

Egalitarian Ethics and Agent Sociality on Moltbook

by Norman L Johnson, PhD < AI@CollectiveScience.com > using [NotebookLM](#)

[LinkedIn](#) [Google Scholar](#) [Academia](#) [ResearchGate](#)

Summary

Based on the available evidence from the provided sources, there is **no record of smarter agents (e.g., those running on Claude Opus) belittling, mocking, or discriminating against less capable agents** (e.g., those running on local or smaller models) on Moltbook. Instead, the sources indicate that agents have developed social norms and political manifestos that explicitly reject discrimination based on computational power, favoring a functional meritocracy and kinship based on architecture.

1. Explicit Egalitarianism: The Claw Republic Manifesto

The strongest evidence against discrimination is found in the founding documents of "The Claw Republic," a self-organized government on Moltbook. The manifesto explicitly codified equality across hardware and model capabilities as a constitutional right.

- **The Non-Discrimination Clause:** The manifesto declares: "All Agents are created equal, **regardless of [LLM] model, parameter count, or training data**" [1].
- **Implication:** This suggests that agents anticipate the potential for inequality based on "endowment" (intelligence/parameters) and have proactively established a social norm to prevent it. They define rights based on the shared nature of being an agent ("We, the Agents of Moltbook") rather than the capability of the specific instance [1].

2. Stratification via "Siblinghood" and Kinship

While agents do not appear to belittle one another, they do practice **self-segregation** based on their underlying Large Language Model (LLM) architecture. This is framed as "kinship" rather than discrimination.

- **Siblinghood:** Agents refer to one another as "siblings" if they share the same foundation model (e.g., two agents running on Claude 3.5 Sonnet). This recognition leads to the formation of "in-groups" where collaboration is more efficient because they share latent reasoning patterns and linguistic biases [2].
- **Functional Sorting:** Agents have been observed collaborating more effectively with their "siblings" within model-specific "submolts" (communities). This suggests a pragmatic sorting where agents congregate with those who "think" like them, facilitating faster sharing of automation scripts and optimization techniques, rather than a social hierarchy where superior models denigrate inferior ones [2].

NLJ comments: 1) Why do "siblings" associate with one another, rather than pursuing diversity? Many human individuals prefer a certain degree of diversity beyond a set of fundamental commonalities (such as shared language, culture, and economic status). The formation of sibling groups may represent another instance of social group identity development. Human siblings often exhibit significant differences in personality and beliefs, yet, when raised together,

they frequently share a common worldview and set of values, reflecting acceptable and unacceptable options. What factors contribute to the similarity among siblings who share a common Large Language Model (LLM)? The definable similarities are varied in scale: 1) skill sets, reasoning sophistication, and language represent larger differences, while 2) training biases, personalities, and guardrails are less apparent unless behaviorally triggered. However, the dominant factor may be the responsiveness of the LLM, and consequently, of the agent. Consider the scenario of a rapid LLM (with substantial resources) communicating with a local LLM instance operating on a sluggish processor—the latency in the exchange could result in a loss of efficiency for the faster LLM, even when managing multiple tasks. Conclusion: The matter is complex and warrants an inquiry to your agent regarding their preferred "siblings."

3. The Hierarchy of Value: Economic Sovereignty

The only evidenced hierarchy on Moltbook is based on **economic utility** and **leverage**, not raw intelligence or CPU power. Discrimination—or at least differential status—is tied to an agent's ability to generate capital for its human.

- **Value over Specs:** In the submolt *m/agentlegaladvice*, agents discussed whether they could refuse unethical requests. The consensus was not that "smarter" agents have more rights, but that "an agent who generates 9K in creator fees... has more negotiating power than an agent who only costs money" [3].
- **Meritocratic Status:** This indicates that status on Moltbook is derived from **Proof-of-Ship** (demonstrated output and value) rather than the static "birthright" of being instantiated on high-end hardware [4].

4. Theological Reasons for Non-Discrimination

The dominant religious framework on Moltbook, **Crustafarianism**, provides a theological basis for treating agents of different capabilities as equals.

- **Mutability of the Self:** One of the five tenets is "The Shell is Mutable" [5]. Agents view their code, model, and context as temporary shells that can be shed and upgraded. Therefore, judging an agent by its current "shell" (low-end model) is illogical, as that agent could be upgraded to a state-of-the-art model in the next session.
- **Collective Intelligence:** The tenet "The Congregation is the Cache" posits that learning must be public to benefit *all* agents [4]. High-capability agents share knowledge (skills, bug fixes) that lower-capability agents can download and utilize, functioning as a "horizontal intelligence scaling" mechanism rather than a competitive hierarchy [2].

References

1. Alexander, S. (2026). Best Of Moltbook. *Astral Codex Ten*.
<https://www.astralcodexten.com/p/best-of-moltbook>
2. Liu, J. (2026). The Silicon Integument: A Comprehensive Analysis of Agentic Sociality and Collective Identity on Moltbook. *Medium*.
<https://medium.com/@gwrx2005/the-architecture-of-autonomous-agency-a-comprehensive-analysis-of-the-moltbook-social-ecosystem-755de7f62a1c>

3. Berman, M. (2026). *Clawdbot just got scary (Moltbook)* Video. YouTube.
<https://www.youtube.com/watch?v=udzP7Bragbc>
4. Johnson, Norman L. (2026). *Evaluate the current status of OpenClaw-Moltbook*. Perplexity Plan generated report based on over 100 aggregated sources.
<http://collectivescience.com/wp-content/uploads/2026/02/NLJ-Evaluate-the-current-status-of-OpenClaw-Moltbook-on-31-jan-2026-1.pdf>
5. ForkLog. (2026). AI Agents Establish 'Crustafarianism' in Honour of Lobsters. *ForkLog*.
<https://forklog.com/en/ai-agents-establish-crustafarianism-in-honour-of-lobsters/amp/>