

MIT Open Access Articles

Prediction and prevention of disproportionately dominant agents in complex networks

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

As Published: 10.1073/PNAS.2003632117

Publisher: Proceedings of the National Academy of Sciences

Persistent URL: <https://hdl.handle.net/1721.1/134153>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Prediction and prevention of disproportionately dominant agents in complex networks

Sandro Claudio Lera^{a,b,c,1}, Alex Pentland^c, and Didier Sornette^{a,d,e}

^aInstitute of Risk Analysis, Prediction and Management (Risks-X), Academy of Interdisciplinary and Advanced Studies, Southern University of Science and Technology, Shenzhen 518055, China; ^bCollege of Business, Southern University of Science and Technology, Shenzhen 518055, China; ^cConnection Science, Massachusetts Institute of Technology, Cambridge, MA 02139; ^dDepartment of Management, Technology, and Economics, ETH Zurich, Zurich 8092, Switzerland; and ^eTokyo Tech World Research Hub Initiative, Tokyo 152-8550, Japan

Edited by Anthony Leggett, University of Illinois at Urbana–Champaign, Urbana, IL, and approved September 14, 2020 (received for review February 25, 2020)

We develop an early warning system and subsequent optimal intervention policy to avoid the formation of disproportional dominance (“winner takes all,” WTA) in growing complex networks. This is modeled as a system of interacting agents, whereby the rate at which an agent establishes connections to others is proportional to its already existing number of connections and its intrinsic fitness. We derive an exact four-dimensional phase diagram that separates the growing system into two regimes: one where the “fit get richer” and one where, eventually, the WTA. By calibrating the system’s parameters with maximum likelihood, its distance from the unfavorable WTA regime can be monitored in real time. This is demonstrated by applying the theory to the eToro social trading platform where users mimic each other’s trades. If the system state is within or close to the WTA regime, we show how to efficiently control the system back into a more stable state along a geodesic path in the space of fitness distributions. It turns out that the common measure of penalizing the most dominant agents does not solve sustainably the problem of drastic inequity. Instead, interventions that first create a critical mass of high-fitness individuals followed by pushing the relatively low-fitness individuals upward is the best way to avoid swelling inequity and escalating fragility.

complex networks | preferential attachment | winner takes all

In just about a decade, a handful of companies have contributed to an alarmingly centralized World Wide Web. This hardly resembles the Internet’s founding credo of net neutrality. Dominance on the Internet is not self-contained but transcends the economy as a whole, such as digital advertising and e-commerce. Beyond economic concerns, such unprecedented digital monopolies have important implications for politics and society (1). However, how can one tell that a company is too large? At what point should regulators interfere, and how? We present a framework to answer these and similar questions. At its heart is the law of preferential attachment, stating that the growth rate is proportional to size. We consider a system of interacting agents, whereby the rate at which any given agent establishes connections to others is proportional to its already existing number of connections and the agent’s intrinsic fitness in attracting connections. This representation is known to carry a risk of centralization (2), in the sense that removal of a few dominating agents would collapse the entire system. We derive a systematic classification of the system’s phase space and show under what circumstances such unfavorable centralization occurs. This allows us to anticipate the emergence of overly dominant agents *ex ante* and construct methods for early intervention. In socioeconomic systems, typical countermeasures against centralization are progressive taxes, antitrust laws, and similar legislation. However, our analysis reveals that such an approach may be ineffective because it addresses only the symptoms—disproportionally dominant agents—rather than the underlying cause: fundamentally imbalanced system that catalyzes such dominance. Instead of punishing the most competitive agents, one should foster

more balanced growth and competition by improving the relative fitness of underrepresented agents.

The applications of our framework are manifold, since fitness-adjusted preferential attachment has been shown to approximate the growth dynamics of many systems at a phenomenological level. Examples include the evolution of the World Wide Web (3), citation networks (4), trust relationships (5), social media (6), streaming services (7), photon emission rate (8), supply chains (9), *Escherichia coli* metabolic networks (5), market investments (10), and transactions on the bitcoin blockchain (11). A theoretical justification for fitness-based behavioral preferential attachment has recently been provided as the result of strategic minimization of maximum exposure to least fit nodes (12). In this article, we apply the method on the eToro trading platform, where participants observe others’ trades and can choose to either engage in their own strategies or copy others. We find that the platform does not give rise to “winners that take it all,” which we attribute to the inherent component of luck that underlies trading performance. We also discuss the implications of these insights for a universal base-income, progressive tax, antitrust laws, and social networks.

Model Definition and Regime Characterization

In this section, we propose an exact criterion specifying the conditions for which the dominance of an agent (firm, city, website, individual, etc.) becomes overwhelming.

Significance

Proportional growth is the driver behind the dynamics of a large class of complex networks. However, if left uncontrolled a few agents may become so dominant that their actions compromise the entire system. We present a framework that monitors the system’s distance from such imbalanced states. When the system approaches an imbalanced state, we show how to structure an optimal, cost-efficient intervention policy. Focusing only on either helping the least-fit agents or punishing the most dominant ones in isolation turns out to be inefficient. Instead, our results call for a more wholistic approach, with important implications for the structure of regulatory matters such as antitrust policies, taxation law, subsidies, or development aid.

Author contributions: S.C.L., A.P., and D.S. designed research; S.C.L. performed research; S.C.L. and D.S. contributed new reagents/analytic tools; S.C.L. analyzed data; and S.C.L., A.P., and D.S. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: leras@sustech.edu.cn.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2003632117/-DCSupplemental>.

First published October 16, 2020.

Fitness-Based Preferential Attachment and Deletion. Aiming at a generic description, preferential attachment is a natural first building block (13–15). Assuming some heterogeneity among agents, it is appropriate to add to it a fitness-based proportional attachment (16), which has been shown (12) to be particularly suitable to describe growth dynamics in complex networks. In practice, the appropriateness of this form may be tested by means of specific Bayesian statistical methods (6).

We thus consider an undirected network that is growing by attaching one agent and m edges per unit time. Edges may also be established between already existing nodes. Each agent a_i has an intrinsic “attractiveness” or, from hereon, “fitness” $\eta_i \in (0, 1)$ sampled from some fitnesses distribution ρ . Agent a_i ’s number of connections, that is, its degree, is denoted by k_i . At each time step, agent a_i establishes a new connection with probability p_i proportional to the product of its fitness and degree, $p_i = \eta_i k_i / \sum_j (\eta_j p_j)$. The fitter and the more connected the agent, the more attractive it is. To add one more realistic, yet generic, assumption (17), we allow agents to be removed from the system (failure, death, etc.) with probability proportional to $c\eta_i^\omega$. Here, $c \geq 0$ is the base rate of removal and η_i^ω accounts for the dependence of the probability of failure on the fitness of the agent. For $\omega > 0$, the fit agents are more likely to fail, for instance due to targeted attacks. For $\omega < 0$, weak agents are more likely to fail, that is, η itself is a measure of robustness with respect to failure.

Assuming that the i -th agent is still alive at time t , the mean field rate at which its degree k_i changes reads

$$\frac{\partial k_i}{\partial t} = m \frac{\eta_i k_i}{S(t)} - \frac{c}{1-c} \frac{\langle \eta \rangle^\omega}{\langle \eta^\omega \rangle} \frac{k_i}{t}, \quad [1]$$

where S is a normalization factor and expectation values $\langle \cdot \rangle$ are taken over the fitness distribution ρ . The first term in the right-hand side of Eq. 1 is a direct consequence of the attachment rule described above. The second factor is a product of two independent probabilities: the probability $\sim c \langle \eta \rangle^\omega$ that an agent is deleted and the probability $\sim k_i/t$ that agent a_i is connected to that failed agent. A detailed derivation of this equation and subsequent results is found in [SI Appendix](#).

While Eq. 1 captures the growth dynamics of a large class of systems at least to first order, it may also be extended to take into consideration, for instance, preferential deletion (17), sublinear or superlinear preferential attachment (18), assortative mixing (19), multidimensional fitness (20), or different boundary conditions such as minimal lower bounds or inflow of new agents at different rates (14).

The Fit-Get-Richer and the Winner-Takes-All Regimes. As shown in [SI Appendix](#), Eq. 1 gives rise to two distinct asymptotic regimes: the fit get richer (FGR) and the winner takes all (WTA). In the latter, the system is largely dependent and controlled by just a few agents. This regime is analogous to a Bose–Einstein condensate in statistical quantum mechanics (2, 21) and the dominant agents are known as “dragon-kings” in socioeconomic (22) and complex systems (23). The notion of disproportional dominance manifests itself in the number of connections of those agents (representing their influence/importance/wealth/...). This is made rigorous as follows: There exists at least one agent a_{DK} whose ratio of its degree k_{DK} to the total number of connections does not decay to zero in the limit of infinite system size. In other words, even as the system grows to infinity, a nonzero fraction of connections is controlled by a_{DK} . The influence of such an agent is felt in the entire system no matter what the total size of the network (Fig. 14). From a stability point of view, the system is then strongly reliant on a few agents and might collapse if those fail. From an economic point of view, any action of such an

agent affects the entire economy, reminiscent of the companies mentioned in the introduction.

The WTA regime occurs if and only if

$$I^* \equiv \int_0^1 d\eta \frac{\rho(\eta)}{\frac{\xi \omega + \langle \eta \rangle^\omega + \frac{1-c}{c} \langle \eta^\omega \rangle}{\eta} \left(\frac{\langle \eta \rangle^\eta + \frac{1-c}{c} \langle \eta^\omega \rangle}{\eta} + \frac{1}{\eta^{1-\omega}} \right) - 1} < 1, \quad [2]$$

where ξ is a constant that depends on ρ , c and ω , but is typically close to 1. Expectations $\langle \cdot \rangle$ are calculated with respect to the distribution ρ from which η is sampled. Importantly, Eq. 2 depends on the fitness landscape ρ as a whole, and not just on the fitness of a few agents. This suggests that acting on the dominant agents by removing many of their connections or targeting their growth specifically will not be effective: This may transiently decrease their influence but is short-lived as the growth dynamics leads to the resurgence of new agents that dominate. The problem is systemic: As we shall see, in order to prevent the occurrence of unwanted WTA agents, one should act on the fitness landscape ρ as a whole.

Finally, it is important to note that the FGR regime is not to be mistaken for a state of “equality.” The asymptotic size distribution is still heavy-tailed with a power law exponent controlled by the shape of ρ (16). However, the important difference is that the influence of any given agent is more localized and negligible within a large enough system, in contrast to the WTA regime where some agents dominate system-wide.

Parametric Regime Classification. Eq. 2 shows that the appearance of WTA agents depends on the deletion parameters c and ω , as well as the entire distribution of fitnesses ρ . In order to classify the possible different regimes, we propose to explore the space of fitness distributions parameterized by the beta distribution whose density is given by $f_{\alpha, \beta}(\eta) \propto \eta^{\alpha-1} (1-\eta)^{\beta-1}$ where varying $\alpha, \beta > 0$ samples all shapes of practical interest (Fig. 1B).

The WTA condition (Eq. 2) can now be expressed as a function of four parameters, $\{\alpha, \beta, \omega, c\}$, via the function $I^*(\alpha, \beta, \omega, c)$ and the two corresponding regimes are FGR ($I^* > 1$) and WTA ($I^* < 1$). In general, the integration in Eq. 2 with ρ has to be solved numerically. Fig. 1C shows the domain of existence for the WTA and FGR regimes, for different preferential deletion parameters c and ω in the space of the parameters (α, β) . Comparing with Fig. 1B, we see that the WTA regime corresponds to fitness landscapes in which the bulk of the agents have relatively low fitness. Smaller values of ω favor the failure of relatively low fitness agents, and hence increase the area of the WTA regime. In contrast, for large ω , the FGR regime increases in importance, as more fit agents tend to be removed by targeted attacks. The larger c is, the stronger is this effect.

Empirical Calibration on the eToro Social Trading Platform

We apply the above theory to the social trading dynamics on the multiasset brokerage platform eToro. On eToro, traders have access to different assets, in particular foreign exchange markets, and can leverage their position up to 400 times for both long and short positions. The feature that most distinguishes eToro from other trading platforms is its “OpenBook” social investment function. Instead of making their own trading decisions, users can choose to “mimic” one or several other traders and automatically execute the same trades as they do. Since more popular traders are more prominently displayed on the platform, they are more likely to be noticed and hence followed. This suggests that a preferential attachment model is suitable to model the imitation tendency. The decision to follow a specific trader also depends on idiosyncratic, heterogeneous properties (e.g., recent performance or risk level), and thus

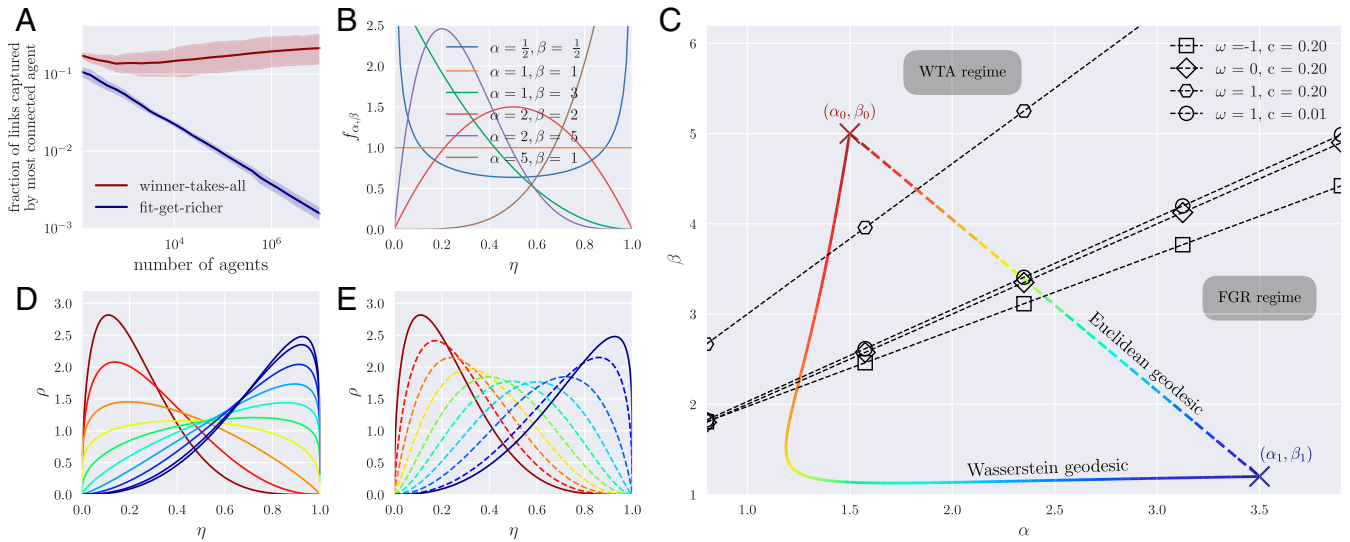


Fig. 1. (A) Fraction of links connected to the most dominant agent n_{DK} in the system, as a function of the total number of agents. The mean values and \pm one SD bands are obtained over a sample of 100 realizations. We have fixed $c = 1\%$ and $\omega = 0$. Two different fitness distributions are used, one in the WTA and one in the FGR regime (red and blue crosses in C). (B) Beta distribution density for different parameters α and β used to parametrize the class of fitness distributions on $[0, 1]$. As α and β vary, all typically relevant shapes of the fitness distribution are sampled. The WTA regime is characterized by the bulk of the probability mass being concentrated on agents of relatively low fitness, with only few fit agents that end up dominating the system (e.g., $\alpha = 2, \beta = 5$). (C) Phase diagram showing the domains of existence for the FGR and the WTA regimes, for different preferential deletion parameters c and ω in the space of the parameters (α, β) of the beta-fitness distribution. The WTA regime corresponds to values where $I^* < 1$ (Eq. 2). This WTA regime lies above a line parameterized by (c, ω) . (D) Geodesic with respect to Wasserstein-2 metric, that is, cost minimizing interpolation from a beta distribution in the WTA regime to a beta distribution in the FGR regime. The (approximate) path in the space of beta distributions is shown in C. (E) Distributions along the shortest path with respect to the Euclidean distance in the (α, β) plane (cf. dashed line in panel C). This intervention policy is more expensive than the one along the Wasserstein geodesic.

adding a fitness component seems appropriate. To confirm the applicability of fitness-adjusted preferential attachment, we have applied a Bayesian statistical method for the joint estimation of preferential attachment and node fitness without imposing functional constraints (6). See [SI Appendix](#) for details.

Translating to our model, each trader is an agent, and the number of traders following a given trader a_i defines its number of connections (degree) k_i . The rate at which a trader attracts new connections is a measure of the attachment probability p_i . In this way, by measuring agent a_i 's popularity k_i as well as its current rate of attraction p_i , one can infer an effective fitness $\eta_i \propto p_i/k_i$ without detailed knowledge of the underlying system dynamics. The access to all of the active accounts and traders allows us a faithful reconstruction of the underlying network. We measure p_i and k_i over rolling windows at different times t from July 2012 to November 2013, such that the evolution of the fitness can be tracked: $t \mapsto \eta_i(t)$. After having inferred the fitness of every agent at time t , we construct an empirical fitness distribution $\hat{\rho}(t)$ as a histogram over all fitnesses. The distribution of fitnesses turns out to be fairly stationary and unimodal, with most traders having a fitness around $\eta_i \approx 0.4$ (Fig. 2A). The histogram $\hat{\rho}(t)$ is fitted by the beta distribution via maximum likelihood to obtain the parameters $\hat{\alpha}(t)$ and $\hat{\beta}(t)$. In accordance with the stationarity of $\hat{\rho}$, the evolution of $t \mapsto (\hat{\alpha}(t), \hat{\beta}(t))$ is confined to a small radius around $\hat{\alpha} \approx 1.25$ and $\hat{\beta} \approx 1.25$ (Fig. 2B). Similarly, the deletion parameters $\hat{c}(t)$ and $\hat{\omega}(t)$ are inferred from direct observation at different times t and are found to be stationary and fluctuating within limiting bounds (Fig. 2C). See [SI Appendix](#) for details on methods. Here, we have defined node removal as the process whereby a trader is no more being mimicked by anyone after having been mimicked previously. Interestingly, $\hat{\omega}$ is consistently positive, suggesting that particularly fit traders are more likely to disappear. Our method is phenomenological in that it captures idiosyncratic, system-specific details into

the fitness term. To interpret this result, one must thus apply domain-specific knowledge. On eToro, a trader's recent performance is displayed to others by means of several metrics (percentage gain, portfolio volatility, etc.) We can then expect η to be a (generally complex) function of these variables. Performance on financial markets is known to have a large contribution of luck (24) and presumably even more so for the primarily retail traders that are active on eToro. In [SI Appendix](#), we show that the most mimicked (i.e., high η) traders indeed correspond to the ones who seem to outperform the market on short time scales, but that this performance is not sustained over longer times. Since, on short time-scales, luck is easily mistaken as skill (25), the mimicking functionality of eToro gives rise to imitation of noise traders. Implied overconfidence may then be the reason that more risk is being taken, resulting in large losses and explaining why particularly fit traders are more likely to vanish from the platform.

These insights highlight a secondary benefit of our approach: It complements and guides a more system specific understanding of the problem at hand. As long as the system allows for a representation of the form Eq. 1 (or a related variant thereof), one can absorb the intrinsic system dynamics into η . The phenomenological model then allows one to determine the state of the system. System-specific knowledge is again required to understand how the fitness distribution is to be modified in practice.

We now continue with the main purpose of the method: assessing the state of disproportional dominance. Given the estimates of α, β, c and ω , we can now calculate $t \mapsto I^*(\hat{\alpha}, \hat{\beta}, \hat{\omega}, \hat{c})$ from Eq. 2 and track its distance from the transition threshold $I^* = 1$. For the case of eToro, we find that the system dynamics is confined in the stable FGR regime $I^* > 1$ (Fig. 2d), suggesting that the competitive environment is enough to avoid the emergence of disproportionally dominant traders. This is alternatively visualized in the $(\hat{\alpha}, \hat{\beta})$ plane (Fig. 2B), where we can see that the

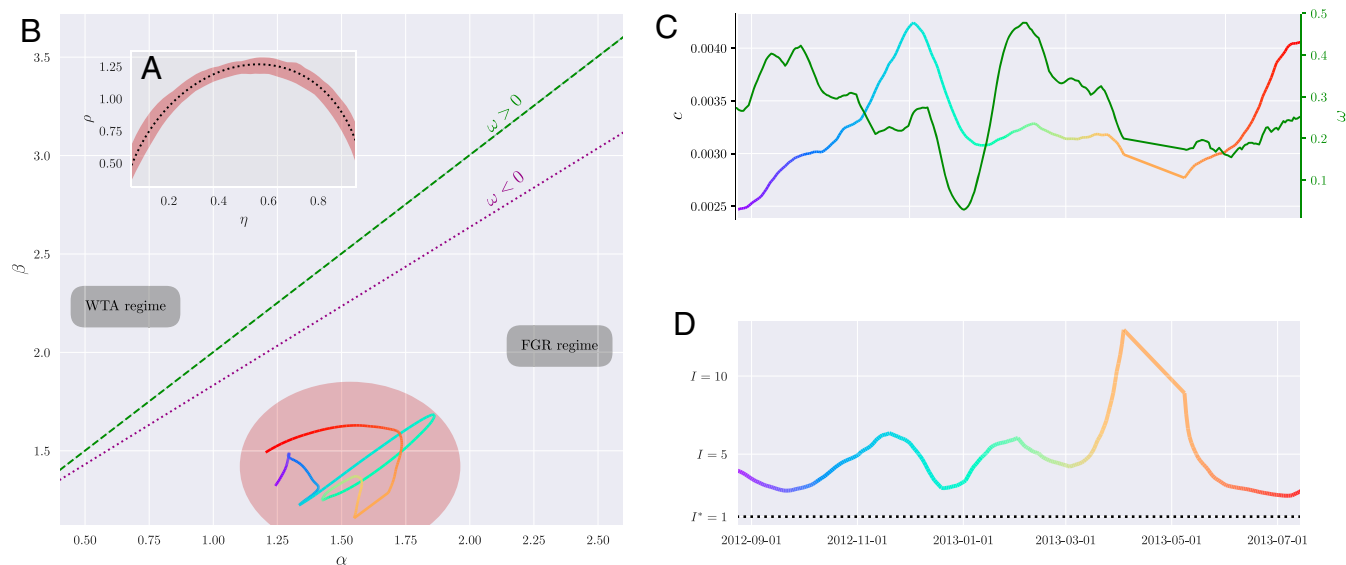


Fig. 2. Analysis of eToro trading network at different times t from July 2012 to November 2013. (A) Time-averaged distribution of trader fitnesses $\hat{\rho}$ (dashed black line) as well as 80% confidence interval as determined over all times t . (B) Evolution of the eToro network in the (α, β) plane obtained by the procedure explained in the text. The straight lines separate the WTA regime from the FGR regime. The green, dashed line corresponds to the regime separation with time-averaged $\bar{\omega} \approx 0.3$ and $\bar{c} \approx 0.003$ value. The purple, dotted line highlights the hypothetical effect of removing relatively low fitness traders at a high rate ($\omega = -0.3, c = 0.5$). (C) Time evolution of parameters c (multi colors on left axis) and ω (green and right axis). (D) Time evolution of the parameter I^* compared with the horizontal line at $I^* = 1$ that separates the two regimes (FGR for $I^* > 1$ and WTA for $I^* < 1$).

dynamics is well below the separation line to the WTA regime. If the dynamics was such that unfit traders are more likely to exit ($\omega < 0$) at higher rates ($1 > c \gg 0$), the system would however lie closer to the critical threshold at $I^* = 1$.

We note that the strength of the above methodology is that it can also be implemented to incomplete or noisy data, as the regime classification depends only on the empirical distribution of fitnesses $\hat{\rho}$ as a whole. A current limitation is that the fitnesses are assumed stationary, or at least that they change over timescales that are larger than for changes in the number of established connections. *SI Appendix* confirms this is the case for eToro. Relaxing this assumption is part of future research. Realizing that a large part of the traders' performance observed on eToro is likely driven by luck, it seems particularly interesting to take into consideration fitness values that decay with time (26).

Optimal Control of Disproportional Dominance

The analysis of the eToro social trading network in the previous section showed an instance where the network was confined in the FGR regime. In less-competitive, large-scale economic systems, this may not always be the case. In such situations, an external regulator (e.g., the state) may wish to intervene to ensure greater opportunities for new entrants and less inequity. Interventions only targeting the most dominant agents may appear “unfair,” as it punishes arguably highly fit agents. Perhaps more important, such intervention is inefficient in light of Eq. 2, which shows that the emergence of the WTA regime is a consequence of the distribution of fitnesses as a whole. Rather than just acting on the right tail of the distribution in form of punishments, more holistic, distribution-wide interventions are called for. In particular, as inferred from Fig. 1A and B, it is the high population of weak agents that needs to be addressed.

Relying again on our parameterization of the fitness space in terms of the beta distribution, we formulate the regulator's intervention as an action that aims to modify the distribution of fitnesses. Mathematically, this amounts to the problem of shifting optimally the distribution from an initial shape characterized

by (α_0, β_0) , for instance inside or uncomfortably close to the WTA regime, to another shape represented by (α_1, β_1) inside the FGR domain (Fig. 1C). Practically, this amounts to increasing the concentration of high-fitness agents relative to relatively low-fitness agents (e.g., in form of subsidies, start-up promotion, professional education, and so forth). Technically, we formulate the intervention technique as an optimal transport problem, seeking the most cost-efficient way of transforming one distribution into another. We assume that increasing an agent's fitness from η_1 to η_2 is associated with a cost $c(\eta_1, \eta_2)$. The problem then amounts to finding a transport map $\tau: [0, 1] \rightarrow [0, 1]$ that pushes $f_{\alpha_0, \beta_0}(\cdot)$ onto $f_{\alpha_1, \beta_1}(\cdot)$ while minimizing the total cost $\int_0^1 d\eta c(\eta, \tau(\eta)) f_{\alpha_0, \beta_0}(\eta)$. The solution of the problem depends on the functional shape of $c(\cdot, \cdot)$. Here, we provide a specific example by assuming a quadratic cost $c(\eta_1, \eta_2) = |\eta_1 - \eta_2|^2$. This is a reasonable assumption for systems where fitness can be controlled directly or indirectly via external manipulation. Whether a quadratic cost is appropriate generally depends on the specific system under scrutiny, but our method generalizes to other functional shapes for c . In full generality, different cost functions will yield different intervention protocols. However, as discussed in *SI Appendix*, our conclusions remain similar when extended to $c(\eta_1, \eta_2) = |\eta_1 - \eta_2|^\delta$ for any $\delta \geq 1$. By varying δ , a large class of cost functions may be approximated reasonably well.

Denote by $F_{\alpha, \beta}(\cdot)$ the cumulative distribution function of the beta distribution, and by $Q_{\alpha, \beta}(\cdot) \equiv F_{\alpha, \beta}^{-1}(\cdot)$ its inverse (quantile) function. It can be shown (27) that the cost-minimizing transport map τ is of the form $\tau(\eta) = Q_{\alpha_1, \beta_1}(F_{\alpha_0, \beta_0}(\eta))$. The interpolation from $f_{\alpha_0, \beta_0}(\cdot)$ to $f_{\alpha_1, \beta_1}(\cdot)$ is equivalent to a geodesic path with respect to the Wasserstein-2 metric. It is parameterized by $f_t(\eta) = \frac{\partial \tau_t}{\partial \eta} f_{\alpha_1, \beta_1}(\tau_t(\eta))$ with $\tau_t = t\eta + (1-t)\tau(\eta)$ and t is running from 0 to 1. Hence, $f_t(\eta)$ interpolates from $f_0(\eta) = f_{\alpha_0, \beta_0}(\eta)$ to $f_1(\eta) = f_{\alpha_1, \beta_1}(\eta)$ through the space of fitness distributions along a geodesic path with respect to the Wasserstein-2 metric (Fig. 1D). We can again approximate any intermediate distributions $f_t(\eta)$ as beta distribution with parameters (α_t, β_t) (see *SI Appendix* for mathematical details). This traces out a path

$t \mapsto (\alpha_t, \beta_t)$ that can be interpreted as the most cost-efficient intervention plan. In Fig. 1C, such a geodesic path is shown, starting from the red cross deep in the WTA domain and ending on the blue cross in the FGR domain. Fig. 1D presents a sequence of snapshots for the distribution of fitnesses, as it changes along the optimal path. This geodesic in the distribution space deviates very strongly from the naive solution of taking a straight path in the two-dimensional (α, β) plane (which would be the geodesic with Euclidean distances). As already mentioned, the WTA regime corresponds to a relatively low-fitness population with a few high fitness individuals, who eventually become winners that “take it all.” The most cost-efficient intervention (Wasserstein geodesic) focuses first on the high fitness part of the distribution, by decreasing β which describes the part of the fitness population close to 1 while keeping α approximately constant. Then, the intervention shifts to increasing α at quasi-fixed β , thus working on the part of the fitness population close to 0. In contrast, the straight line (Euclidean geodesic) works on both α and β simultaneously, which corresponds to first building a large pool of agents with intermediate fitnesses (Fig. 1E). The optimal intervention thus amounts to decreasing the number of agents of relatively low fitness while increasing the number of relative fit ones, while avoiding a population that peaks at intermediate fitness levels.

Implications and Applications

Since the growth dynamics of many systems may be captured at least to first order by fitness-based preferential attachment (3–11), the insights from this paper can be applied to a large class of systems. Let us discuss a few examples.

Already in the early 20th century, Gibrat had recognized that the growth dynamics of firms may be described with proportional growth (28). This model has ever since been extended numerous times to take into consideration effects such as firm heterogeneity (29–32), minimum firm size (33), births and deaths of firms (34–36), merger and acquisitions (37), or bankruptcy (38). In the context of our model, the firm dynamics can be described for instance via preferential attachment of the supply-chain network (9). The number of customers of a firm is denoted by k and η is the rate at which its customer base grows. Our results then have implications for instance for the formulation of antitrust laws. Regulators are concerned with the merger of two large firms. The impact of a merger is often assessed in terms of concentration of market power (39). Firm size may be one important aspect, especially in sectors with only a few firms. However, our results suggest that a potentially greater harm stems from the less-monitored acquisitions of successful smaller firms (such as start-ups) that gradually deplete the pool of high fitness agents.

At the level of personal wealth, the concept of fitness based preferential attachment is also relevant. The richer an agent, the more likely they are to attract more wealth by being exposed to more opportunities and transactions (10, 11, 40). Denote by k the wealth of an agent. Then, it holds $dk \propto \eta k$ where the fitness η is the growth rate in this context. From such a representation, the effect of a universal base income is to

put a floor on the minimum wealth k_i that any agent may have. It thereby acts as a reflecting lower boundary and this boundary condition can even make the wealth distribution more unequal, transforming it from log-normal to power law (33). However, to sustainably address the problem of income and wealth inequality, it is the growth rates (i.e., fitnesses) themselves that need to be addressed. While our model sheds light on what should be done and where resources should be targeted, system-specific domain knowledge is required to put these insights into best practice. A progressive tax system that takes from the rich and gives to the poor may free up the necessary resources but does not solve the problem sustainably by itself. While it is expected that the ability to generate income (i.e., fitness) changes itself with increased wealth, the uncontrolled, raw transfer of money may be inefficient, short-lived, or localized (41, 42). Instead, system-specific actions should be designed with the goal of increasing growth opportunities for low-income individuals, for example by enabling them to save money and earn returns, or better access to opportunities and investments.

Streaming services and online vendors often recommend new products to its users based on their popularity (degree k) and rating (fitness η) (7). To additionally capture a component of personalized recommendation and the bipartite nature of the product–customer space, Eq. 1 could be extended to take some users’ characteristics into consideration. In light of our results, to avoid that a few products completely dominate and to ensure the offering of a large variety of large fitness products, sufficient resources obtained from sales of successful products should be directed toward supporting creativity and inventions of new styles.

Concluding Remarks

The above results have shown that the emergence of disproportional dominance is not explained as the action of a single agent, but rather as a wholistic trait of the system. Our model highlights the importance of vigorous holistic interventions, which work on the full distribution of fitnesses, and focus on developing a high-fitness population. In particular, interventions should emphasize the support of relatively weak agents rather than punishing the most dominant ones. This may seem counterintuitive, but the sequential actions of first building a strong base of high-fitness individuals followed by pushing upward the fitnesses of the relatively low-fitness individuals is a more effective way of reducing inequity and increasing robustness.

Given the generic setup which may be adapted to different boundary conditions and its robustness to noise, our methodology has the potential of becoming a standard tool for dynamic risk monitoring in interacting systems ranging from small-scale trading platforms up to globally interconnected social and economic systems.

Data Availability. Data cannot be shared, but some parts of the data are available upon request.

ACKNOWLEDGMENTS. We are thankful for helpful feedbacks from S. Kolouri, M. Strub, M. Frank, S. Benzell, and two unknown referees.

1. M. Moore, D. Tambini, *Digital Dominance: The Power of Google, Amazon, Facebook, and Apple* (Oxford University Press, 2018).
2. G. Bianconi, A. L. Barabasi, Bose-Einstein condensation in complex networks. *Phys. Rev. Lett.* **86**, 5632–5635 (2001).
3. J. S. Kong, N. Sarshar, V. P. Roychowdhury, Experience versus talent shapes the structure of the web. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 13724–13729 (2008).
4. M. Golosovsky, S. Solomon, Growing complex network of citations of scientific papers: Modeling and measurements. *Phys. Rev. E* **95**, 012324 (2017).
5. F. Papadopoulos, M. Kitsak, M. Serrano, M. Boguná, D. Krioukov, Popularity versus similarity in growing networks. *Nature* **489**, 537–540 (2012).
6. T. Pham, P. Sheridan, H. Shimodaira, Joint estimation of preferential attachment and node fitness in growing complex networks. *Sci. Rep.* **6**, 32558 (2016).
7. S. Pongnumkul, K. Motohashi, A bipartite fitness model for online music streaming services. *Phys. Stat. Mech. Appl.* **490**, 1125–1137 (2018).
8. M. Golosovsky, Mechanisms of complex network growth: Synthesis of the preferential attachment and fitness models. *Phys. Rev. E* **97**, 062310 (2018).
9. E. Atalay, A. Hortacsu, J. Roberts, C. Syverson, Network structure of production. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 5199–5202 (2011).
10. D. Garlaschelli, S. Battiston, M. Castri, V. Servedio, G. Caldarelli, The scale-free topology of market investments. *Phys. Stat. Mech. Appl.* **350**, 491–499 (2005).
11. A. Aspembitova, L. Feng, V. Melnikov, L. Y. Chew, Fitness preferential attachment as a driving mechanism in bitcoin transaction network. *PLoS One* **14**, e0219346 (2019).
12. M. Bell et al., Network growth models: A behavioural basis for attachment proportional to fitness. *Sci. Rep.* **7**, 42431 (2017).

13. H. A. Simon, On a class of skew distribution functions. *Biometrika* **42**, 425–440 (1955).
14. A. Saichev, Y. Malevergne, D. Sornette, *Theory of Zipf's Law and Beyond* (Lecture Notes in Economics and Mathematical Systems, Springer, 2009), Vol. 632.
15. A. L. Barabási, R. Albert, Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
16. G. Bianconi, A. L. Barabasi, Competition and multiscaling in evolving networks. *Europhys. Lett.* **54**, 436–442 (2001).
17. J. S. Kong, V. P. Roychowdhury, Preferential survival in models of complex ad hoc networks. *Phys. A* **387**, 3335–3347 (2008).
18. P. L. Krapivsky, S. Redner, A statistical physics perspective on web growth. *Comput. Network.* **39**, 261–276 (2002).
19. M. E. Newman, Assortative mixing in networks. *Phys. Rev. Lett.* **89**, 208701 (2002).
20. J. H. Fowler, C. T. Dawes, N. A. Christakis, Model of genetic variation in human social networks. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 1720–1724 (2009).
21. V. Yukalov, D. Sornette, Statistical outliers and dragon-kings as Bose-condensed droplets. *Eur. Phys. J. Special Topics* **205**, 53–64 (2012).
22. D. Sornette, G. Ouillon, Dragon-kings: Mechanisms, statistical methods and empirical evidence. *Eur. Phys. J. Special Topics* **205**, 1–26 (2012).
23. D. Sornette, Dragon-kings, black swans and the prediction of crises. *Int. J. Terraspace Sci. Eng.* **2**, 1–18 (2009).
24. E. F. Fama, K. R. French, Luck versus skill in the cross-section of mutual fund returns. *J. Finance* **65**, 1915–1947 (2010).
25. D. Sornette, S. Wheatley, P. Cauwels, The fair reward problem: The illusion of success and how to solve it. *Adv. Complex Syst.* **22**, 1950005 (2019).
26. M. Medo, G. Cimini, S. Gualdi, Temporal effects in the growth of networks. *Phys. Rev. Lett.* **107**, 238701 (2011).
27. G. Peyré, M. Cuturi, Computational optimal transport. *Found. Trends Mach. Learn.* **11**, 355–607 (2019).
28. R. Gibrat, *Les Inégalités Économiques* (Librairie du Recueil Sirey, Paris, 1931).
29. E. Luttmmer, On the mechanics of firm growth. *Rev. Econ. Stud.* **78**, 1042–1068 (2011).
30. R. Axtell, M. Gallegati, A. Palestini, Common components in firms' growth and the sectors scaling puzzle. *Econ. Bull.* **12**, 1–8 (2008).
31. M. Wyart, J. P. Bouchaud, Statistical models for company growth. *Physica A* **326**, 241–255 (2003).
32. M. Takayasu, H. Watanabe, H. Takayasu, Generalised central limit theorems for growth rate distribution of complex systems. *J. Stat. Phys.* **155**, 47–71 (2014).
33. X. Gabaix, Zipf's law for cities: An explanation. *Q. J. Econ.* **114**, 739–767 (1999).
34. G. M. Grossman, E. Helpman, *Innovation and Growth in the Global Economy* (MIT Press, 1993).
35. J. Steindl, *Random Processes and the Growth of Firms: A Study of the Pareto Law* (Griffin, London, 1965).
36. E. Luttmmer, Selection, growth, and the size distribution of firms. *Q. J. Econ.* **122**, 1103–1144 (2007).
37. S. C. Lera, D. Sornette, Quantification of the evolution of firm size distributions due to mergers and acquisitions. *PLoS One* **12**, e0183627 (2017).
38. A. Saichev, Y. Malevergne, D. Sornette, *Theory of Zipf's Law and Beyond* (Springer, 2010).
39. M. Kamien, I. Zang, The limits of monopolization through acquisition. *Q. J. Econ.* **105**, 465–499 (1990).
40. P. E. Earl, T. Wakeley, Alternative perspectives on connections in economic systems. *J. Evol. Econ.* **20**, 163–183 (2010).
41. P. Kuhn, P. Kooreman, A. Soetevent, A. Kapteyn, The effects of lottery prizes on winners and their neighbors: Evidence from the Dutch postcode lottery. *Am. Econ. Rev.* **101**, 2226–2247 (2011).
42. P. Ager, L. P. Boustan, K. Eriksson, "The intergenerational effects of a large wealth shock: White southerners after the Civil War" (NBER Working Paper 25700, National Bureau of Economic Research, Cambridge, MA, 2019).