

Print Edition

Wednesday 7 October 2026


EDITION NO. 76 • 2 PIECES • 7 PAGES

IN THIS EDITION

- | | | |
|----|---|---|
| 1. | Will agents have network effects? • @andrewchen | 1 |
| 2. | Social science and common sense • The Update | 6 |
-

Will agents have network effects?

BY ANDREW CHEN • @ANDREWCHEN • 6 OCTOBER 2026



We are seeing a new wave of general purpose consumer-friendly agents, and it's causing many thousands of AI startups to re-evaluate their place in the market. Some questions:

- Will these new horizontal players take over completely?
- Will it be winner take all?
- Can products that focus vertical use cases survive?

The new wave of products - Muse, Town, Instinct, Grokbot, etc - come from taking the power and magic of Openclaw/Hermes but wrapping it in polished, secure, consumer-friendly experience. This is only the beginning, and soon every major tech company will throw their hat in the ring and thousands of startups that have been working on vertical agents in X (travel, shopping, family, work, SMB, etc etc) have big existential questions to answer.

The simple case against winner-take-all dynamics: These agents are tools, not networks. They inherently provide a solo experience, and although packaged as chat experiences, they don't have network effects the

way that Whatsapp or phone networks do. They're more like email clients - I can use Superhuman and someone else uses Outlook and someone else is on Gmail, but in the end we just send email to each other. In that metaphor, if a Muse agent encounters a Instinct agent, they will just figure out how to talk to each other by building APIs, or using email/messaging/etc., so that there's no advantage to everyone running the same thing. Further, you can swap out the underlying LLM - as we saw in the Claude vs Clawd-bot saga - things will still work. The memory and underlying context are just text files, and AI are great at porting things from one system to another. So where's the moat?

Just because they are tools today doesn't mean they can't evolve into networks. As I've written about in the past, instead of discussing "network effects" as a vague/amorphous concept, let's instead overlay it against core KPIs each product has to build against:

- Acquisition network effects: More users** = lower CAC, higher virality, more signups
- Engagement effects: More users = higher retention, deeper usage, more frequency
- Monetization effects: More users = higher ARPU, higher conversion, more share of wallet

(**PS. and by "more users" of course I actually mean, of course, higher network density with more interconnection within/across networks. Not scale effects)

So looking at these sub-categories of network effects, you might ask: What are the features that you'd build to actually create network effects in agents? How does this category become winner-take-all?

Acquisition network effects

In the pre-AI world, network effects in customer acquisition were driven by users taking actions that then bring in even more users - whether that's sending an invite, sharing a piece of content, or getting mentioned in a comment. Those users would then join the network, repeat the same actions, creating more viral loops that would throw off more users. In the post-AI world, both users and agents can initiate these loops. **Agentic viral loops** are about to be a thing. An agent might suggest, "hey, do you want me to help organize a New Years celebration with family" and then might pull together a group chat with your siblings, in-laws, and parents. It might build a microsite for the event and ask people to sign up and RSVP, creating an opportunity to generate a signup. Or for work, you might finish a slide deck and your agent asks if you want to pull together a meeting to present it to your team. When it does that, it might create a prep doc, take notes, and host it all on a microsite that engaged others.

These types of agentic viral loops need a few things to work: They need to be able to plug into your contacts, and to have enough context to know when to loop which people into which projects/events/activities. They need access to communication channels like email/messaging/otherwise, specifically where they can reach potential users, as this serves as the viral

substrate for propagation. The magic will be how to get all of these things without being creepy, and to create enough trust to reach out to friends/colleagues/etc on your behalf, without being too push. No one wants an agent that says, “hey Andrew would like to get lunch next week” followed by “and sign up for this agent to agree!” - it has to be better and more subtle than that.

Many viral loops in the pre-AI era, particularly in the workplace, were built on a content creation loop - you use Google Slides to make a thing, you share a link, and then people would view it. Eventually they might make their own. Many products like Figma, spreadsheets, Notion, Canvas, are all built on this loop, and so are YouTube, Instagram, etc. In the agentic world, you might start by talking to your agent about interior decoration ideas, and it might ultimately build you a Pinterest-like board and make it easily shareable with others. It might build you a spreadsheet mini-app with a budget, a schedule, etc., and again, make it easily shareable. It might volunteer to do this inside a group chat with your friends, and if they want to view it, they should sign up and get their own agent too.

If the content creation loop works for agentic virality, I think it means there’s a big incentive for every agent to volunteer to make visually appealing, shareable content, do this often. AI is so good at codegen that creating a one-off disposable artifact will probably be more useful than not. And there’s a big incentive to drive virality by building/owning the content yourself as the agent. In the interior decoration example, it would be more viral to create a standalone version hosted by the agent, rather than building an actual Pinterest board, or to build a self-contained mini-spreadsheet rather than creating a GSheet. A properly tuned agentic viral loop will do the former because it helps spread the agent, rather than helping spread Pinterest/GSheets. There’s an incentive towards building walled garden assets rather than building/sharing other products.

Agentic viral loops are limited by different constraints - and thus, lead to a different type of viral factor - than pre-AI viral loops. As I wrote in my viral loop braindump a while back, traditionally getting your viral factor >1 has to do with: 1) reducing the friction (length/focus) for each step of the loop, 2) sharing/inviting enough people, and 3) building on a high-conversion marketing channel. For agentic viral loops, theoretically there’s no friction because your agent can just share things for you. It can pick who to share/invite, so if appropriate it should publish content to millions of people so there’s no limitation there. However, agentic viral loops are naturally limited by the fact that people don’t want their agents running around shilling themselves - the 4th time your sister gets the message “hey I’m X, andrew’s agent, and he wants you to try using an agent too!” you’ll get an annoyed text and you’ll stop using that agent. And of course, the marketing channel itself will become limited if everyone starts to filter/block agents, or even better, agents will filter/block other agents as a form of protection. Yet it’s clear that some types of agents will push the boundaries here, and some - perhaps free offerings targeted at less sophisticated segments of the market - will thrive in the same way that banner/popup-laden websites survive.

This might all sound like horrific spamminess to you, but I think it might be done tastefully. If an agent is high-retention and high-usage, you'll have many shots on goal to gradually invite your friends onto the same platform as you. The first shot might come from an invite loop that says your colleague/friend wants to set up a place to share photos/coordinate calendars/etc., and even if you say no, over time, your friend might share very useful microsites for projects or events or research that's relevant, and you might decide that you want to sign up to leave a comment, but then eventually you might try the other functionality as well.

Engagement network effects

It's easy to imagine that in a world where everyone by default has super intelligent agents that product engagement in agents would go up. The simple argument is that the ubiquity of agents would lead to more successful task completion, making agents more useful and applicable to more problems. Routine tasks like scheduling something or summarizing/sharing notes from a meeting, would happen instantly if it ran through agents and you removed humans completely. As agents become more successful over time, not only with your own tasks but in coordinating larger and more complex multi-user decisions, you'd end up using them more.

However, this does not answer the question of network effects. The question there is, does your agent get better when other people are all using the SAME agent? As I said earlier, the argument for NO is if each type of agent runs on a frontier AI model, can talk to other agents (all different types) in real-time via email/messaging/APIs/connectors/etc. Then there's not much leverage. If you're just organizing a meeting or a dinner party or meeting the agents just ask each other in real-time - and even if they are heterogeneous - they can use email and calendars and figure it out.

But some tasks do benefit from everyone running the same agent, because you get the benefits of centralization: Trust, speed, identity, discovery, etc. Some examples:

- "Who are 10 people in San Francisco that I should meet?" - agents can't just ask each other to find an optimal solution
- "Find me a VP of Engineering who isn't publicly looking for a new job, but would leave for a fast-growing startup that pays X in stock" - there's a mix of private and public info here, and an agent of a passive job-seeker might lie if the asking agent is not trustworthy
- "Who's a trustworthy seller of X that I should buy from?" - reputation might be gamed, and it might be better to have a global source of truth
- "Help me pick out the best dinner in the next 5 minutes that my friends and all their +1s would like, incorporating allergies/schedules/locations" - real-time solution incorporating a bunch of shared context
- ... and even more so, if you imagine that there may be agents that lie, exaggerate, or otherwise to benefit their humans

For the above examples, you could imagine Muse/Town/Instinct/whatever building a shared network+context layer that understands the

network, its relation to public/private data, acts as a system of record for identity/trust, does multi-sided marketplace matching, etc. Then it would allow only its agents access to this platform, and agents outside of the network wouldn't have access. If you're part of this network, and are contributing data, your agent gets more useful over time. The network gets more powerful over time.

The counterpoint to this counterpoint is that perhaps new startups will fill the void as a series of agent-facing point solutions, making it possible for people running agents of all times to interface with each other. Perhaps Muse creates a proprietary ecommerce marketplace that only Muse agents can use, but perhaps there will just be a next-gen eBay for any agent to sell to any other agent, and this service maintains central reputation/identity at that layer. In other words, perhaps the centralization will still happen, but just at a layer above agents. Most likely both will happen - every large tech company adds chat, commerce, search, etc, but also there are pure play versions of each of these. But this feels like more fragmentation than a true winner-take-all.

Some parting thoughts

My argument in this writeup is simple: Agents don't inherently have network effects. If all we build are super-intelligent assistants that can use tools and talk to each other, then the category might look surprisingly fragmented. Your agent can talk to my agent, just as Gmail can send email to Outlook, and there's no particular reason we need to use the same one. Better models, UX, integrations, memory, and distribution might create very large companies, but those are not network effects.

But the next gen of agents have a big opportunity design network effects in their interactions. On acquisition, they can turn the things they create - documents, events, websites, research, group chats, plans - into viral objects that naturally pull other people in. And on engagement, they can create shared networks of identity, reputation, relationships, private context, and intent that make the agent increasingly useful as more of the people around you join.

There's many ways for these network effects to instantiate. If we're lucky, things will look more like the open web - mostly decentralized, with bits of centralization to make certain tasks smoother. Identity could become an open protocol. Reputation could live in marketplaces. Commerce could happen through an agent-native eBay. Professional discovery could happen through an agent-native LinkedIn. Agents could remain interchangeable clients sitting on top of a constellation of networks and marketplaces, much as browsers and email clients do today. In that world, enormous network effects emerge around agents without producing a winner-take-all market for the agents themselves.

My guess is that we'll see a fight over exactly this boundary. Every major horizontal agent will try to pull valuable network functionality inside its walls: host the artifact instead of creating a Google Doc, execute the transaction instead of sending you to Amazon, make the introduction instead of searching LinkedIn, create the event instead of using Partiful. Every

time it does this successfully, it converts somebody else's network effect into its own. Meanwhile, every incumbent network and a new generation of startups will have the opposite incentive: make their networks universally accessible to every agent so that no horizontal agent can recreate and capture them.

This is why I don't think the interesting question is simply, "Will agents have network effects?" Of course some will. The much more consequential question is: Where will the network effects live? If they live primarily in open protocols and specialized networks, we may end up with thousands of agents competing on intelligence, personality, UX, specialization, and price, all interoperating with the same underlying ecosystem. Vertical agents can thrive in this world. But if a few horizontal agents successfully accumulate identity, relationships, private context, reputation, latent intent, and distribution - and keep those assets proprietary - then switching agents starts to mean leaving your network behind. Next couple years will be very fun to see how this plays out.

Social science and common sense

BY STEFAN SCHUBERT · THE UPDATE · 6 OCTOBER 2026

Priming studies suggested that coming across particular images or words (e.g. "wrinkle") could have surprising effects on behaviour (e.g. slower walking pace). Many laypeople found this hard to believe, but Daniel Kahneman told them off in no uncertain terms:

...disbelief is not an option. The results are not made up, nor are they statistical flukes. You have no choice but to accept that the major conclusions of these studies are true.

This has not aged well, as many of the priming studies subsequently failed to replicate. Kahneman's injunction might hold for physics (where rigorous evidence often establishes seemingly bizarre results), but social science is different: our research methods are weaker and our intuitive understanding stronger. Disbelief is often very much an option.

Common sense shapes social science

Similarly, common sense is devalued within social science. There's an implicit ideal that we should follow our research methods mechanically, without appealing to common sense. But our methods simply aren't good enough for that to work, and our common-sense knowledge is too valuable to ignore. We must apply common sense at every stage of the research process, from deciding what to study to making sense of the results. In practice, most social scientists do this, but they should be more open and explicit about it.

Social science shapes common sense

There is also the reverse mistake: overlooking the influence of social science on common sense. These days, most well-informed people know that the prices of goods are determined by supply and demand, not simply by how much labour it took to produce them. But it's easy to forget that this is thanks to economics. Common sense often absorbs research insights, and our failure to see this makes us underestimate social science's achievements.

The root of the problem is that we view common sense as more passive than it really is – reflecting our general neglect of human agency. In fact, common sense challenges, complements, and learns from social science.