashing" in green disclosure, where profit-driven companies deceive investors by mimicking the investment behaviours of companies genuinely involved in green initiatives.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

At the same time, many scholars believe that only relying on a single form of disclosure is insufficient to significantly enhance the allocation efficiency of ESG investments [1]. Fairfax [6] introduced the concept of "dynamic disclosure," arguing that the modern disclosure system operates as an evolving feedback loop in which voluntary and mandatory disclosures are inherently connected and mutually reinforcing. Companies are not only required to comply with compulsory disclosure obligations but also need to enhance transparency through voluntary disclosures.

Additionally, many ESG experts have examined the impact of green disclosure on companies and investors by constructing virtual models [11]. Martin and Moser [15]simulated a market to analyze how investors decide based on managers disclosing or withholding information about green investments. They found that while disclosing affects costs, investors respond positively to the social benefits, showing that financial gains and social impact influence managerial decisions.

Lastly, the ongoing development of algorithmic game theory is proving to be a valuable tool for economic analysis, especially in the green economics field [7]. Cui et al. [3] proposed an evolutionary model using multi-agent games to explore agent sustainability within the green finance framework. Their results show that agent full participation reduces information asymmetry, ensures efficient resource allocation, and promotes green production. Enhanced regulation and lower production costs further improve the efficiency of the green financial market. Pástor et al. [17] developed an equilibrium model to examine the influence of ESG standards on asset pricing. The result shows that ESG preferences increase green firms' market value and lower their capital costs. While green assets typically have lower expected returns, they outperform polluting assets as ESG factors strengthen.

3 GREEN DISCLOSURE MODEL

Our model builds upon Gupta and Starmans's dynamic green disclosure framework for a single firm while taking a bottom-up approach to design a virtual ESG fund market. In this market, agents make decisions on energy upgrades based on localized information. We begin by outlining the ESG market structure and the available strategy set, followed by a detailed demonstration of the game dynamics.

3.1 ESG Markets

We compared the 2024 reports from the MSCI Europe ESG Leaders Index and the BlackRock ESG Strategic Growth Fund [8, 10] and found that large-scale ESG funds typically invest in 150 to 300 firms. Based on this, we selected 150 companies from various categories as model agents. We believe that this selection is sufficient to evaluate the impact of different disclosure models on the overall market.

We construct an ESG fund model consisting of a set N of firms, divided into three categories (or types) of agents, as follows: the d category that represents so called "dirty" firms, who are still not working on green activities; the Lg category that refers to low-green firms; and the Hg category for high-green firms, with the degree of green energy transformation increasing from d to Hg. The set N can then be written as $N = N_d \cup N_{Lg} \cup N_{Hg}$, where N_d, N_{Lg}, N_{Hg} are the respective subsets of agents corresponding to each type. Throughout the paper, we shall slightly abuse the notation and use $N_{(\cdot)}$ to also indicate the size of (i.e., the number of agents in) the respective set, and we initially set $N_d = N_{Lg} = N_{Hg} = 50$, summing up to $N = 150$.

3.1.1 *Firm Profits.*

In the simulation of the model, our focus is on analyzing the interactions between agents at different stages of green transformation within the ESG market. To simplify the analysis, we assume that all agents' operating profit π is the same for all stages, with a random noise disturbance added to account for variability. For each agent i, the profit $\tilde{\pi_i}^t$ at stage t, is expressed as follows:

$$\tilde{\pi_i}^t = \pi_i^t + \epsilon,$$

where π_t represents the initial profit of the agent in the model, and ϵ represents a random disturbance.

3.1.2 *Upgrade Cost.*

In the ESG market, all dirty and low-green firms can choose to upgrade to a higher green level. At each time stage t, if a firm $i \in N_d \cup N_{Lg}$ decides to upgrade one step up, the associated upgrade costs are $c_{i,d}^t$ and $c_{i,Lg}^t$, respectively. We posit that there is a correlation between the upgrade cost and a firm's profit, expressed as:

$$c_{i,d}^t, c_{i,Lg}^t = \alpha \cdot \tilde{\pi_i}^t.$$

Here, α represents the proportional coefficient between the upgrade cost and the firm's profit, which measures the degree to which the upgrade cost changes with the enterprise profit.

3.1.3 *Investor Green Valuation.*

The firm types influence the valuation of these agents by ESG funds. We use λ to represent investors' valuation of firms within green finance: $\lambda = \{\lambda_d, \lambda_{Lg}, \lambda_{Hg}\}$. Since the degree of green transformation follows a progression from low to high, the valuation relationship between different types of firms is expressed as follows:

$$\lambda_{Hg} > \lambda_{Lg} > \lambda_d.$$

We can thus calculate the difference between λ of different company types:

$$\Delta\lambda_{\mathrm{I}} = \lambda_{Lg} - \lambda_d,$$
$$\Delta\lambda_{\mathrm{II}} = \lambda_{Hg} - \lambda_{Lg}.$$

In ESG funds, firms receive varying investment allocations based on their green ratings, with higher-rated green firms capturing a larger share. Specifically, each firm occupies a portion of the ESG fund, but the proportion depends on its green rating. Given N_{Hg} high-green, N_{Lg} low-green, and N_d dirty firms, their market shares are distributed as follows:

$$q_i^t(\mathbf{x}^t) = \frac{\lambda(\theta_i(\mathbf{x}^t))}{N_{Hg}(\mathbf{x}^t)\lambda_{Hg} + N_{Lg}(\mathbf{x}^t)\lambda_{Lg} + N_d(\mathbf{x}^t)\lambda_d} \cdot \frac{I}{\sum_{j=1}^{N} \lambda(\theta_j(\mathbf{x}^t))}. \tag{1}$$

$q_i^t(\mathbf{x}^t)$ represents the market share of firm i. Here, $\mathbf{x}^t = (x_1^t, \ldots, x_n^t)$ is the strategy profile at time t, where each x_j^t indicates the strategy chosen by firm j. A strategy x_j^t specifies whether a firm chooses to maintain its current status or upgrade to a greener type. This choice influences the firm's type θ_j.

$$x_j^t = \begin{cases} 1, & \text{if } j \text{ chooses to upgrade} \\ 0, & \text{otherwise} \end{cases}$$

In the first part of the formula, the numerator $\lambda(\theta_i(\mathbf{x}^t))$ denotes the green valuation of firm i, where $\theta_i(\mathbf{x}^t)$ represents the type of firm (high-level green, low-level green, or dirty) as determined by the strategy profile $\mathbf{x}^t$. The denominator, $N_{Hg}(\mathbf{x}^t)\lambda_{Hg} + N_{Lg}(\mathbf{x}^t)\lambda_{Lg}+N_d(\mathbf{x}^t)\lambda_d$, is the weighted sum of the green valuations of all firms in the market, reflecting the relative competitiveness of different firm types based on the strategies chosen. The second part of the formula, $\frac{I}{\sum_{j=1}^{N}\lambda(\theta_j(\mathbf{x}^t))}$, determines how the total market investment I is allocated based on the green valuations $\lambda(\theta_j(\mathbf{x}^t))$ of all firms j.

Therefore, each agent needs to consider the potential impact of other agents when deciding to upgrade. As the number of green enterprises increases, the overall green valuation of the market rises. However, since the total investment I remains constant, an individual firm's market share $q_i^t(\mathbf{x}^t)$ may decrease due to increased competition from other green firms.

3.1.4 Payoff. At the beginning of each stage t, the agent can choose a strategy, whether to maintain its current status or to upgrade to a greener type. We use the expected payoff of that stage to represent its payoff, denoted as u_i^t for each agent i:

$$\begin{aligned} u_i^t &= \beta \cdot \left(\tilde{\pi}_i^t + q_i^t(\mathbf{x}^t)\cdot I - C_i^t(x_i, \theta_i(\mathbf{x}^t))\right) \\ &= \beta \cdot \Big(\tilde{\pi}_i^t + q_i^t(\mathbf{x}^t)\cdot I \\ &\quad - x_i \cdot \left[\mathbb{1}_{\theta_i(\mathbf{x}^t)=d}\cdot c_{i,d}^t + \mathbb{1}_{\theta_i(\mathbf{x}^t)=Lg}\cdot c_{i,Lg}^t\right]\Big). \end{aligned} \tag{2}$$

Here, β represents the discount factor within the current stage, used to adjust the expected payoff to reflect its present value. In the cost function $C_i^t(x_i,\theta_i(\mathbf{x}^t))$, $\theta_i(\mathbf{x}^t)$ represents the firm type, determined by the strategy profile $\mathbf{x}^t$. The indicator functions $\mathbb{1}_{\theta_i(\mathbf{x}^t)=d}$ and $\mathbb{1}_{\theta_i(\mathbf{x}^t)=Lg}$ are equal to 1 when firm i is a dirty firm or a low-level green firm, respectively, and 0 otherwise.

3.2 Green Disclosure Policies

When developing green disclosure policies in the ESG market, both strict and lax disclosure approaches are considered. The choice of disclosure method can significantly influence a firm's green transformation process.

Strict disclosure requires firms to fully disclose their progress in green transformation, enabling investors to differentiate accurately between high-green, low-green, and dirty firms. Under strict disclosure, λ_d, λ_{Lg}, and λ_{Hg} are clearly distinguished. The core advantage of this method is that it enhances transparency, enabling investors to allocate capital more effectively. As a result, it encourages all firms that have not yet reached advanced green levels to pursue greener transformations more actively.

On the other hand, lax disclosure takes a different approach by grouping "low green" and "high green" firms into a single category without distinguishing their specific levels of green performance. As a result, investors cannot fully assess firms' actual environmental performance in the short term. A key feature of lax disclosure is that it permits a degree of "greenwashing," where the green valuation (λ_g) of all green-stage firms is averaged between λ_{Lg} and λ_{Hg}:

$$\lambda_g = \frac{\lambda_{Hg}+\lambda_{Lg}}{2}.$$

The advantage of this approach is that it helps firms in the early stages of transformation, which face high costs, overcome initial bottlenecks. Artificially misvaluing these firms to "subsidize" them encourages greater adoption of green technology.

4 EMPIRICAL GAMES

In this ESG fund game, the strategy set $\mathcal{S}$ is defined as $\mathcal{S} = \{S_1, S_2\}$, where $S_1 = 0$ represents maintaining the current green level, and $S_2 = 1$ represents upgrading by one level, i.e., from dirty to low-green ($d \rightarrow$ Lg) or from low-green to high-green (Lg $\rightarrow$ Hg). This strategy set includes two choices: maintaining or upgrading. It applies across agent types under both disclosure forms. Under strict disclosure, dirty and low-green agents can upgrade, while high-green agents cannot due to reaching the highest level. Under lax disclosure, only dirty agents participate, as low-green and high-green agents are already at the highest level.

If each firm can independently decide whether to participate in the upgrade by selecting a strategy from the strategy space, there will be 2^{N_d} and $2^{N_d+N_{Lg}}$ different strategy configurations for the lax and strict disclosure versions of the game, respectively. Analyzing games of this size can be extremely challenging computationally. This complexity arises because, in asymmetric games, each player can independently choose their strategy, resulting in an enormous strategy space. To address this issue, we adopt an empirical game theory approach, which reduces the game's size by exploiting symmetries between firms. In the ESG fund game, firms are differentiated by their green valuations and upgrade costs, allowing us to construct distinct empirical games based on different disclosure policies.

4.1 Lax Disclosure Empirical Games

We begin by analyzing the ESG model using the lax form of the game at time step t. In this lax form, only dirty firms participate.

In this scenario, we can simplify the original game into a *Quasi-symmetric game*, where the payoffs for players are independent of the specific arrangement of other players. This transformation significantly reduces the game's complexity and size [24]. In other words, each player's payoff depends only on the number of players adopting each strategy. As a result, we can represent the payoff profiles using a strategy combination vector that tracks the number of players for each strategy. This approach reduces the number of strategy profiles to $\binom{50+|S|-1}{|S|-1} = 51$. Compared to the original game's vast number of possible strategy combinations (2^{50}), this represents a significant reduction in the game strategy space.

To store the payoffs of agents in quasi-symmetric games, we utilize a *heuristic payoff table* (HPT) [21], where the payoff for each strategy is recorded based on the number of players and their payoffs employing that strategy. While firms may have the same expected profits (since they are drawn from the same noise distribution), the actual profits of individual agents can vary in each simulation. This variability means that agents might receive different payoffs depending on the specific strategy combination in the game. As a result, we cannot directly rely on the HPT to compute the equilibrium since the game retains asymmetries [20].

To resolve this issue, we employ the Empirical Game-Theoretic Analysis (EGTA) method [23], which calculates the average payoff for players using each strategy within a given strategy profile and considers this as the general payoff for that strategy.

Specifically, $n^k = \{n_1^k, n_2^k, \dots, n_v^k\}$ is a vector representing the number of players, where n_j^k represents the number of players who choose strategy S_j. Here, v denotes the total number of strategies, and N_d represents the total number of players in the system. We impose the constraint that the sum of the players choosing different strategies $S_1, S_2, \dots, S_v$ must equal N_d, i.e., $\sum_{j=1}^{v} n_j^k = N_d$. This condition ensures that the total number of players is consistent.

The strategy profile $s^k = \{s_1^k, s_2^k, \dots, s_v^k\}$ describes the strategy combination chosen by all players in the game, where each $s_j^k \subseteq N_d$ represents the set of players choosing strategy S_j. The number of players selecting strategy S_j must satisfy $|s_j^k| = n_j^k$, ensuring consistency with the previously defined n_j^k. Additionally, the strategy profile satisfies two conditions: all players should select one strategy, and each can choose only one strategy. Thus, the set of all players N_d is decomposed into disjoint sets of strategies s_j^k that satisfy $\bigcup_{j=1}^{v} s_j^k = N_d$ and $s_i^k \cap s_j^k = \emptyset$ for $i \neq j$.

While the vector n^k of players is fixed for each simulation, the specific players assigned to each strategy are chosen randomly. Accordingly, we performed $M = 10,000$ simulations, each time computing the payoff $p_i^{(m)}(s^k)$ for strategy S_i, where $m \in \{1, 2, \dots, M\}$. The payoff $p_i^{(m)}(s^k)$ for strategy S_i is the average of the individual payoffs of all agents choosing that strategy, calculated as follows:

$$p_i^{(m)}(s^k) = \frac{1}{n_i^k} \sum_{a=1}^{n_i^k} u_a^{(m)},$$

where, as from above, n_i^k represents the number of agents who choose strategy S_i, and $u_a^{(m)}$ is the individual payoff of agent a who chooses strategy S_i in the m^{th} simulation. Through multiple simulations, we can compute the expected payoff of strategy S_i as:

$$\bar{p}_i(s^k) = \frac{1}{M} \sum_{m=1}^{M} \left(\frac{1}{n_i^k} \sum_{a=1}^{n_i^k} u_a^{(m)} \right).$$

Using this method, we average the individual payoffs of each agent choosing the strategy across multiple simulations, thus obtaining the overall payoff of strategy S_i under different strategy profiles s^k.

Therefore, we denote the heuristic payoff table (HPT) as $\mathcal{M} = (\mathcal{N}, \mathcal{U})$, where $\mathcal{N}$ is the strategy distribution matrix of dimension $\binom{10+|S|-1}{|S|-1} \times |S|$, and $\mathcal{U}$ is the corresponding payoff matrix. The entry n_i^k in $\mathcal{N}$ represents the number of players who choose strategy S_i in strategy profile s^k, where $i \in \{1, 2\}$. Meanwhile, the entry $\mathcal{U}_{i,k}$ in $\mathcal{U}$ represents the average payoff for a player choosing strategy S_i in profile s^k. The payoff $\mathcal{U}_{i,k}$ can be expressed as:

$$\mathcal{U}_{i,k} = \begin{cases} \bar{p}_i(s^k) & \text{if } n_i^k > 0, \\ 0 & \text{otherwise.} \end{cases}$$

4.2 Strict Disclosure Empirical Games

In contrast, the strict disclosure game involves two types of agents: dirty and low-green. Thus, we consider two distinct types of players, both subject to the same profit noise distribution but with different upgrade costs and green valuations. Specifically, this strictly disclosed game includes 50 dirty and 50 low-green agents.

To address the complexity of the original asymmetric game, we apply the role symmetric game analysis (RSA) method, which reduces the strategy space by partially transforming the game into a symmetric structure [24]. In this simplified role symmetric game, players are divided into two distinct *roles* based on their types: $\{d, Lg\}$. Here, N_d represents the 50 dirty-type players, and N_{Lg} represents the 50 low-green players. Within each role, the payoff for each strategy is determined by averaging the payoffs of all players who adopt that strategy in the same role.

We define the role r_i of player i as belonging to $\{r_d, r_{Lg}\}$, indicating whether player i is of the dirty type (r_d) or low-green (r_{Lg}). We extend the heuristic payoff table (HPT) as $\mathcal{M} = (\mathcal{N}^d \times \mathcal{N}^{Lg}, \mathcal{U}^d \times \mathcal{U}^{Lg})$. Here, $\mathcal{N}^d \times \mathcal{N}^{Lg}$ represents the strategy profile matrix with dimensions $\binom{50+|S|-1}{|S|-1}^2 \times 2|S|$, where $\mathcal{N}^d$ is the strategy distribution matrix for dirty agents and $\mathcal{N}^{Lg}$ is the strategy distribution matrix for low-green agents. Correspondingly, $\mathcal{U}^d \times \mathcal{U}^{Lg}$ is the payoff matrix of the same dimensions.

Similarly, for a strategy combination $n_r^k = (n_{r,1}^k, n_{r,2}^k, \dots, n_{r,k}^k)$, where $n_{r,j}^k$ represents the number of players of type r who choose strategy S_j in n_r^k, we estimate the payoffs of each strategy through M random simulations. Finally, after multiple simulations, we obtain the average payoff of agents of type r who choose strategy S_j in the strategy profile s_r^k:

$$\bar{p}_j^r(s_r^k) = \frac{1}{M} \sum_{m=1}^{M} \left(\frac{1}{n_{r,j}^k} \sum_{a=1}^{n_{r,j}^k} u_a^{r,(m)} \right).$$

The entry $n_{r,j}^k$ in $\mathcal{N}^r$, where $r \in \{d, Lg\}$, describes the number of players in role r who choose strategy S_j, $j \in \{1, 2\}$, in strategy profile s_r^k. Meanwhile, the entry $\mathcal{U}_{j,k}^r$ in $\mathcal{U}^r$, where $r \in \{d, Lg\}$, represents the average payoff of a player in role r who adopts strategy S_j in profile s_r^k. The value $\mathcal{U}_{k,j}^r$ is calculated as:

$$\mathcal{U}_{j,k}^r = \begin{cases} \bar{p}_j^r(s_r^k) & \text{if } n_{r,j}^k > 0, \\ 0 & \text{otherwise.} \end{cases}$$

5 EXPERIMENTAL SETTING

In the ESG fund model, since defining precise standards for corporate green ratings and upgrade costs is challenging, we opt to use a series of randomly generated model parameters for simulation. The upgrade cost coefficient α is selected from the set $\{0, 0.05, \cdots, 0.45, 0.5\}$, and the green valuation gap Δ (where $\Delta = \Delta\lambda_{\mathrm{I}} = \Delta\lambda_{\mathrm{II}}$) is chosen from $\{10, 30, 50, 70, 90\}$. At the start of the simulation, we generate an initial state of the ESG fund market with $N = 150$ firms, following the model introduced in [9] (with fixed parameters: $N_d = N_{Lg} = N_{Hg} = 50, \pi = 10, \epsilon \sim \mathcal{N}(0, 0.1^2), \lambda_d = 10, I = 1500, \beta = 0.9$).

Specifically, we generate random profits $\tilde{\pi}_0$, as well as the corresponding upgrade costs c, green valuations λ and market share q for all agents based on the selected parameters α and Δ. This process ultimately yields the initial state $\mathbf{v}_0$ of the ESG fund market.

We divide the game into two stages. In the first stage, for each generated initial ESG fund market state, we construct a heuristic payoff table based on the initial state, where the set of players is limited to the number of firms participating under different disclosure policies. For each strategy space s^k, we update the agents' market share in this strategy space and calculate the average payoff $\mathcal{U}_{i,k}$ for the players who select strategy S_i. Then, we apply the α-Rank algorithm [16] to compute the equilibrium of the green upgrade game under the chosen disclosure method and obtain the stationary distribution π_t of the game.

In the second stage, based on the equilibrium results from the first stage, we update the market parameters (such as the number of firm types N_d, N_{Lg}, N_{Hg}) and recalculate the market share and average payoff $\mathcal{U}_{i,k}$. Using the new heuristic payoff table, we again apply the α-Rank algorithm to calculate the game's stationary distribution π_{t+1} and output the final distribution of agent types. For each combination of parameters, we repeated the above experimental process 10,000 times using different initial market states to further reduce the impact of randomness.

To compare the impact of different disclosure methods on corporate strategies within the model, we designed three distinct disclosure approaches:

- **Strict disclosure model**: In both time steps, all firm types are clearly distinguished, offering the highest level of market transparency.
- **Lax disclosure model**: In both time steps, all green firms (low green and high green) are treated as homogeneous, reducing the model to two types of agents.
- **Mix disclosure model**: The lax disclosure form is used in the first stage, transitioning to the strict disclosure form in the second stage.

The entire process of the game is shown in the pseudo-code (see Algorithm 1).

5.1 α-Rank Algorithm

We consider N agents, each denoted by i, having access to a set of strategies of size k_i. We refer to the strategy set for agent i by $\mathcal{S}_i = \{s_{i,1}, \ldots, s_{i,k_i}\}$, $k_i = |\mathcal{S}_i|$, receives a payoff $M_i : \prod_{i=1}^{N} S_i \rightarrow \mathbb{R}$. A joint strategy profile is a set of policies for all participating agents in the joint strategy set, $\mathcal{S}_{Joint} = \prod_{i=1}^{N} S_i = \{s_{1,j_1}, \ldots, s_{N,j_N}\}$, with $s_{i,j_i} \in \mathcal{S}_i$ and $j_i \in \{1, \ldots, k_i\}$. Similarly, the joint payoff in the game is $M_{Joint} = \prod_{i=1}^{N} M_i$.

Each element in the transition probability matrix of a Markov chain represents the probability of an agent switching from one strategy to another, and this switching tendency is related to the reward obtained [16]. Consider any two joint strategy profiles $a = \{s_{i,a}, s_{-i}\}$ and $b = \{s_{i,b}, s_{-i}\}$ that differ in only one individual strategy for the i^{th} agent, where $s_{i,a} \neq s_{i,b}$. For $M_i\left(s_{i,a}, \boldsymbol{s}_{-i}\right) \neq M_i\left(s_{i,b}, \boldsymbol{s}_{-i}\right)$, we can calculate the probability, $\rho_{s_{i,a},s_{i,b}}\left(\boldsymbol{s}_{-i}\right)$, that one copy of agent i with strategy $s_{i,a}$ invades the population with all other agents (in that population m) playing $s_{i,b}$:

$$\rho_{s_{i,a},s_{i,b}}\left(\boldsymbol{s}_{-i}\right) = \frac{1 - e^{-\alpha\left(M_i\left(s_{i,a},\boldsymbol{s}_{-i}\right) - M_i\left(s_{i,b},\boldsymbol{s}_{-i}\right)\right)}}{1 - e^{-m\alpha\left(M_i\left(s_{i,a},\boldsymbol{s}_{-i}\right) - M_i\left(s_{i,b},\boldsymbol{s}_{-i}\right)\right)}}$$

and $1/m$ if $M_i\left(s_{i,a}, \boldsymbol{s}_{-i}\right) = M_i\left(s_{i,b}, \boldsymbol{s}_{-i}\right)$. α is the ranking strength, and $\alpha \geq 0$.

Further we can calculate the Markov probability transition matrix $P_{a,b}$:

$$P_{a,b} = \begin{cases} \frac{1}{\sum_{l=1}^{N}(|S_l|-1)} \rho_{s_{i,a},s_{i,b}}\left(\boldsymbol{s}_{-i}\right) & \text{if } a \neq b \\ 1 - \sum_{b \neq a} P_{a,b} & \text{if } a = b \\ 0 & \text{otherwise .} \end{cases}$$

The goal in α-Rank is to establish an ordering in policy profiles dependent on the evolutionary stability of each joint strategy. In other words, higher-ranked strategies are prevalent in populations with higher average survival time. Formally, such a notion can be easily derived as the limiting vector $\pi = \lim_{t\rightarrow\infty}\left[P^T\right]^t \boldsymbol{\pi}_0$ of Markov chain when evolving from an initial distribution $\boldsymbol{\pi}_0$.

Finally, we calculate the agent ranking corresponding to the smoothly distributed ordered quality. The quality of each agent's stationary distribution constitutes its "score."

Algorithm 1 Empirical Game procedure

Input: ESG fund market initial state $\mathbf{v}_0$, Disclosure form
Parameter: Strategy space $\mathcal{S}$
Output: N^{π}

```
for Step in {1,2} do
    for each entry n^k of N do
        for m ∈ {1, 2, ··· , M} do
            Randomly select a strategy profile s^k, such that |s_i^k| =
            n_i^k for all i ∈ {1, 2}, where n_i^k is the number of agents
            choosing strategy S_i
            Update the set of different types of agents N^k in the
            model according to s^k
            Update the market share of each agent using (1)
            Calculate the payoff of each agent in the set of agents
            s_i^k that selects strategy S_i using (2)
            Calculate the average payoff under different strategies
            to get p_i^(m)(s^k)
        end for
        Set U_{i,k} = Σ_{m=1}^{M} p_i^(m)(s^k)/M, ∀i ∈ {1, 2}
    end for
    Compute equilibrium π of the game induced by U using
    α-Rank
    Update agent type set N^π according to n^π for next step
end for
return N^π
```

6 RESULTS

In this section, we present the experimental results for various scenarios and evaluate the changes in the number of different agent types in the two-step disclosure model by analyzing the game equilibrium outcomes.

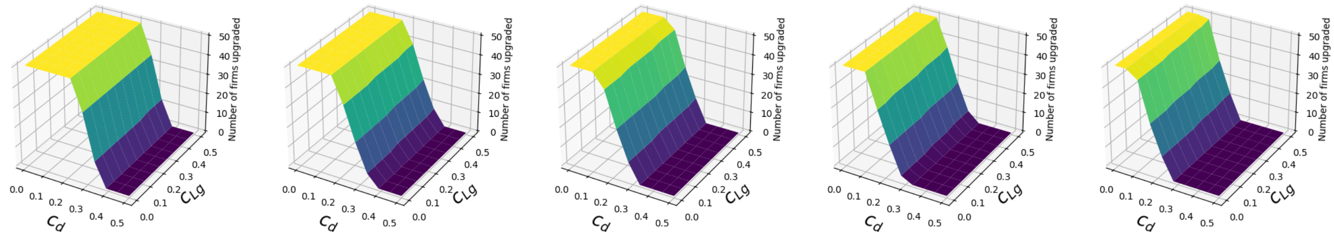

(a) $\Delta = 10$ **(b)** $\Delta = 30$ **(c)** $\Delta = 50$ **(d)** $\Delta = 70$ **(e)** $\Delta = 90$

Figure 1: Dirty agents' equilibrium upgrade choices under different c_d, c_{Lg} and Δ.

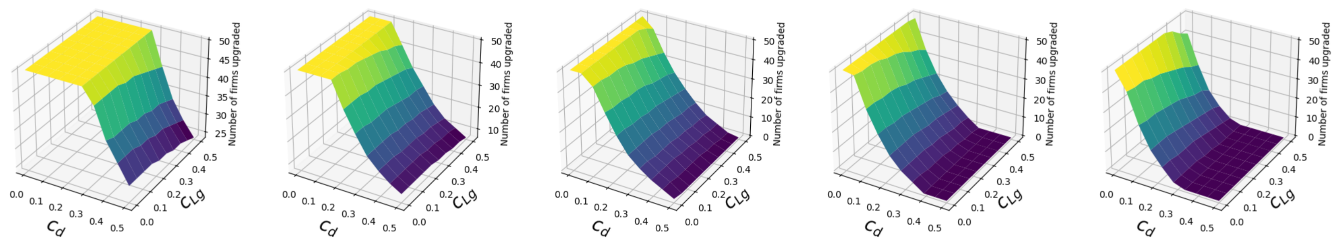

(a) $\Delta = 10$ **(b)** $\Delta = 30$ **(c)** $\Delta = 50$ **(d)** $\Delta = 70$ **(e)** $\Delta = 90$

Figure 2: Low-green agents' equilibrium upgrade choices under different c_d, c_{Lg} and Δ.

6.1 Equilibrium

6.1.1 Strict Disclosure Game. We analyze the number of agents that upgrade under different scenarios in Algorithm 1. Specifically, Figure 1 illustrates the number of dirty firms that opt for an upgrade in the strict disclosure game, considering different upgrade cost parameters c_d, c_{Lg} and green valuation parameters Δ.

From Figure 1, we observe that dirty firms' strategy choices are significantly influenced by the cost parameter c_d but almost unaffected by changes in c_{Lg}. As c_d increases, dirty agents shift from universally choosing the upgrade strategy to completely abandoning it. Moreover, as the green valuation difference Δ increases, dirty agents are less likely to upgrade, possibly because a sizeable green valuation gap reduces the market share of low-green agents, leading dirty firms to avoid upgrading in the final equilibrium.

On the other hand, Figure 2 shows the number of low-green firms that choose to upgrade under different scenarios. Similar to dirty firms, low-green agents are also susceptible to changes in c_d. However, as c_{Lg} increases, there is a slight decrease in the number of upgrades. The change in Δ shows the equilibrium trend between the two different types of agents. We found that low-green agents are more sensitive to changes in Δ than dirty agents. Specifically, as Δ increases, the strategy choices of low-green agents change more drastically than those of dirty agents. When $\Delta = 10$, low-green agents opt to upgrade in most scenarios, but when $\Delta = 90$, there is almost no situation where all low-green agents choose to upgrade.

In order to compare the impact of changes in Δ on the two types of agents in more detail, we selected the scenarios where $c_d = 0.3$ and $c_{Lg} = 0.3$ for further analysis. The specific results are shown in the following figures:

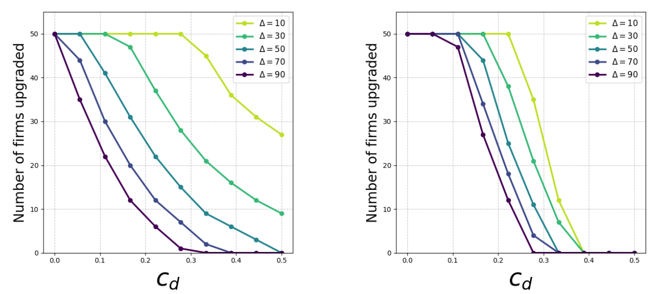

(a) Low-green agents. **(b) Dirty agents.**

Figure 3: Equilibrium upgrade strategy changes for different agents when $c_{Lg} = 0.3$.

From Figure 3, we can observe that when $c_{Lg} = 0.3$, as c_d increases, the equilibrium strategy trends for dirty and low-green agents differ significantly. First, we notice that the strategy of low-green agents experiences a wider range of changes, possibly because their payoff is not directly impacted by c_d, leading to a smoother adjustment in strategy. Additionally, it is evident that as Δ increases, for each level of upgrade cost c_d, low-green agents become progressively less inclined to upgrade. When $\Delta = 10$, even as c_d increases, some agents continue to opt for upgrading. However, when Δ increases to 90, the number of agents choosing to upgrade declines rapidly as c_d increases, eventually dropping to zero.

For dirty agents, Figure 3b shows that the number choosing the upgrade strategy drops to zero within a narrow cost range. Suppose we divide Δ into specific intervals. In that case, agents in the same interval begin to forgo upgrading at the same c_d, with the proportion of those upgrading gradually decreasing to zero at another constant c'_d. This indicates that the same Δ range primarily affects the rate of decline in the upgrade proportion of dirty agents, while c_d and c'_d remain fixed. A broader range of upgrade costs gives agents more flexibility in adapting to green transformations.

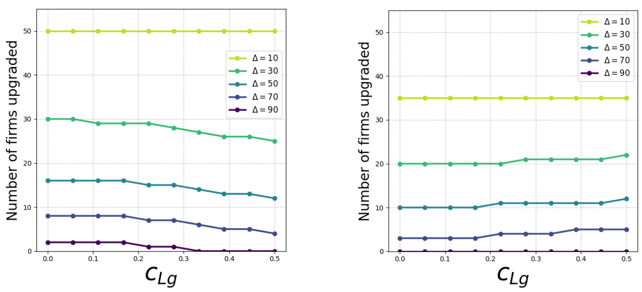

(a) Low-green agents. **(b) Dirty agents.**

Figure 4: Equilibrium upgrade strategy changes for different agents when $c_d = 0.3$.

Figures 4a and 4b illustrate the effect of c_{Lg} on the agent equilibrium strategy when $c_d = 0.3$. We observe that c_{Lg} has minimal impact on both dirty and low-green agents. As c_{Lg} increases, the proportion of low-green agents upgrading decreases slightly, whereas the proportion of dirty agents upgrading increases marginally. Additionally, across different values of Δ, the number of low-green agents upgrading remains higher than that of dirty agents, confirming the conclusion drawn from Figure 2.

6.1.2 Lax Disclosure Game. Next, we analyze the equilibrium under the lax disclosure game. Experiments show that dirty agents tolerate higher upgrade costs in this case. When $c_d = 0.5$, all dirty agents upgraded across all Δ scenarios, suggesting that conflating low-green and high-green valuations incentivizes dirty firm upgrades. To capture dynamic changes, we increased the upgrade costs, with results shown in the figure:

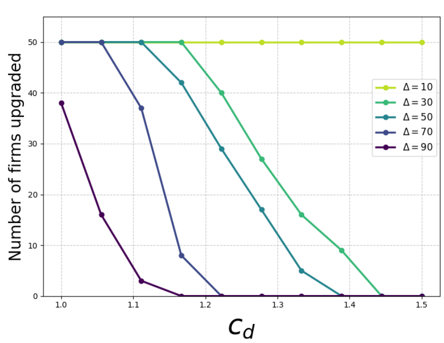

Figure 5: Equilibrium upgrade strategy changes for dirty agents under lax disclosure form.

In Figure 5, we can observe that when the range of c_d is between [1, 1.5], the strategy distribution of dirty agents changes significantly. As Δ increases, the number of agents choosing to upgrade decreases rapidly. When $\Delta = 90$, only a tiny part of agents opt to upgrade their green level. Overall, under the lax disclosure form, dirty agents are more willing to incur into higher upgrade costs, and their strategy distribution is more sensitive to changes in Δ.

6.2 Comparative Analysis of Disclosure Policies

To better understand the impact of different disclosure forms on various agent types in the ESG market, we conducted experiments using a two-stage ESG market model with different disclosure parameters. This allowed us to analyze how disclosure forms at each stage influence agent strategy selection. To compare the effects, we used parameter sets with varying strategy distributions across different disclosure forms for a more intuitive analysis.

For a more detailed analysis, we selected specific parameter sets for comparison across three scenarios: 1) $\{\Delta = 50, c_d = 0.3, c_{Lg} = 0.3\}$, 2) $\{\Delta = 90, c_d = 0.3, c_{Lg} = 0.3\}$, 3) $\{\Delta = 50, c_d = 0.4, c_{Lg} = 0.4\}$. Through these three experiments, we can comprehensively analyze the impact of changes in c_d, c_{Lg} and Δ on the distribution of agent upgrade strategies in a two-stage model. The equilibrium results of these experiments can assist regulators in selecting appropriate disclosure policies to more effectively promote the green transformation of enterprises.

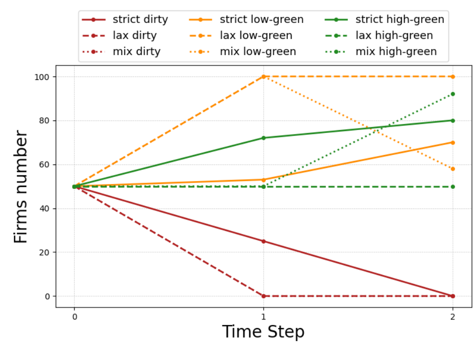

Figure 6: Equilibrium agent distribution in Scenario 1.

Figure 6 shows the experimental results of parameter set 1. In the continuous lax disclosure model, we found that although all dirty agents upgraded to low-green agents, since the disclosure model does not differentiate between low-green and high-green agents, the number of different agent types in the second stage remained unchanged, and no further green level upgrades were achieved. On the other hand, in the continuous strict disclosure model, we observed that as the process progresses, all dirty agents eventually complete the upgrade, and the number of low-green and high-green agents steadily increases over the two stages, eventually reaching 70 and 80, respectively.

Meanwhile, in the mixed disclosure model, we found that not only all dirty agents completed the upgrade in the first stage, but also some low-green agents who were still willing to upgrade to high-green agents in the second stage. In this model, the number of high-green agents reaches its highest point. These experimental results indicate that under this parameter set, all disclosure models can facilitate the upgrade of all dirty agents, but the mixed disclosure model yields the best final outcome.

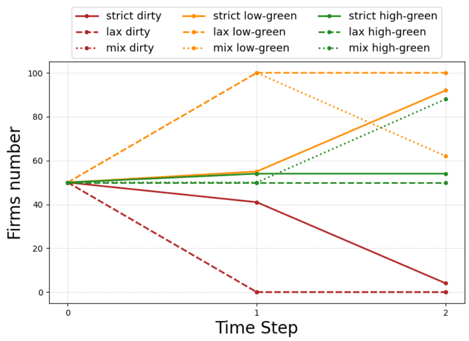

Figure 7: Equilibrium agent distribution in Scenario 2.

Increasing Δ to the value in parameter set 2 yields the results in Figure 7. We find that adjusting Δ does not affect the continuous lax disclosure model, as all dirty agents still upgrade first, with no changes in the second stage. In the strict disclosure model, the upgrade ratio for all agents significantly drops in the first stage. In the second stage, most dirty agents upgrade to low-green, but few low-green agents upgrade to high-green. This shows that increasing Δ lowers the initial upgrade ratio and reduces the likelihood of low-green agents upgrading.

On the other hand, in the mixed disclosure model, all dirty agents upgrade in the first stage, and low-green agents are strongly willing to upgrade in the second stage. However, compared to parameter set 1, the number of high-green agents decreases.

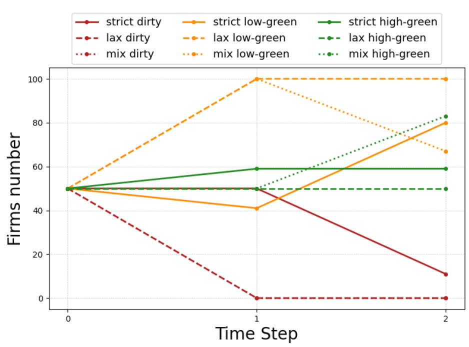

Figure 8: Equilibrium agent distribution in Scenario 3.

Finally, we increased the agents' upgrade costs for the experiment, and the results are shown in Figure 8. In the continuous strict disclosure model, we found that all dirty agents chose not to upgrade in the first stage, and only a small number of low-green agents opted to upgrade. In the second stage, some dirty agents upgraded to low-green agents, but almost no low-green agents advanced to high-green in this stage.

These experimental results suggest that when upgrade costs increase, only a tiny proportion of dirty agents choose to upgrade in the first stage, and only a part of the high-green agents are willing to upgrade in the second stage. Although the number of high-green agents decreased slightly in the mixed disclosure model compared to the previous experiments, this model still demonstrated the best overall outcome.

By comparing the above three sets of experiments, we found that the lax disclosure form encourages all dirty agents to upgrade. However, increases in green valuation gaps and upgrade costs lead to a reduction in upgrade efficiency in the ESG market. Dirty agents are more sensitive to upgrade costs, while low-green agents are more affected by changes in Δ. Finally, by comparing the three different disclosure models, we conclude that the mixed disclosure format helps market agents transition to higher green levels, enhances the sustainability of the ESG market, and provides a viable approach for governments and regulators when designing policies.

7 CONCLUSIONS

This paper presents an agent-based model in the ESG fund market to analyze agents' dynamics under different forms of green disclosure. Using the EGTA method, we simplify the payoff matrix of an asymmetric game and examine how different parameters and disclosure forms influence agents' strategies and upgrade efficiency at equilibrium.

Our findings show that under strict disclosure, dirty and low-green agents are more sensitive to the dirty firm's upgrade cost c_d, while changes in the low-green firm's cost c_{Lg} have little effect on equilibrium. The distribution of low-green firms' strategy varies across c_d, while dirty firms experience sharp shifts in strategy within a narrower c_d range. A larger green valuation gap leads both types of agents to adopt more conservative upgrade strategies.

We found that agents tolerate higher upgrade costs under the lax disclosure form than under the strict one. Simulations show that all agents upgrade within the original c_d range, but strategy distributions vary with different green valuations as costs rise. These findings help regulators determine the best disclosure form to improve allocation efficiency in the ESG market.

Finally, our experiments reveal differences in the number of agents across different disclosure scenarios. Most dirty agents successfully upgraded their green levels, indicating that green disclosure policies can effectively incentivize dirty firms to undertake their initial green upgrades. The lax and strict disclosure forms only performed well in the first and second stages, respectively. On the other hand, the mixed disclosure form demonstrated consistently strong performance across both stages.

Overall, our findings offer valuable insights for policymakers, especially ESG regulators, in setting future standards for green disclosure policies. Additionally, future research could explore how disclosure standards impact firms' financial performance and sustainability, industry-specific differences in green disclosure, and ways to balance transparency and authenticity to regulate "Greenwashing" and support corporate green transformation. Meanwhile, we will collaborate with academic and industry experts in green finance to validate and enhance our model's applicability. In dynamic policy scenarios, we plan to work with practitioners to collect and analyze data on green disclosure implementation and its long-term evolution, strengthening the model's real-world relevance.

REFERENCES

[1] Cyrus Aghamolla and Byeong-Je An. 2023. Mandatory vs. voluntary ESG disclosure, efficiency, and real effects. *Nanyang Business School Research Paper* 22-41 (2023).

[2] Wenni Anggita, Ari Agung Nugroho, et al. 2022. Carbon Emission Disclosure And Green Accounting Practices On The Firm Value. *Jurnal Akuntansi* 26, 3 (2022), 464–481.

[3] Herui Cui, Ruoyao Wang, and Haoran Wang. 2020. An evolutionary analysis of green finance sustainability based on multi-agent game. *Journal of Cleaner Production* 269 (2020), 121799.

[4] Blanca L Delgado-Márquez, Luis E Pedauga, and Eulogio Cordón-Pozo. 2017. Industries regulation and firm environmental disclosure: A stakeholders' perspective on the importance of legitimation and international activities. *Organization & Environment* 30, 2 (2017), 103–121.

[5] Benedikt Erforth. 2020. *The future of European development banking: What role and place for the European Investment Bank?* Number 11/2020. Discussion Paper.

[6] Lisa M Fairfax. 2022. Dynamic disclosure: An exposé on the mythical divide between voluntary and mandatory ESG disclosure. *Tex. L. Rev.* 101 (2022), 273.

[7] Bao Feng, Chunfeng Feng, and Songhui Zhao. 2023. Green Supply Chain Finance Credit Market under Government Regulation: An Evolutionary Game Theory Analysis. *Polish Journal of Environmental Studies* 32, 5 (2023), 3999–4010.

[8] BlackRock Investment Funds. 2024. BlackRock ESG Strategic Growth Fund Factsheet.

[9] Deeksha Gupta and Jan Starmans. 2024. Dynamic Green Disclosure Requirements. *Available at SSRN 4557187* (2024).

[10] MSCI Inc. 2024. MSCI Europe ESG Leaders Index Factsheet.

[11] Roman Inderst and Marcus M Opp. 2022. Socially Optimal Eligibility Criteria for ESG Funds. *Swedish House of Finance Research Paper Forthcoming* (2022).

[12] Habib Zaman Khan, Sudipta Bose, Benedict Sheehy, and Ali Quazi. 2021. Green banking disclosure, firm value and the moderating role of a contextual factor: Evidence from a distinctive regulatory setting. *Business Strategy and the Environment* 30, 8 (2021), 3651–3670.

[13] Blake LeBaron. 2006. Agent-based computational finance. *Handbook of computational economics* 2 (2006), 1187–1233.

[14] Thomas P Lyon and Jay P Shimshack. 2015. Environmental disclosure: Evidence from Newsweek's green companies rankings. *Business & Society* 54, 5 (2015), 632–675.

[15] Patrick R Martin and Donald V Moser. 2016. Managers' green investment disclosures and investors' reaction. *Journal of Accounting and Economics* 61, 1 (2016), 239–254.

[16] Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M Czarnecki, Marc Lanctot, Julien Perolat, and Remi Munos. 2019. α-rank: Multi-agent evaluation by evolution. *Scientific reports* 9, 1 (2019), 9937.

[17] L'uboš Pástor, Robert F Stambaugh, and Lucian A Taylor. 2021. Sustainable investing in equilibrium. *Journal of financial economics* 142, 2 (2021), 550–571.

[18] Egle Samanidou, Elmar Zschischang, Dietrich Stauffer, and Thomas Lux. 2007. Agent-based models of financial markets. *Reports on Progress in Physics* 70, 3 (2007), 409.

[19] Sebastian Steuer and Tobias H Tröger. 2022. The role of disclosure in green finance. *Journal of Financial Regulation* 8, 1 (2022), 1–50.

[20] Karl Tuyls, Julien Perolat, Marc Lanctot, Edward Hughes, Richard Everett, Joel Z Leibo, Csaba Szepesvári, and Thore Graepel. 2020. Bounds and dynamics for empirical game theoretic analysis. *Autonomous Agents and Multi-Agent Systems* 34 (2020), 1–30.

[21] William E Walsh, David C Parkes, and Rajarshi Das. 2004. Choosing samples to compute heuristic-strategy Nash equilibrium. In *Agent-Mediated Electronic Commerce V. Designing Mechanisms and Systems: AAMAS 2003 Workshop, AMEC 2003, Melbourne, Australia, July 15, 2003, Revised Selected Papers 5*. Springer, 109–123.

[22] Michael P Wellman. 2006. Methods for empirical game-theoretic analysis. In *AAAI*. 1552–1556.

[23] Michael P Wellman. 2016. Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems* 30 (2016), 1175–1189.

[24] Michael P Wellman, Karl Tuyls, and Amy Greenwald. 2024. Empirical game-theoretic analysis: A survey. *arXiv preprint arXiv:2403.04018* (2024).

[25] Yue Wu, Kaifu Zhang, and Jinhong Xie. 2020. Bad greenwashing, good greenwashing: Corporate social responsibility and information transparency. *Management Science* 66, 7 (2020), 3095–3112.

[26] Yi Yang, Fan Yang, and Xinxin Zhao. 2022. The impact of the quality of environmental information disclosure on financial performance: The moderating effect of internal and external stakeholders. *Environmental Science and Pollution Research* 29, 45 (2022), 68796–68814.