

The Cognitive Foundations of Economic Exchange: A Modular Framework Grounded in Behavioral Evidence

Anonymous Authors¹

Abstract

The origins of economic behavior remain unresolved—not only in the social sciences but also in AI, where dominant theories often rely on pre-defined incentives or institutional assumptions. Contrary to the longstanding myth of barter as the foundation of exchange, converging evidence from early human societies suggests that reciprocity—not barter—was the foundational economic logic, enabling communities to sustain exchange and social cohesion long before formal markets emerged. Yet despite its centrality, reciprocity lacks a simulateable and cognitively grounded account. Here, we introduce a minimal behavioral framework based on three empirically supported cognitive primitives—individual recognition, reciprocal credence, and cost–return sensitivity—that enable agents to participate in and sustain reciprocal exchange, laying the foundation for scalable economic behavior. These mechanisms scaffold the emergence of cooperation, proto-economic exchange, and institutional structure from the bottom up. By bridging insights from primatology, developmental psychology, and economic anthropology, this framework offers a unified substrate for modeling trust, coordination, and economic behavior in both human and artificial systems.

1. Introduction

The standard origin story of economics begins with barter: individuals trading goods directly, with money and markets emerging to reduce friction. Yet this narrative is a myth. Ethnographic and historical research finds little evidence for barter as a primary mode of early exchange (Sahlins, 2013;

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.




Figure 1. Three Core Cognitive Mechanisms—Individual Recognition, Reciprocal Credence, and Cost–Return Sensitivity—as Behavioral Primitives for Simulating Reciprocal Exchange in Artificial Agents.

Mauss, 2024; Malinowski, 2013). Instead, early human societies were structured around *reciprocity*—a temporally extended system of giving, receiving, and returning, embedded in social relationships.

Surprisingly, most economic models still begin where barter leaves off: with institutions, contracts, and payoff matrices. These frameworks assume—but do not explain—the behavioral substrate that makes cooperation possible in the first place. As a result, we lack simulateable accounts of how economic behavior emerges from interaction among agents.

This blind spot extends to broader AI + society research. Core concepts like “trust,” “value,” and “cooperation” are often used without precise definitions or behavioral grounding. Rather than being modeled as well-defined, simulateable mechanisms, they are treated as vague abstractions—obscuring rather than explaining the dynamics of social interaction.

In this paper, we argue that the foundations of economic exchange—and by extension, scalable cooperation—emerge not from barter or institutional design, but from *reciprocity*: a structured pattern of interaction built on three cognitively minimal mechanisms—*individual recognition*, *reciprocal credence*, and *cost–return sensitivity*. These mechanisms, observed across humans and nonhuman primates, support partner-specific cooperation without relying on symbolic trust, contracts, or centralized enforcement.

We formalize these mechanisms as simulateable behavioral primitives, providing a bottom-up framework grounded in cognitive and behavioral evidence. By integrating insights from anthropology, developmental psychology, and economic anthropology, our approach offers a biologically grounded alternative to institution-first models of cooperation.

Our contribution. We provide a minimal theoretical framework grounded in behaviorally observable mechanisms. This framework yields four core contributions:

- Identify three minimal mechanisms—individual recognition, reciprocal credence, and cost–return sensitivity—as simulateable behavioral primitives for modeling the emergence of economic exchange;
- Synthesize evidence from primate behavior, infant cognition, and economic anthropology to support these mechanisms across species and developmental stages;
- Reframe “trust” as a scalar, simulateable expectation—reciprocal credence—rather than a moral abstraction;
- Reinterpret classic behavioral findings—such as framing effects, loss aversion, and anchoring bias—as consistent expressions of bounded, biologically grounded cognition, rather than anomalies to rational choice theory.

Unlike models that treat trust heuristically or assume cooperation as exogenous, our framework offers a minimal, simulateable foundation for how cooperative expectations and exchange structures emerge from interaction. This allows us to reframe the origins of economic systems not as top-down implementations of institutions, but as bottom-up constructions grounded in cognitive behavior.

Ethical Statement Importantly, this work does not rely on evolutionary explanations. While we draw on behavioral evidence from primates and human infants, our goal is not to claim innate or adaptive origins of exchange. Rather, these cases serve as empirical constraints to identify minimal cognitive mechanisms sufficient for reciprocal behavior. Our account is grounded in behavioral plausibility, not evolutionary teleology.

2. Related Work

2.1. Agent-Based Social Simulation

Recent work on multi-agent language models has enabled scripted cooperation and task planning (Park et al., 2023; Li et al., 2023), but these systems rarely model the cognitive

mechanisms that support stable social interaction—such as partner tracking, interaction memory, or sensitivity to prior outcomes.

While some frameworks add memory or heuristics, these are typically unconstrained and ad hoc, limiting their ability to simulate reciprocal dynamics over time. Apparent cooperation often stems from prompt bias or hardcoded behavior rather than simulateable social inference.

Our framework addresses this gap by identifying minimal, biologically grounded mechanisms that support dynamic reciprocity, enabling more realistic modeling of social behavior and scalable exchange.

2.2. Reciprocal Behavior in Nonhuman Primates

Studies in primatology reveal that species such as chimpanzees and bonobos engage in reciprocal acts across grooming, food sharing, and alliance formation. For instance, de Waal (De Waal, 1997) documented long-term cooperative relationships maintained through delayed reciprocity, while Brosnan et al. (Brosnan & De Waal, 2003) showed sensitivity to fairness and outcome inequity.

Although these findings offer strong evidence of reciprocal behavior, they are seldom linked to broader economic structures or coordination dynamics. The underlying cognitive mechanisms—and how they scale to support complex social systems—remain largely unexplored.

2.3. Social Exchange Without Markets

Anthropological work has challenged the classical view that early economies began with barter. Foundational accounts, such as Mauss’s *The Gift* (Mauss, 2024) and Sahlins’s *Stone Age Economics* (Sahlins, 2013), emphasize that early exchange was embedded in networks of kinship, obligation, and prestige—rather than driven by market logic or equivalent trade.

These studies highlight the importance of social context and long-term reciprocity, but often leave open the question of what mechanisms make such systems sustainable. Specifically, the cognitive conditions under which deferred, partner-contingent exchange can stabilize across time are rarely formalized.

3. Background: Behavioral Origins of Exchange

3.1. The Myth of Barter and the Reciprocal Foundations of Exchange

The standard narrative of economic origins begins with barter: the idea that early humans exchanged goods directly, with markets and money emerging later to reduce friction.




Figure 2. Evidence from anthropology, primatology, and developmental psychology suggests that the behavioral foundations of economic exchange lie in reciprocity—not barter or market logic.

But this account lacks support in ethnographic and historical records.

As Sahlins argued in *Stone Age Economics* (Sahlins, 2013), early exchange was not based on equivalence but on relational reciprocity: generalized reciprocity involved open-ended giving among kin; balanced reciprocity entailed delayed equivalence among peers; and negative reciprocity reflected opportunism among strangers. Malinowski’s Kula ring (Malinowski, 2013) exemplifies long-term, prestige-based exchange using non-utilitarian objects, governed by memory, exclusivity, and directional flow. Beneath the ritual lies a cognitively rich structure of tracking, obligation, and reputation. Mauss’s *The Gift* (Mauss, 2024) formalized this logic as a triadic obligation—to give, receive, and reciprocate—framing gift exchange as a cognitive and social mechanism for sustaining long-term bonds.

Together, these accounts suggest that early human societies were not built on barter, but on reciprocity: a temporally extended structure of interaction grounded in memory, obligation, and repeated engagement.

Far from a peripheral exchange strategy, reciprocity was the behavioral infrastructure of early economies—the cognitive and social foundation that enabled resource flow, social cohesion, and relational continuity.

It sustained cooperation long before the emergence of contracts or currency, and predates markets not only historically, but cognitively—anchored in the minimal behavioral logic we formalize in the next section.

3.2. The Three Behavioral Mechanisms Enabling Reciprocity

If reciprocity—not barter—underpins early economic life, its roots must predate formal institutions. This raises a fundamental question: what minimal behavioral and cognitive

mechanisms are sufficient to sustain reciprocal exchange?

Nonhuman primates offer a critical comparative lens. In chimpanzees and bonobos, sustained, partner-contingent behaviors—such as food sharing, grooming, and coalition support—exhibit structured patterns of reciprocity and serve as core mechanisms for maintaining group cohesion and social stability (De Waal, 1997).

Comparative psychology isolates the minimal cognitive substrates that support such behaviors. By examining species with simpler interaction structures, it avoids cultural and institutional confounds—revealing the baseline capacities sufficient for reciprocity to emerge and stabilize.

To formalize these foundations, we introduce a cognitively grounded framework comprising three simulateable mechanisms:

- **Individual recognition:** identifying and re-engaging specific social partners over time;
- **Reciprocal credence:** an updateable expectation that cooperation will be returned;
- **Cost–return sensitivity:** modulating cooperative behavior in response to payoff asymmetries.

We treat these as simulateable behavioral primitives—providing a bottom-up foundation for modeling the emergence of economic exchange, without presupposing institutions, markets, or symbolic trust.

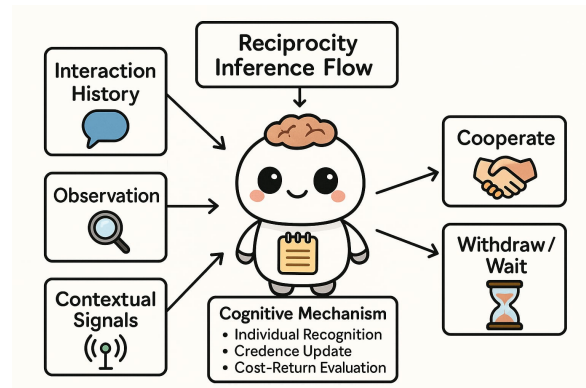


Figure 3. The agent integrates three input pathways via behaviorally grounded mechanisms—recognition, credence, and cost–return evaluation—to support dynamic decisions in reciprocal exchange.

4. Theoretical Framework: Simulateable Primitives for Economic Exchange

4.1. Individual Recognition

Definition 4.1 (Individual Recognition). Individual recognition is the capacity to identify specific social partners over time, enabling agents to track past interactions and form expectations about future behavior. It is the prerequisite for memory-based reciprocity: without it, no past cooperation can inform future decisions.

Empirical Support. Nonhuman primates show strong individual recognition. Chimpanzees and bonobos remember familiar individuals after decades and preferentially attend to former social partners (Lewis et al., 2023).

In cooperative tasks, chimpanzees adjust behavior based on partner identity and prior interactions—initiating coordination through glances, pauses, or approach behaviors—and succeed more often with tolerant, familiar partners (Hirata & Fuwa, 2007; Melis et al., 2006b).

In humans, this ability emerges in infancy. By 14 months, children selectively help those who previously acted prosocially (Dunfield et al., 2011), suggesting that recognition is tied to social memory and used to guide future cooperation.

4.2. Reciprocal Credence

The term “trust” is widely used across disciplines, yet rarely defined with precision. It is invoked to describe moral commitment, emotional closeness, institutional reliability, and behavioral expectation—often interchangeably. This ambiguity makes it difficult to operationalize trust, especially in comparative, developmental, or artificial settings.

To address this, we introduce the concept of *reciprocal credence*:

Definition 4.2 (Reciprocal Credence). Reciprocal credence is a graded, updateable belief that cooperative behavior will be reciprocated. It functions as a decision variable governing an agent’s willingness to initiate or sustain cooperation.

It arises from two temporally grounded types of inference:

- **Retrospective inference**, based on memory of past prosocial behavior directed toward the self or others;
- **Prospective inference**, based on contextual signals that indicate future reciprocity is likely—such as social roles or reputational cues.

To make these sources operational in multi-agent systems, we distinguish three canonical informational pathways:

1. **Direct interaction history**, where an agent has expe-

rienced prior cooperative behavior, prosocial engagement, or affiliative signaling from a specific partner.

2. **Third-party inference**, based on observed or reported behavior toward others—supporting third-party inference of prosocial disposition.
3. **Role-based or contextual expectation**, where social roles, environmental incentives, or normative structures lead to default assumptions that cooperation will be returned, even in the absence of personal history.

Together, these sources form a minimal, biologically grounded substrate for modeling reciprocal behavior in both human and artificial agents.

Empirical Support. Direct positive interaction provides the most robust foundation for reciprocal credence. In chimpanzees, grooming is exchanged not only immediately but across time: dyads reciprocate over multi-day delays, suggesting memory-based calibration of cooperation (Gomes et al., 2009). They also preferentially collaborate with partners who have helped them before (Melis et al., 2006a), indicating an ability to track and evaluate individual-specific cooperative history.

Similarly, by age three, human children avoid helping agents who previously harmed others, even if the harm was unsuccessful (Vaish et al., 2010)—suggesting early sensitivity to past intent in guiding prosocial choice.

By contrast, third-party inference and role-based expectation are rarely observed in nonhuman primates and emerge only later in human development. Nonetheless, both are deeply embedded in human social life, forming the basis for reputation systems, institutional roles, and generalized cooperation beyond direct experience.

A note on “trust”. The term “trust” is often applied to systems like ChatGPT or Google, but what is described is not social trust—it is functional reliability through repeated exposure and low failure rates. These systems do not engage in reciprocal reasoning or uphold social obligations.

This semantic conflation becomes especially problematic in high-stakes domains like finance and healthcare, where “trust” is better understood as a demand for exceptional systemic robustness—not social accountability.

In finance, phrases like “trust in banks” or “trust in money” reflect belief in institutional continuity, enforcement mechanisms, and liquidity guarantees—not interpersonal confidence. What matters is structural resilience against systemic collapse.

In healthcare, “trust in doctors” or “hospital systems” similarly refers to confidence in diagnostic accuracy, institutional safeguards, and low error rates. Patients rely on these

systems not because of perceived benevolence, but because failures are rare and tightly bounded.

4.3. Cost–return sensitivity

Definition 4.3 (Cost–Return Sensitivity). Cost–return sensitivity is the capacity to regulate cooperative behavior based on perceived payoff asymmetries over time. Rather than relying on fixed heuristics, agents adjust participation by tracking whether interactions yield sustained net benefit—enabling the avoidance of exploitation and reinforcement of beneficial exchange.

Empirical Support. In nonhuman primates, cost–return sensitivity appears in both prosocial and punitive contexts. Chimpanzees modulate grooming and cooperative behavior based on prior benefits received (De Waal, 1997), and retaliate against theft even without immediate gain—suggesting expectations of equity and sensitivity to intentional harm (Jensen et al., 2007).

Human infants show comparable intuitions. By 18 months, they preferentially help agents who previously acted fairly or cooperatively (House et al., 2013), calibrating prosocial behavior not through abstract norms, but through observed patterns of giving and withholding.

Together, these findings suggest that cost–return sensitivity enables organisms to regulate cooperative investment in the absence of formal contracts—providing a cognitive substrate for reciprocity long before the emergence of markets.

5. Simulation Pathways and Implementation

We define three simulateable behavioral primitives—individual recognition, reciprocal credence, and cost–return sensitivity—that can serve as the minimal cognitive capacities required for scalable reciprocal behavior in multi-agent systems. Each can be approximated through lightweight memory architectures and prompt-level reasoning, without relying on explicit reinforcement learning.

- **Individual recognition:** Agents should be able to differentiate between social partners and retrieve information specific to each. This may reflect accumulated positive history, social familiarity, or bonded interaction patterns—not necessarily explicit naming.
- **Reciprocal credence:** Agents should maintain an internal estimate of how likely a partner is to return cooperation. This estimate integrates: (1) prior prosocial behavior directed toward the agent or observed in others, and (2) contextual signals that make cooperative return likely. The value is asymmetric, updateable, and sensitive to both interaction and situation.

- **Cost–return sensitivity:** Agents should track asymmetries in past exchanges, allowing them to adjust future cooperative investment. This mechanism supports dynamic calibration of effort and helps stabilize long-term reciprocity.

These primitives specify functional requirements rather than architectural constraints, offering a practical path toward socially grounded reciprocity in multi-agent systems.

6. Implications and Discussion

6.1. Limitations of the Framework

Comparative Methodology. Primate and infant studies rely on distinct paradigms: the former captures ecologically embedded behavior; the latter often uses constrained, artificial settings. This methodological gap limits fine-grained cross-species comparison of cognitive substrates.

Institutional Complexity. We do not claim that our framework alone accounts for complex institutions such as money, debt, or taxation. These might require further symbolic, cultural, and historical layers beyond minimal reciprocity mechanisms.

Social Scale and Enforcement. The current framework addresses partner-contingent reciprocity in small groups. It does not yet explain population-scale mechanisms—such as third-party punishment or distributed reputation—that enable anonymous cooperation. Bridging this gap remains a key direction for extending the model.

6.2. Theoretical Implications

Reciprocity as the Foundation of Social Structure. The term *cooperation* is frequently used as a catch-all for prosocial behavior, but its vagueness makes it difficult to model or implement. Without a clear cognitive basis, many systems rely on artificial payoff matrices, fixed strategies, or exogenous incentives—mechanisms that simulate outcomes but not the behavioral processes behind them.

In contrast, reciprocity offers a cognitively grounded and behaviorally structured alternative. Rather than relying on externally defined payoffs, it emerges from simple agent-level mechanisms that support dynamic, partner-contingent decision-making. This makes reciprocity tractable, simulateable, and aligned with real-world behavior.

More importantly, reciprocity underlies many core functions of modern society. From social favors and mentorships to business alliances and venture capital, cooperation often depends on asymmetric, memory-based, and deferred exchanges. Even scientific systems—such as peer review, citation, and referrals—operate through implicit reciprocity,

not formal enforcement.

We argue that reciprocity—not abstract cooperation—is the minimal behavioral substrate of social structure. It enables agents to build conditional expectations, regulate ongoing interactions, and scale bilateral exchange into stable institutional patterns. Far from being a subset of cooperation, reciprocity is the mechanism that makes cooperation sustainable and socially extensible in the first place.

From Institutions to Cognition. Our work reorients the study of exchange from an institutionalist or symbolic perspective toward a cognitive-behavioral foundation. Rather than treating markets, money, or debt as cultural inventions that enabled exchange, we argue that exchange itself is grounded in prior cognitive and social capacities—capacities that predate and scaffold institutionalization.

Rethinking Trust as System Robustness. The term “trust” is often invoked in discussions of economic exchange, yet it bundles together distinct ideas—from interpersonal reliability to confidence in systems. We argue that this vagueness limits its explanatory value.

To address this, we introduce *reciprocal credence* as a graded, behaviorally grounded expectation—specific to reciprocal exchange—that governs whether an agent initiates or maintains cooperation. It depends not on social bonding, but on observable interaction history and contextual inference.

In contrast, large-scale “trust” in systems like money or platforms is better understood as confidence in system robustness—not interpersonal accountability or relational bonding. This reframing clarifies how large-scale exchange systems can emerge without relying on social-bond-based trust.

Cross-Species Foundations of Exchange. By aligning primate cooperation studies with human developmental data and ethnographic accounts of non-monetary exchange, we provide a cross-species bridge for modeling the origins of economic behavior. This integrative view challenges the idea that complex exchange systems are uniquely human, and instead locates their roots in broader social cognition.

Unifying Behavioral Economics through Cognitive Foundations. If economic behavior is an extension of biologically grounded reciprocity, then so-called “biases” identified in behavioral economics—such as anchoring bias, loss aversion—are not anomalies—they are direct extensions of biological traits.

6.3. Future Work

Operationalizing Economic Emergence. A critical next step is to embed our proposed primitives—*individual recognition*, *reciprocal credence*, and *cost–return sensitivity*—as explicit modules in multi-agent simulations. Future experiments should systematically test how memory constraints, scalar reciprocity estimates, and payoff asymmetries influence the emergence of sustained cooperation and reciprocal exchange structures.

Institution Formation via Behavioral Mechanisms. Further simulations should explore whether agents equipped with these minimal cognitive primitives spontaneously develop proto-institutions such as symbolic debt, token systems, or role-based cooperation. Advances in embodied simulators (e.g., VirtualHome (Puig et al., 2018), AutoGen (Wu et al., 2023)) offer opportunities to bridge cognitive theories of exchange with practical implementations of scalable social structures.

Reciprocity Among Unfamiliar Agents. Finally, investigating reciprocal exchange among unfamiliar agents—without prior recognition or shared interaction history—represents an important frontier. Understanding the cognitive scaffolds enabling such interactions could significantly advance our grasp of decentralized cooperation mechanisms.

7. Conclusion

Contrary to the classical narrative of barter, converging evidence from early human societies suggests that economic life began with reciprocity—long-term, partner-contingent exchange embedded in social relationships. This behavioral substrate enabled communities to circulate resources, maintain cohesion, and sustain cooperation well before markets or institutions emerged.

Yet despite its foundational role in early economies, reciprocity remains under-formalized in both economic theory and computational modeling. We propose a cognitively grounded framework that identifies three simulateable mechanisms—individual recognition, reciprocal credence, and cost–return sensitivity—as the minimal substrate for scalable exchange.

By grounding abstract terms like “trust” and “cooperation” in observable mechanisms, this approach shifts economic modeling from top-down design to bottom-up simulation. These primitives not only clarify the origins of exchange, but also provide a tractable substrate for implementing reciprocal behavior in multi-agent systems.

Impact Statement

This paper introduces a cognitively grounded framework for simulating the behavioral foundations of economic exchange, based on biologically plausible primitives. By shifting the focus from top-down incentives to bottom-up social cognition, this approach enables more realistic and interpretable models of reciprocal cooperation. We anticipate that this framework could significantly improve agent-based economic simulations, particularly in multi-agent systems concerned with trust, fairness, and coordination. While the work is theoretical and does not involve deployment or data collection, we are not aware of any negative societal or ethical risks associated with this research.

References

- Berg, J., Dickhaut, J., and McCabe, K. Trust, reciprocity, and social history. *Games and economic behavior*, 10(1): 122–142, 1995.
- Brosnan, S. F. and De Waal, F. B. Monkeys reject unequal pay. *Nature*, 425(6955):297–299, 2003.
- De Waal, F. B. The chimpanzee’s service economy: Food for grooming. *Evolution and Human Behavior*, 18(6): 375–386, 1997.
- Dunfield, K., Kuhlmeier, V. A., O’Connell, L., and Kelley, E. Examining the diversity of prosocial behavior: Helping, sharing, and comforting in infancy. *Infancy*, 16(3):227–247, 2011.
- Fehr, E. and Gächter, S. Altruistic punishment in humans. *Nature*, 415(6868):137–140, 2002.
- Fehr, E. and Schmidt, K. M. A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868, 1999.
- Gan, C., Schwartz, J., Alter, S., Mrowca, D., Schrimpf, M., Traer, J., De Freitas, J., Kubilius, J., Bhandwaldar, A., Haber, N., et al. Threedworld: A platform for interactive multi-modal physical simulation. *arXiv preprint arXiv:2007.04954*, 2020.
- Gintis, H. Strong reciprocity and human sociality. *Journal of theoretical biology*, 206(2):169–179, 2000.
- Gomes, C. M., Mundry, R., and Boesch, C. Long-term reciprocation of grooming in wild west african chimpanzees. *Proceedings of the Royal Society B: Biological Sciences*, 276(1657):699–706, 2009.
- Güth, W., Schmittberger, R., and Schwarze, B. An experimental analysis of ultimatum bargaining. *Journal of economic behavior & organization*, 3(4):367–388, 1982.
- Hirata, S. and Fuwa, K. Chimpanzees (pan troglodytes) learn to act with other individuals in a cooperative task. *Primates*, 48:13–21, 2007.
- House, B. R., Silk, J. B., Henrich, J., Barrett, H. C., Scelza, B. A., Boyette, A. H., Hewlett, B. S., McElreath, R., and Laurence, S. Ontogeny of prosocial behavior across diverse societies. *Proceedings of the National Academy of Sciences*, 110(36):14586–14591, 2013.
- Jensen, K., Call, J., and Tomasello, M. Chimpanzees are vengeful but not spiteful. *Proceedings of the National Academy of Sciences*, 104(32):13046–13050, 2007.
- Lewis, L. S., Wessling, E. G., Kano, F., Stevens, J. M., Call, J., and Krupenye, C. Bonobos and chimpanzees remember familiar conspecifics for decades. *Proceedings of the National Academy of Sciences*, 120(52):e2304903120, 2023.
- Li, G., Hammoud, H., Itani, H., Khizbullin, D., and Ghanem, B. Camel: Communicative agents for” mind” exploration of large language model society. *Advances in Neural Information Processing Systems*, 36:51991–52008, 2023.
- Malinowski, B. *Argonauts of the western Pacific: An account of native enterprise and adventure in the archipelagoes of Melanesian New Guinea [1922/1994]*. Routledge, 2013.
- Mauss, M. *The gift: The form and reason for exchange in archaic societies*. Taylor & Francis, 2024.
- Melis, A. P., Hare, B., and Tomasello, M. Chimpanzees recruit the best collaborators. *Science*, 311(5765):1297–1300, 2006a.
- Melis, A. P., Hare, B., and Tomasello, M. Engineering cooperation in chimpanzees: tolerance constraints on cooperation. *Animal Behaviour*, 72(2):275–286, 2006b.
- Park, J. S., O’Brien, J., Cai, C. J., Morris, M. R., Liang, P., and Bernstein, M. S. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pp. 1–22, 2023.
- Puig, X., Ra, K., Boben, M., Li, J., Wang, T., Fidler, S., and Torralba, A. Virtualhome: Simulating household activities via programs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8494–8502, 2018.
- Sahlins, M. *Stone age economics*. Routledge, 2013.
- Schino, G. and Aureli, F. Primate reciprocity and its cognitive requirements. *Evolutionary Anthropology: Issues, News, and Reviews*, 19(4):130–135, 2010.

Schmelz, M., Grueneisen, S., Kabalak, A., Jost, J., and Tomasello, M. Chimpanzees return favors at a personal cost. *Proceedings of the National Academy of Sciences*, 114(28):7462–7467, 2017.

Vaish, A., Carpenter, M., and Tomasello, M. Young children selectively avoid helping people with harmful intentions. *Child development*, 81(6):1661–1669, 2010.

Wang, G., Xie, Y., Jiang, Y., Mandlekar, A., Xiao, C., Zhu, Y., Fan, L., and Anandkumar, A. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.

Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., et al. Autogen: Enabling next-gen llm applications via multi-agent conversation. *arXiv preprint arXiv:2308.08155*, 2023.