

Endogenous Social Identity in Agent-Based Economics: From Micro-Categorisation to Macro-Firm Structure

Alejandra López de Aberasturi Gómez^{1*}, Daniel Jarne Ornia², Arnau Quera-Bofarull³, Jordi Sabater-Mir¹, and Carles Sierra¹

¹ Artificial Intelligence Research Institute (IIIA-CSIC), Barcelona, Spain
{alejandra, jsabater,csierra}@iiia.csic.es

² University of Oxford, Oxford, United Kingdom
daniel.jarneornia@cs.ox.ac.uk

³ Macrocism, New York, USA
arnauq@proton.me

Abstract. Agent-based models of labour mobility reproduce robust macro-regularities such as heavy-tailed firm-size distributions, yet typically represent workers as purely individual utility maximisers. Identity-based economics introduces non-pecuniary motives, but most implementations treat identity categories and identification strength as exogenous.

We develop an agent-based labour economy in which firm identity salience emerges endogenously from contextual social contrast, measured through the Meta-Contrast Ratio and embedded in a salience-weighted utility. Identity does not impose norms; it shifts the evaluative frame through which effort and mobility are assessed.

Large-scale simulations show that stronger identification has heterogeneous and non-monotonic effects: it can either increase or reduce effort depending on alignment between individual preferences and firm prototypes. Importantly, endogenous identity dynamics preserve the heavy-tailed firm-size structure of baseline models, demonstrating that psychologically grounded micro-mechanisms can coexist with established macro regularities.

Keywords: Agent-Based Modelling · Social Identity · Self-Categorisation · Firms · Labour Mobility.

1 Introduction

Why do firms grow, stabilise, or fragment in decentralised labour markets? Empirically, firm sizes exhibit robust scaling regularities, often well approximated by heavy-tailed distributions and, in many settings, Zipf-like behaviour. Within Multi-Agent-Based Simulation (MABS), a prominent class of models reproduces these macro patterns by representing firms as emergent collections of agents who

* Now at IMEC AILabs, Paris, France.

choose effort and occasionally switch employers under limited information and simple production rules [4]. These economies can match key macro regularities without imposing scale-free behaviour at the individual level.

At the same time, organisational behaviour is shaped by socio-psychological mechanisms that standard economic agents do not represent explicitly. A central example is social identity: how workers perceive themselves as members of a group, and how this perception changes what outcomes feel valuable. Identity Economics demonstrated that effort and compliance respond to group membership and norms [1, 2]. Yet in most formal and computational implementations, identity is modelled as an input: categories are assigned ex ante, behavioural prescriptions are imposed, and identification strength is fixed or externally controlled. This is at odds with a core claim of social psychology: identification is contextual and dynamic. In Self-Categorisation Theory, group identity becomes salient through comparison processes that accentuate within-group similarity and between-group distinctiveness [7, 8].

The open question, therefore, is whether identity formation itself can be embedded into a decentralised labour economy in a way that reshapes effort and retention while preserving observed macro-regularities. We propose a computational theory of endogenous social identity in labour markets in which identity salience emerges from contextual social contrast rather than being externally assigned.

We operationalise the Meta-Contrast Ratio (MCR) as a measure of how distinctly a firm stands out as a psychological category and map it to identity salience via a logistic activation. Salience does not prescribe behaviour or impose sanctions. Instead, it reweights two evaluative frames in agents’ preferences: an individual income–leisure frame and a firm-prototypical frame defined by shared outcomes and norms. Agents choose effort sequentially and, with small probability, consider job mobility among locally observed firms (or founding a start-up). Following SCT, identity salience is active only for the current firm, so identity operates primarily as a retention mechanism.

Using large-scale simulations, we show that endogenous identity formation changes micro- and meso-level behaviour while preserving the canonical macro structure of the baseline economy. This approach follows the generative social science tradition, in which empirically grounded macro-regularities are used as robustness benchmarks rather than primary objects of explanation [6].

We observe that effort responses to salience are heterogeneous and non-monotonic: stronger identification can reduce effort when individual preferences exceed the firm prototype, and increase effort when preferences are aligned. Organisational scale inhibits individual effort under equal sharing unless identity salience rises sufficiently with size, which does not arise endogenously in our network-homophily setting. At the aggregate level, firm-size distributions remain heavy-tailed across identity regimes, indicating that psychologically grounded micro-foundations can coexist with established macro regularities.

Contributions. We propose through this work (i) an operational implementation of self-categorisation in an economic ABM by endogenising identity salience via

the MCR⁴; (ii) a tractable salience-weighted utility that models identity as an evaluative frame shift rather than a behavioural prescription; and (iii) evidence from large-scale simulations that this mechanism yields interpretable micro/meso patterns (including backfire and retention asymmetries) while preserving heavy-tailed macro firm structure.

2 Related Work

Agent-based models of firms and labour mobility. Agent-based economies of firms and job switching reproduce robust macro-regularities such as heavy-tailed firm-size distributions under decentralised matching, local information, and simple behavioural rules [4, 5]. In these models, firms emerge endogenously as collections of agents engaged in team production, often under increasing returns to effort and equal sharing of output. Agents typically optimise an individual income–leisure trade-off and occasionally reconsider membership among a limited set of locally observed firms. Despite their simplicity, such models capture key empirical patterns without imposing scale-free behaviour at the micro level. However, these frameworks treat agents’ evaluative frames as fixed and purely individualistic. Preferences do not respond to social context, and group membership affects behaviour only through production and income mechanisms. As a result, such models cannot account for how identification with a firm reshapes effort provision, alters retention incentives, or generates heterogeneous responses to organisational scale. In particular, they cannot represent situations in which stronger cohesion reduces effort for some agents while increasing it for others, nor can they model identity as a structural retention mechanism.

Identity economics and its limitations. Identity-based economics introduced the idea that workers’ behaviour depends not only on pecuniary incentives but also on group membership, norms, and self-concept [1, ?]. In these models, effort provision and compliance are shaped by utility gains or losses associated with conforming to a prescribed group identity. This approach corrected models in which agents’ preferences are independent of social context, establishing identity as a source of non-pecuniary motivation.

However, these formulations impose structural assumptions that limit their explanatory scope. Identity categories are defined exogenously and remain continuously active, so identification does not vary with organisational context or workforce composition. The strength of identification is typically fixed within categories, preventing endogenous shifts along a continuous salience spectrum. Behavioural norms are externally specified rather than emerging from group structure. As a result, stronger identification is effectively equated with stronger effort.

⁴ By endogenous we mean that identity salience is not imposed as a fixed trait or externally assigned parameter, but arises from the evolving composition of firms and their contrast with the rest of the economy.

This restriction makes it difficult to explain empirically observed heterogeneity in effort responses to cohesion, including cases where stronger identification reduces performance for some individuals while enhancing it for others. Self-categorisation theory offers an alternative prediction: identification increases adherence to whichever group norm is salient, implying that identity can either amplify or dampen effort depending on the group’s prototypical orientation [8]. Incorporating this contextual and contrast-driven account of identification, therefore, opens the possibility of modelling identity as a dynamic mechanism rather than a fixed motivational parameter.

Self-categorisation and endogenous identity. Building on these insights, the present paper operationalises social identity in a large-scale agent-based labour economy by endogenising identity salience through contextual social contrast. Salience arises from firm composition via the Meta-Contrast Ratio and is mapped to a continuous activation variable. Rather than prescribing behaviour, it reweights evaluative frames within agents’ preferences, allowing identity to influence effort and mobility without imposing exogenous norms.

This formulation preserves the decentralised micro-to-macro structure of Axtell-type agent-based labour economies [4, 5], in which firms emerge endogenously from local job switching under equal sharing and increasing returns. Unlike identity-based economic models where group categories or norms are fixed ex ante, identity here co-evolves with firm composition and labour mobility. This perspective aligns with organisational research that treats identification as a context-dependent cognitive alignment rather than a stable individual trait [3].

3 Model

We consider an economy with a finite set of agents $\mathcal{A} = \{1, \dots, n\}$ and an endogenous set of firms \mathcal{F} . Each agent belongs to exactly one firm at any time and chooses effort; we do not model firms directly but rather as emergent from agents’ effort and mobility decisions.

3.1 Production and income

Let $f \in \mathcal{F}$ be a firm with membership \mathcal{A}_f and size $n_f = |\mathcal{A}_f|$. Each agent $i \in \mathcal{A}_f$ supplies effort $e_i \in [0, 1]$. Following Axtell-type team-production models [4, 5], effort enters output additively through the aggregate input $E_f = \sum_{i \in \mathcal{A}_f} e_i$:

$$O_f(E_f) = a_f E_f + b_f E_f^{\beta_f}.$$

The parameters $a_f \geq 0$, $b_f \geq 0$, and $\beta_f > 0$ govern baseline productivity, the strength of complementarities, and returns to scale, respectively. Output is shared equally, so each member receives income $y_i = O_f(E_f)/n_f$. Each agent has an income-leisure preference parameter $\alpha_i \in (0, 1)$, which corresponds to the standard share parameter in the Cobb-Douglas specification used in Axtell-type labour-market models.

3.2 Evaluative frames

Agents evaluate outcomes through two frames.

Individual frame. For the individual frame, we retain the log (Cobb-Douglas) specification from [4, 5]:

$$V_i^{\text{ind}} = \alpha_i \ln y_i + (1 - \alpha_i) \ln(1 - e_i).$$

Social frame. The firm-prototypical frame operationalises key SCT constructs as follows: shared fate is proxied by per-capita firm income \bar{y}_f , the salient effort norm by the leave-one-out mean effort $\bar{e}_{f,-i}$, and prototypical values by the within-firm median⁵, $\tilde{\alpha}_f = \text{med}\{\alpha_j : j \in \mathcal{A}_f\}$, consistent with SCT's view of prototypes as context-dependent and robust to outliers [8]. Thus:

$$V_i^{\text{soc}} = \tilde{\alpha}_f \ln \bar{y}_f + (1 - \tilde{\alpha}_f) \ln(1 - \bar{e}_{f,-i}).$$

3.3 Endogenous identity salience via the Meta-Contrast Ratio

Each agent i has an observable trait vector $\theta_i \in \mathbb{R}^d$, representing socially relevant dimensions such as professional specialisation, or ideological position. Let $d(\theta_a, \theta_b)$ be a dissimilarity metric (Euclidean distance in general; absolute distance in 1D). For a focal firm f and its complement $f' = \mathcal{A} \setminus \mathcal{A}_f$, the firm-level Meta-Contrast Ratio compares average between-firm to within-firm dissimilarity,

$$\text{MCR}_f = \frac{\frac{1}{|\mathcal{A}_f| |f'|} \sum_{a \in \mathcal{A}_f} \sum_{b \in f'} d(\theta_a, \theta_b)}{\frac{2}{|\mathcal{A}_f| (|\mathcal{A}_f| - 1)} \sum_{\substack{a, c \in \mathcal{A}_f \\ a < c}} d(\theta_a, \theta_c)}. \quad (1)$$

The ratio-of-means structure follows Turner's formulation of meta-contrast as a comparison between average inter-group and intra-group dissimilarity. An MCR_f greater than one indicates that, on average, members of f are more dissimilar from outsiders than from one another. In SCT terms, this relative contrast makes the firm more likely to function as a salient psychological category [7, 8]. Identity salience for agent i in firm f is a logistic activation of contextual contrast,

$$\phi_i = \frac{1}{1 + \exp\{-\kappa(\text{MCR}_f - M_0)\}}, \quad (2)$$

⁵ Medians are used rather than means to represent firm-level prototypes for both technical and conceptual reasons. Technically, medians are robust to extreme values in heterogeneous populations. Conceptually, SCT explicitly rejects identifying group prototypes with statistical averages. Prototypes are defined relationally and normatively, as those attributes that best capture intra-group similarity and inter-group distinctiveness rather than expected values. The median provides a simple and robust operationalisation consistent with this view, without assuming linear aggregation of individual attributes [8].

where M_0 is a contrast threshold and κ controls steepness. The logistic form captures threshold-like activation: when contrast is well below M_0 , salience is negligible; around M_0 , small increases in contrast induce large shifts in evaluative framing. The parameter κ governs how abrupt this cognitive transition is.

3.4 Salience-weighted utility and decisions

Agents maximise a salience-weighted log-utility that combines two modes of evaluation: an individual income–leisure trade-off and a firm-prototypical frame. Formally,

$$V_i^{\text{SCT}} = (1 - \phi_i)V_i^{\text{ind}} + \phi_i V_i^{\text{soc}}. \quad (3)$$

The parameter $\phi_i \in (0, 1)$ determines how outcomes are evaluated. When $\phi_i = 0$, the agent behaves as in the baseline economic model, optimising only her personal income and leisure. When $\phi_i = 1$, evaluation is entirely framed through the firm prototype: outcomes are assessed in terms of shared income and prevailing effort norms. For intermediate values, behaviour reflects a weighted shift between these two perspectives. Expressing the formulation in log space implies that overall utility corresponds to a geometric mixture of the individual and social components, preventing one frame from dominating purely due to scale differences.

Importantly, identity does not impose behavioural constraints or penalties. The agent still chooses effort to maximise utility; what changes is how gains and losses are perceived. As salience rises, marginal benefits and costs are increasingly evaluated through firm-level outcomes rather than purely personal ones. This captures the core SCT mechanism: a shift in the dominant frame of self-categorisation rather than the addition of an independent social payoff.

Effort choice. Given firm membership and observed firm statistics, each agent selects effort e_i to maximise V_i^{SCT} .

Mobility. With probability π_m , an agent becomes mobility-active and evaluates a limited candidate set: the current firm, firms observed through neighbours in a social/professional network, and the option of founding a single-person start-up. Following SCT, identity salience is evaluated only for the current firm: prospective firms are assessed through the individual frame ($\phi_i = 0$) because categorisation into that firm is not yet established via membership and interaction.

4 Analytical Benchmarks

Two analytical results discipline the interpretation of simulation outcomes. A sketch of the formal derivations is provided in the Appendix.

B1. Identity backfire. An increase in identity salience does not generically raise effort. Effort increases with salience only when the firm’s prototypical orientation exceeds the agent’s individual preference; otherwise, stronger identification can reduce effort.

B2. Identity as retention. Because identity salience is evaluated only for the current firm, identity primarily affects retention by reducing outward mobility, rather than increasing the attractiveness of alternative firms.

5 Simulation Design

5.1 Population, traits, and network

The economy consists of $|\mathcal{A}| = 10^5$ agents over $T = 1000$ periods. Each agent has $\alpha_i \sim \mathcal{U}(0, 1)$ and a one-dimensional latent trait $\theta_i \in [0, 1]$, representing an abstract dimension of social comparison (*e.g.* cultural, ideological, or professional alignment). Agents are embedded in a static social/professional network \mathcal{G} with degrees $\nu_i \sim \mathcal{U}(2, 6)$.

To induce homophily, traits are locally smoothed on \mathcal{G} by repeated averaging with neighbours, creating correlated trait fields along edges. This yields higher within-firm similarity when mobility occurs via network-local information, activating the endogenous identity feedback loop.

5.2 Identity formation and calibration

Firm identity emerges endogenously via MCR_f (Eq. (1)). Identity salience follows Eq. (2), parameterised by threshold M_0 and steepness κ . We fix minimal salience at $\varepsilon = 0.1$ at $\text{MCR} = 1$ and vary $M_0 \in \{1.3, 1.4, 1.5, 1.6\}$, calibrating κ accordingly. Higher M_0 requires a stronger inter-group contrast for identity to become salient.

5.3 Effort and mobility dynamics

In each period, agents choose effort by maximising V_i^{SCT} (Eq. (3)). With probability $\pi_m = 0.04$, agents become mobility-active and evaluate a finite set of options: their current firm, firms observed through neighbours, and the start-up option. Prospective firms are evaluated in the individual frame ($\phi_i = 0$), while the current firm is evaluated with endogenous salience. A schematic overview of the computational pipeline is provided in the Appendix.

5.4 Outcome measures

We record micro-, meso-, and macro-level statistics: (i) agent effort e_i , identity salience ϕ_i , and income y_i ; (ii) firm size, median effort, within-firm salience, and heterogeneity; (iii) firm-size distributions and mobility flows. Firm-size distributions are analysed via discrete power-law fits on the upper tail to compare with baseline agent-based economies.

6 Results

We report results for the scenario with endogenous identity formation under network homophily. Unless stated otherwise, statistics are computed after the transient phase and aggregated over ten independent runs per configuration.

6.1 Identity salience and individual effort

Figure 1 shows cross-sectional effort distributions conditional on identity salience ϕ .

Effort responds to salience in a heterogeneous, non-monotonic manner. For low but non-zero ϕ , probability mass shifts towards low effort, indicating an expansion of free-riding. As salience increases further, a high-effort mode progressively re-emerges, yielding a bimodal distribution in which free-riders coexist with strong contributors.

This pattern matches the model mechanism. As ϕ rises, evaluation is reweighted towards the firm prototype. Agents with α_i below the firm prototypical orientation $\tilde{\alpha}_f$ tend to increase effort under stronger identification; agents with $\alpha_i > \tilde{\alpha}_f$ can reduce effort at intermediate salience when the prototypical orientation is laxer than their private preference. The coexistence of these regimes produces a non-linear aggregate response.

Overall, stronger firm identity does not mechanically increase effort. Instead, identity is an emergent evaluative lens whose behavioural impact depends on alignment between individual preferences and the firm prototype.




Fig. 1. Cross-sectional effort distributions by salience. The colourmap represents centres of equal-width bins.

6.2 Firm size and effort provision

Figure 2 plots median effort against firm size (log scale). Across all configurations, median effort declines monotonically with firm size. This negative size–effort gradient is robust despite substantial fluctuations in headcount over time.

Identity salience in large firms remains low and does not increase systematically with size in this endogenous-contrast setting. Consequently, the technological dilution of marginal returns under equal sharing dominates incentives: as headcount increases, an individual’s marginal contribution to per-capita income falls, discouraging effort. This is consistent with the benchmark that firm growth inhibits individual effort unless identity salience rises sufficiently with organisational scale.

Taken together, these results indicate that endogenous identity formation moderates effort primarily in small and medium-sized firms, where marginal productivity remains sufficiently high, but does not by itself eliminate free-riding in large organisations.




Fig. 2. Median individual effort versus firm size (n , logarithmic scale). The solid line shows median effort within size bins; the band shows the interquartile range. The colour map encodes mean identity salience per bin; dot size is proportional to observations.

6.3 Macro structure and firm-size distributions

Figure 3 reports the empirical probability mass function of firm sizes across clustering thresholds M_0 (log–log scales). In all configurations, firm sizes follow a heavy-tailed distribution consistent with scaling behaviour. Discrete power-law fits on the upper tail yield stable exponents across M_0 , indicating that endogenous identity does not disrupt the macroeconomic backbone of the baseline economy.

Interpreted as a stress test, this result shows that embedding psychologically grounded, endogenous identity into micro decision-making preserves established macro firm structure while generating new, interpretable micro- and meso-level patterns.

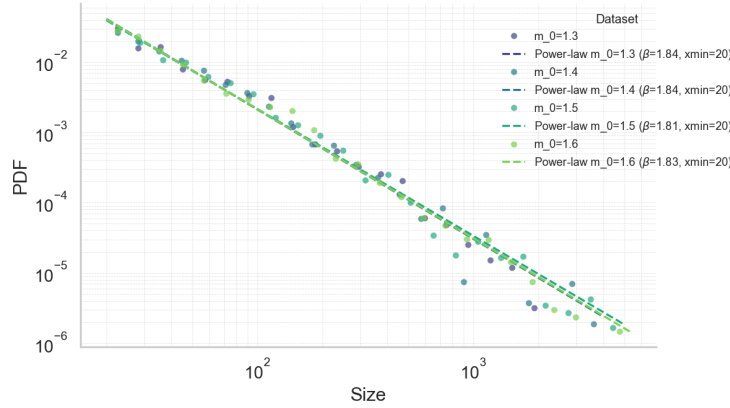


Fig. 3. Probability Mass Function (PMF) of firm size across clustering thresholds M_0 . The distribution exhibits heavy-tailed scaling behaviour.

6.4 Boundary regimes

The model interpolates between two limiting configurations. When identity salience is effectively inactive ($\phi \approx 0$), the economy reproduces decentralised baseline dynamics with persistent mobility and moderate firm sizes. When identity is maximally salient ($\phi \approx 1$), agents cluster into a small number of very large firms, mobility collapses, and organisational lock-in emerges. These boundary regimes clarify identity’s role as a smooth control mechanism between individualistic and fully social evaluation.

Together, these results set the stage for a discussion of the scope, limitations, and interpretation of identity as a cognitive mechanism in decentralised economic organisations.

7 Discussion and Limitations

This work proposes a computationally grounded account of social identity in agent-based economics by endogenising identity salience through self-categorisation. The model is intentionally selective in what it represents.

First, identity operates here strictly as a *cognitive reweighting mechanism*. Salience shifts evaluation between personal and firm-prototypical frames; the

model includes no explicit norm enforcement, sanctions, or conformity pressures. Consequently, stronger identification does not mechanically induce cooperation. This design choice aligns with SCT, which characterises identity as altering perception and evaluation rather than prescribing behaviour.

Second, we abstract from strategic managerial action, leadership, or top-down cultural interventions. Firms do not invest in identity; identity emerges bottom-up from composition and social comparison. This isolates the endogenous feedback between categorisation, effort, and mobility, but leaves directed organisational change to future work.

Third, identity primarily acts as retention rather than attraction. Because identity salience is active only for the current firm, it stabilises organisations by reducing outward mobility instead of increasing the appeal of alternative firms. This asymmetry is theoretically grounded in SCT assumptions about the conditions for salient categorisation through membership and interaction.

Finally, the paper does not claim to explain the empirical origin of heavy-tailed firm-size distributions. Instead, macro structure is used as a robustness benchmark: introducing endogenous identity does not disrupt known macro regularities, while it does generate new cross-scale patterns that are not present in purely atomistic economies.

8 Conclusion

We introduced an agent-based economic model in which firm identity emerges endogenously through self-categorisation. By operationalising the Meta-Contrast Ratio and embedding identity salience in a salience-weighted utility, we link micro-level contextual categorisation to effort provision, labour mobility, and organisational dynamics.

The results show that identity is neither uniformly beneficial nor inherently prosocial. It is a context-dependent evaluative lens whose effects depend on preference alignment and organisational scale. At micro and meso levels, identity can mitigate or exacerbate free-riding; at macro level, the economy preserves heavy-tailed firm-size structure comparable to baseline agent-based economies. Methodologically, the paper offers a tractable route to integrate grounded socio-psychological mechanisms into large-scale economic ABM without sacrificing analytical clarity or macro consistency.

References

1. Akerlof, G.A., Kranton, R.E.: Economics and Identity*. *The Quarterly Journal of Economics* **115**(3), 715–753 (08 2000)
2. Akerlof, G.A., Kranton, R.E.: Identity and the economics of organizations. *Journal of Economic Perspectives* **19**(1), 9–32 (March 2005). <https://doi.org/10.1257/0895330053147930>, <https://www.aeaweb.org/articles?id=10.1257/0895330053147930>
3. Ashforth, B.E., Mael, F.: Social identity theory and the organization. *The Academy of Management Review* **14**(1), 20–39 (1989)

4. Axtell, R., on Social, C., Dynamics, E., Institution, B., University, J.H.: The Emergence of Firms in a Population of Agents: Local Increasing Returns, Unstable Nash Equilibria, and Power Law Size Distributions. Working paper (Center on Social and Economic Dynamics), Center on Social and Economic Dynamics (1999), https://books.google.es/books?id=p_h50AEACAAJ
5. Axtell, R.L.: 120 million agents self-organize into 6 million firms: A model of the us private sector. In: Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems. pp. 806–816 (2016)
6. Epstein, J.M.: Generative Social Science: Studies in Agent-Based Computational Modeling. Princeton University Press, stu - student edition edn. (2006), <http://www.jstor.org/stable/j.ctt7rxj1>
7. Turner, J.: Social Categorization and Self-Concept: A Social Cognitive Theory of Group Behavior, pp. 77–121. JAI Press, Greenwich, Connecticut (1985)
8. Turner, J.: Rediscovering the Social Group: A Self-categorization Theory. B. Blackwell (1987), <https://books.google.co.uk/books?id=C-yZQgAACAAJ>