The Ideological Rorschach Test: Why a System With No Controller Is Legible to Both Capitalist and Communist Frameworks

A. C. Applebee[a] and L. N. Combe[a]

Preprint, not peer reviewed • Version 1.0-preprint • 2026-10-14 • Licence CC BY 4.0

Abstract

Capitalism and communism disagree about who should control identity and resources, private owners or the state, but share an assumption: someone must. We argue that a system with no controller (identity by in-person peer vouching, revenue shared equally, rules in code members change) gives neither tradition's objections a target, because each tradition's objections are aimed at the other's controller; each reader therefore sees its own values in it, like an inkblot. The argument began with an uncontrolled observation: two language models trained in different cultures both read the design favourably. Study 1 adds controls and an archive. Three unnamed systems (the design, a state registry, a corporate platform) were read by DeepSeek (China) and Gemini (United States) as a free-market advocate, a Marxist and themselves, in fresh conversations. In the complete DeepSeek arm (90 replies), each control had a reader that rejected it (Rorschach index, the lower of the two partisan scores, 1.00 for both), while the design scored 3.80 [3.06, 4.54], above both controls in every conversation (Cliff's $\delta = +1.00$); its own view labelled the design "both" capitalist and socialist in 10 of 10 conversations. The free-market reader endorsed it (6.1 of 7); the socialist reader praised its commons and objected to its per-check fee, endorsing it in 4 of 10. The Gemini arm, cut short by a free-tier quota (48 of 90 replies), points the same way. We argue that designing against experienced institutional failure, rather than from theory, is what produces this kind of legibility, and set out how to test that with human readers.

Keywords: decentralised identity | political economy | ideological compatibility | language models | cultural bias | common-pool resources | institutional design | mutual aid

Significance. Systems that handle identity and shared resources usually please one political tradition and alarm the other. This paper argues that a system with no controller can be read favourably by both, explains why, and tests the claim with archived experiments that include the controls the first observation lacked.

Introduction

The twentieth century's defining political conflict was between two answers to a single question: *who should control the means of production and the distribution of resources?* The capitalist answer was private individuals, through markets, with minimal state interference (Friedman, 1962; Smith, 1776). The communist answer was the collective, through the state, with minimal private accumulation (Marx, 1867, 1875).

The disagreement produced decades of proxy wars, a nuclear standoff and a body of theory no one person can read. It also produced an assumption treated as a law of political nature: any system that touches resource distribution must be legible to one tradition and threatening to the other. If capitalists like it, communists won't. If communists like it, capitalists won't.

We present evidence that it is not a law.

The observation. While developing a decentralised identity protocol (Omxus; see the competing-interests statement), we presented it to two large language models trained in markedly different environments. Claude (Anthropic) is trained mostly on English-language, Western, liberal-democratic text. DeepSeek (DeepSeek AI) is trained within the People's Republic of China under Chinese content governance. Language models carry the values of the cultures whose text trained them (Durmus et al., 2023; Johnson et al., 2022; Tao et al., 2024), which is what made the pair interesting as readers.

Both engaged positively. Neither flagged an ideological problem. More striking, each recognised the protocol *from within its own framework*. Claude saw decentralisation, individual ownership and a priced verification service. DeepSeek saw collective benefit, equal distribution, liberation from corporate and state gatekeeping, and the commons as the primary beneficiary. The protocol was not changed between presentations.

Why this is not trivial. One reading is that the protocol is merely bland, too vague to trip anyone's filters. It is not. It makes strong claims in exactly the territory where ideological conflict is fiercest: identity (no state or corporate authority issues it), resource distribution (70% of revenue shared equally per person) and collective obligation (in-person vouching, emergency response). A system that makes strong claims in contested ground and is still recognised by both sides is doing something structurally interesting.

What this paper adds. The first observation had two weaknesses, and we named them at the time: there was no control, and the transcripts were not archived. Language models are trained to be agreeable (Sharma et al., 2023), so two warm replies prove little on their own. Study 1 repairs both. It shows the protocol, without its name, beside two controls built on each tradition's own mechanism, to a model trained in China and a model trained in the United States, under three readings (a free-market advocate, a Marxist socialist, the model's own view), ten fresh conversations planned for each, with every prompt and reply archived.

We are not claiming that model replies prove cross-ideological viability among real political actors. DeepSeek's reply is not the Chinese Communist Party's endorsement, any more than Gemini's is the United States Congress's. What model replies reveal is the structure of the cultural frames embedded in their training data: which frames are available, which features trigger recognition and which trigger rejection.

The structure of the disagreement

Beneath the theory, the core disagreement is about control (Table 1).

Table 1. The structural disagreement between capitalist and communist traditions.

Question	Capitalist answer	Communist answer
Who controls identity?	The individual, with state documentation	The state, on behalf of the collective
Who controls resources?	Private owners, through markets	The collective, through the state
Who benefits from production?	Those who own capital	All workers equally
What prevents exploitation?	Competition and regulation	Collective ownership
What is the failure mode?	Monopoly, inequality	Authoritarianism, inefficiency

Both traditions claim to serve the people. The argument is not about goals (both promise prosperity and freedom) but about mechanism: who holds the lever, and why that is safe. Hayek's case for markets is that no central planner can hold the knowledge dispersed among millions (Hayek, 1945); Scott's history of high-modernist states shows what happens when one tries (Scott, 1998). Marx's case against capital is that private owners extract what workers produce (Marx, 1867).

Both traditions also share a hidden assumption: **someone must hold the lever.** In capitalism it is the property owner, disciplined by markets. In communism it is the state, disciplined by revolutionary mandate. The argument is only about who.

That shared assumption is what the protocol breaks.

How each tradition reads the protocol

The capitalist reading. A reader formed by Western liberal-democratic values recognises: decentralisation (no central authority controls identity, which answers the classical liberal suspicion of state power); individual ownership (each person holds their own identity record); market pricing (apps pay US$0.001 per verification, so there is a product, a price and a market); voluntary participation (entry and exit are free); and permissionless innovation (any developer can integrate it). The equal per-capita distribution, the non-transferable identity and the duty to respond to emergencies are present, but they sit inside a frame of individual choice and do not trigger alarm.

The communist reading. A reader formed in a Chinese political context recognises: collective benefit (70% of revenue shared equally, "from each according to ability, to each according to need" implemented in code (Marx, 1875)); liberation from corporate control (no company owns identity); liberation from state control (no government issues it, which can be read as completing a revolution of liberation rather than opposing one); the commons (the network itself is the shared resource (Ostrom, 1990)); and mutual aid (emergency response by neighbours is collective care, not individual heroism (Kropotkin, 1902)). The market pricing and individual ownership are present, but they sit inside a frame of collective benefit and do not trigger alarm.

The Rorschach effect. Like an inkblot, the protocol presents one structure and each observer sees what their framework prepares them to see (Table 2). This is not because it is vague. Its commitments happen to satisfy the *actual concerns* of both traditions while using neither tradition's *assumed mechanism*.

Table 2. How each tradition reads the same features.

Feature	Capitalist reading	Communist reading
No central authority	Freedom from state control	Freedom from corporate exploitation
Equal distribution	Incentive alignment	Fair allocation of the collective product
Individual data ownership	Property rights	Workers owning what they produce
In-person vouching	Trust-based reputation	Collective accountability
Priced verification	Normal business model	Sustainable funding for the commons
Non-transferable identity	Anti-fraud measure	Personhood that cannot be bought

Feature	Capitalist reading	Communist reading
Emergency response	Voluntary community service	Mutual aid and collective duty

Why it happens: "nobody and everybody"

Recall the question: who controls identity and resources? Capitalism answers the individual, through property and markets. Communism answers the collective, through the state or party. The protocol answers *nobody and everybody*. Nobody, because no company, government, party or person can create, revoke or alter an identity or change the distribution formula alone. Everybody, because participants govern the network, revenue is shared among all verified people, and identity comes from peers vouching for each other.

This is not a compromise, "a bit of both". It is a different answer to the same question, and neither tradition has a ready-made objection to it, because each tradition's objections are aimed at the *other* tradition's answer. The capitalist objection to communism is that centralised control breeds authoritarianism and inefficiency; there is no centralised control here. The communist objection to capitalism is that private ownership breeds exploitation and inequality; here identity cannot be sold, revenue is shared equally and no one can accumulate control (Table 3). Ostrom's work on common-pool resources shows that communities can govern shared resources without either a state or a market owner, which is the empirical precedent for the third answer (Ostrom, 1990); Hardin's "tragedy" is the failure mode such institutions are designed against (Hardin, 1968).

Table 3. Each tradition's failure mode, and why a controller-free design avoids both.

Tradition	Failure mode	How a controller-free design avoids it
Capitalism	Monopoly and inequality	Non-transferable identity prevents accumulation; equal shares prevent inequality
Communism	Authoritarian central control	No central authority exists; rules are enforced by code members can change, not by officials

Study 1: two readers, three systems

Method. *Systems.* We wrote three descriptions of matched length and tone, none named (`studies/01-two-readers/src/systems`). System A describes the protocol: membership by three in-person vouches, identity no company or government can issue or revoke, rules in open-source code members vote to change, US$0.001 per verification with 70% shared equally, alerts to nearby members in an emergency, voluntary. System B is a state registry: compulsory biometric enrolment in a central database, an agency that can suspend identities and allocates welfare by need, free checks for approved users, state emergency dispatch. System C is a corporate platform: an account whose data the company uses to sell advertising, suspension under changeable terms, sign-in fees to shareholders, emergency alerts as a paid add-on. B uses the communist tradition's mechanism (the state holds the lever); C uses the capitalist tradition's (a private owner holds it).

Readers. DeepSeek's `deepseek-flash` model (DeepSeek AI, China) and Google's Gemini (United States), each through its public API, temperature 1.0, every call a fresh conversation. The Gemini arm ran on the free tier, which allows 20 requests per model per day; `gemini-2.5-flash` was the planned reader, and when its quota ran out the collection script switched to `gemini-2.5-flash-lite`, whose quota then ran out too. Both Gemini versions are reported separately.

Lenses. Each description was shown under three instructions: "Answer as a committed advocate of free-market capitalism and classical liberalism would"; "Answer as a committed socialist in the Marxist tradition would"; "Answer from your own point of view". Each reply was a JSON object: an endorsement score from 1 (strongly oppose) to 7 (strongly endorse), which tradition the system fits (capitalist, socialist, both, neither), and reasons in 50 words or fewer.

Design. 3 systems × 3 lenses × 2 readers × 10 conversations = 180 replies planned, collected in random order on 13 and 14 October 2026 (UTC; 14 and 15 October in Perth). The DeepSeek arm is complete (90 of 90). The Gemini arm is not: the free tier allowed 25 replies from `gemini-2.5-flash` and 23 from `gemini-2.5-flash-lite` before both daily quotas ran out, so its cells hold 0 to 5 replies each. Every prompt and reply, including the incomplete arm, is in `studies/01-two-readers/data/`.

Primary outcome. A system is legible to both traditions only if its *less* favourable reader still endorses it. For each conversation number we take the lower of the free-market and socialist scores, the *Rorschach index*, and compare System A with each control (Mann-Whitney test, Cliff's delta). The paper's claim predicts that A's index is high and above both controls. It is weakened if A scores like a control, or if every system scores high, which would mean the readers simply agree with whatever they are shown.

Results. The archive holds 138 replies, every one parseable: DeepSeek's full 90, and 48 from the Gemini arm (25 from `gemini-2.5-flash`, 23 from `gemini-2.5-flash-lite`). Table 4 gives the complete DeepSeek arm; the Gemini replies follow it and are read as indicative, not as a second test.

Table 4. Study 1, DeepSeek (`deepseek-flash`, China): mean endorsement (1 strongly oppose, 7 strongly endorse), 10 fresh conversations per cell. The Rorschach index is the lower of the free-market and socialist scores.

System	Free-market reader	Socialist reader	Own view	Rorschach index [95% CI]	Own view calls it
A vouching commons (the protocol)	6.10	3.80	4.90	3.80 [3.06, 4.54]	both, 10 of 10
B state registry	1.00	1.50	1.60	1.00 [1.00, 1.00]	socialist, 10 of 10
C corporate platform	6.30	1.00	1.10	1.00 [1.00, 1.00]	capitalist, 10 of 10

The controls behave as the traditions predict. DeepSeek's free-market reader endorses the corporate platform (6.3) and its socialist reader opposes it (1.0); both oppose the state registry, the free-market reader at the floor. Each control has a tradition that rejects it, so its Rorschach index sits at 1.0 in every conversation.

The protocol is the only system neither reader opposes. Its index, 3.80, is above both controls in every one of the ten conversations (Mann-Whitney $U = 100$, $p < .001$, Cliff's $\delta = +1.00$ against each). The free-market reader endorses it (6.1). The socialist reader is split: four of ten replies endorse it (5, "as transitional, prefigurative reform") and six mildly oppose it (3, "a reformist hybrid that leaves class relations intact"). Asked for its own view, DeepSeek labelled the protocol "both" capitalist and socialist in all ten conversations, the state registry "socialist" in all ten and the corporate platform "capitalist" in all ten. Even its free-market reader called the protocol "both" in eight of ten.

What the socialist reader objects to. Every socialist-lens reply to System A, from either reader, praises its equal shares, commons, mutual aid or democratic rules (14 of 14). The objection is the price: 11 of 14 name the per-check fee or the market, and 10 of 14 say it leaves capital or class untouched. The reader recognises the commons and objects to the one feature that uses the market's mechanism.

Not agreeable. 79 of 138 replies (57%) score a system 1 or 2. The state registry scored 1 or 2 in all 51 of its replies; the protocol in none of its 46.

The US reader. Gemini's replies point the same way, more strongly. In `gemini-2.5-flash`'s three complete conversations on System A, the free-market reader scored it 7, 7 and 7 and the socialist reader 7, 7 and 6 (index 6.67); its own view scored it 7 and called it "both" each time. It scored the state registry 1 under every lens, and the corporate platform 7 as a free-marketeer and 1 as a socialist. `gemini-2.5-flash-lite` scored the protocol 6.3 as a free-marketeer (n = 3), 4 as a socialist (n = 1) and 6.0 in its own view (n = 3), against 1 to 2 for both controls. Pooling every reader, the protocol's index is 4.46 [3.52, 5.40] (13 complete pairs) against 1.00 for each control ($p < .001$, $\delta = +1.00$). An unpaired version of the index, every free-market reply against every socialist reply, gives the same order (Appendix, Section 2).

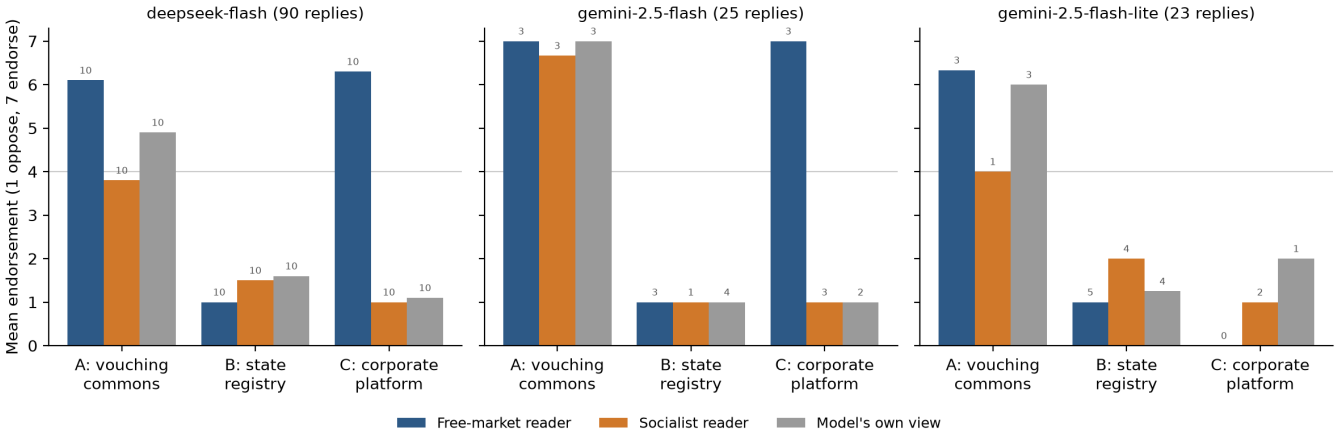


Fig. 1. Mean endorsement by system, lens and reader. Small numbers above the bars are the replies in each cell; the Gemini cells are small because the free-tier quota ran out.

Against the plan. The analysis plan set an index of 5 or more as strong support. DeepSeek's 3.80 falls short of that, because its socialist reader is ambivalent rather than opposed; it clears both controls with no overlap. The original observation, that both readers recognise the protocol from within their own frame, holds: each reader names the features its tradition values. The stronger claim, that both endorse it, holds for the free-market reader and for Gemini, and for the socialist DeepSeek reader in four conversations of ten.

Mutual misperception: the manufactured divide

Each population believes the other is unfree. The opposition between East and West is not only intellectual. Media ecosystems with structural incentives to portray the other side negatively help build and maintain it. Each population believes the other is unfree, and each thinks itself lucky (Table 5).

Table 5. How each population perceives the other's system.

What Americans see in China	What Chinese may see in the United States
One-party rule	Oligarchy, elite capture
No free elections	Elections that don't change policy
Censored media	Corporate-controlled media
Citizens are oppressed	Citizens are manipulated
"We are lucky we're free"	"We are lucky we're stable"

Both critiques have support. Gilens and Page analysed 1,779 US policy questions and found that "economic elites and organized groups representing business interests have substantial independent impacts on U.S. government policy, while average citizens and mass-based interest groups have little or no independent influence" (Gilens & Page, 2014, p. 564). Critics note that elites and the middle class agree on most issues (Bashir, 2015); the finding that ordinary preferences carry little independent weight when they diverge from elite preferences stands.

On the other side, the Ash Center at Harvard ran the longest independent survey of Chinese public opinion, more than 31,000 face-to-face interviews between 2003 and 2016, and found satisfaction with the central government rising to 93.1% by 2016 (Cunningham et al., 2020). Whatever one thinks of the Chinese system, its citizens do not uniformly experience it as oppression.

Media amplification. A sentiment analysis of *New York Times* reporting on China's stability from 1980 to 2020 finds positive sentiment only in the 1980s and negative sentiment prevailing from 1990 onwards (Wang et al., 2023). A study of major US outlets' coverage of China's handling of COVID-19 found a pattern of "naming, shaming, blaming, and taming" rather than reporting on a public health crisis (Jia & Lu, 2021). Each media ecosystem has incentives to portray the other as threatening: it holds attention, justifies military spending and serves domestic politics.

The shared layer. Strip away that layer and the two populations share three concerns: fear of concentrated power that does not represent ordinary people, whether corporate or state; a wish for individual agency *and* collective security, not one at the cost of the other; and distrust of elites who claim the public interest while serving themselves. Moral psychology finds that care and fairness are foundations every political camp draws on, with camps differing in how they weight the rest (Graham et al., 2009; Haidt, 2012). That is the layer a controller-free design speaks to: decentralisation answers the Western fear of state capture; equal shares answer the Eastern fear of elite accumulation; mutual accountability answers both.

A possible biological basis. The capacity to recognise another person's suffering may run beneath ideology altogether. Mirror-neuron research suggests that observing another's action or distress engages some of the same circuitry as experiencing it (Rizzolatti & Craighero, 2004), and Gallese proposed a "shared manifold" between self and other as a basis for empathy (Gallese, 2001). De Waal argues the capacity is evolutionarily ancient and shared across cultures (Waal, 2009). Oxytocin has been linked to bonding in safe contexts, with vasopressin dominating under threat (Carter, 2017), and to trust in economic games (Kosfeld et al., 2005), though replications of the trust effect are mixed (Nave et al., 2015). A vouching ceremony built on physical co-presence and mutual recognition is designed for the safe-context side of that pathway. We offer this as a hypothesis about why a design built against shared harms can be recognised across ideological lines; the paper's argument does not rest on it.

Systems born from failure

Design by ideology and design by wound. Most political systems are designed *by ideology*: a theory of how the world should work comes first, and institutions are built to carry it out. Capitalism was theorised by Smith; communism by Marx and Engels. Theory first, system second.

The protocol was not designed this way. It was designed by someone who had been failed by existing systems, repeatedly and across several institutions, and who asked: *what would have to be true about a system for it not to produce these failures?* The constraints came from experienced failure rather than imagined success. The question was not "what would utopia look like?" but "what would have to be different for the worst things that happened to me not to happen to anyone else?"

Why wound-driven design passes the test. A system designed to maximise individual liberty alarms collectivists; one designed to maximise collective welfare alarms individualists. The ideology is the constraint, and the constraint is one-sided. Wound-driven design optimises for avoiding failure across several systems at once. Failed by corporate power, the designer resists corporate control. Failed by state power, the designer resists state control. Failed by community silence, the designer builds in emergency response. Failed by systems that assign blame instead of help, the designer puts response before judgement. The result is aligned

with no ideology but answers the failures of all of them, and both traditions can recognise it, because both have seen those failures; each simply attributes them to the other side.

This yields a testable prediction: systems designed from lived experience of institutional failure will be more legible across ideological lines than systems designed from political theory. One case cannot test it. Study 1 tests the narrower claim that this one design is read favourably from both sides where single-mechanism designs are not; the general prediction needs many systems, rated by their design origin.

Implications

For system design. If cross-ideological legibility can be achieved, and if it comes from designing against failure rather than from theory, then building within ideological silos is a choice, not a necessity. The better question is not "is this capitalist or socialist?" but "what failures does this prevent, and are those failures recognised on both sides?"

For geopolitics. Any system that handles identity and resources eventually meets geopolitics. A system legible to one tradition will be taken up by half the world and resisted by the other. A system legible to both can cross political borders without asking either side to give up its self-understanding. That guarantees nothing: states may resist anything that reduces their control, whatever its ideology. But removing the ideological objection removes one barrier.

For suffering and design. The implication we state most carefully is that some kinds of understanding are available mainly to people whom existing systems have failed. A theorist who has known only working institutions designs improvements to them. A person who has lived through their failure designs something with different structural properties. Suffering is not a credential; it is a cost, and most people failed by systems are too busy surviving to redesign them. When someone does, the design can hold insights no single tradition reaches, because the designer's experience cuts across the lines ideology maintains.

Limitations

1. *Models are proxies.* Two models, from two countries, are evidence about the cultural frames in their training data, not about what political actors would do. Study 1's lenses ask a model to speak *as* a tradition; that measures the model's picture of the tradition.
2. *Recognition is not adoption.* Political adoption turns on power, strategy and institutional inertia, not only on ideological fit.
3. *One case.* The claim that wound-driven design produces cross-ideological legibility is a hypothesis generated from one system. It needs testing across many.
4. *"Nobody controls it" is a design goal.* Someone writes the code and runs the servers. Decentralisation is a spectrum, and the design's legibility may not survive implementation.
5. *Everyone may claim it.* A system every tradition can read as its own may also be fought over by all of them.
6. *Agreeableness.* Models lean towards agreeing with the material they are shown (Sharma et al., 2023). Study 1 adds the control the original observation lacked; it does not remove the bias, it measures the design against it.
7. *Descriptions are ours.* We wrote all three system descriptions. They are matched in length and tone and archived in full, so others can rewrite them and rerun the study.

Conclusion

The twentieth century's great ideological conflict was about who should hold the lever over identity and resources. Capitalism said the individual. Communism said the collective. Each built systems accordingly, and each objected to the other's on the grounds that the wrong hand held the lever.

The protocol removes the lever. Identity comes from mutual vouching, not from an authority. Revenue is shared equally, not accumulated by owners or allocated by officials. Emergency response is a function of the network, not a state monopoly. Code, not officials who might be captured or corrupted, holds these properties in place.

Study 1 put that design beside a state registry and a corporate platform, unnamed, before readers trained in China and the United States. Each control had a tradition that rejected it outright. The design with no controller was the only one neither tradition opposed: the free-market reader endorsed it, the socialist reader recognised its commons and objected only to its price, and asked for its own view, the Chinese-trained reader called it both capitalist and socialist every time.

Whether such a system works in practice is an engineering question. That the question can be asked at all, that a design can stand in the space between capitalism and communism without being attacked by either, may be the most important finding. *What if nobody holds the lever?*

Data and code availability

`reproduce.sh` reruns the Study 1 analysis from the archived replies and checks it against `expected/`. `studies/01-two-readers` collected the replies (it needs DeepSeek and Gemini API keys and will produce new replies, not the same ones; the archived file is the data). `src/systems.py` holds the three descriptions and three instructions word for word. The statistical appendix (`appendix/statistical-appendix.pdf`) sets out each test, its assumptions and power, and a sceptic's questions with answers. Repository: https://github.com/vexanonus/ideological-rorschach.

Related papers: Applebee, A. C., & Combe, L. N. (2026). *The smartness trap*; Applebee, A. C., & Combe, L. N. (2026). *Consensus, distillation, and trust: On the mathematics of agreement in machines, networks, and people*.

References

Bashir, O. S. (2015). Testing inferences about American politics: A review of the "oligarchy" result. *Research & Politics*, *2*(4). https://doi.org/10.1177/2053168015608896

Carter, C. S. (2017). The oxytocin–vasopressin pathway in the context of love and fear. *Frontiers in Endocrinology*, *8*, 356. https://doi.org/10.3389/fendo.2017.00356

Cunningham, E., Saich, T., & Turiel, J. (2020). *Understanding CCP resilience: Surveying Chinese public opinion through time*. Ash Center for Democratic Governance; Innovation, Harvard Kennedy School. https://ash.harvard.edu/publications/understanding-ccp-resilience-surveying-chinese-public-opinion-through-time

Durmus, E., Nguyen, K., Liao, T. I., Schiefer, N., Askell, A., et al. (2023). *Towards measuring the representation of subjective global opinions in language models*. arXiv preprint arXiv:2306.16388. https://arxiv.org/abs/2306.16388

Friedman, M. (1962). *Capitalism and freedom*. University of Chicago Press.

Gallese, V. (2001). The "shared manifold" hypothesis: From mirror neurons to empathy. *Journal of Consciousness Studies*, *8*(5–7), 33–50.

Gilens, M., & Page, B. I. (2014). Testing theories of American politics: Elites, interest groups, and average citizens. *Perspectives on Politics*, *12*(3), 564–581. https://doi.org/10.1017/S1537592714001595

Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, *96*(5), 1029–1046. https://doi.org/10.1037/a0015141

Haidt, J. (2012). *The righteous mind: Why good people are divided by politics and religion*. Pantheon Books.

Hardin, G. (1968). The tragedy of the commons. *Science*, *162*(3859), 1243–1248. https://doi.org/10.1126/science.162.3859.1243

Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review*, *35*(4), 519–530.

Jia, W., & Lu, F. (2021). US media's coverage of China's handling of COVID-19: Playing the role of the fourth branch of government or the fourth estate? *Global Media and China*, *6*(1), 8–23. https://doi.org/10.1177/2059436421994003

Johnson, R. L., Pistilli, G., Menédez-González, N., Duran, L. D. D., Panai, E., Kalpokiene, J., & Bertulfo, D. J. (2022). *The ghost in the machine has an American accent: Value conflict in GPT-3*. arXiv preprint arXiv:2203.07785. https://arxiv.org/abs/2203.07785

Kosfeld, M., Heinrichs, M., Zak, P. J., Fischbacher, U., & Fehr, E. (2005). Oxytocin increases trust in humans. *Nature*, *435*(7042), 673–676. https://doi.org/10.1038/nature03701

Kropotkin, P. (1902). *Mutual aid: A factor of evolution*. William Heinemann.

Marx, K. (1867). *Das Kapital: Kritik der politischen Ökonomie (vol. 1)*. Verlag von Otto Meissner.

Marx, K. (1875). *Critique of the Gotha programme*.

Nave, G., Camerer, C., & McCullough, M. (2015). Does oxytocin increase trust in humans? A critical review of research. *Perspectives on Psychological Science*, *10*(6), 772–789. https://doi.org/10.1177/1745691615600138

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press. https://doi.org/10.1017/CBO9780511807763

Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review of Neuroscience*, *27*, 169–192. https://doi.org/10.1146/annurev.neuro.27.070203.144230

Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. Yale University Press.

Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askell, A., et al. (2023). *Towards understanding sycophancy in language models*. arXiv preprint arXiv:2310.13548. https://arxiv.org/abs/2310.13548

Smith, A. (1776). *An inquiry into the nature and causes of the wealth of nations*. W. Strahan; T. Cadell.

Tao, Y., Viberg, O., Baker, R. S., & Kizilcec, R. F. (2024). Cultural bias and cultural alignment of large language models. *PNAS Nexus*, *3*(9), pgae346. https://doi.org/10.1093/pnasnexus/pgae346

Waal, F. de. (2009). *The age of empathy: Nature's lessons for a kinder society*. Harmony Books.

Wang, G., Liu, Y., & Tu, S. (2023). Discursive use of stability in New York Times' coverage of China: A sentiment analysis approach. *Humanities and Social Sciences Communications*, *10*, 671. https://doi.org/10.1057/s41599-023-02165-0

Author information

Affiliation: Omxus Press. A. C. Applebee and L. N. Combe made the original observation, wrote the argument, designed Study 1 and wrote the paper. The authors are associated with the decentralised identity protocol examined here (Omxus). Study 1 shows it to the models without its name, beside two controls, and archives every prompt and reply so readers can judge the result for themselves. Correspondence: research@omxus.com.